
Chainer Documentation

Release 4.3.0

Preferred Networks, inc. and Preferred Infrastructure, inc.

Jul 19, 2018

1	Chainer at a Glance	3
1.1	Mushrooms – tasty or deathly?	3
1.2	Initialization	5
1.3	Trainer Structure	6
1.4	Dataset	7
1.5	Iterator	7
1.6	Model	8
1.7	Optimizer	9
1.8	Updater	9
1.9	Extensions	10
1.10	Inference	11
2	Installation	15
2.1	Recommended Environments	15
2.2	Requirements	15
2.3	Install Chainer	16
2.4	Uninstall Chainer	17
2.5	Upgrade Chainer	17
2.6	Reinstall Chainer	17
2.7	Run Chainer with Docker	17
2.8	FAQ	18
3	Guides	19
3.1	Define-by-Run	19
3.2	Variables and Derivatives	19
3.3	Links	21
3.4	Define your own function	23
3.5	Creating Models	32
3.6	Optimizer	33
3.7	Trainer	35
3.8	Trainer Extensions	35
3.9	Using GPU(s) in Chainer	39
3.10	Type Checks	45
3.11	Serializers – saving and loading	49
3.12	Customize your own logging	50
4	Neural Net Examples	53

4.1	MNIST using Trainer	53
4.2	MNIST with a Manual Training Loop	61
4.3	Convolutional Network for Visual Recognition Tasks	69
4.4	Recurrent Nets and their Computational Graph	76
4.5	RNN Language Models	81
4.6	Word2Vec: Obtain word embeddings	91
4.7	Write a Sequence to Sequence (seq2seq) Model	99
5	Reference	115
5.1	Variable and Parameter	115
5.2	Functions	133
5.3	Link and Chains	272
5.4	Optimizers	612
5.5	Weight Initializers	644
5.6	Training Tools	653
5.7	Datasets	689
5.8	Iterator	713
5.9	Serializers	717
5.10	Utilities	726
5.11	Configuring Chainer	738
5.12	Debug Mode	744
5.13	Visualization of Computational Graph	746
5.14	Caffe Model Support	749
5.15	Assertion and Testing	751
6	API Compatibility Policy	759
6.1	Targeted Versions	759
6.2	Versioning and Backward Compatibility	759
6.3	Breaking the Compatibility	759
6.4	Experimental APIs	760
6.5	Supported Backward Compatibility	760
6.6	Model Format Compatibility	762
6.7	Installation Compatibility	762
7	Contribution Guide	763
7.1	Classification of Contributions	763
7.2	Development Cycle	763
7.3	Issues and Pull Requests	765
7.4	Coding Guidelines	767
7.5	Unit Testing	768
7.6	Documentation	770
8	Tips and FAQs	773
8.1	It takes too long time to compile a computational graph. Can I skip it?	773
8.2	MNIST example does not converge in CPU mode on Mac OS X	773
8.3	How do I fix InvalidType error?	774
8.4	How do I accelerate my model using iDeep on Intel CPU?	775
8.5	My training process gets stuck when using MultiprocessIterator	775
9	Performance Best Practices	777
9.1	Use the Latest Version	777
9.2	Enable Hardware Accelerations	778
9.3	Migrate Data Preprocessing Code from NumPy to CuPy	778
9.4	Avoid Data Transfer	778
9.5	Optimize cuDNN Convolution	779

9.6	Fine-Tune Configuration	779
9.7	Load Datasets Concurrently	780
9.8	Use Multiple GPUs	780
9.9	Use Multiple Nodes	780
10	Upgrade Guide	781
10.1	Chainer v4	781
10.2	Chainer v3	783
10.3	Chainer v2	784
11	Comparison with Other Frameworks	799
11.1	A table for quick comparison	799
11.2	Benchmarks	801
12	License	803
13	Indices and tables	805
	Bibliography	807
	Python Module Index	809

Chainer is a powerful, flexible and intuitive deep learning framework.

- Chainer supports CUDA computation. It only requires a few lines of code to leverage a GPU. It also runs on multiple GPUs with little effort.
- Chainer supports various network architectures including feed-forward nets, convnets, recurrent nets and recursive nets. It also supports per-batch architectures.
- Forward computation can include any control flow statements of Python without lacking the ability of back-propagation. It makes code intuitive and easy to debug.

Chainer at a Glance

Welcome to Chainer!

Chainer is a rapidly growing neural network platform. The strengths of Chainer are:

- Python-based – Chainer is developed in Python, allowing for inspection and customization of all code in python and understandable python messages at run time
- Define by Run – neural networks definitions are defined on-the-fly at run time, allowing for dynamic network changes
- NumPy based syntax for working with arrays, thanks to CuPy implementation
- Fully customizable – since Chainer is pure python, all classes and methods can be adapted to allow for the latest cutting edge or specialized approaches
- Broad and deep support – Chainer is actively used for most of the current approaches for neural nets (CNN, RNN, RL, etc.), aggressively adds new approaches as they're developed, and provides support for many kinds of hardware as well as parallelization for multiple GPUs

1.1 Mushrooms – tasty or deathly?

Let's take a look at a basic program of Chainer to see how it works. For a dataset, we'll work with [Kaggle's edible vs. poisonous mushroom dataset](#), which has over 8,000 examples of mushrooms, labelled by 22 categories including odor, cap color, habitat, etc., in a [mushrooms.csv file](#).

How will Chainer learn which mushrooms are edible and which mushrooms will kill you? Let's see!

1.1.1 Full Code

Here's the whole picture of the code:

```
import matplotlib
matplotlib.use('Agg')

import chainer
import chainer.functions as F
import chainer.links as L
from chainer import datasets
from chainer import training
from chainer.training import extensions

import numpy as np

mushroomsfile = 'mushrooms.csv'
```

```
data_array = np.genfromtxt(
    mushroomsfile, delimiter=',', dtype=str, skip_header=1)
for col in range(data_array.shape[1]):
    data_array[:, col] = np.unique(data_array[:, col], return_inverse=True)[1]

X = data_array[:, 1:].astype(np.float32)
Y = data_array[:, 0].astype(np.int32)[:, None]
train, test = datasets.split_dataset_random(
    datasets.TupleDataset(X, Y), int(data_array.shape[0] * .7))

train_iter = chainer.iterators.SerialIterator(train, 100)
test_iter = chainer.iterators.SerialIterator(
    test, 100, repeat=False, shuffle=False)

# Network definition
class MLP(chainer.Chain):
    def __init__(self, n_units, n_out):
        super(MLP, self).__init__()
        with self.init_scope():
            # the input size to each layer inferred from the layer before
            self.l1 = L.Linear(n_units) # n_in -> n_units
            self.l2 = L.Linear(n_units) # n_units -> n_units
            self.l3 = L.Linear(n_out)  # n_units -> n_out

    def __call__(self, x):
        h1 = F.relu(self.l1(x))
        h2 = F.relu(self.l2(h1))
        return self.l3(h2)

model = L.Classifier(
    MLP(44, 1), lossfun=F.sigmoid_cross_entropy, accfun=F.binary_accuracy)

# Setup an optimizer
optimizer = chainer.optimizers.SGD()
optimizer.setup(model)

# Create the updater, using the optimizer
updater = training.StandardUpdater(train_iter, optimizer, device=-1)

# Set up a trainer
trainer = training.Trainer(updater, (50, 'epoch'), out='result')
```

```

# Evaluate the model with the test dataset for each epoch
trainer.extend(extensions.Evaluator(test_iter, model, device=-1))

# Dump a computational graph from 'loss' variable at the first iteration
# The "main" refers to the target link of the "main" optimizer.
trainer.extend(extensions.dump_graph('main/loss'))

trainer.extend(extensions.snapshot(), trigger=(20, 'epoch'))

# Write a log of evaluation statistics for each epoch
trainer.extend(extensions.LogReport())

# Save two plot images to the result dir
if extensions.PlotReport.available():
    trainer.extend(
        extensions.PlotReport(['main/loss', 'validation/main/loss'],
                               'epoch', file_name='loss.png'))
    trainer.extend(
        extensions.PlotReport(
            ['main/accuracy', 'validation/main/accuracy'],
            'epoch', file_name='accuracy.png'))

# Print selected entries of the log to stdout
trainer.extend(extensions.PrintReport(
    ['epoch', 'main/loss', 'validation/main/loss',
     'main/accuracy', 'validation/main/accuracy', 'elapsed_time']))

# Run the training
trainer.run()

x, t = test[np.random.randint(len(test))]

predict = model.predictor(x[None]).data
predict = predict[0][0]

if predict >= 0:
    print('Predicted Poisonous, Actual ' + ['Edible', 'Poisonous'][t[0]])
else:
    print('Predicted Edible, Actual ' + ['Edible', 'Poisonous'][t[0]])

```

```
Predicted ...
```

If you've worked with other neural net frameworks, some of that code may look familiar. Let's break down what it's doing.

1.1.2 Code Breakdown

1.2 Initialization

Let's start our python program. Matplotlib is used for the graphs to show training progress.

```

#!/usr/bin/env python

import matplotlib
matplotlib.use('Agg')

```

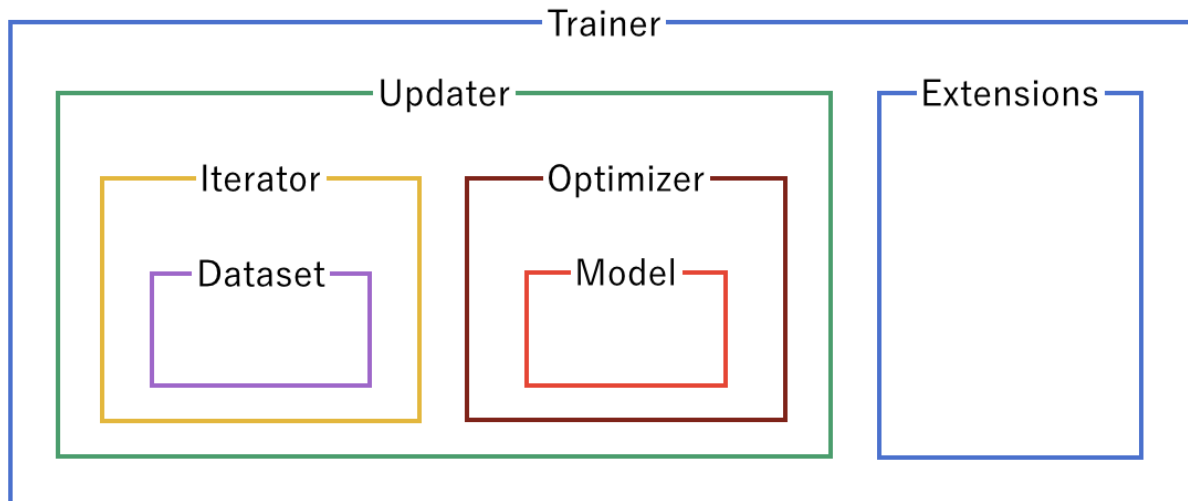
Typical imports for a Chainer program. *links* contain trainable parameters and *functions* do not.

```
import chainer
import chainer.functions as F
import chainer.links as L
from chainer import training
from chainer import datasets
from chainer.training import extensions

import numpy as np
```

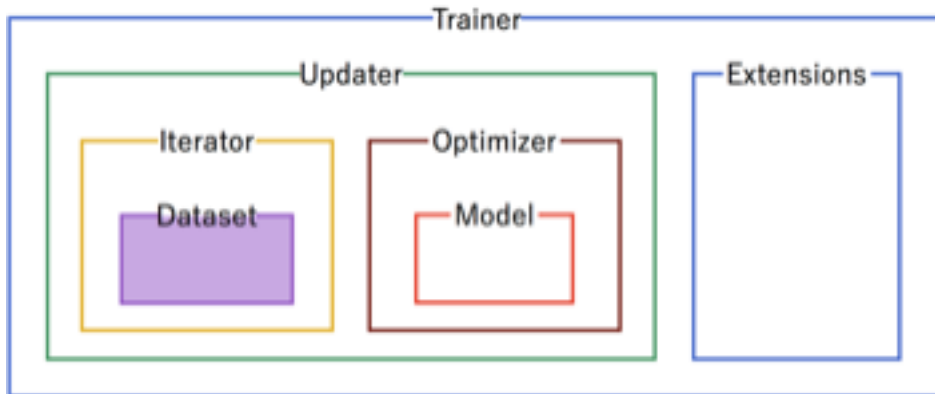
1.3 Trainer Structure

A `chainer.trainer` is used to set up our neural network and data for training. The components of the `chainer.trainer` are generally hierarchical, and are organized as follows:



Each of the components is fed information from the components within it. Setting up the trainer starts at the inner components, and moves outward, with the exception of `extensions`, which are added after the `chainer.trainer` is defined.

1.4 Dataset



Our first step is to format the *datasets*. From the raw `mushrooms.csv`, we format the data into a Chainer *TupleDataset*.

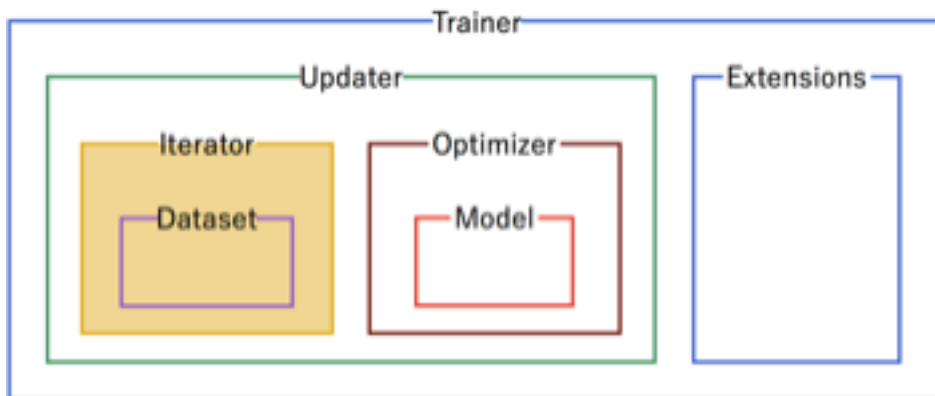
```

data_array = np.genfromtxt(
    'mushrooms.csv', delimiter=',', dtype=str, skip_header=1)
for col in range(data_array.shape[1]):
    data_array[:, col] = np.unique(data_array[:, col], return_inverse=True)[1]

X = data_array[:, 1:].astype(np.float32)
Y = data_array[:, 0].astype(np.int32)[:, None]
train, test = datasets.split_dataset_random(
    datasets.TupleDataset(X, Y), int(data_array.shape[0] * .7))

```

1.5 Iterator



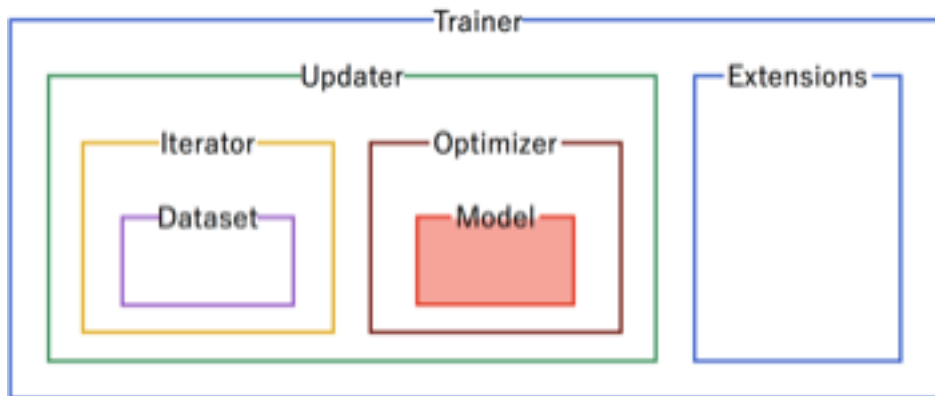
Configure *iterators* to step through batches of the data for training and for testing validation. In this case, we'll use a batch size of 100, no repeating, and shuffling not required since we already shuffled the dataset on reading it in.

```

train_iter = chainer.iterators.SerialIterator(train, 100)
test_iter = chainer.iterators.SerialIterator(
    test, 100, repeat=False, shuffle=False)

```

1.6 Model



Next, we need to define the neural network for inclusion in our model. For our mushrooms, we'll chain together two fully-connected, *Linear*, hidden layers between the input and output layers.

As an activation function, we'll use standard Rectified Linear Units (*relu()*).

```
# Network definition
class MLP(chainer.Chain):

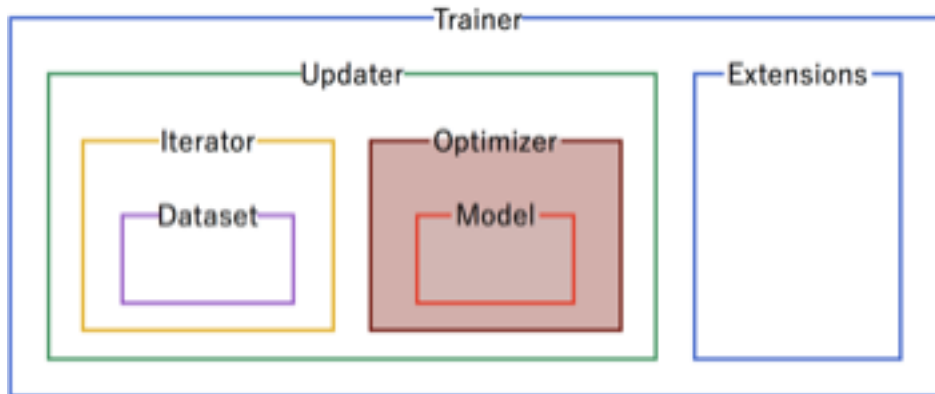
    def __init__(self, n_units, n_out):
        super(MLP, self).__init__()
        with self.init_scope():
            # the input size to each layer inferred from the layer before
            self.l1 = L.Linear(n_units) # n_in -> n_units
            self.l2 = L.Linear(n_units) # n_units -> n_units
            self.l3 = L.Linear(n_out)  # n_units -> n_out

    def __call__(self, x):
        h1 = F.relu(self.l1(x))
        h2 = F.relu(self.l2(h1))
        return self.l3(h2)
```

Since mushrooms are either edible or poisonous (no information on psychedelic effects!) in the dataset, we'll use a Link *Classifier* for the output, with 44 units (double the features of the data) in the hidden layers and a single true/false category for classification.

```
model = L.Classifier(
    MLP(44, 1), lossfun=F.sigmoid_cross_entropy, accfun=F.binary_accuracy)
```

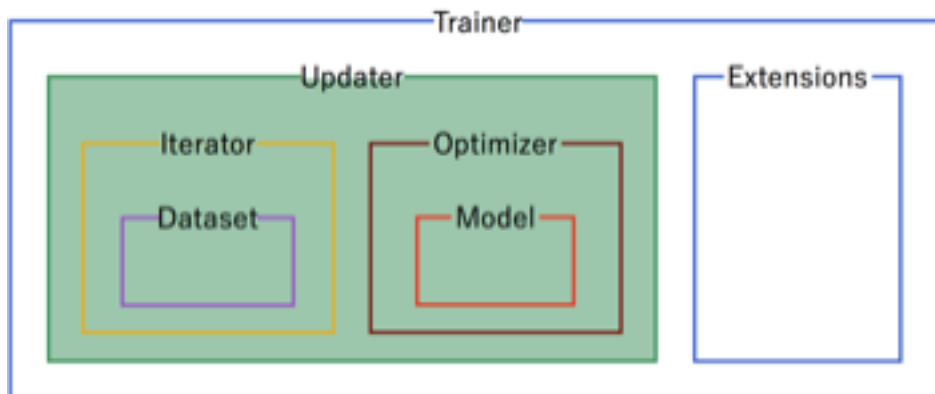
1.7 Optimizer



Pick an optimizer, and set up the model to use it.

```
# Setup an optimizer
optimizer = chainer.optimizers.SGD()
optimizer.setup(model)
```

1.8 Updater



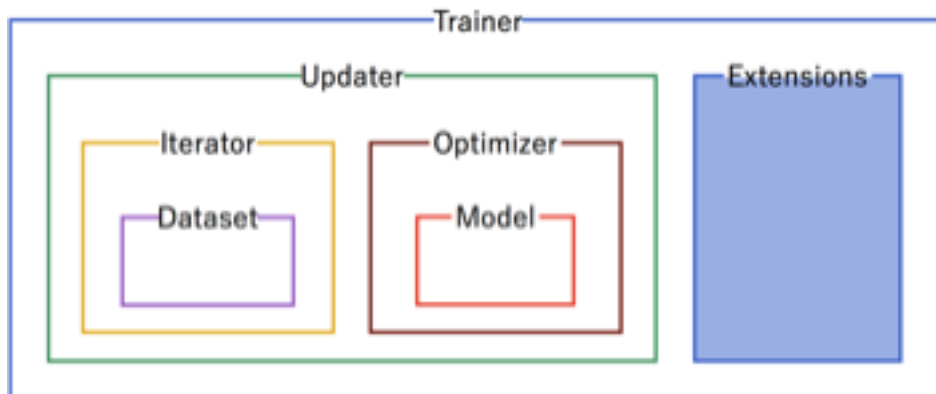
Now that we have the training iterator and optimizer set up, we link them both together into the updater. The updater uses the minibatches from the iterator, and then does the forward and backward processing of the model, and updates the parameters of the model according to the optimizer. Setting the `device=-1` sets the device as the CPU. To use a GPU, set device equal to the number of the GPU, usually `device=0`.

```
# Create the updater, using the optimizer
updater = training.StandardUpdater(train_iter, optimizer, device=-1)
```

Set up the updater to be called after the training batches and set the number of batches per epoch to 100. The learning rate per epoch will be output to the directory `result`.

```
# Set up a trainer
trainer = training.Trainer(updater, (50, 'epoch'), out='result')
```

1.9 Extensions



Use the testing iterator defined above for an *Evaluator* extension to the trainer to provide test scores.

If using a GPU instead of the CPU, set `device` to the ID of the GPU, usually 0.

```
trainer.extend(extensions.Evaluator(test_iter, model, device=-1))
```

Save a computational graph from loss variable at the first iteration. `main` refers to the target link of the main optimizer. The graph is saved in the *Graphviz*'s dot format. The output location (directory) to save the graph is set by the `out` argument of *Trainer*.

```
trainer.extend(extensions.dump_graph('main/loss'))
```

Take a snapshot of the `trainer` object every 20 epochs.

```
trainer.extend(extensions.snapshot(), trigger=(20, 'epoch'))
```

Write a log of evaluation statistics for each epoch.

```
trainer.extend(extensions.LogReport())
```

Save two plot images to the result directory.

```
if extensions.PlotReport.available():
    trainer.extend(
        extensions.PlotReport(['main/loss', 'validation/main/loss'],
                              'epoch', file_name='loss.png'))
    trainer.extend(
        extensions.PlotReport(
            ['main/accuracy', 'validation/main/accuracy'],
            'epoch', file_name='accuracy.png'))
```

Print selected entries of the log to standard output.

```
trainer.extend(extensions.PrintReport(
    ['epoch', 'main/loss', 'validation/main/loss',
     'main/accuracy', 'validation/main/accuracy', 'elapsed_time']))
```

Run the training.

```
trainer.run()
```


1.10 Inference

Once the training is complete, only the model is necessary to make predictions. Let's check that a random line from the test data set and see if the inference is correct:

```
x, t = test[np.random.randint(len(test))]  
  
predict = model.predictor(x[None]).data  
predict = predict[0][0]  
  
if predict >= 0:  
    print('Predicted Poisonous, Actual ' + ['Edible', 'Poisonous'][t[0]])  
else:  
    print('Predicted Edible, Actual ' + ['Edible', 'Poisonous'][t[0]])
```

1.10.1 Output

Output for this instance will look like:

epoch	main/loss	validation/main/loss	main/accuracy	validation/main/accuracy
↪ elapsed_time				
1	0.550724	0.502818	0.733509	0.752821
↪ 0.215426				
2	0.454206	0.446234	0.805439	0.786926
↪ 0.902108				
3	0.402783	0.395893	0.838421	0.835979
↪ 1.50414				
4	0.362979	0.359988	0.862807	0.852632
↪ 2.24171				
5	0.32713	0.329881	0.88	0.874232
↪ 2.83247				
6	0.303469	0.31104	0.892456	0.887284
↪ 3.45173				
7	0.284755	0.288553	0.901754	0.903284
↪ 3.9877				
8	0.26801	0.272033	0.9125	0.907137
↪ 4.54794				
9	0.25669	0.261355	0.920175	0.917937
↪ 5.21672				
10	0.241789	0.251821	0.927193	0.917937
↪ 5.79541				
11	0.232291	0.238022	0.93	0.925389
↪ 6.3055				
12	0.222805	0.22895	0.934035	0.923389
↪ 6.87083				
13	0.21276	0.219291	0.93614	0.928189
↪ 7.54113				
14	0.204822	0.220736	0.938596	0.922589
↪ 8.12495				
15	0.197671	0.207017	0.938393	0.936042
↪ 8.69219				
16	0.190285	0.199129	0.941053	0.934842
↪ 9.24302				
17	0.182827	0.193303	0.944386	0.942695
↪ 9.80991				

(continues on next page)

(continued from previous page)

18	0.176776	0.194284	0.94614	0.934042	└
→ 10.3603					
19	0.16964	0.177684	0.945789	0.945242	└
→ 10.8531					
20	0.164831	0.171988	0.949825	0.947347	└
→ 11.3876					
21	0.158394	0.167459	0.952982	0.949747	└
→ 11.9866					
22	0.153353	0.161774	0.956964	0.949347	└
→ 12.6433					
23	0.148209	0.156644	0.957368	0.951747	└
→ 13.3825					
24	0.144814	0.15322	0.957018	0.955495	└
→ 13.962					
25	0.138782	0.148277	0.958947	0.954147	└
→ 14.6					
26	0.135333	0.145225	0.961228	0.956695	└
→ 15.2284					
27	0.129593	0.141141	0.964561	0.958295	└
→ 15.7413					
28	0.128265	0.136866	0.962632	0.960547	└
→ 16.2711					
29	0.123848	0.133444	0.966071	0.961347	└
→ 16.7772					
30	0.119687	0.129579	0.967193	0.964547	└
→ 17.3311					
31	0.115857	0.126606	0.968596	0.966547	└
→ 17.8252					
32	0.113911	0.124272	0.968772	0.962547	└
→ 18.3121					
33	0.111502	0.122548	0.968596	0.965095	└
→ 18.8973					
34	0.107427	0.116724	0.970526	0.969747	└
→ 19.4723					
35	0.104536	0.114517	0.970877	0.969095	└
→ 20.0804					
36	0.099408	0.112128	0.971786	0.970547	└
→ 20.6509					
37	0.0972982	0.107618	0.973158	0.970947	└
→ 21.2467					
38	0.0927064	0.104918	0.973158	0.969347	└
→ 21.7978					
39	0.0904702	0.101141	0.973333	0.969747	└
→ 22.3328					
40	0.0860733	0.0984015	0.975263	0.971747	└
→ 22.8447					
41	0.0829282	0.0942095	0.977544	0.974947	└
→ 23.5113					
42	0.082219	0.0947418	0.975965	0.969347	└
→ 24.0427					
43	0.0773362	0.0906804	0.977857	0.977747	└
→ 24.5252					
44	0.0751769	0.0886449	0.977895	0.972147	└
→ 25.1722					
45	0.072056	0.0916797	0.978246	0.977495	└
→ 26.0778					
46	0.0708111	0.0811359	0.98	0.979347	└
→ 26.6648					

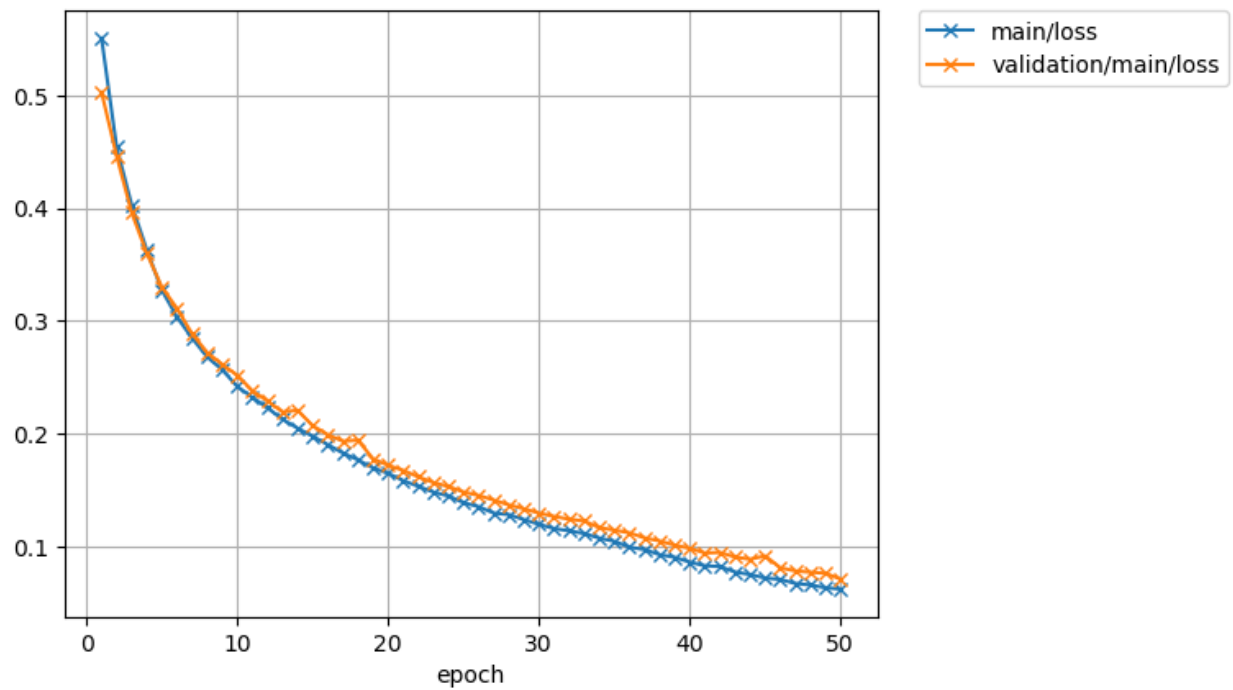
(continues on next page)

(continued from previous page)

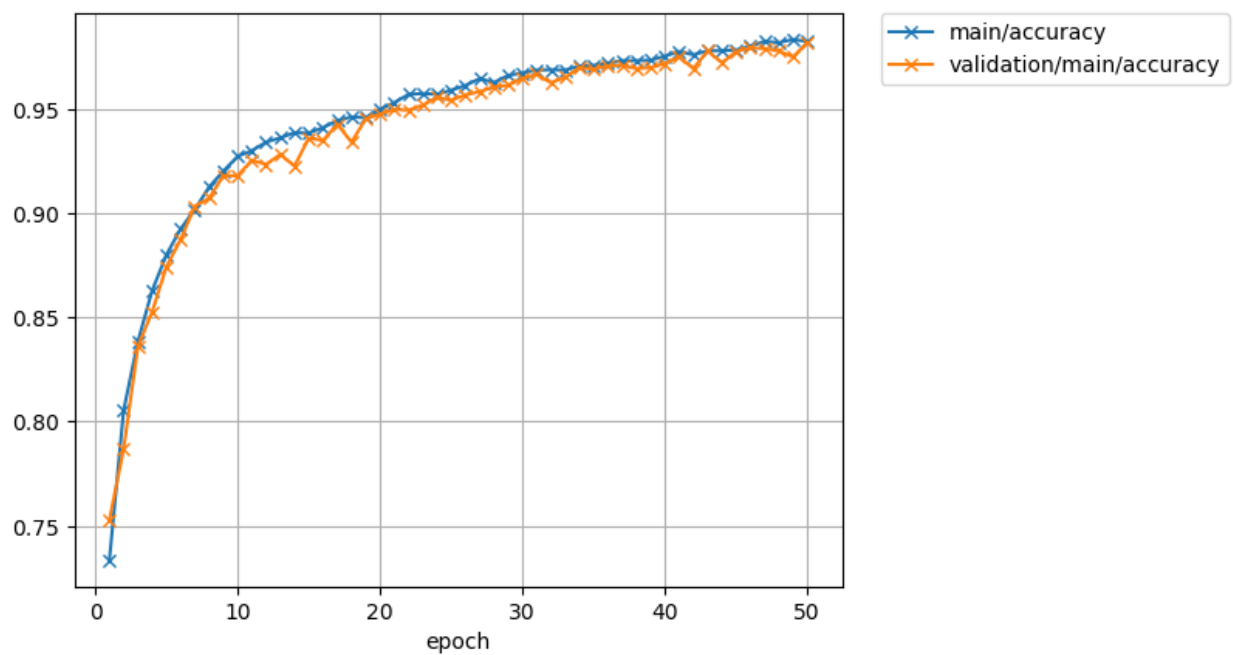
47	0.0671919	0.0783265	0.982456	0.978947	└
→ 27.2929					
48	0.0658817	0.0772342	0.981754	0.977747	└
→ 27.8119					
49	0.0634615	0.0762576	0.983333	0.974947	└
→ 28.3876					
50	0.0622394	0.0710278	0.982321	0.981747	└
→ 28.9067					
Predicted Edible Actual Edible					

Our prediction was correct. Success!

The loss function:



And the accuracy



2.1 Recommended Environments

We recommend the following Linux distributions.

- **Ubuntu** 14.04 / 16.04 LTS (64-bit)
- **CentOS** 7 (64-bit)

Note: We are automatically testing Chainer on all the recommended environments above. We cannot guarantee that Chainer works on other environments including Windows and macOS (especially with CUDA support), even if Chainer may seem to be running correctly.

2.2 Requirements

You need to have the following components to use Chainer.

- **Python**
 - Supported Versions: 2.7.6+, 3.4.3+, 3.5.1+ and 3.6.0+.
- **NumPy**
 - Supported Versions: 1.9, 1.10, 1.11, 1.12, 1.13 and 1.14.
 - NumPy will be installed automatically during the installation of Chainer.

Before installing Chainer, we recommend you to upgrade `setuptools` and `pip`:

```
$ pip install -U setuptools pip
```

2.2.1 Hardware Acceleration Support

You can accelerate performance of Chainer by installing the following optional components.

- **NVIDIA CUDA / cuDNN**
 - CuPy 4.0+
 - See [CuPy Installation Guide](#) for instructions.
- **Intel CPU (experimental)**
 - iDeep 1.0.3+
 - See [Tips and FAQs](#) for instructions.

2.2.2 Optional Features

The following packages are optional dependencies. Chainer can be installed without them, in which case the corresponding features are not available.

- **Image dataset support**
 - pillow 2.3+
 - Run `pip install pillow` to install.
- **HDF5 serialization support**
 - h5py 2.5+
 - Run `pip install h5py` to install.

2.3 Install Chainer

2.3.1 Using pip

We recommend to install Chainer via pip:

```
$ pip install chainer
```

Note: Any optional dependencies (including CuPy) can be added after installing Chainer. Chainer automatically detects the available packages and enables/disables the optional features appropriately.

2.3.2 Using Tarball

The tarball of the source tree is available via `pip download chainer` or from [the release notes page](#). You can install Chainer from the tarball:

```
$ pip install chainer-x.x.x.tar.gz
```

You can also install the development version of Chainer from a cloned Git repository:

```
$ git clone https://github.com/chainer/chainer.git
$ cd chainer
$ pip install .
```

2.3.3 Enable CUDA/cuDNN support

In order to enable CUDA support, you have to install [CuPy](#) manually. If you also want to use cuDNN, you have to install CuPy with cuDNN support. See [CuPy's installation guide](#) to install CuPy. Once CuPy is correctly set up, Chainer will automatically enable CUDA support.

You can refer to the following flags to confirm if CUDA/cuDNN support is actually available.

chainer.backends.cuda.available True if Chainer successfully imports [cupy](#).

chainer.backends.cuda.cudnn_enabled True if cuDNN support is available.

2.4 Uninstall Chainer

Use pip to uninstall Chainer:

```
$ pip uninstall chainer
```

Note: When you upgrade Chainer, pip sometimes install the new version without removing the old one in site-packages. In this case, pip uninstall only removes the latest one. To ensure that Chainer is completely removed, run the above command repeatedly until pip returns an error.

2.5 Upgrade Chainer

Just use pip with -U option:

```
$ pip install -U chainer
```

2.6 Reinstall Chainer

If you want to reinstall Chainer, please uninstall Chainer and then install it. We recommend to use `--no-cache-dir` option as pip sometimes uses cache:

```
$ pip uninstall chainer
$ pip install chainer --no-cache-dir
```

2.7 Run Chainer with Docker

We are providing the official Docker image. Use [nvidia-docker](#) command to run Chainer image with GPU. You can login to the environment with bash, and run the Python interpreter:

```
$ nvidia-docker run -it chainer/chainer /bin/bash
```

Or run the interpreter directly:

```
$ nvidia-docker run -it chainer/chainer /usr/bin/python
```

2.8 FAQ

2.8.1 Warning message “cuDNN is not enabled” appears

You failed to build CuPy with cuDNN. If you don’t need cuDNN, ignore this message. Otherwise, retry to install CuPy with cuDNN. `pip install -vvvv` option helps you. There is no need of re-installing Chainer itself. See [CuPy’s installation guide](#) for more details.

2.8.2 CuPy always raises `cupy.cuda.compiler.CompileException`

See FAQ section of [CuPy’s installation guide](#) for details.

2.8.3 h5py installation failed

If the installation failed with error saying `hdf5.h` is not found, you need to install `libhdf5` first. The way to install it depends on your environment:

```
# Ubuntu 14.04/16.04
$ apt-get install libhdf5-dev

# CentOS 7
$ yum -y install epel-release
$ yum install hdf5-devel
```

Note that `h5py` is not required unless you need HDF5 serialization support.

3.1 Define-by-Run

As mentioned on the top page, Chainer is a flexible framework for neural networks. One major goal is flexibility, so it must enable us to write complex architectures simply and intuitively.

Most existing deep learning frameworks are based on the “**Define-and-Run**” scheme. That is, first a network is defined and fixed, and then the user periodically feeds it with mini-batches of training data. Since the network is statically defined before any forward/backward computation, all the logic must be embedded into the network architecture as *data*. Consequently, defining a network architecture in such systems (e.g. Caffe) follows a declarative approach. Note that one can still produce such a static network definition using imperative languages (e.g. torch.nn, Theano-based frameworks, and TensorFlow).

In contrast, Chainer adopts a “**Define-by-Run**” scheme, i.e., the network is defined dynamically via the actual forward computation. More precisely, Chainer stores the history of computation instead of programming logic. This strategy enables us to fully leverage the power of programming logic in Python. For example, Chainer does not need any magic to introduce conditionals and loops into the network definitions. The Define-by-Run scheme is the core concept of Chainer. We will show in this tutorial how to define networks dynamically.

This strategy also makes it easy to write multi-GPU parallelization, since logic comes closer to network manipulation. We will review such amenities in later sections of this tutorial.

3.2 Variables and Derivatives

In the example code of this tutorial, we assume for simplicity that the following symbols are already imported.

```
import numpy as np
import chainer
from chainer.backends import cuda
from chainer import Function, gradient_check, report, training, utils, Variable
from chainer import datasets, iterators, optimizers, serializers
```

(continues on next page)

(continued from previous page)

```
from chainer import Link, Chain, ChainList
import chainer.functions as F
import chainer.links as L
from chainer.training import extensions
```

As described previously, Chainer uses the “Define-by-Run” scheme, so forward computation itself *defines* the network. In order to start forward computation, we have to set the input array to a `chainer.Variable` object. Here we start with a simple `ndarray` with only one element:

```
>>> x_data = np.array([5], dtype=np.float32)
>>> x = Variable(x_data)
```

A `Variable` object supports basic arithmetic operators. In order to compute $y = x^2 - 2x + 1$, just write:

```
>>> y = x**2 - 2 * x + 1
```

The resulting `y` is also a `Variable` object, whose value can be extracted by accessing the `array` attribute:

```
>>> y.array
array([16.], dtype=float32)
```

Note: `Variable` has two attributes to represent the underlying array: `array` and `data`. There is no difference between the two; both refer to exactly the same object. However it is not recommended to use `.data` because it might be confused with `numpy.ndarray.data` attribute.

What `y` holds is not only the result value. It also holds the history of computation (or computational graph), which enables us to compute its derivative. This is done by calling its `backward()` method:

```
>>> y.backward()
```

This runs *error backpropagation* (a.k.a. *backprop* or *reverse-mode automatic differentiation*). Then, the gradient is computed and stored in the `grad` attribute of the input variable `x`:

```
>>> x.grad
array([8.], dtype=float32)
```

Also we can compute gradients of intermediate variables. Note that Chainer, by default, releases the gradient arrays of intermediate variables for memory efficiency. In order to preserve gradient information, pass the `retain_grad` argument to the `backward` method:

```
>>> z = 2*x
>>> y = x**2 - z + 1
>>> y.backward(retain_grad=True)
>>> z.grad
array([-1.], dtype=float32)
```

All these computations can be generalized to a multi-element array input. While single-element arrays are automatically initialized to `[1]`, to start backward computation from a variable holding a multi-element array, we must set the *initial error* manually. This is done simply by setting the `grad` attribute of the output variable:

```
>>> x = Variable(np.array([[1, 2, 3], [4, 5, 6]], dtype=np.float32))
>>> y = x**2 - 2*x + 1
>>> y.grad = np.ones((2, 3), dtype=np.float32)
>>> y.backward()
```

(continues on next page)

(continued from previous page)

```
>>> x.grad
array([[ 0.,  2.,  4.],
       [ 6.,  8., 10.]], dtype=float32)
```

Note: Many functions taking *Variable* object(s) are defined in the *chainer.functions* module. You can combine them to realize complicated functions with automatic backward computation.

Note: Instead of using *backward()*, you can also calculate gradients of any variables in a computational graph w.r.t. any other variables in the graph using the *chainer.grad()* function.

3.2.1 Higher-Order Derivatives

Variable also supports higher-order derivatives (a.k.a. double backpropagation).

Let's see a simple example. First calculate the first-order derivative. Note that `enable_double_backprop=True` is passed to `y.backward()`.

```
>>> x = chainer.Variable(np.array([[0, 2, 3], [4, 5, 6]], dtype=np.float32))
>>> y = x ** 3
>>> y.grad = np.ones((2, 3), dtype=np.float32)
>>> y.backward(enable_double_backprop=True)
>>> x.grad_var
variable([[ 0., 12., 27.],
          [48., 75., 108.]])
>>> assert x.grad_var.array is x.grad
>>> assert (x.grad == (3 * x**2).array).all()
```

chainer.Variable.grad_var is a *Variable* for *chainer.Variable.grad* (which is an *ndarray*). By passing `enable_double_backprop=True` to `backward()`, a computational graph for the backward calculation is recorded. So, you can start backpropagation from `x.grad_var` to calculate the second-order derivative.

```
>>> gx = x.grad_var
>>> x.cleargrad()
>>> gx.grad = np.ones((2, 3), dtype=np.float32)
>>> gx.backward()
>>> x.grad
array([[ 0., 12., 18.],
       [24., 30., 36.]], dtype=float32)
>>> assert (x.grad == (6 * x).array).all()
```

3.3 Links

In order to write neural networks, we have to combine functions with *parameters* and optimize the parameters. You can use the class *Link* to do this. A *Link* is an object that holds parameters (i.e. optimization targets).

The most fundamental ones are links that behave like regular functions while replacing some arguments by their parameters. We will introduce higher level links, but here think of links as simply functions with parameters.

One of the most frequently used links is the *Linear* link (a.k.a. *fully-connected layer* or *affine transformation*). It represents a mathematical function $f(x) = Wx + b$, where the matrix W and the vector b are parameters. This

link corresponds to its pure counterpart `linear()`, which accepts x, W, b as arguments. A linear link from three-dimensional space to two-dimensional space is defined by the following line:

```
>>> f = L.Linear(3, 2)
```

Note: Most functions and links only accept mini-batch input, where the first dimension of the input array is considered as the *batch dimension*. In the above Linear link case, input must have shape of $(N, 3)$, where N is the mini-batch size.

The parameters of a link are stored as attributes. Each parameter is an instance of *Variable*. In the case of the Linear link, two parameters, W and b , are stored. By default, the matrix W is initialized randomly, while the vector b is initialized with zeros. This is the preferred way to initialize these parameters.

```
>>> f.W.data
array([[ 1.0184761,  0.23103087,  0.5650746 ],
       [ 1.2937803,  1.0782351, -0.56423163]], dtype=float32)
>>> f.b.data
array([0., 0.], dtype=float32)
```

An instance of the Linear link acts like a usual function:

```
>>> x = Variable(np.array([[1, 2, 3], [4, 5, 6]], dtype=np.float32))
>>> y = f(x)
>>> y.data
array([[3.1757617, 1.7575557],
       [8.619507, 7.1809077]], dtype=float32)
```

Note: Sometimes it is cumbersome to compute the dimension of the input space. The linear link and some of (de)convolution links can omit the input dimension in their instantiation and infer it from the first mini-batch.

For example, the following line creates a linear link whose output dimension is two:

```
>>> f = L.Linear(2)
```

If we feed a mini-batch of shape $(2, M)$, the input dimension will be inferred as M , which means `l.W` will be a $2 \times M$ matrix. Note that its parameters are initialized in a lazy manner at the first mini-batch. Therefore, `l` does not have `W` attribute if no data is put to the link.

Gradients of parameters are computed by the `backward()` method. Note that gradients are **accumulated** by the method rather than overwritten. So first you must clear the gradients to renew the computation. It can be done by calling the `cleargrads()` method.

```
>>> f.cleargrads()
```

Note: `cleargrads()` is introduced in v1.15 to replace `zerograds()` for efficiency. `zerograds()` is left only for backward compatibility.

Now we can compute the gradients of parameters by simply calling the backward method and access them via the `grad` property.

```
>>> y.grad = np.ones((2, 2), dtype=np.float32)
>>> y.backward()
```

(continues on next page)

(continued from previous page)

```
>>> f.W.grad
array([[5., 7., 9.],
       [5., 7., 9.]], dtype=float32)
>>> f.b.grad
array([2., 2.], dtype=float32)
```

3.4 Define your own function

In this section, you will learn about the following things:

- How to define a function on variables
- Useful tools to write a function using a GPU
- How to test the function definition

After reading this section, you will be able to:

- Write your own functions
- Define simple kernels in the function definition

In the example code of this tutorial, we assume for simplicity that the following symbols are already imported.

```
import numpy as np
import chainer
from chainer.backends import cuda
from chainer import Function, gradient_check, report, training, utils, Variable
from chainer import datasets, iterators, optimizers, serializers
from chainer import Link, Chain, ChainList
import chainer.functions as F
import chainer.links as L
from chainer.training import extensions
```

3.4.1 Differentiable Functions

Chainer provides a collection of functions in the `chainer.functions` module. It covers typical use cases in deep learning, so many existing works can be implemented with them. On the other hand, deep learning is evolving rapidly and we cannot cover all possible functions to define unseen architectures. So it is important to learn how to define your own functions.

First, suppose we want to define an elementwise function $f(x, y, z) = x * y + z$. While it is possible to implement this equation using a combination of the `*` and `+` functions, defining it as a single function may reduce memory consumption, so it is not *only* a toy example. Here we call this function *MulAdd*.

Let's start with defining *MulAdd* working on the CPU. Any function must inherit the `Function` class. The skeleton of a function looks like:

```
class MulAdd(Function):
    def forward_cpu(self, inputs):
        # do forward computation on CPU
        return some_tuple

    def backward_cpu(self, inputs, grad_outputs):
```

(continues on next page)

(continued from previous page)

```
# do backward computation on CPU
return some_tuple
```

We must implement `forward_cpu()` and `backward_cpu()` methods. The non-self arguments of these functions are tuples of array(s), and these functions must return a tuple of array(s).

Warning: Be careful to return a tuple of arrays even if you have just one array to return.

MulAdd is simple and implemented as follows

```
class MulAdd(Function):
    def forward_cpu(self, inputs):
        x, y, z = inputs
        w = x * y + z
        return w,

    def backward_cpu(self, inputs, grad_outputs):
        x, y, z = inputs
        gw, = grad_outputs

        gx = y * gw
        gy = x * gw
        gz = gw
        return gx, gy, gz
```

As per the warning above, the `forward_cpu` method returns a tuple of single element. Note that all arrays appearing in CPU functions are `numpy.ndarray`. The forward function is straightforward: It unpacks the input tuple, computes the output, and packs it into a tuple. The backward function is a bit more complicated. Recall the rule of differentiation of multiplication. This example just implements the rule. Look at the return values, the function just packs the gradient of each input in same order and returns them.

By just defining the core computation of forward and backward, Function class provides a chaining logic on it (i.e. storing the history of computation, etc.).

Note: Assuming we implement a (forward) function $y = f(x)$ which takes as input the vector $x \in \mathbb{R}^n$ and produces as output a vector $y \in \mathbb{R}^m$. Then the backward method has to compute

$$\lambda_i = \sum_{j=1}^m \frac{\partial y_j}{\partial x_i} \gamma_j \text{ for } i = 1 \dots n$$

where γ is the `grad_outputs`. Note, that the resulting vector λ must have the same shape as the arguments of the forward method.

Now let's define the corresponding GPU methods. You can easily predict that the methods we have to write are named `forward_gpu()` and `backward_gpu()`:

```
class MulAdd(Function):
    def forward_cpu(self, inputs):
        ...

    def backward_cpu(self, inputs, grad_outputs):
        ...
```

(continues on next page)

(continued from previous page)

```

def forward_gpu(self, inputs):
    x, y, z = inputs
    w = x * y + z
    return w,

def backward_gpu(self, inputs, grad_outputs):
    x, y, z = inputs
    gw, = grad_outputs

    gx = y * gw
    gy = x * gw
    gz = gw
    return gx, gy, gz

```

In GPU methods, arrays are of type `cupy.ndarray`. We use arithmetic operators defined for this class. These operators implement the basic elementwise arithmetics.

You may find that the definitions of GPU methods are exactly same as those of CPU methods. In that case, we can reduce them to `forward()` and `backward()` methods

```

class MulAdd(Function):
    def forward(self, inputs):
        x, y, z = inputs
        w = x * y + z
        return w,

    def backward(self, inputs, grad_outputs):
        x, y, z = inputs
        gw, = grad_outputs

        gx = y * gw
        gy = x * gw
        gz = gw
        return gx, gy, gz

```

Since the `cupy.ndarray` class implements many methods of `numpy.ndarray`, we can write these unified methods in most cases.

The `MulAdd` function is used as follows:

```

x = Variable(np.random.uniform(-1, 1, (3, 2)).astype(np.float32))
y = Variable(np.random.uniform(-1, 1, (3, 2)).astype(np.float32))
z = Variable(np.random.uniform(-1, 1, (3, 2)).astype(np.float32))
w = MulAdd()(x, y, z)

```

It looks a bit ugly: we have to explicitly instantiate `MulAdd` before applying it to variables. We also have to be careful that one instance of `MulAdd` must not be used multiple times, since it acts as a node in the computational graph. In Chainer, we often define a thin wrapper Python function that hide the instantiation:

```

def muladd(x, y, z):
    return MulAdd()(x, y, z)

w = muladd(x, y, z)

```

3.4.2 Unified forward/backward methods with NumPy/CuPy functions

CuPy also implements many functions that are compatible to those of NumPy. We can write unified forward/backward methods with them. Consider that we want to write a backprop-able function $f(x, y) = \exp(x) + \exp(y)$. We name it *ExpAdd* here. It can be written straight-forward as follows

```
from chainer.backends import cuda

class ExpAdd(Function):
    def forward_cpu(self, inputs):
        x, y = inputs
        z = np.exp(x) + np.exp(y)
        return z,

    def backward_cpu(self, inputs, grad_outputs):
        x, y = inputs
        gz, = grad_outputs

        gx = gz * np.exp(x)
        gy = gz * np.exp(y)
        return gx, gy

    def forward_gpu(self, inputs):
        cupy = cuda.cupy
        x, y = inputs
        z = cupy.exp(x) + cupy.exp(y)
        return z,

    def backward_gpu(self, inputs, grad_outputs):
        cupy = cuda.cupy
        x, y = inputs
        gz, = grad_outputs

        gx = gz * cupy.exp(x)
        gy = gz * cupy.exp(y)
        return gx, gy

def expadd(x, y):
    return ExpAdd()(x, y)
```

Note: Here we used `cuda.cupy` instead of directly accessing `cupy`. This is because the `cupy` module cannot be imported if the CUDA is not installed. In order to keep the implementation valid in non-CUDA environment, we have to defer the access to the `cupy` module. Note that the `chainer.backends.cuda` module can be imported even if the CUDA is not installed. Of course, the module in such environment is almost useless, but if the interpreter does not run through the code accessing CUDA-dedicated functions, the code is still valid.

The CPU and GPU implementations are almost same, except that `numpy` is replaced by `cupy` in GPU methods. We can unify these functions using the `chainer.backends.cuda.get_array_module()` function. This function accepts arbitrary number of arrays, and returns an appropriate module for them. See the following code

```
class ExpAdd(Function):
    def forward(self, inputs):
        xp = cuda.get_array_module(*inputs)
        x, y = inputs
        z = xp.exp(x) + xp.exp(y)
```

(continues on next page)

(continued from previous page)

```

        return z,

    def backward(self, inputs, grad_outputs):
        xp = cuda.get_array_module(*inputs)
        x, y = inputs
        gz, = grad_outputs

        gx = gz * xp.exp(x)
        gy = gz * xp.exp(y)
        return gx, gy

def expadd(x, y):
    return ExpAdd()(x, y)

```

Note that this code works correctly even if CUDA is not installed in the environment. If CUDA is not found, `get_array_module` function always returns `numpy`. We often use the name `xp` for the variadic module name, which is analogous to the abbreviation `np` for NumPy and `cp` for CuPy.

3.4.3 Write an Elementwise Kernel Function

Let's turn back to the `MulAdd` example.

The GPU implementation of `MulAdd` as shown above is already fast and parallelized on GPU cores. However, it invokes two kernels during each of forward and backward computations. It might hurt performance, since the intermediate temporary arrays are read and written by possibly different GPU cores, which consumes much bandwidth. We can reduce the number of invocations by defining our own kernel. It also reduce the memory consumption.

Most functions only require elementwise operations like `MulAdd`. CuPy provides a useful tool to define elementwise kernels, the `cupy.elementwise.ElementwiseKernel` class, and Chainer wraps it by `cuda.elementwise()` function. Our `MulAdd` implementation can be improved as follows:

```

class MulAdd(Function):
    def forward_cpu(self, inputs):
        ...

    def backward_cpu(self, inputs, grad_outputs):
        ...

    def forward_gpu(self, inputs):
        cupy = cuda.cupy
        x, y, z = inputs
        w = cuda.elementwise(
            'float32 x, float32 y, float32 z',
            'float32 w',
            'w = x * y + z',
            'muladd_fwd')(x, y, z)
        return w,

    def backward_gpu(self, inputs, grad_outputs):
        x, y, z = inputs
        gw, = grad_outputs

        gx, gy = cuda.elementwise(
            'float32 x, float32 y, float32 gw',
            'float32 gx, float32 gy',

```

(continues on next page)

(continued from previous page)

```

'''
    gx = y * gw;
    gy = x * gw;
'''
'muladd_bwd')(x, y, gw)

gz = gw
return gx, gy, gz

```

`chainer.backends.cuda.elementwise()` function accepts the essential implementation of the kernel function, and returns a kernel invocation function (actually, it returns `ElementwiseKernel` object, which is callable). In typical usage, we pass four arguments to this function as follows:

1. Input argument list. This is a comma-separated string each entry of which consists of a type specification and an argument name.
2. Output argument list in the same format as the input argument list.
3. Body of *parallel loop*. We can use the input/output argument names as an element of these arrays.
4. Name of the kernel function, which is shown in debuggers and profilers.

Above code is not compiled on every forward/backward computation thanks to two caching mechanisms provided by `cuda.elementwise()`.

The first one is *binary caching*: `chainer.backends.cuda.elementwise()` function caches the compiled binary in the `$(HOME)/.cupy/kernel_cache` directory with a hash value of the CUDA code, and reuses it if the given code matches the hash value. This caching mechanism is actually implemented in CuPy.

The second one is *upload caching*: Given a compiled binary code, we have to upload it to the current GPU in order to execute it. `chainer.backends.cuda.elementwise()` function memoizes the arguments and the current device, and if it is called with the same arguments for the same device, it reuses the previously uploaded kernel code.

The above `MulAdd` code only works for float32 arrays. The `ElementwiseKernel` also supports the type-variadic kernel definition. In order to define variadic kernel functions, you can use *type placeholder* by placing a single character as type specifier:

```

class MulAdd(Function):
    def forward_cpu(self, inputs):
        ...

    def backward_cpu(self, inputs, grad_outputs):
        ...

    def forward_gpu(self, inputs):
        cupy = cuda.cupy
        x, y, z = inputs
        w = cuda.elementwise(
            'T x, T y, T z',
            'T w',
            'w = x * y + z',
            'muladd_fwd')(x, y, z)
        return w,

    def backward_gpu(self, inputs, grad_outputs):
        x, y, z = inputs
        gw, = grad_outputs

        gx, gy = cuda.elementwise(

```

(continues on next page)

(continued from previous page)

```

    'T x, T y, T gw',
    'T gx, T gy',
    '''
        gx = y * gw;
        gy = x * gw;
    '''
    'muladd_bwd')(x, y, gw)

gz = gw
return gx, gy, gz

```

The type placeholder `T` indicates an arbitrary data type that CuPy supports.

There are more functionalities on user-defined kernels in CuPy. See the [CuPy documentation on user-defined kernels](#) for more details.

3.4.4 Write a function with training/test mode

We sometimes want to make a function behave differently in training and test modes. The training/test mode in Chainer is configured by `chainer.config`. This is a thread-local configuration object, and users can substitute `True` or `False` to its `train` attribute. You can refer to [Configuring Chainer](#) to see how to configure this flag as well as other configuration items.

Here, we just show how to use this flag to make a function support training/test mode. You will need to check the value of the boolean flag `chainer.config.train` and branch appropriately.

For example, consider the following simple dropout function:

```

def dropout(x):
    xp = cuda.get_array_module(x.data)
    mask = 2 * (xp.random.rand(*x.shape) > 0.5).astype(x.dtype)
    return x * mask

```

This function applies dropout to each element and doubles survived elements to preserve the scale. The above implementation applies dropout even in test mode, but it is not a desired behavior. We can fix it as follows:

```

def dropout(x):
    if not chainer.config.train:
        return x

    xp = cuda.get_array_module(x.data)
    mask = 2 * (xp.random.rand(*x.shape) > 0.5).astype(x.dtype)
    return x * mask

```

The function now supports test mode. Note that you usually do not have to implement your own dropout function because `dropout()` is officially provided.

3.4.5 Links that wrap functions

Some functions are meant to be combined with parameters. In such case, it is useful to write a small **link** that wraps the function. We have already seen how to define a chain that wraps other links (by inheriting `Chain` class). Here we study how to define a link that does not hold any other links.

As the first example, suppose that we want to implement elementwise product function between the input array and the parameter array. It can be defined as follows:

```
class EltwiseParamProduct(Link):
    def __init__(self, shape):
        super(EltwiseParamProduct, self).__init__()
        with self.init_scope():
            self.W = chainer.Parameter(initializers.Normal(scale=1.), shape)

    def __call__(self, x):
        return self.W * x
```

For another example, assume we want to define a simple linear layer. It is already defined as [Linear](#), so this is an educational example. The linear layer is divided into two parts: a function and its wrapper link. First, we have to define a function on variables:

```
class LinearFunction(Function):
    def forward(self, inputs):
        x, W, b = inputs
        return x.dot(W.T) + b,

    def backward(self, inputs, grad_outputs):
        x, W, b = inputs
        gy, = grad_outputs

        gx = gy.dot(W)
        gW = gy.T.dot(x)
        gb = gy.sum(axis=0)
        return gx, gW, gb

def linear(x, W, b):
    return LinearFunction()(x, W, b)
```

This function takes three arguments: input, weight, and bias. It can be used as a part of model definition, though is inconvenient since the user have to manage the weight and bias parameters directly. In order to make a convenient module, let's wrap it into a link:

```
class Linear(Link):
    def __init__(self, in_size, out_size):
        super(Linear, self).__init__()
        with self.init_scope():
            self.W = chainer.Parameter(
                initializers.Normal(1. / math.sqrt(in_size)),
                (out_size, in_size))
            self.b = chainer.Parameter(0, (out_size,))

    def __call__(self, x):
        return linear(x, self.W, self.b)
```

This link hides the parameters of the linear layer.

Note: An advanced tip to implement functions: if you want to preserve some information between forward and backward computations (e.g. to cache some arrays), you can store it as attributes. Be careful that it might increase the memory consumption during the whole forward-backward computation. If you want to train very large networks on a GPU with limited memory, it is not recommended to cache arrays between forward and backward. There is one exception for this: caching the output arrays does not change the memory consumption, because they are also held by the output Variable objects.

Warning: You should not assume a one-to-one match of calls of forward and backward. Some users may call backward more than once after one forward call.

3.4.6 Testing Function

In order to isolate the cause of learning failure from implementation bugs, it is important to test function implementations. Chainer provides simple utilities to help writing unit tests. They are defined in the `gradient_check` module.

The most important test utility is the `numerical_grad()` function. This function computes the numerical gradient of given function using finite differences. It can be used as follows

```
x = np.random.randn(4, 3).astype(np.float32)
gy = np.ones((4, 3), dtype=np.float32)
f = lambda: (x * x,)
gx = gradient_check.numerical_grad(f, (x,), (gy,))
```

`f` is a closure that returns a tuple of array(s) computed from input arrays. The second and third arguments of `numerical_grad()` are tuples of input arrays and output gradient arrays, respectively. The code above computes the numerical gradients of `sum(f(x))`, where `sum` indicates the summation over all elements. The summation can be weighted by changing `gy`. `numerical_grad()` function also accepts additional `eps` argument, which indicates the quantization width of finite differences.

Note: `numerical_grad()` function accepts both CPU and GPU arrays. Note that we cannot mix CPU and GPU arrays.

Another utility is `chainer.testing.assert_allclose()` function. This is similar to `numpy.testing.assert_allclose()` function. The difference is that Chainer's version accepts CPU and GPU arrays as inputs. We can mix them in one invocation of `chainer.testing.assert_allclose()`. The default values of optional arguments are also different.

Here is a typical usage of gradient checking utilities. This is a test example of `functions.relu()` function

```
import unittest

from chainer import testing

class TestReLU(unittest.TestCase):
    def test_backward_cpu(self):
        x = Variable(np.random.randn(3, 2).astype(np.float32))
        y = F.relu(x)
        y.grad = np.random.randn(3, 2).astype(np.float32)
        y.backward()

        def f():
            return F.relu(x).data,

        gx, = gradient_check.numerical_grad(f, (x.data,), (y.grad,))
        testing.assert_allclose(gx, x.grad)
```

The first four lines of the test code are simple forward and backward computation of ReLU function. The next two lines compute numerical gradient using the same forward function without backward routine. And at last, we compare these

two results elementwise. Note that the above test code can be easily modified to test GPU version just by replacing CPU arrays to GPU arrays.

In most cases, we do not write the code like the above explicitly because Chainer offers a utility function `chainer.gradient_check.check_backward()` that follows this procedure.

```
import unittest

from chainer import gradient_check

class TestReLU(unittest.TestCase):
    def test_backward_cpu(self):

        def f(x):
            return F.relu(x)

        x = np.random.randn(3, 2).astype(np.float32)
        y_grad = np.random.randn(3, 2).astype(np.float32)

        gradient_check.check_backward(f, x, y_grad, atol=1e-4, rtol=1e-4)
```

You can find many examples of function tests under `tests/chainer_tests/functions_tests` directory.

3.5 Creating Models

In the example code of this tutorial, we assume for simplicity that the following symbols are already imported.

```
import numpy as np
import chainer
from chainer.backends import cuda
from chainer import Function, gradient_check, report, training, utils, Variable
from chainer import datasets, iterators, optimizers, serializers
from chainer import Link, Chain, ChainList
import chainer.functions as F
import chainer.links as L
from chainer.training import extensions
```

Most neural network architectures contain multiple links. For example, a multi-layer perceptron consists of multiple linear layers. We can write complex procedures with parameters by combining multiple links like this:

```
>>> l1 = L.Linear(4, 3)
>>> l2 = L.Linear(3, 2)

>>> def my_forward(x):
...     h = l1(x)
...     return l2(h)
```

Here the `L` indicates the `links` module. A procedure with parameters defined in this way is hard to reuse. More Pythonic way is combining the links and procedures into a class:

```
>>> class MyProc(object):
...     def __init__(self):
...         self.l1 = L.Linear(4, 3)
...         self.l2 = L.Linear(3, 2)
... 
```

(continues on next page)

(continued from previous page)

```
...     def forward(self, x):
...         h = self.l1(x)
...         return self.l2(h)
```

In order to make it more reusable, we want to support parameter management, CPU/GPU migration, robust and flexible save/load features, etc. These features are all supported by the `Chain` class in Chainer. Then, what we have to do here is just define the above class as a subclass of `Chain`:

```
>>> class MyChain(Chain):
...     def __init__(self):
...         super(MyChain, self).__init__()
...         with self.init_scope():
...             self.l1 = L.Linear(4, 3)
...             self.l2 = L.Linear(3, 2)
...
...     def __call__(self, x):
...         h = self.l1(x)
...         return self.l2(h)
```

It shows how a complex chain is constructed by simpler links. Links like `l1` and `l2` are called *child links* of `MyChain`. **Note that `Chain` itself inherits `Link`.** It means we can define more complex chains that hold `MyChain` objects as their child links.

Note: We often define a single forward method of a link by the `__call__` operator. Such links and chains are callable and behave like regular functions of Variables.

Note: In Chainer v1, we could also register the trainable layers (i.e., *Link*s) to the model by putting them to the `__init__()` of `Chain` or registering them via `add_link()`. But as these ways are deprecated in Chainer v2, users are recommended to use the way explained above.

Another way to define a chain is using the `ChainList` class, which behaves like a list of links:

```
>>> class MyChain2(ChainList):
...     def __init__(self):
...         super(MyChain2, self).__init__(
...             L.Linear(4, 3),
...             L.Linear(3, 2),
...         )
...
...     def __call__(self, x):
...         h = self[0](x)
...         return self[1](h)
```

`ChainList` can conveniently use an arbitrary number of links, however if the number of links is fixed like in the above case, the `Chain` class is recommended as a base class.

3.6 Optimizer

In the example code of this tutorial, we assume for simplicity that the following symbols are already imported.

```
import numpy as np
import chainer
from chainer.backends import cuda
from chainer import Function, gradient_check, report, training, utils, Variable
from chainer import datasets, iterators, optimizers, serializers
from chainer import Link, Chain, ChainList
import chainer.functions as F
import chainer.links as L
from chainer.training import extensions
```

From the previous guide on *Creating Models*, let's use the `MyChain` class:

```
>>> class MyChain(Chain):
...     def __init__(self):
...         super(MyChain, self).__init__()
...         with self.init_scope():
...             self.l1 = L.Linear(4, 3)
...             self.l2 = L.Linear(3, 2)
...
...     def __call__(self, x):
...         h = self.l1(x)
...         return self.l2(h)
```

To tune parameters values to minimize loss, etc., we have to optimize them by the `Optimizer` class. It runs a numerical optimization algorithm on a given link. Many algorithms are implemented in the `optimizers` module. Here we use the simplest one, called Stochastic Gradient Descent (SGD):

```
>>> model = MyChain()
>>> optimizer = optimizers.SGD().setup(model)
```

The method `setup()` prepares for the optimization given a link.

Some parameter/gradient manipulations, e.g. weight decay and gradient clipping, can be done by setting *hook functions* to the optimizer. Hook functions are called after the gradient computation and right before the actual update of parameters. For example, we can set weight decay regularization by running the next line beforehand:

```
>>> optimizer.add_hook(chainer.optimizer_hooks.WeightDecay(0.0005))
```

Of course, you can write your own hook functions. It should be a function or a callable object.

There are two ways to use the optimizer. One is using it via *Trainer*, which we will see in the following sections. The other way is using it directly. We here review the latter case. To use the optimizer in an automated fashion, see the *Trainer* guide.

There are two further ways to use the optimizer directly. One is manually computing gradients and then calling the `update()` method with no arguments. Do not forget to clear the gradients beforehand!

```
>>> x = np.random.uniform(-1, 1, (2, 4)).astype(np.float32)
>>> model.cleargrads()
>>> # compute gradient here...
>>> loss = F.sum(model(chainer.Variable(x)))
>>> loss.backward()
>>> optimizer.update()
```

The other way is just passing a loss function to the `update()` method. In this case, `cleargrads()` is automatically called by the `update` method, so the user does not have to call it manually.


```
>>> def lossfun(arg1, arg2):
...     # calculate loss
...     loss = F.sum(model(arg1 - arg2))
...     return loss

>>> arg1 = np.random.uniform(-1, 1, (2, 4)).astype(np.float32)
>>> arg2 = np.random.uniform(-1, 1, (2, 4)).astype(np.float32)
>>> optimizer.update(lossfun, chainer.Variable(arg1), chainer.Variable(arg2))
```

See `Optimizer.update()` for the full specification.

3.7 Trainer

When we want to train neural networks, we have to run *training loops* that update the parameters many times. A typical training loop consists of the following procedures:

1. Iterations over training datasets
2. Preprocessing of extracted mini-batches
3. Forward/backward computations of the neural networks
4. Parameter updates
5. Evaluations of the current parameters on validation datasets
6. Logging and printing of the intermediate results

Chainer provides a simple yet powerful way to make it easy to write such training processes. The training loop abstraction mainly consists of two components:

- **Dataset abstraction.** It implements 1 and 2 in the above list. The core components are defined in the `dataset` module. There are also many implementations of datasets and iterators in `datasets` and `iterators` modules, respectively.
- **Trainer.** It implements 3, 4, 5, and 6 in the above list. The whole procedure is implemented by `Trainer`. The way to update parameters (3 and 4) is defined by `Updater`, which can be freely customized. 5 and 6 are implemented by instances of `Extension`, which appends an extra procedure to the training loop. Users can freely customize the training procedure by adding extensions. Users can also implement their own extensions.

3.8 Trainer Extensions

In this section, you will learn about the following topics:

- How to create your own trainer extension
 - by defining a simple function
 - by defining a function decorated with `@make_extension`
 - by defining a class inherited from `Extension` class

In the example code of this tutorial, we assume for simplicity that the following symbols are already imported.

```
import numpy as np
import chainer
from chainer.backends import cuda
from chainer import Function, gradient_check, report, training, utils, Variable
from chainer import datasets, iterators, optimizers, serializers
from chainer import Link, Chain, ChainList
import chainer.functions as F
import chainer.links as L
from chainer.training import extensions
```

3.8.1 What is trainer Extension?

Extension is a callable object that takes a *Trainer* object as an argument. By adding an *Extension* to a *Trainer* using the *extend()* method, the *Extension* will be called according to the schedule specified by using a trigger object (See the details in *1. trigger*)

The *Trainer* object contains all information used in a training loop, e.g., models, optimizers, updaters, iterators, and datasets, etc. This makes it possible to change settings such as the learning rate of an optimizer.

3.8.2 Write a simple function

You can make a new *Extension* by writing a simple function which takes a *Trainer* object as its argument. For example, when you want to reduce the learning rate periodically during training, an *lr_drop* extension can be written as follows:

```
def lr_drop(trainer):
    trainer.updater.get_optimizer('main').lr *= 0.1
```

Then you can add this function to a *Trainer* object via *extend()* method.

```
trainer.extend(lr_drop, trigger=(10, 'epoch'))
```

It lowers the learning rate every 10 epochs by multiplying 0.1 with the current learning rate.

3.8.3 Write a function decorated with @make_extension

make_extension() is a decorator that adds some attributes to a given function. For example, the simple extension we created above can be written in this form:

```
@training.make_extension(trigger=(10, 'epoch'))
def lr_drop(trainer):
    trainer.updater.get_optimizer('main').lr *= 0.1
```

The difference between the above example and this is whether it has a default trigger or not. In the latter case, *lr_drop()* has its default trigger so that unless another trigger is specified via *extend()* method, the trigger specified in *make_extension()* is used by default. The code below acts the same as the former example, i.e., it reduces the learning rate every 10 epochs.

```
trainer.extend(lr_drop)
```

There are several attributes you can add using the *make_extension()* decorator.

1. trigger

trigger is an object that takes a *Trainer* object as an argument and returns a boolean value. If a tuple in the form (period, unit) is given as a trigger, it will be considered as an *IntervalTrigger* that invokes the extension every period unit. For example, when the given tuple is (10, 'epoch'), the extension will run every 10 epochs.

trigger can also be given to the `extend()` method that adds an extension to a *Trainer* object. The priority of triggers is as follows:

- When both `extend()` and a given *Extension* have triggers, the trigger given to `extend()` is used.
- When None is given to `extend()` as the trigger argument and a given *Extension* has trigger, the trigger given to the *Extension* is used.
- When both trigger attributes in `extend()` and *Extension* are None, the *Extension* will be fired every iteration.

See the details in the documentation of `get_trigger()` for more information.

2. default_name

An *Extension* is kept in a dictionary which is a property in a *Trainer*. This argument gives the name of the *Extension*. Users will see this name in the keys of the snapshot which is a dictionary generated by serialization.

3. priority

As a *Trainer* object can be assigned multiple *Extension* objects, the execution order is defined according to the following three values:

- `PRIORITY_WRITER`: The priority for extensions that write some records to the observation dictionary. It includes cases that the extension directly adds values to the observation dictionary, or the extension uses the `chainer.report()` function to report values to the observation dictionary. Extensions which write something to reporter should go first because other Extensions which read those values may be added.
- `PRIORITY_EDITOR`: The priority for extensions that edit the observation dictionary based on already reported values. Extensions which edit some values of reported ones should go after the extensions which write values to reporter but before extensions which read the final values.
- `PRIORITY_READER`: The priority for extensions that only read records from the observation dictionary. This is also suitable for extensions that do not use the observation dictionary at all. Extensions which read the reported values should be fired after all the extensions which have other priorities, e.g, `PRIORITY_WRITER` and `PRIORITY_EDITOR` because it should read the final values.

See the details in the documentation of *Trainer* for more information.

4. finalizer

You can specify a function which takes a *Trainer* object as an argument to finalize the extension. It is called once at the end of the training loop, i.e., when `run()` has finished.

5. initializer

You can specify a function which takes a *Trainer* object as an argument to initialize the extension. It is called once before the training loop begins.

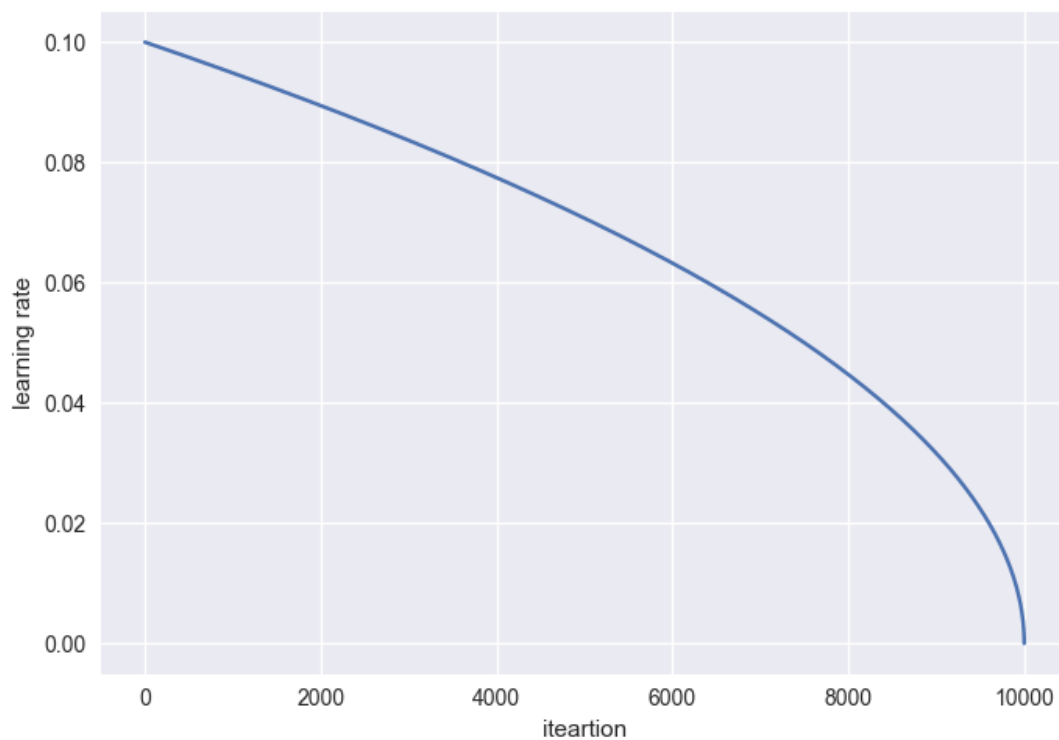
3.8.4 Write a class inherited from the Extension class

This is the way to define your own extension with the maximum degree of freedom. You can keep any values inside of the extension and serialize them.

As an example, let's make an extension that drops the learning rate polynomially. It calculates the learning rate by this equation:

$$\eta = \eta_{\text{init}} \left(1 - \frac{t}{t_{\text{max}}} \right)^{\text{power}}$$

The learning rate will be dropped according to the curve below with power = 0.5:



```
class PolynomialShift(training.Extension):

    def __init__(self, attr, power, stop_trigger, batchsize=None,
                  len_dataset=None):
        self._attr = attr
        self._power = power
        self._init = None
        self._t = 0
        self._last_value = 0

        if stop_trigger[1] == 'iteration':
            self._maxiter = stop_trigger[0]
        elif stop_trigger[1] == 'epoch':
            if batchsize is None or len_dataset is None:
                raise ValueError(
```

(continues on next page)

(continued from previous page)

```

        'When the unit of \'stop_trigger\' is \'epoch\', '
        '\batchsize\' and \'len_dataset\' should be '
        'specified to calculate the maximum iteration.')
    n_iter_per_epoch = len_dataset / float(batchsize)
    self._maxiter = float(stop_trigger[0] * n_iter_per_epoch)

    def initialize(self, trainer):
        optimizer = trainer.updater.get_optimizer('main')
        # ensure that _init is set
        if self._init is None:
            self._init = getattr(optimizer, self._attr)

    def __call__(self, trainer):
        self._t += 1

        optimizer = trainer.updater.get_optimizer('main')
        value = self._init * ((1 - (self._t / self._maxiter)) ** self._power)
        setattr(optimizer, self._attr, value)
        self._last_value = value

    def serialize(self, serializer):
        self._t = serializer('_t', self._t)
        self._last_value = serializer('_last_value', self._last_value)
        if isinstance(self._last_value, np.ndarray):
            self._last_value = np.asscalar(self._last_value)

```

```

stop_trigger = (10000, 'iteration')
trainer.extend(PolynomialShift('lr', 0.5, stop_trigger))

```

This extension `PolynomialShift` takes five arguments.

- `attr`: The name of the optimizer property you want to update using this extension.
- `power`: The power of the above equation to calculate the learning rate.
- `stop_trigger`: The trigger given to the `Trainer` object to specify when to stop the training loop.
- `batchsize`: The training mini-batchsize.
- `len_dataset`: The length of the dataset, i.e., the number of data in the training dataset.

This extension calculates the number of iterations which will be performed during training by using `stop_trigger`, `batchsize`, and `len_dataset`, then stores it as a property `_maxiter`. This property will be used in the `__call__()` method to update the learning rate. The `initialize()` method obtains the initial learning rate from the optimizer given to the `Trainer` object. The `serialize()` method stores or recovers the properties, `_t` (number of iterations) and `_last_value` (the latest learning rate), belonging to this extension.

3.9 Using GPU(s) in Chainer

In the example code of this tutorial, we assume for simplicity that the following symbols are already imported.

```

import numpy as np
import chainer
from chainer.backends import cuda
from chainer import Function, gradient_check, report, training, utils, Variable

```

(continues on next page)

(continued from previous page)

```
from chainer import datasets, iterators, optimizers, serializers
from chainer import Link, Chain, ChainList
import chainer.functions as F
import chainer.links as L
from chainer.training import extensions
```

In this section, you will learn about the following topics:

- Relationship between Chainer and CuPy
- Basics of CuPy
- Single-GPU usage of Chainer
- Multi-GPU usage of model-parallel computing
- Multi-GPU usage of data-parallel computing

After reading this section, you will be able to:

- Use Chainer on a CUDA-enabled GPU
- Write model-parallel computing in Chainer
- Write data-parallel computing in Chainer

3.9.1 Relationship between Chainer and CuPy

Note: From v2.0.0, CuPy is turned into a separate package and repository. Even if you have CUDA installed in your environment, you have to install CuPy separately to use GPUs. See [Working with Custom CUDA Installation](#) for the way to set up CUDA support.

Chainer uses [CuPy](#) as its backend for GPU computation. In particular, the `cupy.ndarray` class is the GPU array implementation for Chainer. CuPy supports a subset of features of NumPy with a compatible interface. It enables us to write a common code for CPU and GPU. It also supports PyCUDA-like user-defined kernel generation, which enables us to write fast implementations dedicated to GPU.

Note: The `chainer.backends.cuda` module imports many important symbols from CuPy. For example, the `cupy` namespace is referred as `cuda.cupy` in the Chainer code. Note that the `chainer.backends.cuda` module can be imported even if CUDA is not installed.

Chainer uses a memory pool for GPU memory allocation. As shown in the previous sections, Chainer constructs and destructs many arrays during learning and evaluating iterations. It is not well suited for CUDA architecture, since memory allocation and release in CUDA (i.e. `cudaMalloc` and `cudaFree` functions) synchronize CPU and GPU computations, which hurts performance. In order to avoid memory allocation and deallocation during the computation, Chainer uses CuPy's memory pool as the standard memory allocator. Chainer changes the default allocator of CuPy to the memory pool, so user can use functions of CuPy directly without dealing with the memory allocator.

3.9.2 Basics of `cupy.ndarray`

See [the document of CuPy](#) for the basic usage of `cupy.ndarray`

CuPy is a GPU array backend that implements a subset of NumPy interface. The `cupy.ndarray` class is in its core, which is a compatible GPU alternative of `numpy.ndarray`. CuPy implements many functions on `cupy.ndarray`

objects. See the reference for the supported subset of NumPy API. Understanding NumPy might help utilizing most features of CuPy. See the NumPy documentation for learning it.

The main difference of `cupy.ndarray` from `numpy.ndarray` is that the content is allocated on the device memory. The allocation takes place on the current device by default. The current device can be changed by `cupy.cuda.Device` object as follows:

```
with cupy.cuda.Device(1):
    x_on_gpu1 = cupy.array([1, 2, 3, 4, 5])
```

Most operations of CuPy is done on the current device. Be careful that it causes an error to process an array on a non-current device.

Chainer provides some convenient functions to automatically switch and choose the device. For example, the `chainer.backends.cuda.to_gpu()` function copies a `numpy.ndarray` object to a specified device:

```
x_cpu = np.ones((5, 4, 3), dtype=np.float32)
x_gpu = cuda.to_gpu(x_cpu, device=1)
```

It is equivalent to the following code using CuPy:

```
x_cpu = np.ones((5, 4, 3), dtype=np.float32)
with cupy.cuda.Device(1):
    x_gpu = cupy.array(x_cpu)
```

Moving a device array to the host can be done by `chainer.backends.cuda.to_cpu()` as follows:

```
x_cpu = cuda.to_cpu(x_gpu)
```

It is equivalent to the following code using CuPy:

```
with x_gpu.device:
    x_cpu = x_gpu.get()
```

Note: The `with` statements in these codes are required to select the appropriate CUDA device. If user uses only one device, these device switching is not needed. `chainer.backends.cuda.to_cpu()` and `chainer.backends.cuda.to_gpu()` functions automatically switch the current device correctly.

Chainer also provides a convenient function `chainer.backends.cuda.get_device_from_id()` and `chainer.backends.cuda.get_device_from_array()` to select a device. The former function accepts an integer or None. When None is given, it returns a *dummy device object*. Otherwise, it returns a corresponding device object. The latter function accepts CuPy array or NumPy array. When a NumPy array is given, it returns a *dummy device object*. Otherwise, it returns a corresponding device object to the give CuPy array. The dummy device object also supports `with` statements like the above example but does nothing. Here are some other examples:

```
cuda.get_device_from_id(1).use()
x_gpu1 = cupy.empty((4, 3), dtype=cupy.float32)

with cuda.get_device_from_id(1):
    x_gpu1 = cupy.empty((4, 3), dtype=cupy.float32)

with cuda.get_device_from_array(x_gpu1):
    y_gpu1 = x_gpu1 + 1
```

Since it accepts NumPy arrays, we can write a function that accepts both NumPy and CuPy arrays with correct device switching:

```
def add1(x):  
    with cuda.get_device_from_array(x):  
        return x + 1
```

The compatibility of CuPy with NumPy enables us to write CPU/GPU generic code. It can be made easy by the `chainer.backends.cuda.get_array_module()` function. This function returns the `numpy` or `cupy` module based on arguments. A CPU/GPU generic function is defined using it like follows:

```
# Stable implementation of log(1 + exp(x))  
def softplus(x):  
    xp = cuda.get_array_module(x)  
    return xp.maximum(0, x) + xp.log1p(xp.exp(-abs(x)))
```

3.9.3 Run Neural Networks on a Single GPU

Single-GPU usage is very simple. What you have to do is transferring *Link* and input arrays to the GPU beforehand. In this subsection, the code is based on *our first MNIST example in this tutorial*.

A *Link* object can be transferred to the specified GPU using the `to_gpu()` method.

This time, we make the number of input, hidden, and output units configurable. The `to_gpu()` method also accepts a device ID like `model.to_gpu(0)`. In this case, the link object is transferred to the appropriate GPU device. The current device is used by default.

If we use `chainer.training.Trainer`, what we have to do is just let the updater know the device ID to send each mini-batch.

```
updater = training.updaters.StandardUpdater(train_iter, optimizer, device=0)  
trainer = training.Trainer(updater, (20, 'epoch'), out='result')
```

We also have to specify the device ID for an evaluator extension as well.

```
trainer.extend(extensions.Evaluator(test_iter, model, device=0))
```

When we write down the training loop by hand, we have to transfer each mini-batch to the GPU manually:

```
model.to_gpu()  
batchsize = 100  
datasize = len(x_train)  
for epoch in range(20):  
    print('epoch %d' % epoch)  
    indexes = np.random.permutation(datasize)  
    for i in range(0, datasize, batchsize):  
        x = Variable(cuda.to_gpu(x_train[indexes[i : i + batchsize]]))  
        t = Variable(cuda.to_gpu(y_train[indexes[i : i + batchsize]]))  
        optimizer.update(model, x, t)
```

3.9.4 Model-parallel Computation on Multiple GPUs

Parallelization of machine learning is roughly classified into two types called “model-parallel” and “data-parallel”. Model-parallel means parallelizations of the computations inside the model. In contrast, data-parallel means parallelizations using data sharding. In this subsection, we show how to use the model-parallel approach on multiple GPUs in Chainer.

Recall the MNIST example. Now suppose that we want to modify this example by expanding the network to 6 layers with 2000 units each using two GPUs. In order to make multi-GPU computation efficient, we only make the two GPUs communicate at the third and sixth layer. The overall architecture looks like the following diagram:

```
(GPU0) input --+--> 11 --> 12 --> 13 --+--> 14 --> 15 --> 16 --+--> output
              |                               |
(GPU1)         +--> 11 --> 12 --> 13 --+--> 14 --> 15 --> 16 --+
```

We can use the above MLP chain as following diagram:

```
(GPU0) input --+--> mlp1 --+--> mlp2 --+--> output
              |               |
(GPU1)         +--> mlp1 --+--> mlp2 --+
```

Let's write a link for the whole network.

```
class ParallelMLP(Chain):
    def __init__(self):
        super(ParallelMLP, self).__init__()
        with self.init_scope():
            # the input size, 784, is inferred
            self.mlp1_gpu0 = MLP(1000, 2000).to_gpu(0)
            self.mlp1_gpu1 = MLP(1000, 2000).to_gpu(1)

            # the input size, 2000, is inferred
            self.mlp2_gpu0 = MLP(1000, 10).to_gpu(0)
            self.mlp2_gpu1 = MLP(1000, 10).to_gpu(1)

    def __call__(self, x):
        # assume x is on GPU 0
        z0 = self.mlp1_gpu0(x)
        z1 = self.mlp1_gpu1(F.copy(x, 1))

        # sync
        h0 = F.relu(z0 + F.copy(z1, 0))
        h1 = F.relu(z1 + F.copy(z0, 1))

        y0 = self.mlp2_gpu0(h0)
        y1 = self.mlp2_gpu1(h1)

        # sync
        y = y0 + F.copy(y1, 0)
        return y # output is on GPU0
```

Recall that the `Link.to_gpu()` method returns the link itself. The `copy()` function copies an input variable to specified GPU device and returns a new variable on the device. The copy supports backprop, which just reversely transfers an output gradient to the input device.

Note: Above code is not parallelized on CPU, but is parallelized on GPU. This is because all the functions in the above code run asynchronously to the host CPU.

An almost identical example code can be found at [examples/mnist/train_mnist_model_parallel.py](https://github.com/chainer/examples/blob/master/mnist/train_mnist_model_parallel.py).

3.9.5 Data-parallel Computation on Multiple GPUs with Trainer

Data-parallel computation is another strategy to parallelize online processing. In the context of neural networks, it means that a different device does computation on a different subset of the input data. In this subsection, we review the way to achieve data-parallel learning on two GPUs.

Suppose again our task is *the MNIST example*. This time we want to directly parallelize the three-layer network. The most simple form of data-parallelization is parallelizing the gradient computation for a distinct set of data. First, define a model and optimizer instances:

```
model = L.Classifier(MLP(1000, 10)) # the input size, 784, is inferred
optimizer = optimizers.SGD()
optimizer.setup(model)
```

Recall that the `MLP` link implements the multi-layer perceptron, and the `Classifier` link wraps it to provide a classifier interface. We used `StandardUpdater` in the previous example. In order to enable data-parallel computation with multiple GPUs, we only have to replace it with `ParallelUpdater`.

```
updater = training.updaters.ParallelUpdater(train_iter, optimizer,
                                           devices={'main': 0, 'second': 1})
```

The `devices` option specifies which devices to use in data-parallel learning. The device with name 'main' is used as the main device. The original model is sent to this device, so the optimization runs on the main device. In the above example, the model is also cloned and sent to GPU 1. Half of each mini-batch is fed to this cloned model. After every backward computation, the gradient is accumulated into the main device, the parameter update runs on it, and then the updated parameters are sent to GPU 1 again.

See also the example code in `examples/mnist/train_mnist_data_parallel.py`.

3.9.6 Data-parallel Computation on Multiple GPUs without Trainer

We here introduce a way to write data-parallel computation without the help of `Trainer`. Most users can skip this section. If you are interested in how to write a data-parallel computation by yourself, this section should be informative. It is also helpful to, e.g., customize the `ParallelUpdater` class.

We again start from the MNIST example. At this time, we use a suffix like `_0` and `_1` to distinguish objects on each device. First, we define a model.

```
model_0 = L.Classifier(MLP(1000, 10)) # the input size, 784, is inferred
```

We want to make two copies of this instance on different GPUs. The `Link.to_gpu()` method runs in place, so we cannot use it to make a copy. In order to make a copy, we can use `Link.copy()` method.

```
model_1 = model_0.copy()
model_0.to_gpu(0)
model_1.to_gpu(1)
```

The `Link.copy()` method copies the link into another instance. *It just copies the link hierarchy*, and does not copy the arrays it holds.

Then, set up an optimizer:

```
optimizer = optimizers.SGD()
optimizer.setup(model_0)
```

Here we use the first copy of the model as *the master model*. Before its update, gradients of `model_1` must be aggregated to those of `model_0`.

Then, we can write a data-parallel learning loop as follows:

```
batchsize = 100
datasize = len(x_train)
for epoch in range(20):
    print('epoch %d' % epoch)
    indexes = np.random.permutation(datasize)
    for i in range(0, datasize, batchsize):
        x_batch = x_train[indexes[i : i + batchsize]]
        y_batch = y_train[indexes[i : i + batchsize]]

        x0 = Variable(cuda.to_gpu(x_batch[:batchsize//2], 0))
        t0 = Variable(cuda.to_gpu(y_batch[:batchsize//2], 0))
        x1 = Variable(cuda.to_gpu(x_batch[batchsize//2:], 1))
        t1 = Variable(cuda.to_gpu(y_batch[batchsize//2:], 1))

        loss_0 = model_0(x0, t0)
        loss_1 = model_1(x1, t1)

        model_0.cleargrads()
        model_1.cleargrads()

        loss_0.backward()
        loss_1.backward()

        model_0.addgrads(model_1)
        optimizer.update()

    model_1.copyparams(model_0)
```

Do not forget to clear the gradients of both model copies! One half of the mini-batch is forwarded to GPU 0, the other half to GPU 1. Then the gradients are accumulated by the `Link.addgrads()` method. This method adds the gradients of a given link to those of the self. After the gradients are prepared, we can update the optimizer in usual way. Note that the update only modifies the parameters of `model_0`. So we must manually copy them to `model_1` using `Link.copyparams()` method.

Note: If the batch size used in one model remain the same, the scale of the gradient is roughly proportional to the number of models, when we aggregate gradients from all models by `chainer.Link.addgrads()`. So you need to adjust the batch size and/or learning rate of the optimizer accordingly.

Now you can use Chainer with GPUs. All examples in the `examples` directory support GPU computation, so please refer to them if you want to know more practices on using GPUs. In the next section, we will show how to define a differentiable (i.e. *backpropable*) function on Variable objects. We will also show there how to write a simple (elementwise) CUDA kernel using Chainer's CUDA utilities.

3.10 Type Checks

In this section, you will learn about the following things:

- Basic usage of type check
- Detail of type information
- Internal mechanism of type check

- More complicated cases
- Call functions
- Typical type check example

After reading this section, you will be able to:

- Write a code to check types of input arguments of your own functions

3.10.1 Basic usage of type check

When you call a function with an invalid type of array, you sometimes receive no error, but get an unexpected result by broadcasting. When you use CUDA with an illegal type of array, it causes memory corruption, and you get a serious error. These bugs are hard to fix. Chainer can check preconditions of each function, and helps to prevent such problems. These conditions may help a user to understand specification of functions.

Each implementation of `Function` has a method for type check, `check_type_forward()`. This function is called just before the `forward()` method of the `Function` class. You can override this method to check the condition on types and shapes of arguments.

`check_type_forward()` gets an argument `in_types`:

```
def check_type_forward(self, in_types):  
    ...
```

`in_types` is an instance of `TypeInfoTuple`, which is a sub-class of `tuple`. To get type information about the first argument, use `in_types[0]`. If the function gets multiple arguments, we recommend to use new variables for readability:

```
x_type, y_type = in_types
```

In this case, `x_type` represents the type of the first argument, and `y_type` represents the second one.

We describe usage of `in_types` with an example. When you want to check if the number of dimension of `x_type` equals to 2, write this code:

```
utils.type_check.expect(x_type.ndim == 2)
```

When this condition is true, nothing happens. Otherwise this code throws an exception, and the user gets a message like this:

```
Traceback (most recent call last):  
...  
chainer.utils.type_check.InvalidType: Expect: in_types[0].ndim == 2  
Actual: 3 != 2
```

This error message means that “ndim of the first argument expected to be 2, but actually it is 3”.

3.10.2 Detail of type information

You can access three information of `x_type`.

- `.shape` is a tuple of ints. Each value is size of each dimension.
- `.ndim` is `int` value representing the number of dimensions. Note that `ndim == len(shape)`
- `.dtype` is `numpy.dtype` representing data type of the value.

You can check all members. For example, the size of the first dimension must be positive, you can write like this:

```
utils.type_check.expect(x_type.shape[0] > 0)
```

You can also check data types with `.dtype`:

```
utils.type_check.expect(x_type.dtype == np.float64)
```

And an error is like this:

```
Traceback (most recent call last):
...
chainer.utils.type_check.InvalidType: Expect: in_types[0].dtype == <class 'numpy.
↪float64'>
Actual: float32 != <class 'numpy.float64'>
```

You can also check kind of dtype. This code checks if the type is floating point

```
utils.type_check.expect(x_type.dtype.kind == 'f')
```

You can compare between variables. For example, the following code checks if the first argument and the second argument have the same length:

```
utils.type_check.expect(x_type.shape[1] == y_type.shape[1])
```

3.10.3 Internal mechanism of type check

How does it show an error message like `"in_types[0].ndim == 2"`? If `x_type` is an object containing `ndim` member variable, we cannot show such an error message because this equation is evaluated as a boolean value by Python interpreter.

Actually `x_type` is a *Expr* objects, and doesn't have a `ndim` member variable itself. *Expr* represents a syntax tree. `x_type.ndim` makes a *Expr* object representing `(getattr, x_type, 'ndim')`. `x_type.ndim == 2` makes an object like `(eq, (getattr, x_type, 'ndim'), 2)`. `type_check.expect()` gets a *Expr* object and evaluates it. When it is `True`, it causes no error and shows nothing. Otherwise, this method shows a readable error message.

If you want to evaluate a *Expr* object, call `eval()` method:

```
actual_type = x_type.eval()
```

`actual_type` is an instance of `TypeInfo`, while `x_type` is an instance of *Expr*. In the same way, `x_type.shape[0].eval()` returns an int value.

3.10.4 More powerful methods

Expr class is more powerful. It supports all mathematical operators such as `+` and `*`. You can write a condition that the first dimension of `x_type` is the first dimension of `y_type` times four:

```
utils.type_check.expect(x_type.shape[0] == y_type.shape[0] * 4)
```

When `x_type.shape[0] == 3` and `y_type.shape[0] == 1`, users can get the error message below:

```
Traceback (most recent call last):
...
chainer.utils.type_check.InvalidType: Expect: in_types[0].shape[0] == in_types[1].
↪ shape[0] * 4
Actual: 3 != 4
```

To compare a member variable of your function, wrap a value with `Variable` to show readable error message:

```
x_type.shape[0] == utils.type_check.Variable(self.in_size, "in_size")
```

This code can check the equivalent condition below:

```
x_type.shape[0] == self.in_size
```

However, the latter condition doesn't know the meaning of this value. When this condition is not satisfied, the latter code shows unreadable error message:

```
chainer.utils.type_check.InvalidType: Expect: in_types[0].shape[0] == 4 # what does
↪ '4' mean?
Actual: 3 != 4
```

Note that the second argument of `utils.type_check.Variable` is only for readability.

The former shows this message:

```
chainer.utils.type_check.InvalidType: Expect: in_types[0].shape[0] == in_size # OK,
↪ `in_size` is a value that is given to the constructor
Actual: 3 != 4 # You can also check actual value here
```

3.10.5 Call functions

How to check summation of all values of shape? *Expr* also supports function call:

```
sum = utils.type_check.Variable(np.sum, 'sum')
utils.type_check.expect(sum(x_type.shape) == 10)
```

Why do we need to wrap the function `numpy.sum` with `utils.type_check.Variable`? `x_type.shape` is not a tuple but an object of *Expr* as we have seen before. Therefore, `numpy.sum(x_type.shape)` fails. We need to evaluate this function lazily.

The above example produces an error message like this:

```
Traceback (most recent call last):
...
chainer.utils.type_check.InvalidType: Expect: sum(in_types[0].shape) == 10
Actual: 7 != 10
```

3.10.6 More complicated cases

How to write a more complicated condition that can't be written with these operators? You can evaluate *Expr* and get its result value with `eval()` method. Then check the condition and show warning message by hand:

```
x_shape = x_type.shape.eval() # get actual shape (int tuple)
if not more_complicated_condition(x_shape):
    expect_msg = 'Shape is expected to be ...'
    actual_msg = 'Shape is ...'
    raise utils.type_check.InvalidType(expect_msg, actual_msg)
```

Please write a readable error message. This code generates the following error message:

```
Traceback (most recent call last):
...
chainer.utils.type_check.InvalidType: Expect: Shape is expected to be ...
Actual: Shape is ...
```

3.10.7 Typical type check example

We show a typical type check for a function.

First check the number of arguments:

```
utils.type_check.expect(in_types.size() == 2)
```

`in_types.size()` returns a *Expr* object representing the number of arguments. You can check it in the same way.

And then, get each type:

```
x_type, y_type = in_types
```

Don't get each value before checking `in_types.size()`. When the number of argument is illegal, `type_check.expect` might output unuseful error messages. For example, this code doesn't work when the size of `in_types` is 0:

```
utils.type_check.expect(
    in_types.size() == 2,
    in_types[0].ndim == 3,
)
```

After that, check each type:

```
utils.type_check.expect(
    x_type.dtype == np.float32,
    x_type.ndim == 3,
    x_type.shape[1] == 2,
)
```

The above example works correctly even when `x_type.ndim == 0` as all conditions are evaluated lazily.

3.11 Serializers – saving and loading

Serializer is a simple interface to serialize or deserialize an object. `Link`, `Optimizer`, and `Trainer` support serialization.

Concrete serializers are defined in the `serializers` module. It supports NumPy NPZ and HDF5 formats.

For example, we can serialize a link object into NPZ file by the `serializers.save_npz()` function:

Assuming we have defined a model:

```
>>> from chainer import serializers
>>> serializers.save_npz('my.model', model)
```

This saves the parameters of `model` into the file `'my.model'` in NPZ format. The saved model can be read back from `my.model` back into `model` by the `serializers.load_npz()` function:

```
>>> serializers.load_npz('my.model', model)
```

Note: Note that only the parameters and the *persistent values* are serialized by this serialization code. Other attributes are not saved automatically. You can register arrays, scalars, or any serializable objects as persistent values by the `Link.add_persistent()` method. The registered values can be accessed by attributes of the name passed to the `add_persistent` method.

The state of an optimizer can also be saved by the same functions:

```
>>> serializers.save_npz('my.state', optimizer)
>>> serializers.load_npz('my.state', optimizer)
```

Note: Note that serialization of optimizer only saves its internal states including number of iterations, momentum vectors of MomentumSGD, etc. It does not save the parameters and persistent values of the target link. We have to explicitly save the target link with the optimizer to resume the optimization from saved states.

Support of the HDF5 format is enabled if the `h5py` package is installed. Serialization and deserialization with the HDF5 format are almost identical to those with the NPZ format; just replace `save_npz()` and `load_npz()` by `save_hdf5()` and `load_hdf5()`, respectively.

3.12 Customize your own logging

In this section, you will learn about the following things:

- What is `chainer.Reporter`?
- How to report logging with `chainer.Reporter`?
- The naming rule for the reported values.

After reading this section, you will be able to:

- Write your own report.

3.12.1 What is Reporter?

`chainer.Reporter` is used to collect values that users want to watch. The reporter object manipulates a dictionary from value names to the actually observed values. We call this dictionary as *observation*.

See the following example:

```
>>> from chainer import Reporter, report, report_scope
>>>
>>> reporter = Reporter()
```

(continues on next page)

(continued from previous page)

```
>>> observer = object() # it can be an arbitrary (reference) object
>>> reporter.add_observer('my_observer:', observer)
>>> observation = {}
>>> with reporter.scope(observation):
...     reporter.report({'x': 1}, observer)
...
>>> observation
{'my_observer:/x': 1}
```

When a value is passed to the `reporter`, an object called `observer` can be optionally attached. In this case, the name of the observer is added as the prefix of the value name. The observer name should be registered beforehand. Using `reporter.scope`, you can select which observation to save the observed values.

There are also a global API `chainer.report()`, which reports observed values with the current reporter object. In this case, *current* means which with statement scope the current code line is in. This function calls the `Reporter.report()` method of the current reporter.

```
>>> observation = {}
>>> with reporter.scope(observation):
...     report({'x': 1}, observer)
...
>>> observation
{'my_observer:/x': 1}
```

3.12.2 Use report in Chain or Link

The most important application of `Reporter` is to report observed values from each `Link` or `Chain` in the training and validation procedures.

But, how to report the observed values from each link or chain? Should we prepare the `Reporter`? No, you only need to call `report()` in chain or link, because `Trainer` and some extensions prepare their own `Reporter` object with the hierarchy of the target link registered as observers. We can use `report()` function inside any links and chains to report the observed values (e.g., training loss, accuracy, activation statistics, etc.).

See the following example:

```
>>> class Classifier(Chain):
...     def __init__(self, predictor):
...         super(Classifier, self).__init__()
...         with self.init_scope():
...             self.predictor = predictor
...
...     def __call__(self, x, t):
...         y = self.predictor(x)
...         loss = F.softmax_cross_entropy(y, t)
...         accuracy = F.accuracy(y, t)
...         report({'loss': loss, 'accuracy': accuracy}, self)
...         return loss
...
... 
```

If the link is named 'main' in the hierarchy (which is the default name of the target link in the `StandardUpdater`), these reported values are named 'main/loss' and 'main/accuracy'. If these values are reported inside the `Evaluator` extension, 'validation/' is added at the head of the link name, thus the item names are changed to 'validation/main/loss' and 'validation/main/accuracy' ('validation' is the default name of the `Evaluator` extension).

3.12.3 Naming rule for the reported values

So, you know almost everything about `Reporter`. However, there is one more thing. It is what is the naming rule for the reported values, especially when the values are reported from a link that is not the root of the link hierarchy.

As we explained in the previous section, the root of links is named as 'main' by the `StandardUpdater` and the names of reported values in the root have the prefix 'main/'. When the values are reported from a link that is not the root of the link hierarchy, the prefix of the names are determined by the link hierarchy, or `namedlinks()`.

See the following example:

```
>>> class MLP(Chain):
...     def __init__(self, n_units, n_out):
...         super(MLP, self).__init__()
...         with self.init_scope():
...             # the size of the inputs to each layer will be inferred
...             self.l1 = L.Linear(None, n_units) # n_in -> n_units
...             self.l2 = L.Linear(None, n_units) # n_units -> n_units
...             self.l3 = L.Linear(None, n_out)   # n_units -> n_out
...
...     def __call__(self, x):
...         h1 = F.relu(self.l1(x))
...         h2 = F.relu(self.l2(h1))
...         y = self.l3(h2)
...         report({'sum_y': F.sum(y)}, self)
...         return y
...
>>> model = Classifier(MLP(100, 10))
>>> for name, observer in model.namedlinks(skipself=True):
...     print(name)
/predictor
/predictor/l1
/predictor/l2
/predictor/l3
```

You can get the parameters of the link hierarchy by `namedlinks()`. In this example, we report 'loss' and 'accuracy' in the root of links, and 'sum_y' in the link of '/predictor'. So, you can access the reported values by 'main/accuracy', 'main/accuracy', and 'main/predictor/sum_y'.

See what we explained is correct:

```
>>> train, test = datasets.get_mnist()
>>> train_iter = iterators.SerialIterator(train, batch_size=100, shuffle=True)
>>> test_iter = iterators.SerialIterator(test, batch_size=100, repeat=False,
↳shuffle=False)
>>> optimizer = optimizers.SGD()
>>> optimizer.setup(model)
>>> updater = training.StandardUpdater(train_iter, optimizer)
>>> trainer = training.Trainer(updater, (1, 'epoch'), out='result')
>>> trainer.extend(extensions.Evaluator(test_iter, model))
>>> trainer.extend(extensions.LogReport())
>>> trainer.extend(extensions.PrintReport(
...     ['epoch', 'main/accuracy', 'main/loss', 'main/predictor/sum_y', 'validation/
↳main/accuracy']))
>>> trainer.run()
epoch      main/accuracy  main/loss  main/predictor/sum_y  validation/main/accuracy
1          0.662317   1.38345   47.9927               0.8498
```

4.1 MNIST using Trainer

In the example code of this tutorial, we assume for simplicity that the following symbols are already imported.

```
import numpy as np
import chainer
from chainer.backends import cuda
from chainer import Function, gradient_check, report, training, utils, Variable
from chainer import datasets, iterators, optimizers, serializers
from chainer import Link, Chain, ChainList
import chainer.functions as F
import chainer.links as L
from chainer.training import extensions
```

By using *Trainer*, you don't need to write the training loop explicitly any more. Furthermore, Chainer provides many useful extensions that can be used with *Trainer* to visualize your results, evaluate your model, store and manage log files more easily.

This example will show how to use the *Trainer* to train a fully-connected feed-forward neural network on the MNIST dataset.

Note: If you would like to know how to write a training loop without using the *Trainer*, please check *MNIST with a Manual Training Loop* instead of this tutorial.

4.1.1 1. Prepare the dataset

Load the MNIST dataset, which contains a training set of images and class labels as well as a corresponding test set.

```
from chainer.datasets import mnist

train, test = mnist.get_mnist()
```

Note: You can use a Python list as a dataset. That's because *Iterator* can take any object as a dataset whose elements can be accessed via `[]` accessor and whose length can be obtained with `len()` function. For example,

```
train = [(x1, t1), (x2, t2), ...]
```

a list of tuples like this can be used as a dataset.

There are many utility dataset classes defined in *datasets*. It's recommended to utilize them in the actual applications.

For example, if your dataset consists of a number of image files, it would take a large amount of memory to load those data into a list like above. In that case, you can use *ImageDataset*, which just keeps the paths to image files. The actual image data will be loaded from the disk when the corresponding element is requested via `[]` accessor. Until then, no images are loaded to the memory to reduce memory use.

4.1.2 2. Prepare the dataset iterations

Iterator creates a mini-batch from the given dataset.

```
batchsize = 128

train_iter = iterators.SerialIterator(train, batchsize)
test_iter = iterators.SerialIterator(test, batchsize, False, False)
```

4.1.3 3. Prepare the model

Here, we are going to use the same model as the one defined in *MNIST with a Manual Training Loop*.

```
class MLP(Chain):

    def __init__(self, n_mid_units=100, n_out=10):
        super(MLP, self).__init__()
        with self.init_scope():
            self.l1 = L.Linear(None, n_mid_units)
            self.l2 = L.Linear(None, n_mid_units)
            self.l3 = L.Linear(None, n_out)

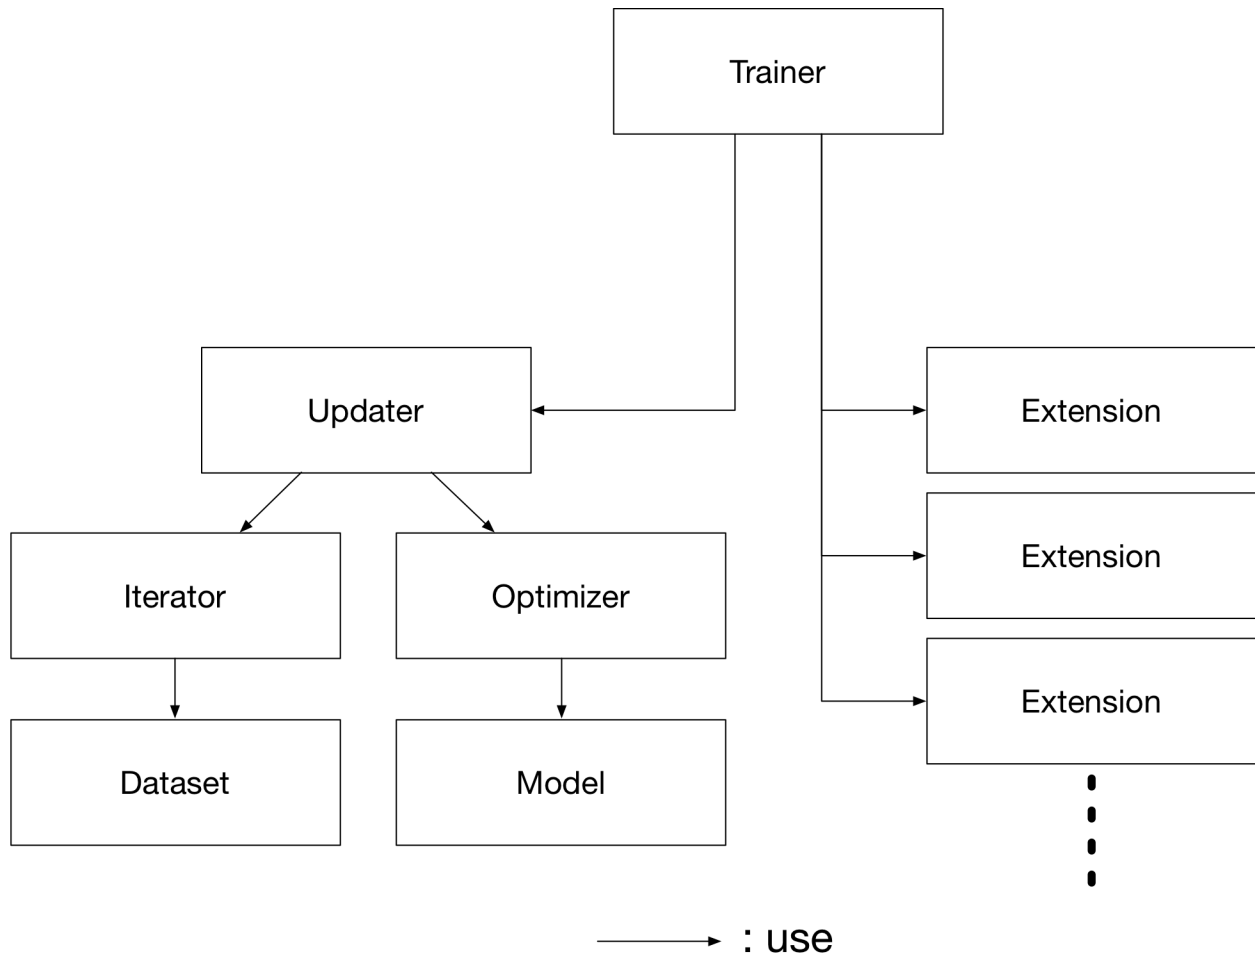
    def __call__(self, x):
        h1 = F.relu(self.l1(x))
        h2 = F.relu(self.l2(h1))
        return self.l3(h2)

gpu_id = 0 # Set to -1 if you use CPU

model = MLP()
if gpu_id >= 0:
    model.to_gpu(gpu_id)
```

4.1.4 4. Prepare the Updater

Trainer is a class that holds all of the necessary components needed for training. The main components are shown below.



Basically, all you need to pass to *Trainer* is an *Updater*. However, *Updater* contains an *Iterator* and *Optimizer*. Since *Iterator* can access the dataset and *Optimizer* has references to the model, *Updater* can access to the model to update its parameters.

So, *Updater* can perform the training procedure as shown below:

1. Retrieve the data from dataset and construct a mini-batch (*Iterator*)
2. Pass the mini-batch to the model and calculate the loss
3. Update the parameters of the model (*Optimizer*)

Now let's create the *Updater* object !

```

max_epoch = 10

# Wrap your model by Classifier and include the process of loss calculation within_
# your model.
# Since we do not specify a loss function here, the default 'softmax_cross_entropy'
# is used.
model = L.Classifier(model)
  
```

(continues on next page)

(continued from previous page)

```
# selection of your optimizing method
optimizer = optimizers.MomentumSGD()

# Give the optimizer a reference to the model
optimizer.setup(model)

# Get an updater that uses the Iterator and Optimizer
updater = training.updaters.StandardUpdater(train_iter, optimizer, device=gpu_id)
```

Note: Here, the model defined above is passed to *Classifier* and changed to a new *Chain*. *Classifier*, which in fact inherits from the *Chain* class, keeps the given *Chain* model in its `predictor` attribute. Once you give the input data and the corresponding class labels to the model by the `()` operator,

1. `__call__()` of the model is invoked. The data is then given to `predictor` to obtain the output `y`.
2. Next, together with the given labels, the output `y` is passed to the loss function which is determined by `lossfun` argument in the constructor of *Classifier*.
3. The loss is returned as a *Variable*.

In *Classifier*, the `lossfun` is set to `softmax_cross_entropy()` as default.

StandardUpdater is the simplest class among several updaters. There are also the *ParallelUpdater* and the *MultiprocessParallelUpdater* to utilize multiple GPUs. The *MultiprocessParallelUpdater* uses the NVIDIA NCCL library, so you need to install NCCL and re-install CuPy before using it.

4.1.5 5. Setup Trainer

Lastly, we will setup *Trainer*. The only requirement for creating a *Trainer* is to pass the *Updater* object that we previously created above. You can also pass a `stop_trigger` to the second trainer argument as a tuple like `(length, unit)` to tell the trainer when to stop the training. The `length` is given as an integer and the `unit` is given as a string which should be either `epoch` or `iteration`. Without setting `stop_trigger`, the training will never be stopped.

```
# Setup a Trainer
trainer = training.Trainer(updater, (max_epoch, 'epoch'), out='mnist_result')
```

The `out` argument specifies an output directory used to save the log files, the image files of plots to show the time progress of loss, accuracy, etc. when you use *PlotReport* extension. Next, we will explain how to display or save those information by using trainer *Extension*.

4.1.6 6. Add Extensions to the Trainer object

The *Trainer* extensions provide the following capabilities:

- Save log files automatically (*LogReport*)
- Display the training information to the terminal periodically (*PrintReport*)
- Visualize the loss progress by plotting a graph periodically and save it as an image file (*PlotReport*)
- Automatically serialize the state periodically (`snapshot()` / `snapshot_object()`)
- Display a progress bar to the terminal to show the progress of training (*ProgressBar*)

- Save the model architecture as a Graphviz's dot file (`dump_graph()`)

To use these wide variety of tools for your training task, pass *Extension* objects to the `extend()` method of your *Trainer* object.

```
from chainer.training import extensions

trainer.extend(extensions.LogReport())
trainer.extend(extensions.snapshot(filename='snapshot_epoch-{}.updater.epoch'))
trainer.extend(extensions.snapshot_object(model.predictor, filename='model_epoch-{}.
    ↳updater.epoch'))
trainer.extend(extensions.Evaluator(test_iter, model, device=gpu_id))
trainer.extend(extensions.PrintReport(['epoch', 'main/loss', 'main/accuracy',
    ↳'validation/main/loss', 'validation/main/accuracy', 'elapsed_time']))
trainer.extend(extensions.PlotReport(['main/loss', 'validation/main/loss'], x_key=
    ↳'epoch', file_name='loss.png'))
trainer.extend(extensions.PlotReport(['main/accuracy', 'validation/main/accuracy'], x_
    ↳key='epoch', file_name='accuracy.png'))
trainer.extend(extensions.dump_graph('main/loss'))
```

LogReport

Collect loss and accuracy automatically every epoch or iteration and store the information under the log file in the directory specified by the `out` argument when you create a *Trainer* object.

snapshot()

The `snapshot()` method saves the *Trainer* object at the designated timing (default: every epoch) in the directory specified by `out`. The *Trainer* object, as mentioned before, has an *Updater* which contains an *Optimizer* and a model inside. Therefore, as long as you have the snapshot file, you can use it to come back to the training or make inferences using the previously trained model later.

snapshot_object()

However, when you keep the whole *Trainer* object, in some cases, it is very tedious to retrieve only the inside of the model. By using `snapshot_object()`, you can save the particular object (in this case, the model wrapped by *Classifier*) as a separate snapshot. *Classifier* is a *Chain* object which keeps the model that is also a *Chain* object as its `predictor` property, and all the parameters are under the `predictor`, so taking the snapshot of `predictor` is enough to keep all the trained parameters.

This is a list of commonly used trainer extensions:

LogReport This extension collects the loss and accuracy values every epoch or iteration and stores in a log file. The log file will be located under the output directory (specified by `out` argument of the *Trainer* object).

snapshot() This extension saves the *Trainer* object at the designated timing (default: every epoch) in the output directory. The *Trainer* object, as mentioned before, has an *Updater* which contains an *Optimizer* and a model inside. Therefore, as long as you have the snapshot file, you can use it to come back to the training or make inferences using the previously trained model later.

snapshot_object() `snapshot()` extension above saves the whole *Trainer* object. However, in some cases, it is tedious to retrieve only the inside of the model. By using `snapshot_object()`, you can save the particular object (in the example above, the model wrapped by *Classifier*) as a separated snapshot. Taking the snapshot of `predictor` is enough to keep all the trained parameters, because *Classifier* (which is a subclass of *Chain*) keeps the model as its `predictor` property, and all the parameters are under this property.

`dump_graph()` This extension saves the structure of the computational graph of the model. The graph is saved in Graphviz dot format under the output directory of the `Trainer`.

`Evaluator` `Iterators` that use the evaluation dataset and the model object are required to use `Evaluator` extension. It evaluates the model using the given dataset (typically it's a validation dataset) at the specified timing interval.

`PrintReport` This extension outputs the specified values to the standard output.

`PlotReport` This extension plots the values specified by its arguments and saves it as a image file.

This is not an exhaustive list of built-in extensions. Please take a look at [Extensions](#) for more of them.

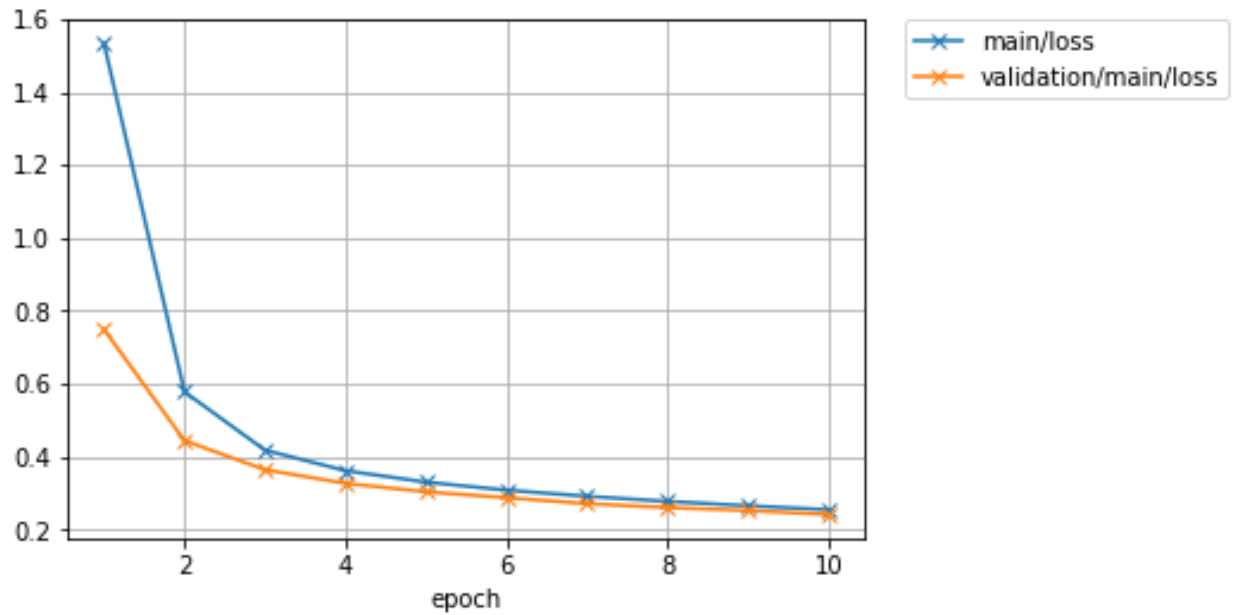
4.1.7 7. Start Training

Just call `run()` method from `Trainer` object to start training.

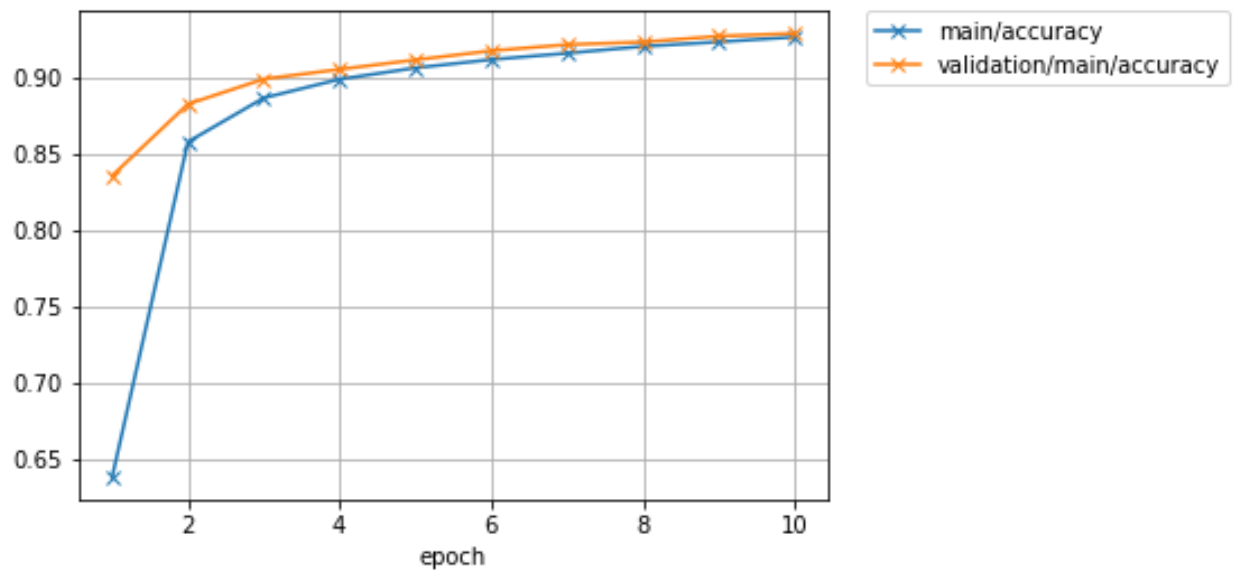
```
trainer.run()
```

epoch	main/loss	main/accuracy	validation/main/loss	validation/main/accuracy	
↪ elapsed_time					
1	1.53241	0.638409	0.74935	0.835839	↪
↪ 4.93409					
2	0.578334	0.858059	0.444722	0.882812	↪
↪ 7.72883					
3	0.418569	0.886844	0.364943	0.899229	↪
↪ 10.4229					
4	0.362342	0.899089	0.327569	0.905558	↪
↪ 13.148					
5	0.331067	0.906517	0.304399	0.911788	↪
↪ 15.846					
6	0.309019	0.911964	0.288295	0.917722	↪
↪ 18.5395					
7	0.292312	0.916128	0.272073	0.921776	↪
↪ 21.2173					
8	0.278291	0.92059	0.261351	0.923457	↪
↪ 23.9211					
9	0.266266	0.923541	0.253195	0.927314	↪
↪ 26.6612					
10	0.255489	0.926739	0.242415	0.929094	↪
↪ 29.466					

Let's see the plot of loss progress saved in the `mnist_result` directory.

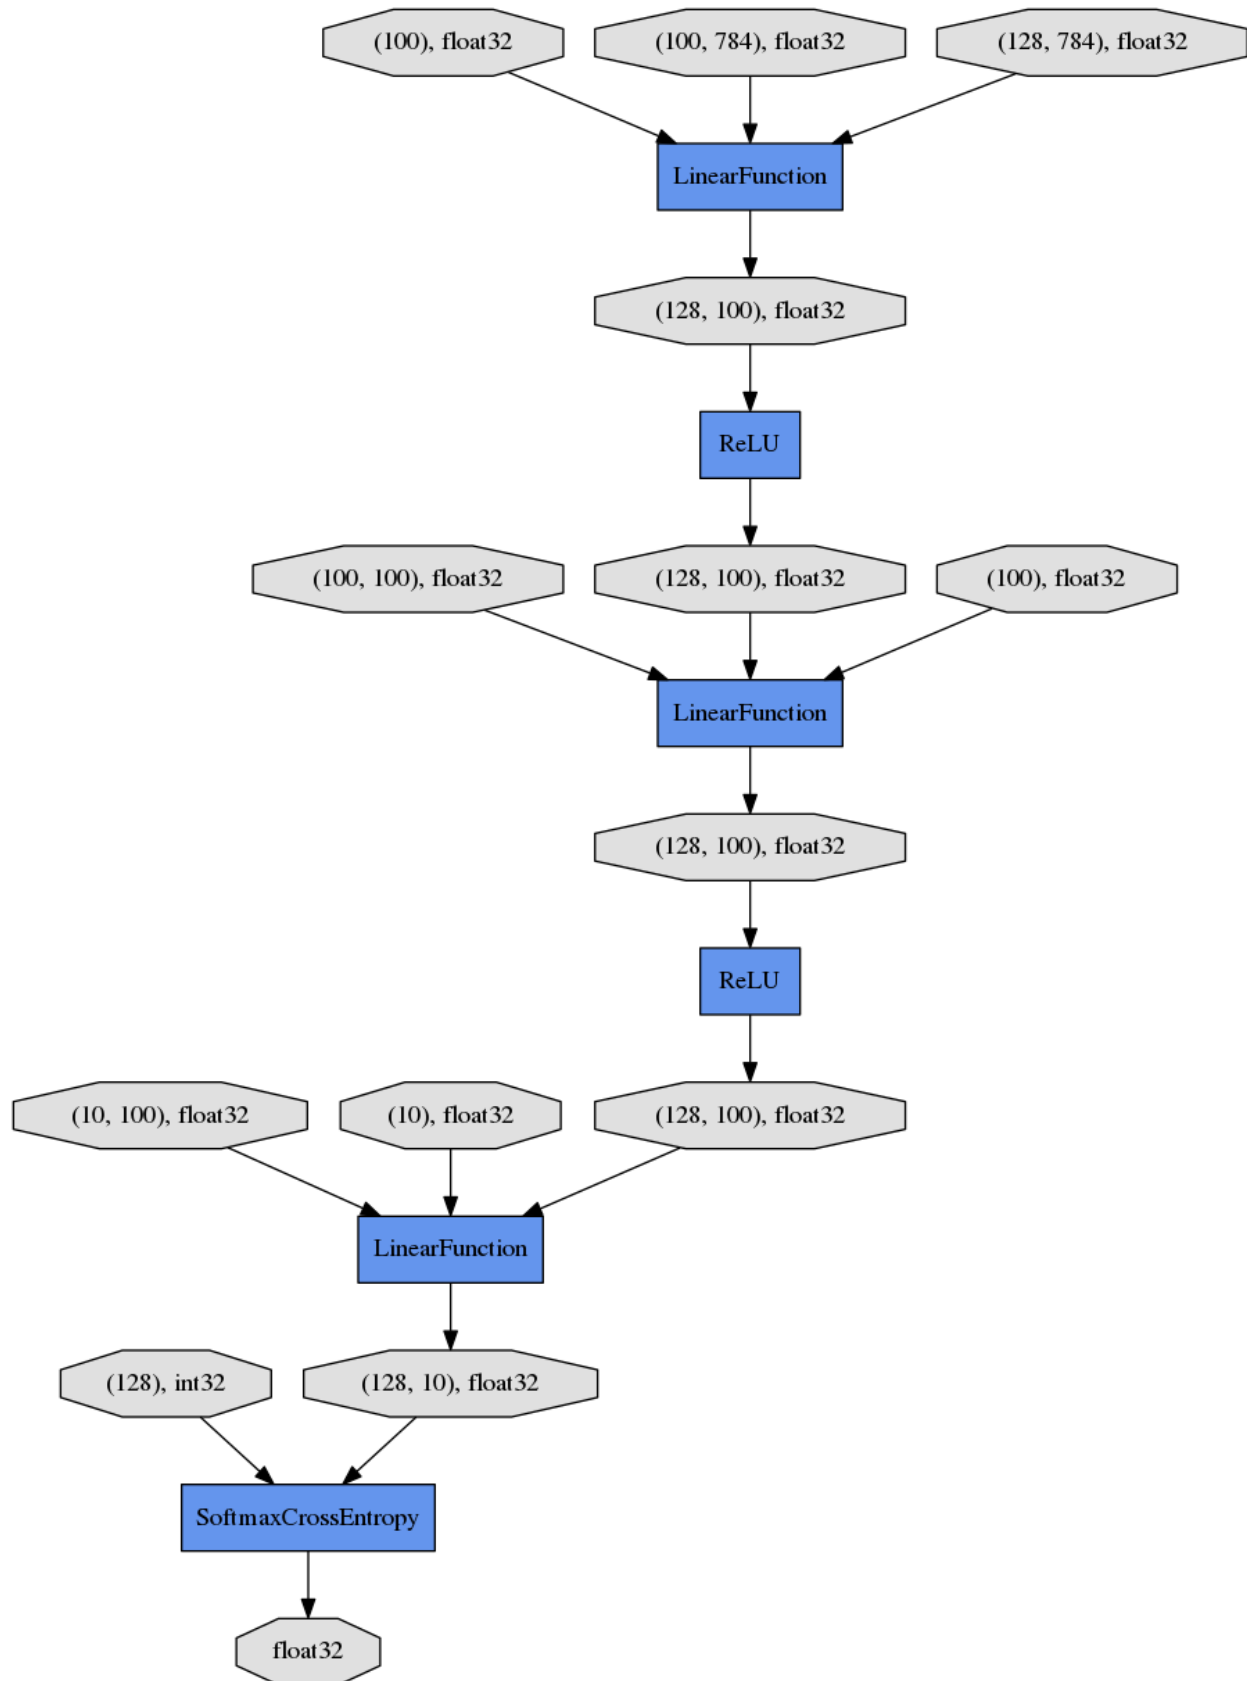


How about the accuracy?



Furthermore, let's visualize the computational graph saved with `dump_graph()` using Graphviz.

```
% dot -Tpng mnist_result/cg.dot -o mnist_result/cg.png
```



From the top to the bottom, you can see the data flow in the computational graph. It basically shows how data and parameters are passed to the *Functions*.

4.1.8 8. Evaluate a pre-trained model

Evaluation using the snapshot of a model is as easy as what explained in the *MNIST with a Manual Training Loop*.

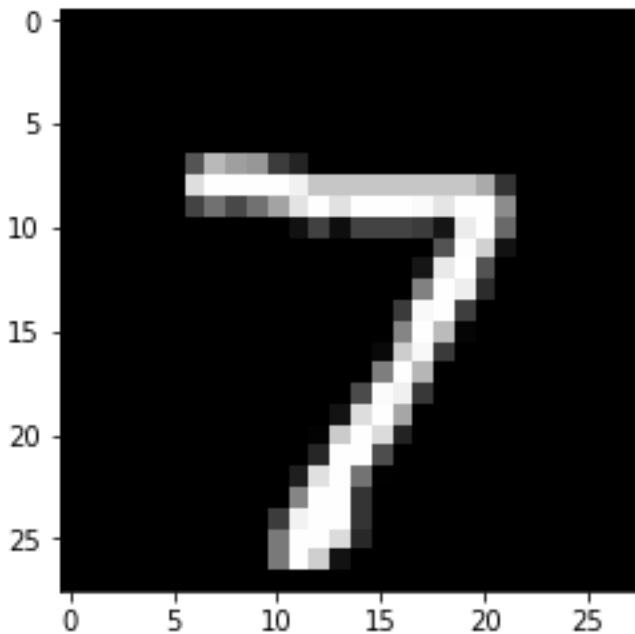
```
import matplotlib.pyplot as plt

model = MLP()
serializers.load_npz('mnist_result/model_epoch-10', model)

# Show the output
x, t = test[0]
plt.imshow(x.reshape(28, 28), cmap='gray')
plt.show()
print('label:', t)

y = model(x[None, ...])

print('predicted_label:', y.data.argmax(axis=1)[0])
```



```
label: 7
predicted_label: 7
```

The prediction looks correct. Success!

4.2 MNIST with a Manual Training Loop

In the example code of this tutorial, we assume for simplicity that the following symbols are already imported.

```
import numpy as np
import chainer
from chainer.backends import cuda
from chainer import Function, gradient_check, report, training, utils, Variable
from chainer import datasets, iterators, optimizers, serializers
from chainer import Link, Chain, ChainList
import chainer.functions as F
import chainer.links as L
from chainer.training import extensions
```

In this tutorial section, we will learn how to train a deep neural network to classify images of hand-written digits in the popular MNIST dataset. This dataset contains 50,000 training examples and 10,000 test examples. Each example is a set of a 28 x 28 greyscale image and a corresponding class label. Since the digits from 0 to 9 are used, there are 10 classes for the labels.

Chainer provides a feature called *Trainer* that can simplify the training procedure of your model. However, it is also good to know how the training works in Chainer before starting to use the useful *Trainer* class that hides the actual processes. Writing your own training loop can be useful for learning how *Trainer* works or for implementing features not included in the standard trainer.

The complete training procedure consists of the following steps:

1. *Prepare a dataset*
2. *Create a dataset iterator*
3. *Define a network*
4. *Select an optimization algorithm*
5. *Write a training loop*
 - (a) Retrieve a set of examples (mini-batch) from the training dataset.
 - (b) Feed the mini-batch to your network.
 - (c) Run a forward pass of the network and compute the loss.
 - (d) Just call the `backward()` method from the loss *Variable* to compute the gradients for all trainable parameters.
 - (e) Run the optimizer to update those parameters.
6. *Save the trained model*
7. *Perform classification by the saved model* and check the network performance on validation/test sets.

4.2.1 1. Prepare a dataset

Chainer contains some built-in functions to use some popular datasets like MNIST, CIFAR10/100, etc. Those can automatically download the data from servers and provide dataset objects which are easy to use.

The code below shows how to retrieve the MNIST dataset from the server and save an image from its training split to make sure the images are correctly obtained.

```
from __future__ import print_function
import matplotlib.pyplot as plt
from chainer.datasets import mnist

# Download the MNIST data if you haven't downloaded it yet
train, test = mnist.get_mnist(withlabel=True, ndim=1)
```

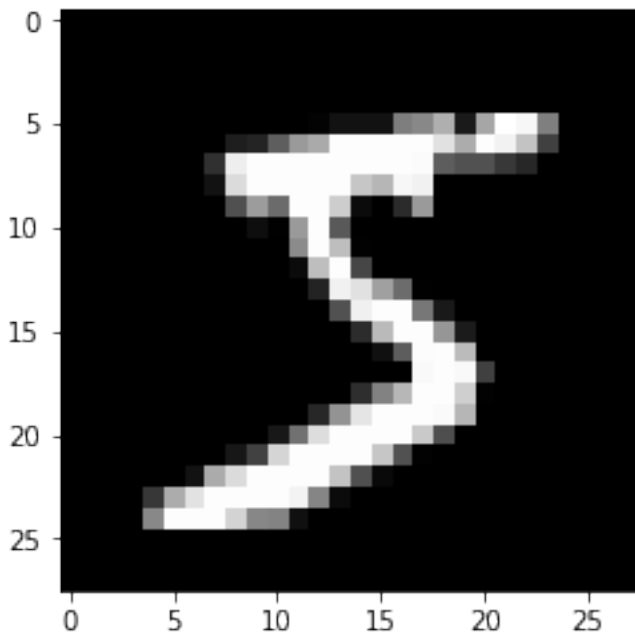
(continues on next page)

(continued from previous page)

```
# Display an example from the MNIST dataset.
# `x` contains the input image array and `t` contains that target class
# label as an integer.
x, t = train[0]
plt.imshow(x.reshape(28, 28), cmap='gray')
plt.savefig('5.png')
print('label:', t)
```

```
label: 5
```

The saved image 5.png will look like:



4.2.2 2. Create a dataset iterator

Although this is an optional step, we'd like to introduce the *Iterator* class that retrieves a set of data and labels from the given dataset to easily make a mini-batch. There are some subclasses that can perform the same thing in different ways, e.g., using multi-processing to parallelize the data loading part, etc.

Here, we use *SerialIterator*, which is also a subclass of *Iterator* in the example code below. The *SerialIterator* can provide mini-batches with or without shuffling the order of data in the given dataset.

All *Iterators* produce a new mini-batch by calling its *next()* method. All *Iterators* also have properties to know how many times we have taken all the data from the given dataset (*epoch*) and whether the next mini-batch will be the start of a new epoch (*is_new_epoch*), and so on.

The code below shows how to create a *SerialIterator* object from a dataset object.

```
from chainer import iterators

# Choose the minibatch size.
batchsize = 128
```

(continues on next page)

(continued from previous page)

```
train_iter = iterators.SerialIterator(train, batchsize)
test_iter = iterators.SerialIterator(test, batchsize,
                                     repeat=False, shuffle=False)
```

Note: `iterators` can take a built-in Python list as a given dataset. It means that the example code below is able to work,

```
train = [(x1, t1), (x2, t2), ...] # A list of tuples
train_iter = iterators.SerialIterator(train, batchsize)
```

where `x1`, `x2`, ... denote the input data and `t1`, `t2`, ... denote the corresponding labels.

Details of `SerialIterator`

- `SerialIterator` is a built-in subclass of `Iterator` that can retrieve a mini-batch from a given dataset in either sequential or shuffled order.
- The `Iterator`'s constructor takes two arguments: a dataset object and a mini-batch size.
- If you want to use the same dataset repeatedly during the training process, set the `repeat` argument to `True` (default). Otherwise, the dataset will be used only one time. The latter case is actually for the evaluation.
- If you want to shuffle the training dataset every epoch, set the `shuffle` argument to `True`. Otherwise, the order of each data retrieved from the dataset will be always the same at each epoch.

In the example code shown above, we set `batchsize = 128` in both `train_iter` and `test_iter`. So, these iterators will provide 128 images and corresponding labels at a time.

4.2.3 3. Define a network

Now let's define a neural network that we will train to classify the MNIST images. For simplicity, we use a three-layer perceptron here. We set each hidden layer to have 100 units and set the output layer to have 10 units, which is corresponding to the number of class labels of the MNIST.

Create your network as a subclass of `Chain`

You can create your network by writing a new subclass of `Chain`. The main steps are twofold:

1. Register the network components which have trainable parameters to the subclass. Each of them must be instantiated and assigned to a property in the scope specified by `init_scope()`:
2. Define a `__call__()` method that represents the actual **forward computation** of your network. This method takes one or more `Variable`, `numpy.array`, or `cupy.array` as its inputs and calculates the forward pass using them.

```
class MyNetwork(Chain):

    def __init__(self, n_mid_units=100, n_out=10):
        super(MyNetwork, self).__init__()
        with self.init_scope():
            self.ll = L.Linear(None, n_mid_units)
```

(continues on next page)

(continued from previous page)

```

        self.l2 = L.Linear(n_mid_units, n_mid_units)
        self.l3 = L.Linear(n_mid_units, n_out)

    def __call__(self, x):
        h = F.relu(self.l1(x))
        h = F.relu(self.l2(h))
        return self.l3(h)

model = MyNetwork()

gpu_id = 0 # Set to -1 if you use CPU
if gpu_id >= 0:
    model.to_gpu(gpu_id)

```

Link, *Chain*, *ChainList*, and those subclass objects which contain trainable parameters should be registered to the model by assigning it as a property inside the `init_scope()`. For example, a *FunctionNode* does not contain any trainable parameters, so there is no need to keep the object as a property of your network. When you want to use `relu()` in your network, using it as a function in `__call__()` works correctly.

In Chainer, the Python code that implements the forward computation itself represents the network. In other words, we can conceptually think of the computation graph for our network being constructed dynamically as this forward computation code executes. This allows Chainer to describe networks in which different computations can be performed in each iteration, such as branched networks, intuitively and with a high degree of flexibility. This is the key feature of Chainer that we call **Define-by-Run**.

4.2.4 4. Select an optimization algorithm

Chainer provides a wide variety of optimization algorithms that can be used to optimize the network parameters during training. They are located in `optimizers` module.

Here, we are going to use the stochastic gradient descent (SGD) method with momentum, which is implemented by *MomentumSGD*. To use the optimizer, we give the network object (typically it's a *Chain* or *ChainList*) to the `setup()` method of the optimizer object to register it. In this way, the *Optimizer* can automatically find the model parameters and update them during training.

You can easily try out other optimizers as well. Please test and observe the results of various optimizers. For example, you could try to change *MomentumSGD* to *Adam*, *RMSprop*, etc.

```

from chainer import optimizers

# Choose an optimizer algorithm
optimizer = optimizers.MomentumSGD(lr=0.01, momentum=0.9)

# Give the optimizer a reference to the model so that it
# can locate the model's parameters.
optimizer.setup(model)

```

Note: In the above example, we set `lr` to 0.01 in the constructor. This value is known as the “learning rate”, one of the most important hyperparameters that need to be adjusted in order to obtain the best performance. The various optimizers may each have different hyperparameters and so be sure to check the documentation for the details.

4.2.5 5. Write a training loop

We now show how to write the training loop. Since we are working on a digit classification problem, we will use `softmax_cross_entropy()` as the loss function for the optimizer to minimize. For other types of problems, such as regression models, other loss functions might be more appropriate. See the [Chainer documentation for detailed information on the various loss functions](#) for more details.

Our training loop will be structured as follows.

1. We will first get a mini-batch of examples from the training dataset.
2. We will then feed the batch into our network by calling it (a `Chain` object) like a function. This will execute the forward-pass code that are written in the `__call__()` method.
3. This will return the network output that represents class label predictions. We supply it to the loss function along with the true (that is, target) values. The loss function will output the loss as a `Variable` object.
4. We then clear any previous gradients in the network and perform the backward pass by calling the `backward()` method on the loss variable which computes the parameter gradients. We need to clear the gradients first because the `backward()` method accumulates gradients instead of overwriting the previous values.
5. Since the optimizer already has a reference to the network, it has access to the parameters and the computed gradients so that we can now call the `update()` method of the optimizer which will update the model parameters.

In addition to the above steps, you might want to check the performance of the network with a validation dataset. This allows you to observe how well it is generalized to new data so far, namely, you can check whether it is overfitting to the training data. The code below checks the performance on the test set at the end of each epoch. The code has the same structure as the training code except that no backpropagation is performed and we also compute the accuracy on the test data using the `accuracy()` function.

The training loop code is as follows:

```
import numpy as np
from chainer.dataset import concat_examples
from chainer.backends.cuda import to_cpu

max_epoch = 10

while train_iter.epoch < max_epoch:

    # ----- One iteration of the training loop -----
    train_batch = train_iter.next()
    image_train, target_train = concat_examples(train_batch, gpu_id)

    # Calculate the prediction of the network
    prediction_train = model(image_train)

    # Calculate the loss with softmax_cross_entropy
    loss = F.softmax_cross_entropy(prediction_train, target_train)

    # Calculate the gradients in the network
    model.cleargrads()
    loss.backward()

    # Update all the trainable parameters
    optimizer.update()
    # ----- until here -----
```

(continues on next page)

(continued from previous page)

```

# Check the validation accuracy of prediction after every epoch
if train_iter.is_new_epoch: # If this iteration is the final iteration of the_
    ↪current epoch

    # Display the training loss
    print('epoch:{:02d} train_loss:{:.04f} '.format(
        train_iter.epoch, float(to_cpu(loss.data))), end='')

    test_losses = []
    test_accuracies = []
    while True:
        test_batch = test_iter.next()
        image_test, target_test = concat_examples(test_batch, gpu_id)

        # Forward the test data
        prediction_test = model(image_test)

        # Calculate the loss
        loss_test = F.softmax_cross_entropy(prediction_test, target_test)
        test_losses.append(to_cpu(loss_test.data))

        # Calculate the accuracy
        accuracy = F.accuracy(prediction_test, target_test)
        accuracy.to_cpu()
        test_accuracies.append(accuracy.data)

    if test_iter.is_new_epoch:
        test_iter.epoch = 0
        test_iter.current_position = 0
        test_iter.is_new_epoch = False
        test_iter._pushed_position = None
        break

    print('val_loss:{:.04f} val_accuracy:{:.04f}'.format(
        np.mean(test_losses), np.mean(test_accuracies)))

```

Output

```

epoch:01 train_loss:0.8072 val_loss:0.7592 val_accuracy:0.8289
epoch:02 train_loss:0.5021 val_loss:0.4467 val_accuracy:0.8841
epoch:03 train_loss:0.3539 val_loss:0.3673 val_accuracy:0.9007
epoch:04 train_loss:0.2524 val_loss:0.3307 val_accuracy:0.9067
epoch:05 train_loss:0.4232 val_loss:0.3076 val_accuracy:0.9136
epoch:06 train_loss:0.3033 val_loss:0.2910 val_accuracy:0.9167
epoch:07 train_loss:0.2004 val_loss:0.2773 val_accuracy:0.9222
epoch:08 train_loss:0.2885 val_loss:0.2679 val_accuracy:0.9239
epoch:09 train_loss:0.2818 val_loss:0.2579 val_accuracy:0.9266
epoch:10 train_loss:0.2403 val_loss:0.2484 val_accuracy:0.9307

```

4.2.6 6. Save the trained model

Chainer provides two types of *serializers* that can be used to save and restore model state. One supports the HDF5 format and the other supports the NumPy NPZ format. For this example, we are going to use the NPZ format to save our model since it is easy to use with NumPy and doesn't need to install any additional dependencies or libraries.

```
serializers.save_npz('my_mnist.model', model)
```

4.2.7 7. Perform classification by the saved model

Let's use the saved model to classify a new image. In order to load the trained model parameters, we need to perform the following two steps:

1. Instantiate the same network as what you trained.
2. Overwrite all parameters in the model instance with the saved weights using the `load_npz()` function.

Once the model is restored, it can be used to predict image labels on new input data.

```
from chainer import serializers

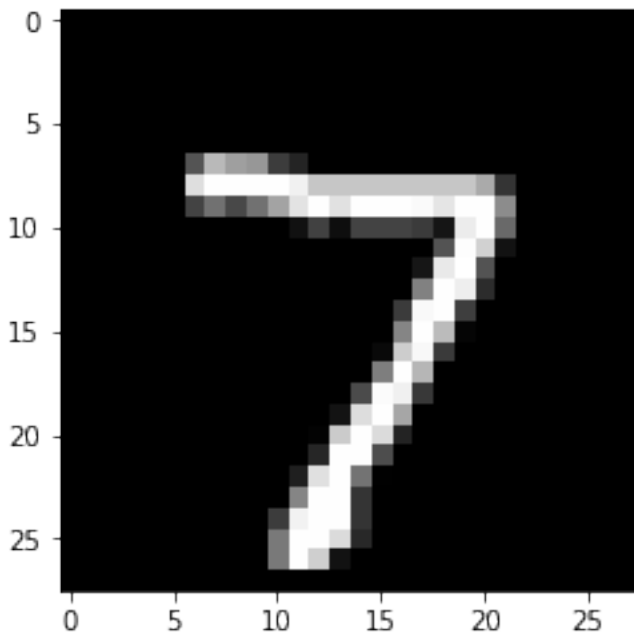
# Create an instance of the network you trained
model = MyNetwork()

# Load the saved parameters into the instance
serializers.load_npz('my_mnist.model', model)

# Get a test image and label
x, t = test[0]
plt.imshow(x.reshape(28, 28), cmap='gray')
plt.savefig('7.png')
print('label:', t)
```

```
label: 7
```

The saved test image looks like:



```

# Change the shape of the minibatch.
# In this example, the size of minibatch is 1.
# Inference using any mini-batch size can be performed.

print(x.shape, end=' -> ')
x = x[None, ...]
print(x.shape)

# Forward calculation of the model by sending X
y = model(x)

# The result is given as Variable, then we can take a look at the contents by the_
# attribute, .data.
y = y.data

# Look up the most probable digit number using argmax
pred_label = y.argmax(axis=1)

print('predicted label:', pred_label[0])

```

```

(784,) -> (1, 784)
predicted label: 7

```

The prediction result looks correct. Yay!

4.3 Convolutional Network for Visual Recognition Tasks

In this section, you will learn how to write

- A small convolutional network with a model class that is inherited from `Chain`,
- A large convolutional network that has several building block networks with `ChainList`.

After reading this section, you will be able to:

- Write your own original convolutional network in Chainer

A convolutional network (ConvNet) is mainly comprised of convolutional layers. This type of network is commonly used for various visual recognition tasks, e.g., classifying hand-written digits or natural images into given object classes, detecting objects from an image, and labeling all pixels of an image with the object classes (semantic segmentation), and so on.

In such tasks, a typical ConvNet takes a set of images whose shape is (N, C, H, W) , where

- N denotes the number of images in a mini-batch,
- C denotes the number of channels of those images,
- H and W denote the height and width of those images,

respectively. Then, it typically outputs a fixed-sized vector as membership probabilities over the target object classes. It also can output a set of feature maps that have the corresponding size to the input image for a pixel labeling task, etc.

In the example code of this tutorial, we assume for simplicity that the following symbols are already imported.

```
import numpy as np
import chainer
from chainer.backends import cuda
from chainer import Function, gradient_check, report, training, utils, Variable
from chainer import datasets, iterators, optimizers, serializers
from chainer import Link, Chain, ChainList
import chainer.functions as F
import chainer.links as L
from chainer.training import extensions
```

4.3.1 LeNet5

Here, let's start by defining LeNet5 [LeCun98] in Chainer. This is a ConvNet model that has 5 layers comprised of 3 convolutional layers and 2 fully-connected layers. This was proposed to classify hand-written digit images in 1998. In Chainer, the model can be written as follows:

```
class LeNet5(Chain):
    def __init__(self):
        super(LeNet5, self).__init__()
        with self.init_scope():
            self.conv1 = L.Convolution2D(
                in_channels=1, out_channels=6, ksize=5, stride=1)
            self.conv2 = L.Convolution2D(
                in_channels=6, out_channels=16, ksize=5, stride=1)
            self.conv3 = L.Convolution2D(
                in_channels=16, out_channels=120, ksize=4, stride=1)
            self.fc4 = L.Linear(None, 84)
            self.fc5 = L.Linear(84, 10)

    def __call__(self, x):
        h = F.sigmoid(self.conv1(x))
        h = F.max_pooling_2d(h, 2, 2)
        h = F.sigmoid(self.conv2(h))
        h = F.max_pooling_2d(h, 2, 2)
        h = F.sigmoid(self.conv3(h))
        h = F.sigmoid(self.fc4(h))
        if chainer.config.train:
            return self.fc5(h)
        return F.softmax(self.fc5(h))
```

A typical way to write your network is creating a new class inherited from `Chain` class. When defining your model in this way, typically, all the layers which have trainable parameters are registered to the model by assigning the objects of `Link` as an attribute.

The model class is instantiated before the forward and backward computations. To give input images and label vectors simply by calling the model object like a function, `__call__()` is usually defined in the model class. This method performs the forward computation of the model. Chainer uses the powerful autograd system for any computational graphs written with `FunctionNodes` and `Links` (actually a `Link` calls a corresponding `FunctionNode` inside of it), so that you don't need to explicitly write the code for backward computations in the model. Just prepare the data, then give it to the model. The way this works is the resulting output `Variable` from the forward computation has a `backward()` method to perform autograd. In the above model, `__call__()` has a `if` statement at the end to switch its behavior by the Chainer's running mode, i.e., training mode or not. Chainer presents the running mode as a global variable `chainer.config.train`. When it's in training mode, `__call__()` returns the output value of the last layer as is to compute the loss later on, otherwise it returns a prediction result by calculating `softmax()`.

Note: In Chainer v1, if a function or link behaved differently in training and other modes, it was common that it held an attribute that represented its running mode or was provided with the mode from outside as an argument. In Chainer v2, it is recommended to use the global configuration `chainer.config.train` to switch the running mode.

If you don't want to write `conv1` and the other layers more than once, you can also write the same model like in this way:

```
from functools import partial

class LeNet5(Chain):
    def __init__(self):
        super(LeNet5, self).__init__()
        net = [('conv1', L.Convolution2D(1, 6, 5, 1))]
        net += [('_sigm1', F.sigmoid)]
        net += [('_mpool1', partial(F.max_pooling_2d, ksize=2, stride=2))]
        net += [('_conv2', L.Convolution2D(6, 16, 5, 1))]
        net += [('_sigm2', F.sigmoid)]
        net += [('_mpool2', partial(F.max_pooling_2d, ksize=2, stride=2))]
        net += [('_conv3', L.Convolution2D(16, 120, 4, 1))]
        net += [('_sigm3', F.sigmoid)]
        net += [('_mpool3', partial(F.max_pooling_2d, ksize=2, stride=2))]
        net += [('_fc4', L.Linear(None, 84))]
        net += [('_sigm4', F.sigmoid)]
        net += [('_fc5', L.Linear(84, 10))]
        net += [('_sigm5', F.sigmoid)]
        with self.init_scope():
            for n in net:
                if not n[0].startswith('_'):
                    setattr(self, n[0], n[1])
        self.forward = net

    def __call__(self, x):
        for n, f in self.forward:
            if not n.startswith('_'):
                x = getattr(self, n)(x)
            else:
                x = f(x)
        if chainer.config.train:
            return x
        return F.softmax(x)
```

Note: You can also use `Sequential` to write the above model more simply. Please note that `Sequential` is an experimental feature introduced in Chainer v4 and its interface may be changed in the future versions.

This code creates a list of pairs of component name (e.g., `conv1`, `_sigm1`, etc.) and all `Links` and functions (e.g., `F.sigmoid`, which internally invokes `FunctionNode`) after calling its superclass's constructor. In this case, components whose name start with `_` are functions (`FunctionNode`), which doesn't have any trainable parameters, so that we don't register (`setattr`) it to the model. Others (`conv1`, `fc4`, etc.) are `Links`, which are trainable layers that hold parameters. This operation can be freely replaced with many other ways because those component names are just designed to select `Links` only from the list `net` easily. The list `net` is stored as an attribute `forward` to refer it in `__call__()`. In `__call__()`, it retrieves all layers in the network from `self.forward` sequentially and gives the input variable or the intermediate output from the previous layer to the current layer. The last part of the `__call__()` to switch its behavior by the training/inference mode is the same as the former way.

Ways to calculate loss

When you train the model with label vector `t`, the loss should be calculated using the output from the model. There also are several ways to calculate the loss:

```
model = LeNet5()

# Input data and label
x = np.random.rand(32, 1, 28, 28).astype(np.float32)
t = np.random.randint(0, 10, size=(32,)).astype(np.int32)

# Forward computation
y = model(x)

# Loss calculation
loss = F.softmax_cross_entropy(y, t)
```

This is a primitive way to calculate a loss value from the output of the model. On the other hand, the loss computation can be included in the model itself by wrapping the model object (*Chain* or *ChainList* object) with a class inherited from *Chain*. The outer *Chain* should take the model defined above and register it with `init_scope()`. *Chain* is actually inherited from *Link*, so that *Chain* itself can also be registered as a trainable *Link* to another *Chain*. Actually, *Classifier* class to wrap the model and add the loss computation to the model already exists. Actually, there is already a *Classifier* class that can be used to wrap the model and include the loss computation as well. It can be used like this:

```
model = L.Classifier(LeNet5())

# Forward & Loss calculation
loss = model(x, t)
```

This class takes a model object as an input argument and registers it to a `predictor` property as a trained parameter. As shown above, the returned object can then be called like a function in which we pass `x` and `t` as the input arguments and the resulting loss value (which we recall is a *Variable*) is returned.

See the detailed implementation of *Classifier* from here: `chainer.links.Classifier` and check the implementation by looking at the source.

From the above examples, we can see that Chainer provides the flexibility to write our original network in many different ways. Such flexibility intends to make it intuitive for users to design new and complex models.

4.3.2 VGG16

Next, let's write some larger models in Chainer. When you write a large network consisting of several building block networks, *ChainList* is useful. First, let's see how to write a VGG16 [Simonyan14] model.

```
class VGG16(chainer.ChainList):
    def __init__(self):
        super(VGG16, self).__init__(
            VGGBlock(64),
            VGGBlock(128),
            VGGBlock(256, 3),
            VGGBlock(512, 3),
            VGGBlock(512, 3, True))

    def __call__(self, x):
        for f in self.children():
```

(continues on next page)

(continued from previous page)

```

        x = f(x)
    if chainer.config.train:
        return x
    return F.softmax(x)

class VGGBlock(chainer.Chain):
    def __init__(self, n_channels, n_convs=2, fc=False):
        w = chainer.initializers.HeNormal()
        super(VGGBlock, self).__init__()
        with self.init_scope():
            self.conv1 = L.Convolution2D(None, n_channels, 3, 1, 1, initialW=w)
            self.conv2 = L.Convolution2D(
                n_channels, n_channels, 3, 1, 1, initialW=w)
            if n_convs == 3:
                self.conv3 = L.Convolution2D(
                    n_channels, n_channels, 3, 1, 1, initialW=w)
            if fc:
                self.fc4 = L.Linear(None, 4096, initialW=w)
                self.fc5 = L.Linear(4096, 4096, initialW=w)
                self.fc6 = L.Linear(4096, 1000, initialW=w)

        self.n_convs = n_convs
        self.fc = fc

    def __call__(self, x):
        h = F.relu(self.conv1(x))
        h = F.relu(self.conv2(h))
        if self.n_convs == 3:
            h = F.relu(self.conv3(h))
        h = F.max_pooling_2d(h, 2, 2)
        if self.fc:
            h = F.dropout(F.relu(self.fc4(h)))
            h = F.dropout(F.relu(self.fc5(h)))
            h = self.fc6(h)
        return h

```

That's it. VGG16 is a model which won the 1st place in [classification + localization task](#) at ILSVRC 2014, and since then, has become one of the standard models for many different tasks as a pre-trained model. This has 16-layers, so it's called "VGG-16", but we can write this model without writing all layers independently. Since this model consists of several building blocks that have the same architecture, we can build the whole network by re-using the building block definition. Each part of the network is consisted of 2 or 3 convolutional layers and activation function ([relu\(\)](#)) following them, and [max_pooling_2d\(\)](#) operations. This block is written as `VGGBlock` in the above example code. And the whole network just calls this block one by one in sequential manner.

4.3.3 ResNet152

How about ResNet? ResNet [\[He16\]](#) came in the following year's ILSVRC. It is a much deeper model than VGG16, having up to 152 layers. This sounds super laborious to build, but it can be implemented in almost same manner as VGG16. In the other words, it's easy. One possible way to write ResNet-152 is:

```

class ResNet152(chainer.Chain):
    def __init__(self, n_blocks=[3, 8, 36, 3]):
        w = chainer.initializers.HeNormal()
        super(ResNet152, self).__init__()

```

(continues on next page)

(continued from previous page)

```

    with self.init_scope():
        self.conv1 = L.Convolution2D(None, 64, 7, 2, 3, initialW=w, nobias=True)
        self.bn1 = L.BatchNormalization(64)
        self.res2 = ResBlock(n_blocks[0], 64, 64, 256, 1)
        self.res3 = ResBlock(n_blocks[1], 256, 128, 512)
        self.res4 = ResBlock(n_blocks[2], 512, 256, 1024)
        self.res5 = ResBlock(n_blocks[3], 1024, 512, 2048)
        self.fc6 = L.Linear(2048, 1000)

    def __call__(self, x):
        h = self.bn1(self.conv1(x))
        h = F.max_pooling_2d(F.relu(h), 2, 2)
        h = self.res2(h)
        h = self.res3(h)
        h = self.res4(h)
        h = self.res5(h)
        h = F.average_pooling_2d(h, h.shape[2:], stride=1)
        h = self.fc6(h)
        if chainer.config.train:
            return h
        return F.softmax(h)

class ResBlock(chainer.ChainList):
    def __init__(self, n_layers, n_in, n_mid, n_out, stride=2):
        super(ResBlock, self).__init__()
        self.add_link(BottleNeck(n_in, n_mid, n_out, stride, True))
        for _ in range(n_layers - 1):
            self.add_link(BottleNeck(n_out, n_mid, n_out))

    def __call__(self, x):
        for f in self.children():
            x = f(x)
        return x

class BottleNeck(chainer.Chain):
    def __init__(self, n_in, n_mid, n_out, stride=1, proj=False):
        w = chainer.initializers.HeNormal()
        super(BottleNeck, self).__init__()
        with self.init_scope():
            self.conv1x1a = L.Convolution2D(
                n_in, n_mid, 1, stride, 0, initialW=w, nobias=True)
            self.conv3x3b = L.Convolution2D(
                n_mid, n_mid, 3, 1, 1, initialW=w, nobias=True)
            self.conv1x1c = L.Convolution2D(
                n_mid, n_out, 1, 1, 0, initialW=w, nobias=True)
            self.bn_a = L.BatchNormalization(n_mid)
            self.bn_b = L.BatchNormalization(n_mid)
            self.bn_c = L.BatchNormalization(n_out)
            if proj:
                self.conv1x1r = L.Convolution2D(
                    n_in, n_out, 1, stride, 0, initialW=w, nobias=True)
                self.bn_r = L.BatchNormalization(n_out)
        self.proj = proj

    def __call__(self, x):

```

(continues on next page)

(continued from previous page)

```

h = F.relu(self.bn_a(self.conv1x1a(x)))
h = F.relu(self.bn_b(self.conv3x3b(h)))
h = self.bn_c(self.conv1x1c(h))
if self.proj:
    x = self.bn_r(self.conv1x1r(x))
return F.relu(h + x)

```

In the `BottleNeck` class, depending on the value of the `proj` argument supplied to the initializer, it will conditionally compute a convolutional layer `conv1x1r` which will extend the number of channels of the input `x` to be equal to the number of channels of the output of `conv1x1c`, and followed by a batch normalization layer before the final ReLU layer. Writing the building block in this way improves the re-usability of a class. It switches not only the behavior in `__class__()` by flags but also the parameter registration. In this case, when `proj` is `False`, the `BottleNeck` doesn't have `conv1x1r` and `bn_r` layers, so the memory usage would be efficient compared to the case when it registers both anyway and just ignore them if `proj` is `False`.

Using nested `Chains` and `ChainList` for sequential part enables us to write complex and very deep models easily.

4.3.4 Use Pre-trained Models

Various ways to write your models were described above. It turns out that VGG16 and ResNet are very useful as general feature extractors for many kinds of tasks, including but not limited to image classification. So, Chainer provides you with the pre-trained VGG16 and ResNet-50/101/152 models with a simple API. You can use these models as follows:

```

from chainer.links import VGG16Layers

model = VGG16Layers()

```

When `VGG16Layers` is instantiated, the pre-trained parameters are automatically downloaded from the author's server. So you can immediately start to use VGG16 with pre-trained weight as a good image feature extractor. See the details of this model here: [chainer.links.VGG16Layers](#).

In the case of ResNet models, there are three variations differing in the number of layers. We have `chainer.links.ResNet50Layers`, `chainer.links.ResNet101Layers`, and `chainer.links.ResNet152Layers` models with easy parameter loading feature. ResNet's pre-trained parameters are not available for direct downloading, so you need to download the weight from the author's web page first, and then place it into the dir `$CHAINER_DATASET_ROOT/pfnet/chainer/models` or your favorite place. Once the preparation is finished, the usage is the same as VGG16:

```

from chainer.links import ResNet152Layers

model = ResNet152Layers()

```

```

Traceback (most recent call last):
OSError: The pre-trained caffemodel does not exist. Please download it from 'https://
↳github.com/KaimingHe/deep-residual-networks', and place it on ...

```

Please see the details of usage and how to prepare the pre-trained weights for ResNet here: [chainer.links.ResNet50Layers](#)

References

4.4 Recurrent Nets and their Computational Graph

In the example code of this tutorial, we assume for simplicity that the following symbols are already imported.

```
import numpy as np
import chainer
from chainer.backends import cuda
from chainer import Function, gradient_check, report, training, utils, Variable
from chainer import datasets, iterators, optimizers, serializers
from chainer import Link, Chain, ChainList
import chainer.functions as F
import chainer.links as L
from chainer.training import extensions
```

In this section, you will learn how to write

- recurrent nets with full backprop,
- recurrent nets with truncated backprop,
- evaluation of networks with few memory.

After reading this section, you will be able to:

- Handle input sequences of variable length
- Truncate upper stream of the network during forward computation
- Use no-backprop mode to prevent network construction

4.4.1 Recurrent Nets

Recurrent nets are neural networks with loops. They are often used to learn from sequential input/output. Given an input stream $x_1, x_2, \dots, x_t, \dots$ and the initial state h_0 , a recurrent net iteratively updates its state by $h_t = f(x_t, h_{t-1})$, and at some or every point in time t , it outputs $y_t = g(h_t)$. If we expand the procedure along the time axis, it looks like a regular feed-forward network except that same parameters are repeatedly used within the network.

Here we learn how to write a simple one-layer recurrent net. The task is language modeling: given a finite sequence of words, we want to predict the next word at each position without peeking the successive words. Suppose there are 1,000 different word types, and that we use 100 dimensional real vectors to represent each word (a.k.a. word embedding).

Let's start from defining the recurrent neural net language model (RNNLM) as a chain. We can use the `chainer.links.LSTM` link that implements a fully-connected stateful LSTM layer. This link looks like an ordinary fully-connected layer. On construction, you pass the input and output size to the constructor:

```
>>> l = L.LSTM(100, 50)
```

Then, call on this instance `l(x)` executes *one step of LSTM layer*:

```
>>> l.reset_state()
>>> x = Variable(np.random.randn(10, 100).astype(np.float32))
>>> y = l(x)
```

Do not forget to reset the internal state of the LSTM layer before the forward computation! Every recurrent layer holds its internal state (i.e. the output of the previous call). At the first application of the recurrent layer, you must reset the internal state. Then, the next input can be directly fed to the LSTM instance:

```
>>> x2 = Variable(np.random.randn(10, 100).astype(np.float32))
>>> y2 = l(x2)
```

Based on this LSTM link, let's write our recurrent network as a new chain:

```
class RNN(Chain):
    def __init__(self):
        super(RNN, self).__init__()
        with self.init_scope():
            self.embed = L.EmbedID(1000, 100) # word embedding
            self.mid = L.LSTM(100, 50) # the first LSTM layer
            self.out = L.Linear(50, 1000) # the feed-forward output layer

    def reset_state(self):
        self.mid.reset_state()

    def __call__(self, cur_word):
        # Given the current word ID, predict the next word.
        x = self.embed(cur_word)
        h = self.mid(x)
        y = self.out(h)
        return y

rnn = RNN()
model = L.Classifier(rnn)
optimizer = optimizers.SGD()
optimizer.setup(model)
```

Here *EmbedID* is a link for word embedding. It converts input integers into corresponding fixed-dimensional embedding vectors. The last linear link *out* represents the feed-forward output layer.

The RNN chain implements a *one-step-forward computation*. It does not handle sequences by itself, but we can use it to process sequences by just feeding items in a sequence straight to the chain.

Suppose we have a list of word variables *x_list*. Then, we can compute loss values for the word sequence by simple for loop.

```
def compute_loss(x_list):
    loss = 0
    for cur_word, next_word in zip(x_list, x_list[1:]):
        loss += model(cur_word, next_word)
    return loss
```

Of course, the accumulated loss is a Variable object with the full history of computation. So we can just call its *backward()* method to compute gradients of the total loss according to the model parameters:

```
# Suppose we have a list of word variables x_list.
rnn.reset_state()
model.cleargrads()
loss = compute_loss(x_list)
loss.backward()
optimizer.update()
```

Or equivalently we can use the *compute_loss* as a loss function:

```
rnn.reset_state()
optimizer.update(compute_loss, x_list)
```

4.4.2 Truncate the Graph by Unchaining

Learning from very long sequences is also a typical use case of recurrent nets. Suppose the input and state sequence is too long to fit into memory. In such cases, we often truncate the backpropagation into a short time range. This technique is called *truncated backprop*. It is heuristic, and it makes the gradients biased. However, this technique works well in practice if the time range is long enough.

How to implement truncated backprop in Chainer? Chainer has a smart mechanism to achieve truncation, called **backward unchaining**. It is implemented in the `Variable.unchain_backward()` method. Backward unchaining starts from the Variable object, and it chops the computation history backwards from the variable. The chopped variables are disposed automatically (if they are not referenced explicitly from any other user object). As a result, they are no longer a part of computation history, and are not involved in backprop anymore.

Let's write an example of truncated backprop. Here we use the same network as the one used in the previous subsection. Suppose we are given a very long sequence, and we want to run backprop truncated at every 30 time steps. We can write truncated backprop using the model defined above:

```
loss = 0
count = 0
seqlen = len(x_list[1:])

rnn.reset_state()
for cur_word, next_word in zip(x_list, x_list[1:]):
    loss += model(cur_word, next_word)
    count += 1
    if count % 30 == 0 or count == seqlen:
        model.cleargrads()
        loss.backward()
        loss.unchain_backward()
        optimizer.update()
```

State is updated at `model()`, and the losses are accumulated to `loss` variable. At each 30 steps, backprop takes place at the accumulated loss. Then, the `unchain_backward()` method is called, which deletes the computation history backward from the accumulated loss. Note that the last state of `model` is not lost, since the RNN instance holds a reference to it.

The implementation of truncated backprop is simple, and since there is no complicated trick on it, we can generalize this method to different situations. For example, we can easily extend the above code to use different schedules between backprop timing and truncation length.

4.4.3 Network Evaluation without Storing the Computation History

On evaluation of recurrent nets, there is typically no need to store the computation history. While unchaining enables us to walk through unlimited length of sequences with limited memory, it is a bit of a work-around.

As an alternative, Chainer provides an evaluation mode of forward computation which does not store the computation history. This is enabled by just calling `no_backprop_mode()` context:

```
with chainer.no_backprop_mode():
    x_list = [Variable(...) for _ in range(100)] # list of 100 words
    loss = compute_loss(x_list)
```

Note that we cannot call `loss.backward()` to compute the gradient here, since the variable created in the no-backprop context does not remember the computation history.

No-backprop context is also useful to evaluate feed-forward networks to reduce the memory footprint.

We can combine a fixed feature extractor network and a trainable predictor network using `no_backprop_mode()`. For example, suppose we want to train a feed-forward network `predictor_func`, which is located on top of another fixed pre-trained network `fixed_func`. We want to train `predictor_func` without storing the computation history for `fixed_func`. This is simply done by following code snippets (suppose `x_data` and `y_data` indicate input data and label, respectively):

```
with chainer.no_backprop_mode():
    x = Variable(x_data)
    feat = fixed_func(x)
y = predictor_func(feat)
y.backward()
```

At first, the input variable `x` is in no-backprop mode, so `fixed_func` does not memorize the computation history. Then `predictor_func` is executed in backprop mode, i.e., with memorizing the history of computation. Since the history of computation is only memorized between variables `feat` and `y`, the backward computation stops at the `feat` variable.

4.4.4 Making it with Trainer

The above codes are written with plain Function/Variable APIs. When we write a training loop, it is better to use `Trainer`, since we can then easily add functionalities by extensions.

Before implementing it on `Trainer`, let's clarify the training settings. We here use Penn Tree Bank dataset as a set of sentences. Each sentence is represented as a word sequence. We concatenate all sentences into one long word sequence, in which each sentence is separated by a special word `<eos>`, which stands for "End of Sequence". This dataset is easily obtained by `chainer.datasets.get_ptb_words()`. This function returns train, validation, and test dataset, each of which is represented as a long array of integers. Each integer represents a word ID.

Our task is to learn a recurrent neural net language model from the long word sequence. We use words in different locations to form mini-batches. It means we maintain B indices pointing to different locations in the sequence, read from these indices at each iteration, and increment all indices after the read. Of course, when one index reaches the end of the whole sequence, we turn the index back to 0.

In order to implement this training procedure, we have to customize the following components of `Trainer`:

- Iterator. Built-in iterators do not support reading from different locations and aggregating them into a mini-batch.
- Update function. The default update function does not support truncated BPTT.

When we write a dataset iterator dedicated to the dataset, the dataset implementation can be arbitrary; even the interface is not fixed. On the other hand, the iterator must support the `Iterator` interface. The important methods and attributes to implement are `batch_size`, `epoch`, `epoch_detail`, `is_new_epoch`, `iteration`, `__next__`, and `serialize`. Following is a code from the official example in the `examples/ptb` directory.

```
from __future__ import division

class ParallelSequentialIterator(chainer.dataset.Iterator):
    def __init__(self, dataset, batch_size, repeat=True):
        self.dataset = dataset
        self.batch_size = batch_size
        self.epoch = 0
        self.is_new_epoch = False
```

(continues on next page)

(continued from previous page)

```

self.repeat = repeat
self.offsets = [i * len(dataset) // batch_size for i in range(batch_size)]
self.iteration = 0

def __next__(self):
    length = len(self.dataset)
    if not self.repeat and self.iteration * self.batch_size >= length:
        raise StopIteration
    cur_words = self.get_words()
    self.iteration += 1
    next_words = self.get_words()

    epoch = self.iteration * self.batch_size // length
    self.is_new_epoch = self.epoch < epoch
    if self.is_new_epoch:
        self.epoch = epoch

    return list(zip(cur_words, next_words))

@property
def epoch_detail(self):
    return self.iteration * self.batch_size / len(self.dataset)

def get_words(self):
    return [self.dataset[(offset + self.iteration) % len(self.dataset)]
            for offset in self.offsets]

def serialize(self, serializer):
    self.iteration = serializer('iteration', self.iteration)
    self.epoch = serializer('epoch', self.epoch)

train_iter = ParallelSequentialIterator(train, 20)
val_iter = ParallelSequentialIterator(val, 1, repeat=False)

```

Although the code is slightly long, the idea is simple. First, this iterator creates `offsets` pointing to positions equally spaced within the whole sequence. The *i*-th examples of mini-batches refer the sequence with the *i*-th offset. The iterator returns a list of tuples of the current words and the next words. Each mini-batch is converted to a tuple of integer arrays by the `concat_examples` function in the standard updater (see the previous tutorial).

Backprop Through Time is implemented as follows.

```

class BPPTUpdater(training.updaters.StandardUpdater):

    def __init__(self, train_iter, optimizer, bprop_len):
        super(BPPTUpdater, self).__init__(train_iter, optimizer)
        self.bprop_len = bprop_len

    # The core part of the update routine can be customized by overriding.
    def update_core(self):
        loss = 0
        # When we pass one iterator and optimizer to StandardUpdater.__init__,
        # they are automatically named 'main'.
        train_iter = self.get_iterator('main')
        optimizer = self.get_optimizer('main')

        # Progress the dataset iterator for bprop_len words at each iteration.
        for i in range(self.bprop_len):

```

(continues on next page)

(continued from previous page)

```

# Get the next batch (a list of tuples of two word IDs)
batch = train_iter.__next__()

# Concatenate the word IDs to matrices and send them to the device
# self.converter does this job
# (it is chainer.dataset.concat_examples by default)
x, t = self.converter(batch)

# Compute the loss at this time step and accumulate it
loss += optimizer.target(chainer.Variable(x), chainer.Variable(t))

optimizer.target.cleargrads() # Clear the parameter gradients
loss.backward() # Backprop
loss.unchain_backward() # Truncate the graph
optimizer.update() # Update the parameters

updater = BPTTUpdater(train_iter, optimizer, bprop_len) # instantiation

```

In this case, we update the parameters on every `bprop_len` consecutive words. The call of `unchain_backward` cuts the history of computation accumulated to the LSTM links. The rest of the code for setting up `Trainer` is almost same as one given in the previous tutorial.

In this section we have demonstrated how to write recurrent nets in Chainer and some fundamental techniques to manage the history of computation (a.k.a. computational graph). The example in the [examples/ptb](#) directory implements truncated backprop learning of a LSTM language model from the Penn Treebank corpus. In the next section, we will review how to use GPU(s) in Chainer.

4.5 RNN Language Models

4.5.1 0. Introduction

The **language model** is modeling the probability of generating natural language sentences or documents. You can use the language model to estimate how natural a sentence or a document is. Also, with the language model, you can generate new sentences or documents.

Let's start with modeling the probability of generating sentences. We represent a sentence as $\mathbf{X} = (\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_T)$, in which \mathbf{x}_t is a one-hot vector. Generally, \mathbf{x}_0 is the one-hot vector of **BOS** (beginning of sentence), and \mathbf{x}_T is that of **EOS** (end of sentence).

A language model models the probability of a word occurrence under the condition of its previous words in a sentence. Let $\mathbf{X}_{[i,j]}$ be $(\mathbf{x}_i, \mathbf{x}_{i+1}, \dots, \mathbf{x}_j)$, the occurrence probability of sentence \mathbf{X} can be represented as follows:

$$P(\mathbf{X}) = P(\mathbf{x}_0) \prod_{t=1}^T P(\mathbf{x}_t | \mathbf{X}_{[0,t-1]})$$

So, the language model $P(\mathbf{X})$ can be decomposed into word probabilities conditioned with its previous words. In this tutorial, we model $P(\mathbf{x}_t | \mathbf{X}_{[0,t-1]})$ with a recurrent neural network to obtain a language model $P(\mathbf{X})$.

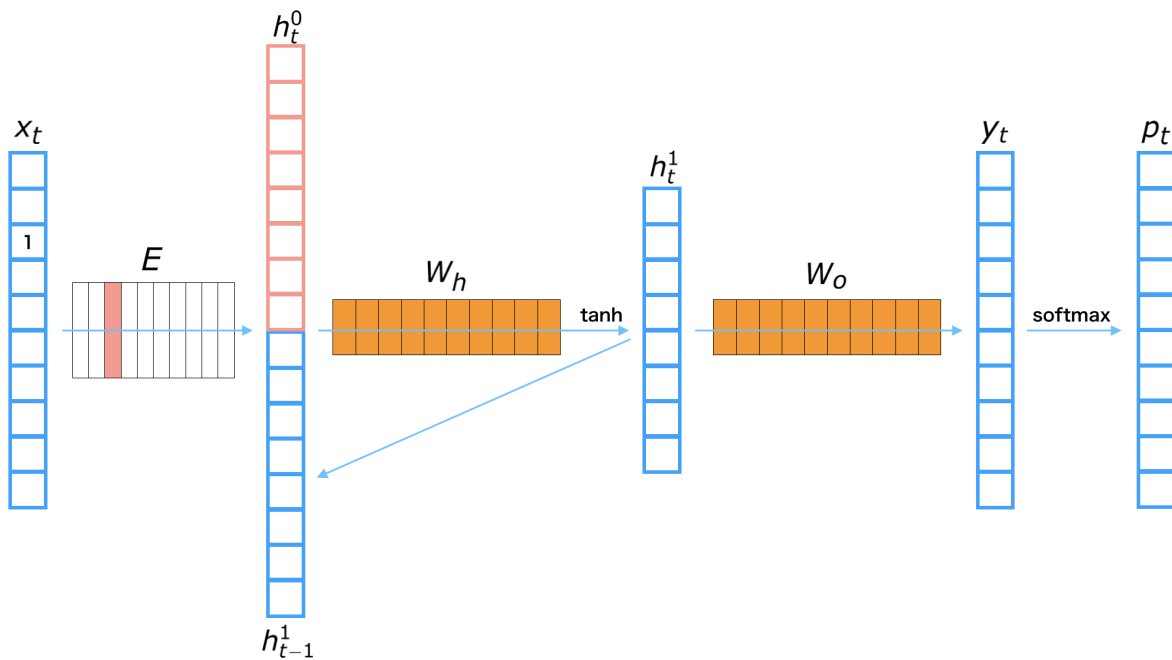
4.5.2 1. Basic Idea of Recurrent Neural Net Language Model

1.1 Recurrent Neural Net Language Model

Recurrent Neural Net Language Model (RNNLM) is a type of neural net language models which contains the RNNs in the network. Since an RNN can deal with the variable length inputs, it is suitable for modeling the sequential data such as sentences in natural language.

We show one layer of an RNNLM with these parameters.

Symbol	Definition
\mathbf{x}_t	the one-hot vector of t -th word
\mathbf{y}_t	the t -th output
$\mathbf{h}_t^{(i)}$	the t -th hidden layer of i -th layer
\mathbf{p}_t	the next word's probability of t -th word
\mathbf{E}	Embedding matrix
\mathbf{W}_h	Hidden layer matrix
\mathbf{W}_o	Output layer matrix



The process to get a next word prediction from i -th input word \mathbf{x}_t

1. Get the embedding vector: $\mathbf{h}_t^{(0)} = \mathbf{E}\mathbf{x}_t$
2. Calculate the hidden layer: $\mathbf{h}_t^{(1)} = \tanh\left(\mathbf{W}_h \begin{bmatrix} \mathbf{h}_t^{(0)} \\ \mathbf{h}_{t-1}^{(1)} \end{bmatrix}\right)$
3. Calculate the output layer: $\mathbf{y}_t = \mathbf{W}_o\mathbf{h}_t^{(1)}$
4. Transform to probability: $\mathbf{p}_t = \text{softmax}(\mathbf{y}_t)$

Note:

- Note that \tanh in the above equation is applied to the input vector in element-wise manner.
- Note that $\begin{bmatrix} \mathbf{a} \\ \mathbf{b} \end{bmatrix}$ denotes a concatenated vector of \mathbf{a} and \mathbf{b} .
- Note that softmax in the above equation converts an arbitrary real vector to a probability vector which the summation over all elements is 1.

1.2 Perplexity (Evaluation of the language model)

Perplexity is the common evaluation metric for a language model. Generally, it measures how well the proposed probability model $P_{\text{model}}(\mathbf{X})$ represents the target data $P^*(\mathbf{X})$. Let a validation dataset be $D = \{\mathbf{X}^{(n)}\}_{n=1}^{|D|}$, which is a set of sentences, where the n -th sentence length is $T^{(n)}$, and the vocabulary size of this dataset is $|\mathcal{V}|$, the perplexity is represented as follows:

$$b^z \text{ s.t. } z = -\frac{1}{|\mathcal{V}|} \sum_{n=1}^{|D|} \sum_{t=1}^{T^{(n)}} \log_b P_{\text{model}}(\mathbf{x}_t^{(n)}, \mathbf{x}_{[a,t-1]}^{(n)})$$

We usually use $b = 2$ or $b = e$. The perplexity shows how much varied the predicted distribution for the next word is. When a language model represents the dataset well, it should show a high probability only for the correct next word, so that the entropy should be high. In the above equation, the sign is reversed, so that smaller perplexity means better model.

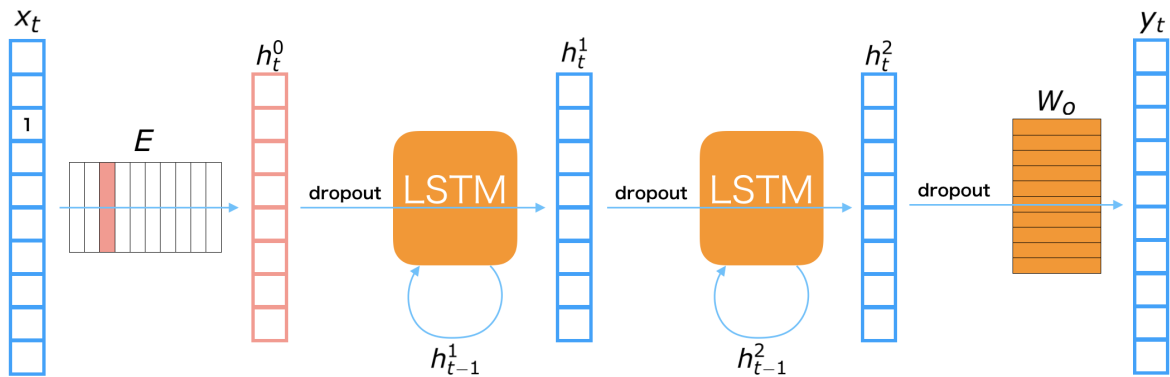
During training, we minimize the below cross entropy:

$$\mathcal{H}(\hat{P}, P_{\text{model}}) = -\hat{P}(\mathbf{X}) \log P_{\text{model}}(\mathbf{X})$$

where \hat{P} is the empirical distribution of a sequence in the training dataset.

4.5.3 2. Implementation of Recurrent Neural Net Language Model

There is an example of RNN language model in the official repository, so we will explain how to implement a RNNLM in Chainer based on that: [examples/ptb](#)

2.1 Model Overview

The RNNLM used in this notebook is depicted in the above figure. The symbols appeared in the figure are defined as follows:

Symbol	Definition
\mathbf{x}_t	the one-hot vector of t -th word
\mathbf{y}_t	the t -th output
$\mathbf{h}_t^{(i)}$	the t -th hidden layer of i -th layer
\mathbf{p}_t	the next word's probability of t -th word
\mathbf{E}	Embedding matrix
\mathbf{W}_h	Hidden layer matrix
\mathbf{W}_o	Output layer matrix

LSTMs (long short-term memory) are used for the connection of hidden layers. A LSTM is one of major recurrent neural net modules. It is designed for remembering the long-term memory, so that it should be able to consider relationships of distant words, such that a word at beginning of sentence and it at the end. We also use **Dropout** before both LSTMs and linear transformations. Dropout is one of regularization techniques for preventing overfitting on training dataset.

2.2 Step-by-step Implementation

2.2.1 Import Package

First, let's import necessary packages.

Listing 1: train_ptb.py

```
import numpy as np

import chainer
import chainer.functions as F
import chainer.links as L
from chainer import training
from chainer.training import extensions
```

2.2.2 Define Training Settings

Define all training settings here.

Listing 2: train_ptb.py

```
parser.add_argument('--batchsize', '-b', type=int, default=20,
                    help='Number of examples in each mini-batch')
parser.add_argument('--bpropflen', '-l', type=int, default=35,
                    help='Number of words in each mini-batch '
                         '(= length of truncated BPTT)')
parser.add_argument('--epoch', '-e', type=int, default=39,
                    help='Number of sweeps over the dataset to train')
parser.add_argument('--gpu', '-g', type=int, default=-1,
                    help='GPU ID (negative value indicates CPU)')
parser.add_argument('--gradclip', '-c', type=float, default=5,
                    help='Gradient norm threshold to clip')
parser.add_argument('--out', '-o', default='result',
```

(continues on next page)

(continued from previous page)

```

        help='Directory to output the result')
parser.add_argument('--resume', '-r', default='',
                    help='Resume the training from snapshot')
parser.add_argument('--test', action='store_true',
                    help='Use tiny datasets for quick tests')
parser.set_defaults(test=False)
parser.add_argument('--unit', '-u', type=int, default=650,
                    help='Number of LSTM units in each layer')
parser.add_argument('--model', '-m', default='model.npz',
                    help='Model file name to serialize')

```

2.2.3 Define Network Structure

An RNNLM written in Chainer is shown below. It implements the model depicted in the above figure.

Listing 3: train_ptb.py

```

class RNNForLM(chainer.Chain):

    def __init__(self, n_vocab, n_units):
        super(RNNForLM, self).__init__()
        with self.init_scope():
            self.embed = L.EmbedID(n_vocab, n_units)
            self.l1 = L.LSTM(n_units, n_units)
            self.l2 = L.LSTM(n_units, n_units)
            self.l3 = L.Linear(n_units, n_vocab)

        for param in self.params():
            param.data[...] = np.random.uniform(-0.1, 0.1, param.data.shape)

    def reset_state(self):
        self.l1.reset_state()
        self.l2.reset_state()

    def __call__(self, x):
        h0 = self.embed(x)
        h1 = self.l1(F.dropout(h0))
        h2 = self.l2(F.dropout(h1))
        y = self.l3(F.dropout(h2))
        return y

```

- When we instantiate this class for making a model, we give the vocabulary size to `n_vocab` and the size of hidden vectors to `n_units`.
- This network uses `chainer.links.LSTM`, `chainer.links.Linear`, and `chainer.functions.dropout` as its building blocks. All the layers are registered and initialized in the context with `self.init_scope()`.
- You can access all the parameters in those layers by calling `self.params()`.
- In the constructor, it initializes all parameters with values sampled from a uniform distribution $U(-1, 1)$.
- The `__call__` method takes an word ID `x`, and calculates the word probability vector for the next word by forwarding it through the network, and returns the output.
- Note that the word ID `x` is automatically converted to a $|\mathcal{V}|$ -dimensional one-hot vector and then multiplied with the input embedding matrix in `self.embed(x)` to obtain an embed vector `h0` at the first line of `__call__`.

2.2.4 Load the Penn Tree Bank Long Word Sequence Dataset

In this notebook, we use Penn Tree Bank dataset that contains number of sentences. Chainer provides an utility function to obtain this dataset from server and convert it to a long single sequence of word IDs. `chainer.datasets.get_ptb_words()` actually returns three separated datasets which are for train, validation, and test.

Let's download and make dataset objects using it:

Listing 4: train_ptb.py

```
# Load the Penn Tree Bank long word sequence dataset
train, val, test = chainer.datasets.get_ptb_words()
```

2.2.5 Define Iterator for Making a Mini-batch from the Dataset

Dataset iterator creates a mini-batch of couple of words at different positions, namely, pairs of current word and its next word. Each example is a part of sentences starting from different offsets equally spaced within the whole sequence.

Listing 5: train_ptb.py

```
class ParallelSequentialIterator(chainer.dataset.Iterator):

    def __init__(self, dataset, batch_size, repeat=True):
        self.dataset = dataset
        self.batch_size = batch_size # batch size
        # Number of completed sweeps over the dataset. In this case, it is
        # incremented if every word is visited at least once after the last
        # increment.
        self.epoch = 0
        # True if the epoch is incremented at the last iteration.
        self.is_new_epoch = False
        self.repeat = repeat
        length = len(dataset)
        # Offsets maintain the position of each sequence in the mini-batch.
        self.offsets = [i * length // batch_size for i in range(batch_size)]
        # NOTE: this is not a count of parameter updates. It is just a count of
        # calls of ``__next__``.
        self.iteration = 0
        # use -1 instead of None internally
        self._previous_epoch_detail = -1.

    def __next__(self):
        # This iterator returns a list representing a mini-batch. Each item
        # indicates a different position in the original sequence. Each item is
        # represented by a pair of two word IDs. The first word is at the
        # "current" position, while the second word at the next position.
        # At each iteration, the iteration count is incremented, which pushes
        # forward the "current" position.
        length = len(self.dataset)
        if not self.repeat and self.iteration * self.batch_size >= length:
            # If not self.repeat, this iterator stops at the end of the first
            # epoch (i.e., when all words are visited once).
            raise StopIteration
        cur_words = self.get_words()
        self._previous_epoch_detail = self.epoch_detail
```

(continues on next page)

(continued from previous page)

```

self.iteration += 1
next_words = self.get_words()

epoch = self.iteration * self.batch_size // length
self.is_new_epoch = self.epoch < epoch
if self.is_new_epoch:
    self.epoch = epoch

return list(zip(cur_words, next_words))

@property
def epoch_detail(self):
    # Floating point version of epoch.
    return self.iteration * self.batch_size / len(self.dataset)

@property
def previous_epoch_detail(self):
    if self._previous_epoch_detail < 0:
        return None
    return self._previous_epoch_detail

def get_words(self):
    # It returns a list of current words.
    return [self.dataset[(offset + self.iteration) % len(self.dataset)]
            for offset in self.offsets]

def serialize(self, serializer):
    # It is important to serialize the state to be recovered on resume.
    self.iteration = serializer('iteration', self.iteration)
    self.epoch = serializer('epoch', self.epoch)
    try:
        self._previous_epoch_detail = serializer(
            'previous_epoch_detail', self._previous_epoch_detail)
    except KeyError:
        # guess previous_epoch_detail for older version
        self._previous_epoch_detail = self.epoch + \
            (self.current_position - self.batch_size) / len(self.dataset)
    if self.epoch_detail > 0:
        self._previous_epoch_detail = max(
            self._previous_epoch_detail, 0.)
    else:
        self._previous_epoch_detail = -1.

```

2.2.6 Define Updater

We use Backpropagation through time (BPTT) for optimize the RNNLM. BPTT can be implemented by overriding `update_core()` method of `StandardUpdater`. First, in the constructor of the `BPTTUpdater`, it takes `bprop_len` as an argument in addition to other arguments `StandardUpdater` needs. `bprop_len` defines the length of sequence T to calculate the loss:

$$\mathcal{L} = - \sum_{t=0}^T \sum_{n=1}^{|\mathcal{V}|} \hat{P}(\mathbf{x}_{t+1}^{(n)}) \log P_{\text{model}}(\mathbf{x}_{t+1}^{(n)} | \mathbf{x}_t^{(n)})$$

where $\hat{P}(\mathbf{x}_t^n)$ is a probability for n -th word in the vocabulary at the position t in the training data sequence.

Listing 6: train_ptb.py

```
class BPTTUpdater(training.updaters.StandardUpdater):

    def __init__(self, train_iter, optimizer, bprop_len, device):
        super(BPTTUpdater, self).__init__(
            train_iter, optimizer, device=device)
        self.bprop_len = bprop_len

    # The core part of the update routine can be customized by overriding.
    def update_core(self):
        loss = 0
        # When we pass one iterator and optimizer to StandardUpdater.__init__,
        # they are automatically named 'main'.
        train_iter = self.get_iterator('main')
        optimizer = self.get_optimizer('main')

        # Progress the dataset iterator for bprop_len words at each iteration.
        for i in range(self.bprop_len):
            # Get the next batch (a list of tuples of two word IDs)
            batch = train_iter.__next__()

            # Concatenate the word IDs to matrices and send them to the device
            # self.converter does this job
            # (it is chainer.dataset.concat_examples by default)
            x, t = self.converter(batch, self.device)

            # Compute the loss at this time step and accumulate it
            loss += optimizer.target(chainer.Variable(x), chainer.Variable(t))

        optimizer.target.cleargrads() # Clear the parameter gradients
        loss.backward() # Backprop
        loss.unchain_backward() # Truncate the graph
        optimizer.update() # Update the parameters
```

2.2.7 Define Evaluation Function (Perplexity)

Define a function to calculate the perplexity from the loss value. If we take e as b in the above definition of perplexity, calculating the perplexity is just to give the loss value to the power of e :

Listing 7: train_ptb.py

```
def compute_perplexity(result):
    result['perplexity'] = np.exp(result['main/loss'])
    if 'validation/main/loss' in result:
        result['val_perplexity'] = np.exp(result['validation/main/loss'])
```

2.2.8 Create Iterator

Here, the code below just creates iterator objects from dataset splits (train/val/test).

Listing 8: train_ptb.py

```
train_iter = ParallelSequentialIterator(train, args.batchsize)
val_iter = ParallelSequentialIterator(val, 1, repeat=False)
test_iter = ParallelSequentialIterator(test, 1, repeat=False)
```

2.2.9 Create RNN and Classification Model

Instantiate RNNLM model and wrap it with `chainer.links.Classifier` because it calculates softmax cross entropy as the loss.

Listing 9: train_ptb.py

```
rnn = RNNForLM(n_vocab, args.unit)
model = L.Classifier(rnn)
model.compute_accuracy = False # we only want the perplexity
```

Note that `Classifier` computes not only the loss but also accuracy based on a given input/label pair. To learn the RNN language model, we only need the loss (cross entropy) in the `Classifier` because we calculate the perplexity instead of classification accuracy to check the performance of the model. So, we turn off computing the accuracy by giving `False` to `model.compute_accuracy` attribute.

2.2.10 Setup Optimizer

Prepare an optimizer. Here, we use `GradientClipping` to prevent gradient explosion. It automatically clips the gradient to be used to update the parameters in the model with given constant `gradclip`.

Listing 10: train_ptb.py

```
optimizer = chainer.optimizers.SGD(lr=1.0)
optimizer.setup(model)
optimizer.add_hook(chainer.optimizer_hooks.GradientClipping(args.gradclip))
```

2.2.11 Setup and Run Trainer

Let's make a trainer object and start the training! Note that we add an `eval_hook` to the `Evaluator` extension to reset the internal states before starting evaluation process. It can prevent to use training data during evaluating the model.

Listing 11: train_ptb.py

```
updater = BPTTUpdater(train_iter, optimizer, args.bprop_len, args.gpu)
trainer = training.Trainer(updater, (args.epoch, 'epoch'), out=args.out)

eval_model = model.copy() # Model with shared params and distinct states
eval_rnn = eval_model.predictor
trainer.extend(extensions.Evaluator(
    val_iter, eval_model, device=args.gpu,
```

(continues on next page)

(continued from previous page)

```

    # Reset the RNN state at the beginning of each evaluation
    eval_hook=lambda _: eval_rnn.reset_state())

interval = 10 if args.test else 500
trainer.extend(extensions.LogReport(postprocess=compute_perplexity,
                                   trigger=(interval, 'iteration')))
trainer.extend(extensions.PrintReport(
    ['epoch', 'iteration', 'perplexity', 'val_perplexity']
), trigger=(interval, 'iteration'))
trainer.extend(extensions.ProgressBar(
    update_interval=1 if args.test else 10))
trainer.extend(extensions.snapshot())
trainer.extend(extensions.snapshot_object(
    model, 'model_iter_{.updater.iteration}'))
if args.resume:
    chainer.serializers.load_npz(args.resume, trainer)

trainer.run()

```

2.2.12 Evaluate the trained model on test dataset

Let's see the perplexity on the test split. *Trainer*'s extension can be used as just a normal function outside of *Trainer*.

Listing 12: train_ptb.py

```

print('test')
eval_rnn.reset_state()
evaluator = extensions.Evaluator(test_iter, eval_model, device=args.gpu)
result = evaluator()
print('test perplexity:', np.exp(float(result['main/loss'])))

```

2.3 Run Example

2.3.1 Training the model

You can train the model with the script: `examples/ptb/train_ptb.py`

```

$ pwd
/root2chainer/chainer/examples/ptb
$ python train_ptb.py --test # run by test mode. If you want to use all data, remove
  ↳ "--test".
Downloading from https://raw.githubusercontent.com/wojzaremba/lstm/master/data/ptb.
  ↳ train.txt...
Downloading from https://raw.githubusercontent.com/wojzaremba/lstm/master/data/ptb.
  ↳ valid.txt...
Downloading from https://raw.githubusercontent.com/wojzaremba/lstm/master/data/ptb.
  ↳ test.txt...
#vocab = 10000
test
test perplexity: 29889.9857364

```


2.3.2 Generating sentences

You can generate the sentence which starts with a word in the vocabulary. In this example, we generate a sentence which starts with the word apple. We use the script in the PTB example of the official repository: [examples/ptb/gentxt.py](#)

```
$ pwd
/root2chainer/chainer/examples/ptb
$ python gentxt.py -m model.npz -p apple
apple a new u.s. economist with <unk> <unk> fixed more than to N the company said who_
→is looking back to
```

4.6 Word2Vec: Obtain word embeddings

4.6.1 0. Introduction

Word2vec is the tool for generating the distributed representation of words, which is proposed by Mikolov et al[1]. When the tool assigns a real-valued vector to each word, the closer the meanings of the words, the greater similarity the vectors will indicate.

Distributed representation means assigning a real-valued vector for each word and representing the word by the vector. When representing a word by distributed representation, we call the **word embeddings**. In this tutorial, we aim at explaining how to get the word embeddings from Penn Tree Bank dataset.

Let's think about what the meaning of word is. Since we are human, we can understand that the words “animal” and “dog” are deeply related each other. But what information will Word2vec use to learn the vectors for words? The words “animal” and “dog” should have similar vectors, but the words “food” and “dog” should be far from each other. How to know the features of those words automatically?

4.6.2 1. Basic Idea

Word2vec learns the similarity of word meanings from simple information. It learns the representation of words from sentences. The core idea is based on the assumption that the meaning of a word is affected by the words around it. This idea follows **distributional hypothesis**[2].

The word we focus on to learn its representation is called **center word**, and the words around it are called **context words**. The window size C determines the number of context words which is considered.

Here, let's see the algorithm by using an example sentence: “**The cute cat jumps over the lazy dog.**”.

- All of the following figures consider “cat” as the center word.
- According to the window size C , you can see that the number of context words is changed.

 : Center Word
 : Context Word

c=0 The cute  jumps over the lazy dog.

c=1 The    over the lazy dog.

c=2      the lazy dog.

4.6.3 2. Main Algorithm

Word2vec, the tool for creating the word embeddings, is actually built with two models, which are called **Skip-gram** and **CBoW**.

To explain the models with the figures below, we will use the following symbols.

Symbol	Definition
$ \mathcal{V} $	The size of vocabulary
D	The size of embedding vector
\mathbf{v}_t	A one-hot center word vector
$V_{t\pm C}$	A set of $2C$ context vectors around \mathbf{v}_t , namely, $\{\mathbf{v}_{t+c}\}_{c=-C}^C \setminus \mathbf{v}_t$
\mathbf{l}_H	An embedding vector of an input word vector
\mathbf{l}_O	An output vector of the network
\mathbf{W}_H	The embedding matrix for inputs
\mathbf{W}_O	The embedding matrix for outputs

Note: Using **negative sampling** or **hierarchical softmax** for the loss function is very common, however, in this tutorial, we will use the **softmax over all words** and skip the other variants for the sake of simplicity.

2.1 Skip-gram

This model learns to predict context words $V_{t\pm C}$ when a center word \mathbf{v}_t is given. In the model, each row of the embedding matrix for input \mathbf{W}_H becomes a word embedding of each word.

When you input a center word \mathbf{v}_t into the network, you can predict one of context words $\hat{\mathbf{v}}_{t+c} \in V_{t\pm C}$ as follows:

1. Calculate an embedding vector of the input center word vector: $\mathbf{l}_H = \mathbf{W}_H \mathbf{v}_t$
2. Calculate an output vector of the embedding vector: $\mathbf{l}_O = \mathbf{W}_O \mathbf{l}_H$
3. Calculate a probability vector of a context word: $\hat{\mathbf{v}}_{t+c} = \text{softmax}(\mathbf{l}_O)$

Each element of the $|\mathcal{V}|$ -dimensional vector $\hat{\mathbf{v}}_{t+c}$ is a probability that a word in the vocabulary turns out to be a context word at position c . So, the probability $p(\mathbf{v}_{t+c}|\mathbf{v}_t)$ can be estimated by a dot product of the one-hot vector \mathbf{v}_{t+c} which represents the actual word at the position c and the output vector $\hat{\mathbf{v}}_{t+c}$.

$$p(\mathbf{v}_{t+c}|\mathbf{v}_t) = \mathbf{v}_{t+c}^T \hat{\mathbf{v}}_{t+c}$$

The loss function to predict all the context words $V_{t\pm C}$ given a center word \mathbf{v}_t is defined as follows:

$$\begin{aligned} L(V_{t\pm C}|\mathbf{v}_t; \mathbf{W}_H, \mathbf{W}_O) &= \sum_{V_{t\pm C}} -\log(p(\mathbf{v}_{t+c} | \mathbf{v}_t)) \\ &= \sum_{V_{t\pm C}} -\log(\mathbf{v}_{t+c}^T \hat{\mathbf{v}}_{t+c}) \end{aligned}$$

2.2 Continuous Bag of Words (CBoW)

This model learns to predict center word \mathbf{v}_t when context words $V_{t\pm C}$ is given. When you give a set of context words $V_{t\pm C}$ to the network, you can estimate the probability of the center word $\hat{\mathbf{v}}_t$ as follows:

1. Calculate a mean embedding vector over all context words: $\mathbf{l}_H = \frac{1}{2C} \sum_{V_{t\pm C}} \mathbf{W}_H \mathbf{v}_{t+c}$
2. Calculate an output vector of the embedding vector: $\mathbf{l}_O = \mathbf{W}_O \mathbf{l}_H$
3. Calculate a probability vector of a center word: $\hat{\mathbf{v}}_t = \text{softmax}(\mathbf{l}_O)$

Each element of the $|\mathcal{V}|$ -dimensional vector $\hat{\mathbf{v}}_t$ is a probability that a word in the vocabulary turns out to be a center word. So, the probability $p(\mathbf{v}_t|V_{t\pm C})$ can be estimated by a dot product of the one-hot vector \mathbf{v}_t which represents the actual center word and the output vector $\hat{\mathbf{v}}_t$.

$$p(\mathbf{v}_t|V_{t\pm C}) = \mathbf{v}_t^T \hat{\mathbf{v}}_t$$

The loss function to predict the center word \mathbf{v}_t given context words $V_{t\pm C}$ is defined as follows:

$$\begin{aligned} L(\mathbf{v}_t|V_{t\pm C}; \mathbf{W}_H, \mathbf{W}_O) &= -\log(p(\mathbf{v}_t | V_{t\pm C})) \\ &= -\log(\mathbf{v}_t^T \hat{\mathbf{v}}_t) \end{aligned}$$

4.6.4 3. Details of Skip-gram

In this tutorial, we mainly explain Skip-gram model because

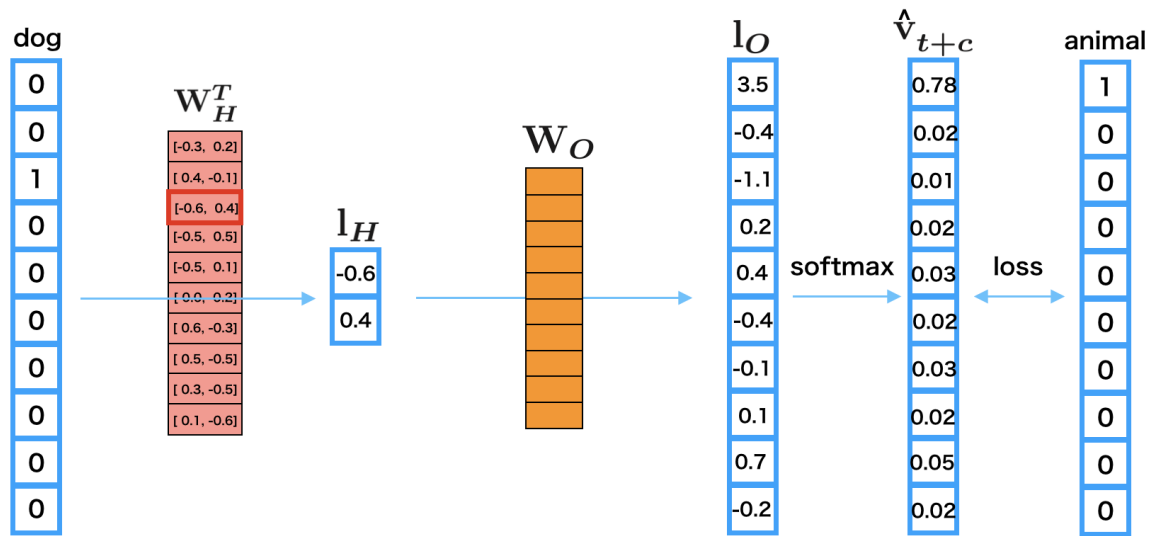
1. It is easier to understand the algorithm than CBoW.
2. Even if the number of words increases, the accuracy is largely maintained. So, it is more scalable.

So, let's think about a concrete example of calculating Skip-gram under this setup:

- The size of vocabulary $|\mathcal{V}|$ is 10.
- The size of embedding vector D is 2.
- Center word is “dog”.
- Context word is “animal”.

Since there should be more than one context word, repeat the following process for each context word.

1. The one-hot vector of “dog” is $[0 \ 0 \ 1 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0]$ and you input it as the center word.
2. The third row of embedding matrix \mathbf{W}_H is used for the word embedding of “dog” \mathbf{l}_H .
3. Then, multiply \mathbf{W}_O with \mathbf{l}_H to obtain the output vector \mathbf{l}_O .
4. Give \mathbf{l}_O to the softmax function to make it a predicted probability vector $\hat{\mathbf{v}}_{t+c}$ for a context word at the position c .
5. Calculate the error between $\hat{\mathbf{v}}_{t+c}$ and the one-hot vector of “animal”; $[1 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0]$.
6. Propagate the error back to the network to update the parameters.



4.6.5 4. Implementation of Skip-gram in Chainer

There is an example of Word2vec in the official repository of Chainer, so we will explain how to implement Skip-gram based on this: [examples/word2vec](#)

4.1 Preparation

First, let's import necessary packages:

Listing 13: train_word2vec.py

```
import argparse
import collections

import numpy as np
import six

import chainer
from chainer.backends import cuda
import chainer.functions as F
import chainer.initializers as I
import chainer.links as L
import chainer.optimizers as O
from chainer import reporter
from chainer import training
from chainer.training import extensions
```

4.2 Define a Skip-gram model

Next, let's define a network for Skip-gram.

Listing 14: train_word2vec.py

```

class SkipGram(chainer.Chain):
    """Definition of Skip-gram Model"""

    def __init__(self, n_vocab, n_units, loss_func):
        super(SkipGram, self).__init__()

        with self.init_scope():
            self.embed = L.EmbedID(
                n_vocab, n_units, initialW=I.Uniform(1. / n_units))
            self.loss_func = loss_func

    def __call__(self, x, contexts):
        e = self.embed(contexts)
        batch_size, n_context, n_units = e.shape
        x = F.broadcast_to(x[:, None], (batch_size, n_context))
        e = F.reshape(e, (batch_size * n_context, n_units))
        x = F.reshape(x, (batch_size * n_context,))
        loss = self.loss_func(e, x)
        reporter.report({'loss': loss}, self)
        return loss

```

Listing 15: train_word2vec.py

```

class SoftmaxCrossEntropyLoss(chainer.Chain):
    """Softmax cross entropy loss function preceded by linear transformation.

    """

    def __init__(self, n_in, n_out):
        super(SoftmaxCrossEntropyLoss, self).__init__()
        with self.init_scope():
            self.out = L.Linear(n_in, n_out, initialW=0)

    def __call__(self, x, t):
        return F.softmax_cross_entropy(self.out(x), t)

```

Note:

- The weight matrix `self.embed.W` is the embedding matrix for input vector `x`.
- The function call `__call__` takes the word ID of a center word `x` and word IDs of context words `contexts` as inputs, and outputs the error calculated by the loss function `loss_func` s.t. `SoftmaxCrossEntropyLoss`.
- Note that the initial shape of `x` and `contexts` are `(batch_size,)` and `(batch_size, n_context)`, respectively.
- The `batch_size` means the size of mini-batch, and `n_context` means the number of context words.

First, we obtain the embedding vectors of contexts by `e = self.embed(contexts)`. Then `F.broadcast_to(x[:, None], (batch_size, n_context))` performs broadcasting of `x` (its shape is `(batch_size,)`) to `(batch_size, n_context)` by copying the same value `n_context` time to fill the second axis, and then the broadcasted `x` is reshaped into 1-D vector `(batchsize * n_context,)` while `e` is reshaped to `(batch_size * n_context, n_units)`. In Skip-gram model, predicting a context word from the center word is the same as predicting the center word from a context word because the center word is always a context

word when considering the context word as a center word. So, we create `batch_size * n_context` center word predictions by applying `self.out` linear layer to the embedding vectors of context words. Then, calculate softmax cross entropy between the broadcasted center word ID `x` and the predictions.

4.3 Prepare dataset and iterator

Let's retrieve the Penn Tree Bank (PTB) dataset by using Chainer's dataset utility `get_ptb_words()` method.

```
train, val, _ = chainer.datasets.get_ptb_words()
counts = collections.Counter(train)
```

Then define an iterator to make mini-batches that contain a set of center words with their context words. `train` and `val` means training data and validation data. Each data contains the list of Document IDs:

```
>>> train
array([ 0,  1,  2, ..., 39, 26, 24], dtype=int32)
>>> val
array([2211, 396, 1129, ..., 108, 27, 24], dtype=int32)
```

Listing 16: `train_word2vec.py`

```
class WindowIterator(chainer.dataset.Iterator):
    """Dataset iterator to create a batch of sequences at different positions.

    This iterator returns a pair of the current words and the context words.
    """

    def __init__(self, dataset, window, batch_size, repeat=True):
        self.dataset = np.array(dataset, np.int32)
        self.window = window # size of context window
        self.batch_size = batch_size
        self._repeat = repeat
        # order is the array which is shuffled `[window, window + 1, ...,
        # len(dataset) - window - 1]`
        self.order = np.random.permutation(
            len(dataset) - window * 2).astype(np.int32)
        self.order += window
        self.current_position = 0
        # Number of completed sweeps over the dataset. In this case, it is
        # incremented if every word is visited at least once after the last
        # increment.
        self.epoch = 0
        # True if the epoch is incremented at the last iteration.
        self.is_new_epoch = False

    def __next__(self):
        """This iterator returns a list representing a mini-batch.

        Each item indicates a different position in the original sequence.
        """
        if not self._repeat and self.epoch > 0:
            raise StopIteration

        i = self.current_position
        i_end = i + self.batch_size
        position = self.order[i:i_end]
```

(continues on next page)

(continued from previous page)

```

w = np.random.randint(self.window - 1) + 1
offset = np.concatenate([np.arange(-w, 0), np.arange(1, w + 1)])
pos = position[:, None] + offset[None, :]
contexts = self.dataset.take(pos)
center = self.dataset.take(position)

if i_end >= len(self.order):
    np.random.shuffle(self.order)
    self.epoch += 1
    self.is_new_epoch = True
    self.current_position = 0
else:
    self.is_new_epoch = False
    self.current_position = i_end

return center, contexts

@property
def epoch_detail(self):
    return self.epoch + float(self.current_position) / len(self.order)

def serialize(self, serializer):
    self.current_position = serializer('current_position',
                                       self.current_position)
    self.epoch = serializer('epoch', self.epoch)
    self.is_new_epoch = serializer('is_new_epoch', self.is_new_epoch)
    if self._order is not None:
        serializer('_order', self._order)

```

- In the constructor, we create an array `self.order` which denotes shuffled indices of `[window, window + 1, ..., len(dataset) - window - 1]` in order to choose a center word randomly from dataset in a mini-batch.
- The iterator definition `__next__` returns `batch_size` sets of center word and context words.
- The code `self.order[i:i_end]` returns the indices for a set of center words from the random-ordered array `self.order`. The center word IDs center at the random indices are retrieved by `self.dataset.take`.
- `np.concatenate([np.arange(-w, 0), np.arange(1, w + 1)])` creates a set of offsets to retrieve context words from the dataset.
- The code `position[:, None] + offset[None, :]` generates the indices of context words for each center word index in `position`. The context word IDs context are retrieved by `self.dataset.take`.

4.4 Prepare model, optimizer, and updater

Listing 17: train_word2vec.py

```
model = SkipGram(n_vocab, args.unit, loss_func)
```

Listing 18: train_word2vec.py

```
optimizer = O.Adam()
optimizer.setup(model)
```

Listing 19: train_word2vec.py

```
train_iter = WindowIterator(train, args.window, args.batchsize)
val_iter = WindowIterator(val, args.window, args.batchsize, repeat=False)

# Set up an updater
updater = training.updaters.StandardUpdater(
    train_iter, optimizer, converter=convert, device=args.gpu)
```

Listing 20: train_word2vec.py

```
trainer = training.Trainer(updater, (args.epoch, 'epoch'), out=args.out)

trainer.extend(extensions.Evaluator(
    val_iter, model, converter=convert, device=args.gpu))
trainer.extend(extensions.LogReport())
trainer.extend(extensions.PrintReport(
    ['epoch', 'main/loss', 'validation/main/loss']))
trainer.extend(extensions.ProgressBar())
trainer.run()
```

4.5 Start training

```
$ pwd
/root2chainer/chainer/examples/word2vec
$ python train_word2vec.py --test # run by test mode. If you want to use all data,
  ↪ remove "--test".
GPU: -1
# unit: 100
Window: 5
Minibatch-size: 1000
# epoch: 20
Training model: skipgram
Output type: hsm

n_vocab: 10000
data length: 100
epoch      main/loss  validation/main/loss
1          4233.75  2495.33
2          1411.14  4990.66
3          4233.11  1247.66
4          2821.66  4990.65
5          4231.94  1247.66
6          5642.04  2495.3
7          5640.82  4990.64
8          5639.31  2495.28
9          2817.89  4990.62
10         1408.03  3742.94
11         5633.11  1247.62
12         4221.71  2495.21
13         4219.3   4990.56
14         4216.57  2495.16
15         4213.52  2495.12
```

(continues on next page)

(continued from previous page)

16	5616.03	1247.55
17	5611.34	3742.78
18	2800.31	3742.74
19	1397.79	2494.95
20	2794.1	3742.66

4.5 Search the similar words

```
$ pwd
/root2chainer/chainer/examples/word2vec
$ python search.py
>> apple
query: apple
compaq: 0.6169619560241699
chip: 0.49579331278800964
retailer: 0.4904134273529053
maker: 0.4684058427810669
computer: 0.4652436673641205
>> animal
query: animal
beauty: 0.5680124759674072
human: 0.5404794216156006
insulin: 0.5365156531333923
cell: 0.5186758041381836
photographs: 0.5077002048492432
```

4.6.6 5. Reference

- [1] Mikolov, Tomas; et al. “Efficient Estimation of Word Representations in Vector Space”. arXiv:1301.3781
- [2] Distributional Hypothesis

4.7 Write a Sequence to Sequence (seq2seq) Model

4.7.1 0. Introduction

The **sequence to sequence (seq2seq) model**^{[1][2]} is a learning model that converts an input sequence into an output sequence. In this context, the **sequence** is a list of symbols, corresponding to the words in a sentence. The seq2seq model has achieved great success in fields such as machine translation, dialogue systems, question answering, and text summarization. All of these tasks can be regarded as the task to learn a model that converts an input sequence into an output sequence.

4.7.2 1. Basic Idea of Seq2seq Model

1.1 Overview of Seq2seq Model

The Notations of Sequence

The seq2seq model converts an input sequence into an output sequence. Let the input sequence and the output sequence be \mathbf{X} and \mathbf{Y} . The i -th element of the input sequence is represented as \mathbf{x}_i , and the j -th element of the output sequence is also represented as \mathbf{y}_j . Generally, each of the \mathbf{x}_i and the \mathbf{y}_j is the one-hot vector of the symbols. For example, in natural language processing(NLP), the one-hot vector represents the word and its size becomes the vocabulary size.

Let's think about the seq2seq model in the context of NLP. Let the vocabulary of the inputs and the outputs be $\mathcal{V}^{(s)}$ and $\mathcal{V}^{(t)}$, all the elements \mathbf{x}_i and \mathbf{y}_j satisfy $\mathbf{x}_i \in \mathbb{R}^{|\mathcal{V}^{(s)}|}$ and $\mathbf{y}_j \in \mathbb{R}^{|\mathcal{V}^{(t)}|}$. The input sequence \mathbf{X} and the output sequence \mathbf{Y} are represented as the following equations:

$$\begin{aligned}\mathbf{X} &= (\mathbf{x}_1, \dots, \mathbf{x}_I) = (\mathbf{x}_i)_{i=1}^I \\ \mathbf{Y} &= (\mathbf{y}_1, \dots, \mathbf{y}_J) = (\mathbf{y}_j)_{j=1}^J\end{aligned}$$

I and J are the length of the input sequence and the output sequence. Using the typical NLP notation, \mathbf{y}_0 is the one-hot vector of *BOS*, which is the virtual word representing the beginning of the sentence, and \mathbf{y}_{J+1} is that of *EOS*, which is the virtual word representing the end of the sentence.

The Notations of Conditional Probability $P(\mathbf{Y}|\mathbf{X})$

Next, let's think about the conditional probability $P(\mathbf{Y}|\mathbf{X})$ generating the output sequence \mathbf{Y} when the input sequence \mathbf{X} is given. The purpose of seq2seq model is modeling the probability $P(\mathbf{Y}|\mathbf{X})$. However, the seq2seq model does not model the probability $P(\mathbf{Y}|\mathbf{X})$ directly. Actually, it models the probability $P(\mathbf{y}_j|\mathbf{Y}_{<j}, \mathbf{X})$, which is the probability of generating the j -th element of the output sequence \mathbf{y}_j given the $\mathbf{Y}_{<j}$ and \mathbf{X} . $\mathbf{Y}_{<j}$ means the output sequence from 1 to $j-1$, or $(\mathbf{y}_j)_{j=1}^{j-1}$. In this notation, you can write the model $P_\theta(\mathbf{Y}|\mathbf{X})$ with the product of $P_\theta(\mathbf{y}_j|\mathbf{Y}_{<j}, \mathbf{X})$:

$$P_\theta(\mathbf{Y}|\mathbf{X}) = \prod_{j=1}^{J+1} P_\theta(\mathbf{y}_j|\mathbf{Y}_{<j}, \mathbf{X})$$

Processing Steps in Seq2seq Model

Now, let's think about the processing steps in seq2seq model. The feature of seq2seq model is that it consists of the two processes:

1. The process that generates the fixed size vector \mathbf{z} from the input sequence \mathbf{X}
2. The process that generates the output sequence \mathbf{Y} from \mathbf{z}

In other words, the information of \mathbf{X} is conveyed by \mathbf{z} , and $P_\theta(\mathbf{y}_j|\mathbf{Y}_{<j}, \mathbf{X})$ is actually calculated by $P_\theta(\mathbf{y}_j|\mathbf{Y}_{<j}, \mathbf{z})$.

First, we represent the process which generating \mathbf{z} from \mathbf{X} by the function Λ :

$$\mathbf{z} = \Lambda(\mathbf{X})$$

The function Λ may be the recurrent neural net such as LSTMs.

Second, we represent the process which generating \mathbf{Y} from \mathbf{z} by the following formula:

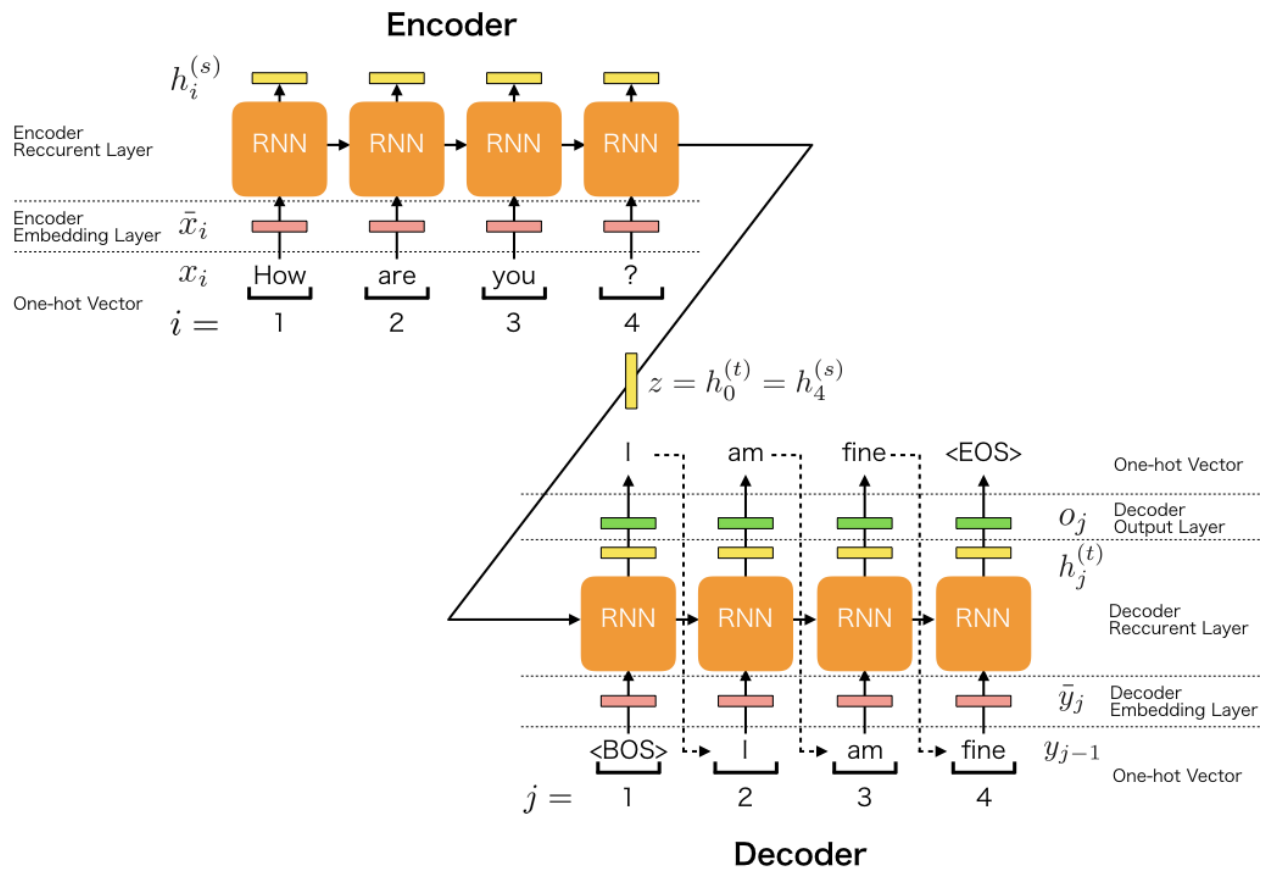
$$\begin{aligned}P_\theta(\mathbf{y}_j|\mathbf{Y}_{<j}, \mathbf{X}) &= \Upsilon(\mathbf{h}_j^{(t)}, \mathbf{y}_j) \\ \mathbf{h}_j^{(t)} &= \Psi(\mathbf{h}_{j-1}^{(t)}, \mathbf{y}_{j-1})\end{aligned}$$

Ψ is the function to generate the hidden vectors $\mathbf{h}_j^{(t)}$, and Υ is the function to calculate the generative probability of the one-hot vector \mathbf{y}_j . When $j = 1$, $\mathbf{h}_{j-1}^{(t)}$ or $\mathbf{h}_0^{(t)}$ is \mathbf{z} generated by $\Lambda(\mathbf{X})$, and \mathbf{y}_{j-1} or \mathbf{y}_0 is the one-hot vector of *BOS*.

1.2 Model Architecture of Seq2seq Model

In this section, we describe the architecture of seq2seq model. To simplify the explanation, we use the most basic architecture. The architecture of seq2seq model can be separated to the five major roles.

1. Encoder Embedding Layer
2. Encoder Recurrent Layer
3. Decoder Embedding Layer
4. Decoder Recurrent Layer
5. Decoder Output Layer



The encoder consists of two layers: the embedding layer and the recurrent layer, and the decoder consists of three layers: the embedding layer, the recurrent layer, and the output layer.

In the explanation, we use the following symbols:

Symbol	Definition
H	the size of the hidden vector
D	the size of the embedding vector
\mathbf{x}_i	the one-hot vector of i -th word in the input sentence
$\bar{\mathbf{x}}_i$	the embedding vector of i -th word in the input sentence
$\mathbf{E}^{(s)}$	Embedding matrix of the encoder
$\mathbf{h}_i^{(s)}$	the i -th hidden vector of the encoder
\mathbf{y}_j	the one-hot vector of j -th word in the output sentence
$\bar{\mathbf{y}}_j$	the embedding vector of j -th word in the output sentence
$\mathbf{E}^{(t)}$	Embedding matrix of the decoder
$\mathbf{h}_j^{(t)}$	the j -th hidden vector of the decoder

1.2.1 Encoder Embedding Layer

The first layer, or the encoder embedding layer converts the each word in the input sentence to the embedding vector. When processing the i -th word in the input sentence, the input and the output of the layer are the following:

- The input is \mathbf{x}_i : the one-hot vector which represents i -th word
- The output is $\bar{\mathbf{x}}_i$: the embedding vector which represents i -th word

Each embedding vector is calculated by the following equation:

$$\bar{\mathbf{x}}_i = \mathbf{E}^{(s)} \mathbf{x}_i$$

$\mathbf{E}^{(s)} \in \mathbb{R}^{D \times |\mathcal{V}^{(s)}|}$ is the embedding matrix of the encoder.

1.2.2 Encoder Recurrent Layer

The encoder recurrent layer generates the hidden vectors from the embedding vectors. When processing the i -th embedding vector, the input and the output of the layer are the following:

- The input is $\bar{\mathbf{x}}_i$: the embedding vector which represents the i -th word
- The output is $\mathbf{h}_i^{(s)}$: the hidden vector of the i -th position

For example, when using the uni-directional RNN of one layer, the process can be represented as the following function $\Psi^{(s)}$:

$$\begin{aligned} \mathbf{h}_i^{(s)} &= \Psi^{(s)}(\bar{\mathbf{x}}_i, \mathbf{h}_{i-1}^{(s)}) \\ &= \tanh \left(\mathbf{W}^{(s)} \begin{bmatrix} \mathbf{h}_{i-1}^{(s)} \\ \bar{\mathbf{x}}_i \end{bmatrix} + \mathbf{b}^{(s)} \right) \end{aligned}$$

In this case, we use the tanh as the activation function.

1.2.3 Decoder Embedding Layer

The decoder embedding layer converts the each word in the output sentence to the embedding vector. When processing the j -th word in the output sentence, the input and the output of the layer are the following:

- The input is \mathbf{y}_{j-1} : the one-hot vector which represents the $(j - 1)$ -th word generated by the decoder output layer

- The output is $\bar{\mathbf{y}}_j$: the embedding vector which represents the $(j - 1)$ -th word

Each embedding vector is calculated by the following equation:

$$\bar{\mathbf{y}}_j = \mathbf{E}^{(t)} \mathbf{y}_{j-1}$$

$\mathbf{E}^{(t)} \in \mathbb{R}^{D \times |\mathcal{V}^{(t)}|}$ is the embedding matrix of the encoder.

1.2.4 Decoder Recurrent Layer

The decoder recurrent layer generates the hidden vectors from the embedding vectors. When processing the j -th embedding vector, the input and the output of the layer are the following:

- The input is $\bar{\mathbf{y}}_j$: the embedding vector
- The output is $\mathbf{h}_j^{(t)}$: the hidden vector of j -th position

For example, when using the uni-directional RNN of one layer, the process can be represented as the following function $\Psi^{(t)}$:

$$\begin{aligned} \mathbf{h}_j^{(t)} &= \Psi^{(t)}(\bar{\mathbf{y}}_j, \mathbf{h}_{j-1}^{(t)}) \\ &= \tanh \left(\mathbf{W}^{(t)} \begin{bmatrix} \mathbf{h}_{j-1}^{(t)} \\ \bar{\mathbf{y}}_j \end{bmatrix} + \mathbf{b}^{(t)} \right) \end{aligned}$$

In this case, we use the tanh as the activation function. And we must use the encoder's hidden vector of the last position as the decoder's hidden vector of first position as following:

$$\mathbf{h}_0^{(t)} = \mathbf{z} = \mathbf{h}_I^{(s)}$$

1.2.5 Decoder Output Layer

The decoder output layer generates the probability of the j -th word of the output sentence from the hidden vector. When processing the j -th embedding vector, the input and the output of the layer are the following:

- The input is $\mathbf{h}_j^{(t)}$: the hidden vector of j -th position
- The output is p_j : the probability of generating the one-hot vector \mathbf{y}_j of the j -th word

$$\begin{aligned} p_j &= P_\theta(\mathbf{y}_j | \mathbf{Y}_{<j}) = \text{softmax}(\mathbf{o}_j) \cdot \mathbf{y}_j \\ &= \text{softmax}(\mathbf{W}^{(o)} \mathbf{h}_j^{(t)} + \mathbf{b}^{(o)}) \cdot \mathbf{y}_j \end{aligned}$$

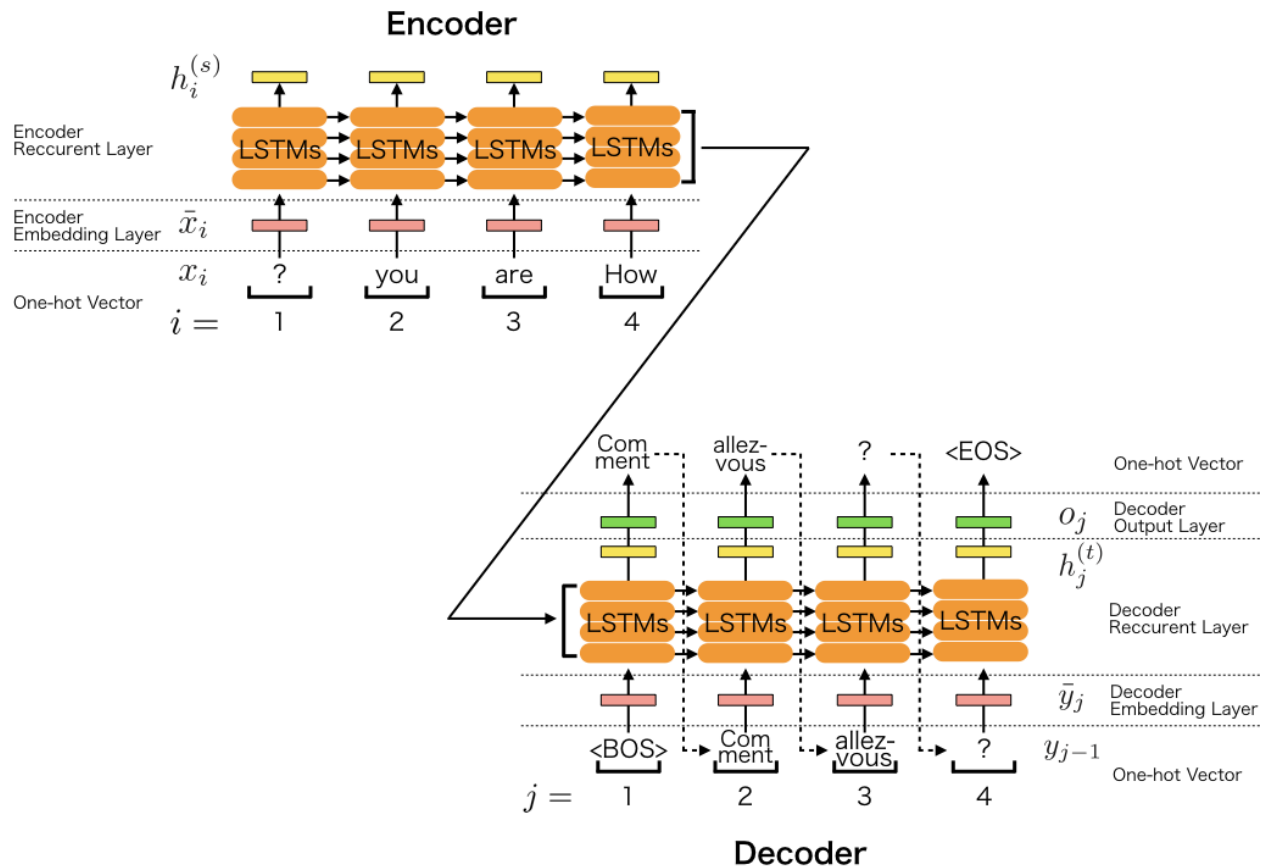
Note: There are a lot of varieties of seq2seq models. We can use the different RNN models in terms of: (1) directionality (unidirectional or bidirectional), (2) depth (single-layer or multi-layer), (3) type (a vanilla RNN, a Long Short-term Memory (LSTM), or a gated recurrent unit (GRU)), and (4) additional functionality (s.t. Attention Mechanism).

4.7.3 2. Implementation of Seq2seq Model

The official Chainer repository includes a neural machine translation example using the seq2seq model. We will now provide an overview of the example and explain its implementation in detail. [chainer/examples/seq2seq](https://github.com/chainer/examples/tree/master/seq2seq)

2.1 Model Overview

In this simple example, an input sequence is processed by a stacked **LSTM-RNN** (long short-term memory recurrent neural networks) and it is encoded as a fixed-size vector. The output sequence is also processed by another stacked LSTM-RNN. At decoding time, an output sequence is generated using argmax.



2.2 Step-by-step Implementation

2.2.1 Import Package

First, let's import necessary packages.

Listing 21: seq2seq.py

```
from nltk.translate import bleu_score
import numpy
import progressbar
import six

import chainer
from chainer.backends import cuda
import chainer.functions as F
import chainer.links as L
```

(continues on next page)

(continued from previous page)

```
from chainer import training
from chainer.training import extensions
```

2.2.2 Define Training Settings

Define all training settings here.

Listing 22: seq2seq.py

```
parser.add_argument('SOURCE', help='source sentence list')
parser.add_argument('TARGET', help='target sentence list')
parser.add_argument('SOURCE_VOCAB', help='source vocabulary file')
parser.add_argument('TARGET_VOCAB', help='target vocabulary file')
parser.add_argument('--validation-source',
                    help='source sentence list for validation')
parser.add_argument('--validation-target',
                    help='target sentence list for validation')
parser.add_argument('--batchsize', '-b', type=int, default=64,
                    help='number of sentence pairs in each mini-batch')
parser.add_argument('--epoch', '-e', type=int, default=20,
                    help='number of sweeps over the dataset to train')
parser.add_argument('--gpu', '-g', type=int, default=-1,
                    help='GPU ID (negative value indicates CPU)')
parser.add_argument('--resume', '-r', default='',
                    help='resume the training from snapshot')
parser.add_argument('--unit', '-u', type=int, default=1024,
                    help='number of units')
parser.add_argument('--layer', '-l', type=int, default=3,
                    help='number of layers')
parser.add_argument('--min-source-sentence', type=int, default=1,
                    help='minimum length of source sentence')
parser.add_argument('--max-source-sentence', type=int, default=50,
                    help='maximum length of source sentence')
parser.add_argument('--min-target-sentence', type=int, default=1,
                    help='minimum length of target sentence')
parser.add_argument('--max-target-sentence', type=int, default=50,
                    help='maximum length of target sentence')
parser.add_argument('--log-interval', type=int, default=200,
                    help='number of iteration to show log')
parser.add_argument('--validation-interval', type=int, default=4000,
                    help='number of iteration to evaluate the model '
                    'with validation dataset')
parser.add_argument('--out', '-o', default='result',
                    help='directory to output the result')
```

2.2.3 Define Network Structure

The Chainer implementation of seq2seq is shown below. It implements the model depicted in the above figure.

Listing 23: seq2seq.py

```
class Seq2seq(chainer.Chain):
```

(continues on next page)

(continued from previous page)

```

def __init__(self, n_layers, n_source_vocab, n_target_vocab, n_units):
    super(Seq2seq, self).__init__()
    with self.init_scope():
        self.embed_x = L.EmbedID(n_source_vocab, n_units)
        self.embed_y = L.EmbedID(n_target_vocab, n_units)
        self.encoder = L.NStepLSTM(n_layers, n_units, n_units, 0.1)
        self.decoder = L.NStepLSTM(n_layers, n_units, n_units, 0.1)
        self.W = L.Linear(n_units, n_target_vocab)

    self.n_layers = n_layers
    self.n_units = n_units

def __call__(self, xs, ys):
    xs = [x[:-1] for x in xs]

    eos = self.xp.array([EOS], numpy.int32)
    ys_in = [F.concat([eos, y], axis=0) for y in ys]
    ys_out = [F.concat([y, eos], axis=0) for y in ys]

    # Both xs and ys_in are lists of arrays.
    exs = sequence_embed(self.embed_x, xs)
    eys = sequence_embed(self.embed_y, ys_in)

    batch = len(xs)
    # None represents a zero vector in an encoder.
    hx, cx, _ = self.encoder(None, None, exs)
    _, _, os = self.decoder(hx, cx, eys)

    # It is faster to concatenate data before calculating loss
    # because only one matrix multiplication is called.
    concat_os = F.concat(os, axis=0)
    concat_ys_out = F.concat(ys_out, axis=0)
    loss = F.sum(F.softmax_cross_entropy(
        self.W(concat_os), concat_ys_out, reduce='no')) / batch

    chainer.report({'loss': loss.data}, self)
    n_words = concat_ys_out.shape[0]
    perp = self.xp.exp(loss.data * batch / n_words)
    chainer.report({'perp': perp}, self)
    return loss

def translate(self, xs, max_length=100):
    batch = len(xs)
    with chainer.no_backprop_mode(), chainer.using_config('train', False):
        xs = [x[:-1] for x in xs]
        exs = sequence_embed(self.embed_x, xs)
        h, c, _ = self.encoder(None, None, exs)
        ys = self.xp.full(batch, EOS, numpy.int32)
        result = []
        for i in range(max_length):
            eys = self.embed_y(ys)
            eys = F.split_axis(eys, batch, 0)
            h, c, ys = self.decoder(h, c, eys)
            cys = F.concat(ys, axis=0)
            wy = self.W(cys)
            ys = self.xp.argmax(wy.data, axis=1).astype(numpy.int32)
            result.append(ys)

```

(continues on next page)

(continued from previous page)

```

# Using `xp.concatenate(...)` instead of `xp.stack(result)` here to
# support NumPy 1.9.
result = cuda.to_cpu(
    self.xp.concatenate([self.xp.expand_dims(x, 0) for x in result]).T)

# Remove EOS taggs
outs = []
for y in result:
    inds = numpy.argwhere(y == EOS)
    if len(inds) > 0:
        y = y[:inds[0, 0]]
    outs.append(y)
return outs

```

- In Seq2seq, three functions are defined: the constructor `__init__`, the function call `__call__`, and the function for translation `translate`.

Listing 24: seq2seq.py

```

def __init__(self, n_layers, n_source_vocab, n_target_vocab, n_units):
    super(Seq2seq, self).__init__()
    with self.init_scope():
        self.embed_x = L.EmbedID(n_source_vocab, n_units)
        self.embed_y = L.EmbedID(n_target_vocab, n_units)
        self.encoder = L.NStepLSTM(n_layers, n_units, n_units, 0.1)
        self.decoder = L.NStepLSTM(n_layers, n_units, n_units, 0.1)
        self.W = L.Linear(n_units, n_target_vocab)

    self.n_layers = n_layers
    self.n_units = n_units

```

- When we instantiate this class for making a model, we give the number of stacked lstms to `n_layers`, the vocabulary size of the source language to `n_source_vocab`, the vocabulary size of the target language to `n_target_vocab`, and the size of hidden vectors to `n_units`.
- This network uses `chainer.links.NStepLSTM`, `chainer.links.EmbedID`, and `chainer.links.Linear` as its building blocks. All the layers are registered and initialized in the context with `self.init_scope()`.
- You can access all the parameters in those layers by calling `self.params()`.
- In the constructor, it initializes all parameters with values sampled from a uniform distribution $U(-1, 1)$.

Listing 25: seq2seq.py

```

def __call__(self, xs, ys):
    xs = [x[:-1] for x in xs]

    eos = self.xp.array([EOS], numpy.int32)
    ys_in = [F.concat([eos, y], axis=0) for y in ys]
    ys_out = [F.concat([y, eos], axis=0) for y in ys]

    # Both xs and ys_in are lists of arrays.
    exs = sequence_embed(self.embed_x, xs)
    eys = sequence_embed(self.embed_y, ys_in)

```

(continues on next page)

(continued from previous page)

```

batch = len(xs)
# None represents a zero vector in an encoder.
hx, cx, _ = self.encoder(None, None, exs)
_, _, os = self.decoder(hx, cx, eys)

# It is faster to concatenate data before calculating loss
# because only one matrix multiplication is called.
concat_os = F.concat(os, axis=0)
concat_ys_out = F.concat(ys_out, axis=0)
loss = F.sum(F.softmax_cross_entropy(
    self.W(concat_os), concat_ys_out, reduce='no')) / batch

chainer.report({'loss': loss.data}, self)
n_words = concat_ys_out.shape[0]
perp = self.xp.exp(loss.data * batch / n_words)
chainer.report({'perp': perp}, self)
return loss

```

- The `__call__` method takes sequences of source language's word IDs `xs` and sequences of target language's word IDs `ys`. Each sequence represents a sentence, and the size of `xs` is mini-batch size.
- Note that the sequences of word IDs `xs` and `ys` are converted to a vocabulary-size one-hot vectors and then multiplied with the embedding matrix in `sequence_embed` to obtain embedding vectors `exs` and `eys`.

Listing 26: seq2seq.py

```

def sequence_embed(embed, xs):
    x_len = [len(x) for x in xs]
    x_section = numpy.cumsum(x_len[:-1])
    ex = embed(F.concat(xs, axis=0))
    exs = F.split_axis(ex, x_section, 0)
    return exs

```

- `self.encoder` and `self.decoder` are the encoder and the decoder of the seq2seq model. Each element of the decoder output `os` is $h_{[1:J]}^{(t)}$ in the figure above.
- After calculating the recurrent layer output, the loss `loss` and the perplexity `perp` are calculated, and the values are logged by `chainer.report`.

Note: It is well known that the seq2seq model learns much better when the source sentences are reversed. The paper[1] says that “While the LSTM is capable of solving problems with long term dependencies, we discovered that the LSTM learns much better when the source sentences are reversed (the target sentences are not reversed). By doing so, the LSTM’s test perplexity dropped from 5.8 to 4.7, and the test BLEU scores of its decoded translations increased from 25.9 to 30.6.” So, at the first line in the `__call__`, the input sentences are reversed `xs = [x[::-1] for x in xs]`.

Listing 27: seq2seq.py

```

def translate(self, xs, max_length=100):
    batch = len(xs)
    with chainer.no_backprop_mode(), chainer.using_config('train', False):
        xs = [x[::-1] for x in xs]
        exs = sequence_embed(self.embed_x, xs)
        h, c, _ = self.encoder(None, None, exs)

```

(continues on next page)

(continued from previous page)

```

ys = self.xp.full(batch, EOS, numpy.int32)
result = []
for i in range(max_length):
    eys = self.embed_y(ys)
    eys = F.split_axis(eys, batch, 0)
    h, c, ys = self.decoder(h, c, eys)
    cys = F.concat(ys, axis=0)
    wy = self.W(cys)
    ys = self.xp.argmax(wy.data, axis=1).astype(numpy.int32)
    result.append(ys)

# Using `xp.concatenate(...)` instead of `xp.stack(result)` here to
# support NumPy 1.9.
result = cuda.to_cpu(
    self.xp.concatenate([self.xp.expand_dims(x, 0) for x in result]).T)

# Remove EOS taggs
outs = []
for y in result:
    inds = numpy.argwhere(y == EOS)
    if len(inds) > 0:
        y = y[:inds[0, 0]]
    outs.append(y)
return outs

```

- After the model learned the parameters, the function `translate` is called to generate the translated sentences `outs` from the source sentences `xs`.
- So as not to change the parameters, the codes for the translation are nested in the scope `chainer.no_backprop_mode()` and `chainer.using_config('train', False)`.

2.2.4 Load French-English Corpus from WMT15 Dataset

In this tutorial, we use French-English corpus from WMT15 [website](#) that contains 10^9 documents. We must prepare additional libraries, dataset, and parallel corpus. To understand the pre-processing, see [2.3.1 Requirements](#).

After the pre-processing the dataset, let's make dataset objects:

Listing 28: seq2seq.py

```

# Load pre-processed dataset
source_ids = load_vocabulary(args.SOURCE_VOCAB)
target_ids = load_vocabulary(args.TARGET_VOCAB)
train_source = load_data(source_ids, args.SOURCE)
train_target = load_data(target_ids, args.TARGET)
assert len(train_source) == len(train_target)

train_data = [
    (s, t)
    for s, t in six.moves.zip(train_source, train_target)
    if (args.min_source_sentence <= len(s) <= args.max_source_sentence and
        args.min_target_sentence <= len(t) <= args.max_target_sentence)]

train_source_unknown = calculate_unknown_ratio(
    [s for s, _ in train_data])

```

(continues on next page)

(continued from previous page)

```

train_target_unknown = calculate_unknown_ratio(
    [t for _, t in train_data])

print('Source vocabulary size: %d' % len(source_ids))
print('Target vocabulary size: %d' % len(target_ids))
print('Train data size: %d' % len(train_data))
print('Train source unknown ratio: %.2f%%' % (train_source_unknown * 100))
print('Train target unknown ratio: %.2f%%' % (train_target_unknown * 100))

target_words = {i: w for w, i in target_ids.items()}
source_words = {i: w for w, i in source_ids.items()}

```

- This code uses utility functions below:

Listing 29: seq2seq.py

```

def load_vocabulary(path):
    with open(path) as f:
        # +2 for UNK and EOS
        word_ids = {line.strip(): i + 2 for i, line in enumerate(f)}
    word_ids['<UNK>'] = 0
    word_ids['<EOS>'] = 1
    return word_ids

```

Listing 30: seq2seq.py

```

def load_data(vocabulary, path):
    n_lines = count_lines(path)
    bar = progressbar.ProgressBar()
    data = []
    print('loading...: %s' % path)
    with open(path) as f:
        for line in bar(f, max_value=n_lines):
            words = line.strip().split()
            array = numpy.array([vocabulary.get(w, UNK)
                                for w in words], numpy.int32)
            data.append(array)
    return data

```

Listing 31: seq2seq.py

```

def calculate_unknown_ratio(data):
    unknown = sum((s == UNK).sum() for s in data)
    total = sum(s.size for s in data)
    return unknown / total

```

2.2.5 Define Evaluation Function (Bleu Score)

BLEU[3] (bilingual evaluation understudy) is the evaluation metric for the quality of text which has been machine-translated from one natural language to another.

Listing 32: seq2seq.py

```

class CalculateBleu(chainer.training.Extension):

    trigger = 1, 'epoch'
    priority = chainer.training.PRIORITY_WRITER

    def __init__(
        self, model, test_data, key, batch=100, device=-1, max_length=100):
        self.model = model
        self.test_data = test_data
        self.key = key
        self.batch = batch
        self.device = device
        self.max_length = max_length

    def __call__(self, trainer):
        with chainer.no_backprop_mode():
            references = []
            hypotheses = []
            for i in range(0, len(self.test_data), self.batch):
                sources, targets = zip(*self.test_data[i:i + self.batch])
                references.extend([t.tolist() for t in targets])

                sources = [
                    chainer.dataset.to_device(self.device, x) for x in sources]
                ys = [y.tolist()
                    for y in self.model.translate(sources, self.max_length)]
                hypotheses.extend(ys)

        bleu = bleu_score.corpus_bleu(
            references, hypotheses,
            smoothing_function=bleu_score.SmoothingFunction().method1)
        chainer.report({self.key: bleu})

```

2.2.6 Create Iterator

Here, the code below just creates iterator objects.

Listing 33: seq2seq.py

```

train_iter = chainer.iterators.SerialIterator(train_data, args.batchsize)

```

2.2.7 Create RNN and Classification Model

Instantiate Seq2seq model.

Listing 34: seq2seq.py

```
model = Seq2seq(args.layer, len(source_ids), len(target_ids), args.unit)
```

2.2.8 Setup Optimizer

Prepare an optimizer. We use `chainer.optimizers.Adam`.

Listing 35: seq2seq.py

```
optimizer = chainer.optimizers.Adam()
optimizer.setup(model)
```

2.2.9 Setup and Run Trainer

Let's make a trainer object.

Listing 36: seq2seq.py

```
updater = training.updaters.StandardUpdater(
    train_iter, optimizer, converter=convert, device=args.gpu)
trainer = training.Trainer(updater, (args.epoch, 'epoch'), out=args.out)
trainer.extend(extensions.LogReport(
    trigger=(args.log_interval, 'iteration')))
trainer.extend(extensions.PrintReport(
    ['epoch', 'iteration', 'main/loss', 'validation/main/loss',
     'main/perp', 'validation/main/perp', 'validation/main/bleu',
     'elapsed_time']),
    trigger=(args.log_interval, 'iteration'))
```

Setup the trainer's extension to see the BLEU score on the test data.

Listing 37: seq2seq.py

```
test_source = load_data(source_ids, args.validation_source)
test_target = load_data(target_ids, args.validation_target)
assert len(test_source) == len(test_target)
test_data = list(six.moves.zip(test_source, test_target))
test_data = [(s, t) for s, t in test_data if 0 < len(s) and 0 < len(t)]
test_source_unknown = calculate_unknown_ratio(
    [s for s, _ in test_data])
test_target_unknown = calculate_unknown_ratio(
    [t for _, t in test_data])

print('Validation data: %d' % len(test_data))
print('Validation source unknown ratio: %.2f%%' %
      (test_source_unknown * 100))
print('Validation target unknown ratio: %.2f%%' %
      (test_target_unknown * 100))

@chainer.training.make_extension()
```

(continues on next page)

(continued from previous page)

```
def translate(trainer):
    source, target = test_data[numpy.random.choice(len(test_data))]
    result = model.translate([model.xp.array(source)])[0]

    source_sentence = ' '.join([source_words[x] for x in source])
    target_sentence = ' '.join([target_words[y] for y in target])
    result_sentence = ' '.join([target_words[y] for y in result])
    print('# source : ' + source_sentence)
    print('# result : ' + result_sentence)
    print('# expect : ' + target_sentence)

trainer.extend(
    translate, trigger=(args.validation_interval, 'iteration'))
trainer.extend(
    CalculateBleu(
        model, test_data, 'validation/main/bleu', device=args.gpu),
    trigger=(args.validation_interval, 'iteration'))
```

Let's start the training!

Listing 38: seq2seq.py

```
trainer.run()
```

2.3 Run Example

2.3.1 Requirements

Before running the example, you must prepare additional libraries, dataset, and parallel corpus.

- See the detail description: [chainer/examples/seq2seq/README.md](#)

2.3.1 Training the model

You can train the model with the script: [chainer/examples/seq2seq/seq2seq.py](#)

```
$ pwd
/root2chainer/chainer/examples/seq2seq
$ python seq2seq.py --gpu=0 giga-fren.preprocess.en giga-fren.preprocess.fr \
vocab.en vocab.fr \
--validation-source newstest2013.preprocess.en \
--validation-target newstest2013.preprocess.fr > log
100% (22520376 of 22520376) |#####| Elapsed Time: 0:09:20 Time: 0:09:20
100% (22520376 of 22520376) |#####| Elapsed Time: 0:10:36 Time: 0:10:36
100% (3000 of 3000) |#####| Elapsed Time: 0:00:00 Time: 0:00:00
100% (3000 of 3000) |#####| Elapsed Time: 0:00:00 Time: 0:00:00
epoch      iteration  main/loss  validation/main/loss  main/perp  validation/main/
→perp     validation/main/bleu  elapsed_time
0          200          171.449                991.556
→          85.6739
```

(continues on next page)

(continued from previous page)

0	400	143.918	183.594	└
↪		172.473		
0	600	133.48	126.945	└
↪		260.315		
0	800	128.734	104.127	└
↪		348.062		
0	1000	124.741	91.5988	└
↪		436.536		
...				

Note: Before running the script, be careful the locale and the python's encoding. Please setup them to use utf-8 encoding.

2.3.1 Validate the model

While you are training the model, you can get the validation results:

```
...
# source : We knew the Government had tried many things , like launching <UNK> with
↪<UNK> or organising speed dating evenings .
# result : Nous savions que le gouvernement avait <UNK> plusieurs fois , comme le
↪<UNK> <UNK> , le <UNK> ou le <UNK> <UNK> .
# expect : Nous savions que le gouvernement avait tenté plusieurs choses comme lancer
↪des parfums aux <UNK> ou organiser des soirées de <UNK>
...
```

4.7.4 3. Reference

- [1] [Sequence to Sequence Learning with Neural Networks](#)
- [2] [Learning Phrase Representations using RNN Encoder–Decoder for Statistical Machine Translation](#)
- [3] [BLEU](#)

5.1 Variable and Parameter

<code>chainer.Variable</code>	Array with a structure to keep track of computation.
<code>chainer.as_variable</code>	Converts an array or a variable into <i>Variable</i> .
<code>chainer.Parameter</code>	Parameter variable that can be registered to a link.
<code>chainer.variable.VariableNode</code>	Node in the backward computational graph representing a variable.

5.1.1 chainer.Variable

class `chainer.Variable` (*data=None*, *, *name=None*, *grad=None*, *requires_grad=True*)

Array with a structure to keep track of computation.

Every variable holds a data array of type either `numpy.ndarray` or `cupy.ndarray`.

A variable object holds a data array and a *VariableNode* object of a computational graph. If the variable is constructed by the user, the node is *root* and does not hold any parent. If the variable is constructed by a *FunctionNode* object (i.e., by calling functions under `chainer.functions` or user-defined functions), or by using operators (see the list below), the node holds a reference to its parent called *creator_node*. This reference is used in backpropagation to backtrack the graph.

Users can disable (resp. enable) this chaining behavior by calling `no_backprop_mode()` (resp. `force_backprop_mode()`). In the former context, a variable never creates a computational graph, whereas in the latter context, it is forced to create.

Note: The following operators are defined for variable(s).

- Indexing: `a[slices]` (`__getitem__()`)
- Addition: `a + b` (`__add__()`, `__radd__()`)
- Subtraction: `a - b` (`__sub__()`, `__rsub__()`)

- Multiplication: `a * b` (`__mul__()`, `__rmul__()`)
 - Division: `a / b` (`__div__()`, `__rdiv__()`, `__truediv__()`, `__rtruediv__()`)
 - Floor Division: `a // b` (`__floordiv__()`, `__rfloordiv__()`)
 - Exponentiation: `a ** b` (`__pow__()`, `__rpow__()`)
 - Matrix Multiplication: `a @ b` (`__matmul__()`, `__rmatmul__()`)
 - Negation (Arithmetic): `- a` (`__neg__()`)
 - Absolute value: `abs(a)` (`__abs__()`)
-

Warning: `volatile` argument is not supported anymore since v2. Instead, use `chainer.no_backprop_mode()`.

Parameters

- **data** (`numpy.ndarray` or `cupy.ndarray`) – Initial data array.
- **name** (`str`) – Name of the variable.
- **grad** (`numpy.ndarray` or `cupy.ndarray`) – Initial gradient array.
- **requires_grad** (`bool`) – Boolean indicating whether `grad` will be set in backward calculation.

Methods

`__getitem__` (`slices`)

Extract elements from array with specified shape, axes and offsets.

Parameters

- **x** (`Variable` or `numpy.ndarray` or `cupy.ndarray`) – A variable to be sliced.
- **slices** (`int`, `slice`, `Ellipsis`, `None`, `integer array-like`, `boolean array-like` or `tuple of them`) – An object to specify the selection of elements.

Returns A `Variable` object which contains sliced array of `x`.

Note: It only supports types that are supported by CUDA's `atomicAdd` when an integer array is included in `slices`. The supported types are `numpy.float32`, `numpy.int32`, `numpy.uint32`, `numpy.uint64` and `numpy.ulonglong`.

Note: It does not support `slices` that contains multiple boolean arrays.

Note: See NumPy document for details of [indexing](#).

Example

```

>>> x = np.arange(12).reshape((2, 2, 3))
>>> x
array([[[ 0,  1,  2],
        [ 3,  4,  5]],

       [[ 6,  7,  8],
        [ 9, 10, 11]]])
>>> F.get_item(x, 0)
variable([[0, 1, 2],
          [3, 4, 5]])
>>> F.get_item(x, (0, 0, slice(0, 2, 1))) # equals x[0, 0, 0:2:1]
variable([0, 1])
>>> F.get_item(x, (Ellipsis, 2)) # equals x[..., 2]
variable([[ 2,  5],
          [ 8, 11]])
>>> F.get_item(x, (1, np.newaxis, 1, 0)) # equals x[1, None, 1, 0]
variable([9])

```

__len__()

Returns the first dimension of the data array.

Returns Number of the first dimension of the data array.

Return type `int`

__copy__()

addgrad(*var*)

Accumulates the gradient array from given source variable.

This method adds the gradient of a given variable to the gradient of this variable. The accumulation is even done across the host and different devices. If this variable has uninitialized data/grad arrays, this method initializes it with the shape of the given variable and then accumulates the gradient.

Parameters *var* (`Variable`) – Source variable.

backward(*retain_grad=False, enable_double_backprop=False, loss_scale=None*)

Runs error backpropagation (a.k.a. backprop) from this variable.

On backprop, `FunctionNode.backward()` is called on each `FunctionNode` object appearing in the backward graph starting from this variable. The backward graph is represented by backward references from variable nodes to their creators, and from function nodes to their input variable nodes. The backprop stops at all root nodes. Some function nodes set `None` as gradients of some inputs, where further backprop does not take place at such inputs.

This method uses `grad` as the initial error array. User can manually set a gradient array before calling this method. If the shape of `data` is `()` (i.e., it is scalar) and `grad` is `None`, then this method automatically complements 1.0 as the initial error. This is useful on starting backprop from some scalar loss value.

From v3, this method supports *differentiable backprop* (a.k.a. double backprop, grad of grads). To enable it, pass `enable_double_backprop=True`.

Parameters

- **retain_grad** (`bool`) – If `True`, the gradient arrays of all intermediate variables are kept. Otherwise, `grad` of the intermediate variables are set to `None` on appropriate timing, which may reduce the maximum memory consumption.

In most cases of training some models, the purpose of backprop is to compute gradients of parameters, not of all variables, and therefore it is recommended to set this flag `False`.

- **enable_double_backprop** (*bool*) – (Added in v3.0) If `True`, computational trace of the whole backpropagation procedure is recorded to the computational graph so that one can further do backpropagation from the resulting gradients. Note that enabling it results in larger memory consumption needed to store the gradients w.r.t intermediate variables that are required for the second gradient computation.
- **loss_scale** (*float*) – Loss scaling factor. Loss scaling is a useful technique to mitigate vanishing gradient issue that tends to happen when low precision data type like `float16` is used during training. If you set loss scaling factor, gradients of loss values are to be multiplied by the factor before backprop starts. The factor is propagated to whole gradients in a computational graph along the backprop. The gradients of parameters are divided by the factor just before the parameters are to be updated.

cleargrad ()

Clears the gradient array.

copydata (*var*)

Copies the data array from given source variable.

This method copies the data array from given variable to this variable. The copy is done even if the arrays reside on different devices, including across the host and a GPU device. If this variable has an uninitialized data array, this method initializes it by the data array of the given variable. Similarly, if the given variable has an uninitialized data array, this method initializes it by the data array of this variable (`self`). If both are uninitialized, this method does nothing.

Parameters **var** (*Variable*) – Source variable.

debug_print ()

Display a summary of the stored data and location of the Variable

reshape (**shape*)

Returns a variable of a different shape and the same content.

See also:

[`chainer.functions.reshape\(\)`](#) for full documentation,

retain_data ()

Lets the corresponding variable node keep the underlying array.

set_creator (*gen_func*)

Notifies the variable that the given function is its creator.

Parameters **gen_func** (*Function*) – Function object that creates this variable as one of its outputs.

set_creator_node (*fnode*)

Notifies the variable that the given node is its creator.

Parameters **fnode** (*FunctionNode*) – Function node that has this variable as an output.

summary ()

to_cpu ()

Copies the data and gradient arrays to CPU.

to_gpu (*device=None*)

Copies the data and gradient arrays to specified GPU.

Parameters **device** – Target device specifier. If omitted, the current device is used.

to_intel64 ()

Copies the data and gradient arrays to intel64 specific ndarray.

If the array is not suited for intel64, it will be converted to `numpy.ndarray`.

transpose (*axes)

Permute the dimensions of an input variable without copy.

See also:

`chainer.functions.transpose()` for full documentation.

unchain ()

Deletes the reference to the creator of this variable.

This method deletes the reference to the creator from the corresponding variable node. Unlike `unchain_backward()`, it does not backtrack the graph.

This method is equivalent to `self.creator_node = None`.

unchain_backward ()

Deletes references between variable nodes and functions backward.

After this method completes, intermediate variable nodes and functions that are not referenced from anywhere are deallocated by reference count GC. Also this variable itself deletes the reference to its creator function from the node, i.e. the node becomes root in the computation graph. It indicates that backprop after unchaining stops at this variable. This behavior is useful to implement truncated BPTT.

zerograd ()

Initializes the gradient array by zeros.

Note that the gradient variable is unchained from the computational graph by this method because this operation breaks the backprop validity.

Deprecated since version v1.15: Use `cleargrad()` instead.

__eq__ (other)

This operator is not defined for Variable.

__ne__ (other)

This operator is not defined for Variable.

__lt__ (other)

This operator is not defined for Variable.

__le__ (other)

This operator is not defined for Variable.

__gt__ (other)

This operator is not defined for Variable.

__ge__ (other)

This operator is not defined for Variable.

__nonzero__ ()

This operator is not defined for Variable.

__bool__ ()

This operator is not defined for Variable.

__neg__ ()

Element-wise negation.

Returns Output variable.

Return type *Variable*

`__abs__()`
Element-wise absolute.

Returns Output variable.

Return type *Variable*

`__add__()`
Element-wise addition.

Returns Output variable.

Return type *Variable*

`__radd__()`
Element-wise addition.

Returns Output variable.

Return type *Variable*

`__sub__(rhs)`
Element-wise subtraction.

Returns Output variable.

Return type *Variable*

`__rsub__(rhs)`
Element-wise subtraction.

Returns Output variable.

Return type *Variable*

`__mul__(rhs)`
Element-wise multiplication.

Returns Output variable.

Return type *Variable*

`__rmul__(rhs)`
Element-wise multiplication.

Returns Output variable.

Return type *Variable*

`__div__(rhs)`
Element-wise division

Returns Output variable.

Return type *Variable*

`__truediv__(rhs)`
Element-wise division

Returns Output variable.

Return type *Variable*

`__rdiv__(rhs)`
Element-wise division.

Returns Output variable.

Return type *Variable*

`__rtruediv__` (*rhs*)
Element-wise division.

Returns Output variable.

Return type *Variable*

`__floordiv__` (*rhs*)
Element-wise floor division.

Returns Output variable.

Return type *Variable*

`__rfloordiv__` (*rhs*)
Element-wise floor division.

Returns Output variable.

Return type *Variable*

`__pow__` (*rhs*)
Element-wise power function.

Returns Output variable.

Return type *Variable*

`__rpow__` (*rhs*)
Element-wise power function.

Returns Output variable.

Return type *Variable*

`__matmul__` (*rhs*)
Matrix multiplication.

Returns Output variable.

Return type *Variable*

`__rmatmul__` (*rhs*)
Matrix multiplication.

Returns Output variable.

Return type *Variable*

Attributes

T
Transposition of this variable.

array
The underlying data array.

It is either `numpy.ndarray` or `cupy.ndarray` object, or None if the variable is in an uninitialized state.

creator
Function implementation that created this variable.

When this variable has been created by an old-style function (i.e., it is implemented as a subclass of *Function*), this property returns that *Function* object.

When this variable has been created by a new-style function (i.e., it is implemented as a subclass of *FunctionNode* class), this property returns that node object.

creator_node

FunctionNode object that created this variable.

This property has a setter to which *None* can be set. Setting *None* to this property is equivalent to call *unchain()*; it purges the variable from the function that created this variable.

The setter also accepts the original *FunctionNode* object that created this variable. For example, you can once set *None* to this property and then set the original value again.

Note: Setting an irrelevant *FunctionNode()* object does not emit any error immediately, whereas the behavior is undefined. Do not set a *FunctionNode()* object that did not create this variable object.

data

The underlying data array (equivalent to *array*).

Note that using this attribute directly is discouraged; use *array* instead. Using *array*, you can find an error earlier when your code mixes up *Variable* and *ndarray* because *ndarray* does not have an attribute *.array* while it has *.data*.

dtype**grad**

Gradient array of this variable.

Note that this property returns the underlying array of the gradient variable instead of the gradient variable itself; to get/set gradient variable, use *grad_var* instead.

grad_var

Gradient variable.

label

Short text that represents the variable.

name**ndim****node****rank****requires_grad**

It indicates that *grad* will be set in backward calculation.

shape**size****xp**

Array module for this variable.

Depending on which of CPU/GPU this variable is on, this property returns *numpy* or *cupy*.

5.1.2 chainer.as_variable

`chainer.as_variable(obj)`

Converts an array or a variable into *Variable*.

This is a convenient function to get a *Variable* object transparently from a raw array or a variable.

Note that this function should only be used for type consistency (i.e., to enforce the return value of an API having type *Variable*). The `requires_grad` flag is kept as is; if `obj` is a raw array, the newly created variable has `requires_grad = False`. In order to make a variable w.r.t. which you want to compute the gradient, you should use *Variable* directly.

Parameters `obj` (*numpy.ndarray* or *cupy.ndarray* or *Variable*) – An array or a variable that you want to convert to *Variable*.

Returns A variable converted from `obj`. If `obj` is a raw array, this is a new *Variable* object that wraps the array. If `obj` is already a *Variable* object, this function returns `obj` as is.

Return type *Variable*

5.1.3 chainer.Parameter

`class chainer.Parameter(initializer=None, shape=None, name=None)`

Parameter variable that can be registered to a link.

Parameter is a subclass of *Variable*. It almost behaves as same as a usual variable except that a parameter can be registered to a *Link* object just by assigning it to an attribute of the link within an `init_scope()` context.

Parameter also supports an initialization by an initializer. It can have two initializers: one for the data array, and the other for the gradient array. The initializer only specifies the way of filling the elements of these arrays, and the shape information is specified at the initialization point.

When a link that the parameter has been registered to is passed to an *GradientMethod*, an update rule is set to the parameter. This update rule specifies how to update the data array of the parameter using its gradient array.

Parameters

- **initializer** (*Initializer* or *numpy.ndarray* or *cupy.ndarray*) – Initializer of the data array. If `shape` is given, this initializer is immediately used to initialize the data array. Otherwise, if it is an array, it is immediately used as the data array, and otherwise the data array is left uninitialized and will be initialized by this initializer in `initialize()`. It can also be a scalar, in which case the data array will be filled by this scalar. Note that float32 is used in this case.
- **shape** (*int* or *tuple of int* or *None*) – Shape of the parameter. If it is *None*, the initialization is deferred to the call of `initialize()`.
- **name** (*str*) – Name of the parameter.

Variables

- **initializer** – Initializer of the data array. It is used for initializing the data array of an uninitialized variable.
- **update_rule** – *UpdateRule* instance that updates this variable as a parameter. This argument is set to `update_rule`.

Methods

`__getitem__` (*slices*)

Extract elements from array with specified shape, axes and offsets.

Parameters

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – A variable to be sliced.
- **slices** (*int, slice, Ellipsis, None, integer array-like, boolean array-like or tuple of them*) – An object to specify the selection of elements.

Returns A *Variable* object which contains sliced array of x.

Note: It only supports types that are supported by CUDA's atomicAdd when an integer array is included in slices. The supported types are `numpy.float32`, `numpy.int32`, `numpy.uint32`, `numpy.uint64` and `numpy.ulonglong`.

Note: It does not support slices that contains multiple boolean arrays.

Note: See NumPy document for details of indexing.

Example

```
>>> x = np.arange(12).reshape((2, 2, 3))
>>> x
array([[[ 0,  1,  2],
        [ 3,  4,  5]],

       [[ 6,  7,  8],
        [ 9, 10, 11]])
>>> F.get_item(x, 0)
variable([[0, 1, 2],
          [3, 4, 5]])
>>> F.get_item(x, (0, 0, slice(0, 2, 1))) # equals x[0, 0, 0:2:1]
variable([0, 1])
>>> F.get_item(x, (Ellipsis, 2)) # equals x[..., 2]
variable([[ 2,  5],
          [ 8, 11]])
>>> F.get_item(x, (1, np.newaxis, 1, 0)) # equals x[1, None, 1, 0]
variable([9])
```

`__len__` ()

Returns the first dimension of the data array.

Returns Number of the first dimension of the data array.

Return type `int`

`__copy__` ()

addgrad (*var*)

Accumulates the gradient array from given source variable.

This method adds the gradient of a given variable to the gradient of this variable. The accumulation is even done across the host and different devices. If this variable has uninitialized data/grad arrays, this method initializes it with the shape of the given variable and then accumulates the gradient.

Parameters **var** (*Variable*) – Source variable.

backward (*retain_grad=False, enable_double_backprop=False, loss_scale=None*)

Runs error backpropagation (a.k.a. backprop) from this variable.

On backprop, *FunctionNode.backward()* is called on each *FunctionNode* object appearing in the backward graph starting from this variable. The backward graph is represented by backward references from variable nodes to their creators, and from function nodes to their input variable nodes. The backprop stops at all root nodes. Some function nodes set *None* as gradients of some inputs, where further backprop does not take place at such inputs.

This method uses *grad* as the initial error array. User can manually set a gradient array before calling this method. If the shape of *data* is *()* (i.e., it is scalar) and *grad* is *None*, then this method automatically complements 1.0 as the initial error. This is useful on starting backprop from some scalar loss value.

From v3, this method supports *differentiable backprop* (a.k.a. double backprop, grad of grads). To enable it, pass *enable_double_backprop=True*.

Parameters

- **retain_grad** (*bool*) – If *True*, the gradient arrays of all intermediate variables are kept. Otherwise, *grad* of the intermediate variables are set to *None* on appropriate timing, which may reduce the maximum memory consumption.

In most cases of training some models, the purpose of backprop is to compute gradients of parameters, not of all variables, and therefore it is recommended to set this flag *False*.

- **enable_double_backprop** (*bool*) – (*Added in v3.0*) If *True*, computational trace of the whole backpropagation procedure is recorded to the computational graph so that one can further do backpropagation from the resulting gradients. Note that enabling it results in larger memory consumption needed to store the gradients w.r.t intermediate variables that are required for the second gradient computation.
- **loss_scale** (*float*) – Loss scaling factor. Loss scaling is a useful technique to mitigate vanishing gradient issue that tends to happen when low precision data type like float16 is used during training. If you set loss scaling factor, gradients of loss values are to be multiplied by the factor before backprop starts. The factor is propagated to whole gradients in a computational graph along the backprop. The gradients of parameters are divided by the factor just before the parameters are to be updated.

cleargrad ()

Clears the gradient array.

copydata (*var*)

Copies the data array from given source variable.

This method copies the data array from given variable to this variable. The copy is done even if the arrays reside on different devices, including across the host and a GPU device. If this variable has an uninitialized data array, this method initializes it by the data array of the given variable. Similarly, if the given variable has an uninitialized data array, this method initializes it by the data array of this variable (*self*). If both are uninitialized, this method does nothing.

Parameters **var** (*Variable*) – Source variable.

debug_print()

Display a summary of the stored data and location of the Variable

initialize(shape)

Initializes the uninitialized variable.

Uninitialized variable is a variable created with the data array set to None. This method creates and initializes the data array. The shape of the variable can be left unknown until this method is called.

Parameters **shape** (*tuple of int*) – Shape of the data array.

reshape(*shape)

Returns a variable of a different shape and the same content.

See also:

[`chainer.functions.reshape\(\)`](#) for full documentation,

retain_data()

Lets the corresponding variable node keep the underlying array.

set_creator(gen_func)

Notifies the variable that the given function is its creator.

Parameters **gen_func** (`Function`) – Function object that creates this variable as one of its outputs.

set_creator_node(fnode)

Notifies the variable that the given node is its creator.

Parameters **fnode** (`FunctionNode`) – Function node that has this variable as an output.

summary()**to_cpu()**

Copies the data and gradient arrays to CPU.

to_gpu(device=None)

Copies the data and gradient arrays to specified GPU.

Parameters **device** – Target device specifier. If omitted, the current device is used.

to_intel64()

Copies the data and gradient arrays to intel64 specific mdarray.

If the array is not suited for intel64, it will be converted to `numpy.ndarray`.

transpose(*axes)

Permute the dimensions of an input variable without copy.

See also:

[`chainer.functions.transpose\(\)`](#) for full documentation.

unchain()

Deletes the reference to the creator of this variable.

This method deletes the reference to the creator from the corresponding variable node. Unlike [`unchain_backward\(\)`](#), it does not backtrack the graph.

This method is equivalent to `self.creator_node = None`.

unchain_backward()

Deletes references between variable nodes and functions backward.

After this method completes, intermediate variable nodes and functions that are not referenced from anywhere are deallocated by reference count GC. Also this variable itself deletes the reference to its creator function from the node, i.e. the node becomes root in the computation graph. It indicates that backprop after unchaining stops at this variable. This behavior is useful to implement truncated BPTT.

update()

Updates the data array using the gradient and the update rule.

This method updates the parameter using the attached update rule.

zerograd()

Initializes the gradient array by zeros.

Note that the gradient variable is unchained from the computational graph by this method because this operation breaks the backprop validity.

Deprecated since version v1.15: Use `cleargrad()` instead.

__eq__(other)

This operator is not defined for Variable.

__ne__(other)

This operator is not defined for Variable.

__lt__(other)

This operator is not defined for Variable.

__le__(other)

This operator is not defined for Variable.

__gt__(other)

This operator is not defined for Variable.

__ge__(other)

This operator is not defined for Variable.

__nonzero__()

This operator is not defined for Variable.

__bool__()

This operator is not defined for Variable.

__neg__()

Element-wise negation.

Returns Output variable.

Return type *Variable*

__abs__()

Element-wise absolute.

Returns Output variable.

Return type *Variable*

__add__()

Element-wise addition.

Returns Output variable.

Return type *Variable*

__radd__()

Element-wise addition.

Returns Output variable.

Return type *Variable*

`__sub__(rhs)`

Element-wise subtraction.

Returns Output variable.

Return type *Variable*

`__rsub__(rhs)`

Element-wise subtraction.

Returns Output variable.

Return type *Variable*

`__mul__(rhs)`

Element-wise multiplication.

Returns Output variable.

Return type *Variable*

`__rmul__(rhs)`

Element-wise multiplication.

Returns Output variable.

Return type *Variable*

`__div__(rhs)`

Element-wise division

Returns Output variable.

Return type *Variable*

`__truediv__(rhs)`

Element-wise division

Returns Output variable.

Return type *Variable*

`__rdiv__(rhs)`

Element-wise division.

Returns Output variable.

Return type *Variable*

`__rtruediv__(rhs)`

Element-wise division.

Returns Output variable.

Return type *Variable*

`__floordiv__(rhs)`

Element-wise floor division.

Returns Output variable.

Return type *Variable*

`__rfloordiv__ (rhs)`

Element-wise floor division.

Returns Output variable.

Return type *Variable*

`__pow__ (rhs)`

Element-wise power function.

Returns Output variable.

Return type *Variable*

`__rpow__ (rhs)`

Element-wise power function.

Returns Output variable.

Return type *Variable*

`__matmul__ (rhs)`

Matrix multiplication.

Returns Output variable.

Return type *Variable*

`__rmatmul__ (rhs)`

Matrix multiplication.

Returns Output variable.

Return type *Variable*

Attributes

T

Transposition of this variable.

array

The underlying data array.

It is either `numpy.ndarray` or `cupy.ndarray` object, or `None` if the variable is in an uninitialized state.

creator

Function implementation that created this variable.

When this variable has been created by an old-style function (i.e., it is implemented as a subclass of *Function*), this property returns that *Function* object.

When this variable has been created by a new-style function (i.e., it is implemented as a subclass of *FunctionNode* class), this property returns that node object.

creator_node

FunctionNode object that created this variable.

This property has a setter to which `None` can be set. Setting `None` to this property is equivalent to call `unchain()`; it purges the variable from the function that created this variable.

The setter also accepts the original *FunctionNode* object that created this variable. For example, you can once set `None` to this property and then set the original value again.

Note: Setting an irrelevant `FunctionNode()` object does not emit any error immediately, whereas the behavior is undefined. Do not set a `FunctionNode()` object that did not create this variable object.

data

The underlying data array (equivalent to `array`).

Note that using this attribute directly is discouraged; use `array` instead. Using `array`, you can find an error earlier when your code mixes up `Variable` and `ndarray` because `ndarray` does not have an attribute `.array` while it has `.data`.

dtype**grad**

Gradient array of this variable.

Note that this property returns the underlying array of the gradient variable instead of the gradient variable itself; to get/set gradient variable, use `grad_var` instead.

grad_var

Gradient variable.

initializer = None

label

Short text that represents the variable.

name**ndim****node****rank****requires_grad**

It indicates that `grad` will be set in backward calculation.

shape**size****xp**

Array module for this variable.

Depending on which of CPU/GPU this variable is on, this property returns `numpy` or `cupy`.

5.1.4 `chainer.variable.VariableNode`

class `chainer.variable.VariableNode` (*variable*, *name*, ***kwargs*)

Node in the backward computational graph representing a variable.

This object represents a variable node in a computational graph. The node is used in error backpropagation (a.k.a. backprop) to determine which gradient to be passed to each function.

A variable node is held by the corresponding `Variable` object, which is managed by users. `FunctionNode` objects that take the variable as an input also hold references to the variable node.

Note that the node does not hold a reference to the corresponding data array in general. The data array is actually accessible by the node in the following cases.

1. If there exists a `Variable` object that holds a reference to the variable node, the variable node holds a weak reference to the variable object, and thus the data array is accessible via the weak reference.

2. If `retain_data()` is called, the node holds a reference to the data array. It is mainly called by a function that needs the input or output data array in its backprop procedure. See `FunctionNode.retain_inputs()` and `FunctionNode.retain_outputs()` for more details.

Users usually do not need to touch this variable node object. The computational graph is automatically managed by Chainer, and any interface that is beneficial for users is also provided by `Variable`.

Parameters

- **variable** (`Variable`) – The corresponding variable object.
- **name** (`str`) – Name of the variable node.

Variables

- **dtype** – Data type of the data array.
- **shape** – Shape of the data array.
- **name** (`str`) – Name of the variable node.

Methods

`get_variable()`

Returns the corresponding `Variable` object.

`VariableNode` object holds a weak reference of the variable object. If the reference is alive, it is returned by this property. Otherwise, this property creates a new `Variable` object from this node object and returns it.

Returns The variable object that refers this node.

Return type `Variable`

`get_variable_or_none()`

Returns the holding `Variable` object or `None`.

`VariableNode` object holds a weak reference of the variable object. If the reference is alive, it is returned by this property. Otherwise, returns `None`.

Returns The variable object that refers this node.

Return type `Variable`

`retain_data()`

Lets the node hold a reference to the underlying data array.

This method gets the data array of the corresponding variable and keeps it. If the weak reference to the corresponding variable is dead, it raises an error.

`set_creator(creator)`

Sets a `Function` object that created this node.

This method is equivalent to `self.creator = creator`. A `FunctionNode` object can also be passed.

Parameters **creator** (`Function` or `FunctionNode`) – Function that has created this variable.

`set_creator_node(creator_node)`

Sets a `FunctionNode` object that created this node.

This method is equivalent to `self.creator_node = creator_node`. A `Function` object can also be passed, in which case the `Function.node` attribute is used.

Parameters `creator_node` (`FunctionNode` or `Function`) – Function node that has this variable as an output.

`unchain()`

Deletes the reference to the creator of this variable node.

This method is equivalent to `self.creator_node = None`.

Attributes

`creator`

Function object that created this variable node.

When the function is implemented with the old-style API (i.e., it uses `Function` class), this property returns the `Function` object. The object is extracted from the `FunctionAdapter` object, so the returned object is not the function node, but instead the actual implementation of forward and backward procedures.

When the function is implemented with the new-style API (i.e., it uses `FunctionNode` class), this property returns the function node object. In this case, the returned object is same as `creator_node`.

Warning: As of v3.0.0, when the creator is an old-style function, the following code is invalid:

```
creator = v.creator
v.creator = None
...
v.creator = creator
```

The point is that `FunctionNode` objects are used as nodes in the computational graph instead of `Function`, and each `Function` object only holds a *weak reference* to the corresponding `FunctionNode`. Since `creator` returns the `Function` object, the `FunctionNode` object is not kept by preserving `creator`.

The above code should be fixed as follows.

```
creator_node = v.creator_node
v.creator_node = None
...
v.creator_node = creator_node
```

`creator_node`

Function node that has this variable as an output.

See `FunctionNode` for the definition of a function node.

`data`

Data array of the corresponding variable.

If the data is not available, it returns `None`.

`grad`

Gradient array of the corresponding variable.

If the variable is not available, it returns `None`.

`grad_var`

Gradient variable of the corresponding variable.

If the corresponding variable is not available, it return `None`.

label

Short text that represents the variable node.

rank**requires_grad**

It indicates that `grad` will be set in backward calculation.

5.2 Functions

Chainer provides variety of built-in function implementations in `chainer.functions` package. These functions return a `Variable` object or a tuple of multiple `Variable` objects.

Note: Functions implemented in Chainer consists of the following two parts:

- A class that inherits `FunctionNode`, which defines forward/backward computation.
- A “wrapper” function around the class.

APIs listed in this page are “wrapper” of `FunctionNode` implementations. In most cases, you don’t have to use `FunctionNode` classes directly.

For example, `chainer.functions.sum()` is a wrapper function defined as `def sum(...):` in `chainer/functions/math/sum.py`, and it calls its corresponding `FunctionNode` implementation, `Sum`. Some functions may not have the corresponding `FunctionNode` implementation; one example is `chainer.functions.average()`, which is defined in `chainer/functions/math/average.py`, which calls other wrapper functions to calculate average.

If you are implementing your own functions, please see *Define your own function*.

Note: As of v1.5, the concept of parameterized functions are gone, and they are replaced by corresponding `Link` implementations. They are found in the `chainer.links` namespace.

5.2.1 Arithmetic functions

Basic arithmetic operations for `Variables` are implemented as operators. Refer to the Notes section of `Variable` for details.

`chainer.functions.add()` provides better performance when accumulating three or more `Variables` at once.

`chainer.functions.add`

Element-wise addition.

`chainer.functions.add`

`chainer.functions.add(*xs)`

Element-wise addition.

Returns Output variable.

Return type `Variable`

5.2.2 Activation functions

<code>chainer.functions.clipped_relu</code>	Clipped Rectifier Unit function.
<code>chainer.functions.crelu</code>	Concatenated Rectified Linear Unit function.
<code>chainer.functions.elu</code>	Exponential Linear Unit function.
<code>chainer.functions.hard_sigmoid</code>	Element-wise hard-sigmoid function.
<code>chainer.functions.leaky_relu</code>	Leaky Rectified Linear Unit function.
<code>chainer.functions.log_softmax</code>	Channel-wise log-softmax function.
<code>chainer.functions.lstm</code>	Long Short-Term Memory units as an activation function.
<code>chainer.functions.maxout</code>	Maxout activation function.
<code>chainer.functions.prelu</code>	Parametric ReLU function.
<code>chainer.functions.relu</code>	Rectified Linear Unit function.
<code>chainer.functions.selu</code>	Scaled Exponential Linear Unit function.
<code>chainer.functions.sigmoid</code>	Element-wise sigmoid logistic function.
<code>chainer.functions.slstm</code>	S-LSTM units as an activation function.
<code>chainer.functions.softmax</code>	Softmax function.
<code>chainer.functions.softplus</code>	Element-wise softplus function.
<code>chainer.functions.swish</code>	Swish activation function.
<code>chainer.functions.tanh</code>	Elementwise hyperbolic tangent function.
<code>chainer.functions.tree_lstm</code>	TreeLSTM unit as an activation function.

`chainer.functions.clipped_relu`

`chainer.functions.clipped_relu(x, z=20.0)`

Clipped Rectifier Unit function.

For a clipping value $z(> 0)$, it computes

$$\text{ClippedReLU}(x, z) = \min(\max(0, x), z).$$

Parameters

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable. A (s_1, s_2, \dots, s_n) -shaped float array.
- **z** (*float*) – Clipping value. (default = 20.0)

Returns Output variable. A (s_1, s_2, \dots, s_n) -shaped float array.

Return type *Variable*

Example

```
>>> x = np.random.uniform(-100, 100, (10, 20)).astype(np.float32)
>>> z = 10.0
>>> np.any(x < 0)
True
>>> np.any(x > z)
True
>>> y = F.clipped_relu(x, z=z)
>>> np.any(y.data < 0)
False
>>> np.any(y.data > z)
False
```

chainer.functions.crelu

`chainer.functions.crelu(x, axis=1)`
 Concatenated Rectified Linear Unit function.

This function is expressed as follows

$$f(x) = (\max(0, x), \max(0, -x)).$$

Here, two output values are concatenated along an axis.

See: <https://arxiv.org/abs/1603.05201>

Parameters

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable. A (s_1, s_2, \dots, s_N) -shaped float array.
- **axis** (*int*) – Axis that the output values are concatenated along. Default is 1.

Returns Output variable of concatenated array. If the axis is 1, A $(s_1, s_2 \times 2, \dots, s_N)$ -shaped float array.

Return type *Variable*

Example

```
>>> x = np.array([[ -1,  0], [ 2, -3]], np.float32)
>>> x
array([[ -1.,  0.],
       [  2., -3.]], dtype=float32)
>>> y = F.crelu(x, axis=1)
>>> y.data
array([[0.,  0.,  1.,  0.],
       [2.,  0.,  0.,  3.]], dtype=float32)
```

chainer.functions.elu

`chainer.functions.elu(x, alpha=1.0)`
 Exponential Linear Unit function.

For a parameter α , it is expressed as

$$f(x) = \begin{cases} x & \text{if } x \geq 0 \\ \alpha(\exp(x) - 1) & \text{if } x < 0, \end{cases}$$

See: <https://arxiv.org/abs/1511.07289>

Parameters

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable. A (s_1, s_2, \dots, s_N) -shaped float array.
- **alpha** (*float*) – Parameter α . Default is 1.0.

Returns Output variable. A (s_1, s_2, \dots, s_N) -shaped float array.

Return type *Variable*

Example

```
>>> x = np.array([[ -1,  0], [ 2, -3]], np.float32)
>>> x
array([[ -1.,  0.],
       [  2., -3.]], dtype=float32)
>>> y = F.elu(x, alpha=1.)
>>> y.data
array([[ -0.63212055,  0.          ],
       [  2.          , -0.95021296]], dtype=float32)
```

chainer.functions.hard_sigmoid

`chainer.functions.hard_sigmoid(x)`

Element-wise hard-sigmoid function.

This function is defined as

$$f(x) = \begin{cases} 0 & \text{if } x < -2.5 \\ 0.2x + 0.5 & \text{if } -2.5 < x < 2.5 \\ 1 & \text{if } 2.5 < x. \end{cases}$$

Parameters *x* (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable. A (s_1, s_2, \dots, s_N) -shaped float array.

Returns Output variable. A (s_1, s_2, \dots, s_N) -shaped float array.

Return type *Variable*

Example

It maps the input values into the range of $[0, 1]$.

```
>>> x = np.array([-2.6, -1, 0, 1, 2.6])
>>> x
array([-2.6, -1. ,  0. ,  1. ,  2.6])
>>> F.hard_sigmoid(x).data
array([0. , 0.3, 0.5, 0.7, 1. ])
```

chainer.functions.leaky_relu

`chainer.functions.leaky_relu(x, slope=0.2)`

Leaky Rectified Linear Unit function.

This function is expressed as

$$f(x) = \begin{cases} x & \text{if } x \geq 0 \\ ax & \text{if } x < 0, \end{cases}$$

where a is a configurable slope value.

Parameters

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable. A (s_1, s_2, \dots, s_N) -shaped float array.
- **slope** (*float*) – Slope value a .

Returns Output variable. A (s_1, s_2, \dots, s_N) -shaped float array.

Return type *Variable*

Example

```
>>> x = np.array([[ -1,  0], [ 2, -3], [-2,  1]], np.float32)
>>> x
array([[ -1.,  0.],
       [  2., -3.],
       [-2.,  1.]], dtype=float32)
>>> F.leaky_relu(x, slope=0.2).data
array([[ -0.2,  0. ],
       [  2. , -0.6],
       [-0.4,  1. ]], dtype=float32)
```

chainer.functions.log_softmax

`chainer.functions.log_softmax(x)`

Channel-wise log-softmax function.

This function computes its logarithm of softmax along the second axis. Let $c = (c_1, c_2, \dots, c_D)$ be the slice of x along with the second axis. For each slice c , it computes the logarithm of the function $f(c)$ defined as

$$f(c) = \frac{\exp(c)}{\sum_d \exp(c_d)}.$$

This method is theoretically equivalent to `log(softmax(x))` but is more stable.

Note: `log(softmax(x))` may cause underflow when x is too small, because `softmax(x)` may returns 0. `log_softmax` method is more stable.

Parameters **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable. A n -dimensional ($n \geq 2$) float array.

Returns Output variable. A n -dimensional ($n \geq 2$) float array, which is the same shape with x .

Return type *Variable*

See also:

`softmax()`

Example

```
>>> x = np.array([[0, 1, 2], [0, 2, 4]], np.float32)
>>> x
array([[0., 1., 2.],
       [0., 2., 4.]], dtype=float32)
```

(continues on next page)

(continued from previous page)

```
>>> F.log_softmax(x).data
array([[ -2.407606   ,  -1.4076059 ,  -0.4076059 ],
       [-4.1429315 ,  -2.1429315 ,  -0.14293146]], dtype=float32)
>>> np.allclose(F.log_softmax(x).data, F.log(F.softmax(x)).data)
True
```

chainer.functions.lstm

`chainer.functions.lstm(c_prev, x)`

Long Short-Term Memory units as an activation function.

This function implements LSTM units with forget gates. Let the previous cell state `c_prev` and the input array `x`.

First, the input array `x` is split into four arrays `a, i, f, o` of the same shapes along the second axis. It means that `x` 's second axis must have 4 times the `c_prev` 's second axis.

The split input arrays are corresponding to:

- `a` : sources of cell input
- `i` : sources of input gate
- `f` : sources of forget gate
- `o` : sources of output gate

Second, it computes the updated cell state `c` and the outgoing signal `h` as:

$$c = \tanh(a)\sigma(i) + c_{\text{prev}}\sigma(f),$$

$$h = \tanh(c)\sigma(o),$$

where σ is the elementwise sigmoid function. These are returned as a tuple of two variables.

This function supports variable length inputs. The mini-batch size of the current input must be equal to or smaller than that of the previous one. When mini-batch size of `x` is smaller than that of `c`, this function only updates `c[0:len(x)]` and doesn't change the rest of `c`, `c[len(x):]`. So, please sort input sequences in descending order of lengths before applying the function.

Parameters

- **`c_prev`** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Variable that holds the previous cell state. The cell state should be a zero array or the output of the previous call of LSTM.
- **`x`** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Variable that holds the sources of cell input, input gate, forget gate and output gate. It must have the second dimension whose size is four times of that of the cell state.

Returns Two *Variable* objects `c` and `h`. `c` is the updated cell state. `h` indicates the outgoing signal.

Return type `tuple`

See the original paper proposing LSTM with forget gates: [Long Short-Term Memory in Recurrent Neural Networks](#).

See also:

LSTM

Example

Assuming y is the current incoming signal, c is the previous cell state, and h is the previous outgoing signal from an `lstm` function. Each of y , c and h has `n_units` channels. Most typical preparation of x is:

```
>>> n_units = 100
>>> y = chainer.Variable(np.zeros((1, n_units), np.float32))
>>> h = chainer.Variable(np.zeros((1, n_units), np.float32))
>>> c = chainer.Variable(np.zeros((1, n_units), np.float32))
>>> model = chainer.Chain()
>>> with model.init_scope():
...     model.w = L.Linear(n_units, 4 * n_units)
...     model.v = L.Linear(n_units, 4 * n_units)
>>> x = model.w(y) + model.v(h)
>>> c, h = F.lstm(c, x)
```

It corresponds to calculate the input array x , or the input sources a, i, f, o , from the current incoming signal y and the previous outgoing signal h . Different parameters are used for different kind of input sources.

Note: We use the naming rule below.

- **incoming signal** The formal input of the formulation of LSTM (e.g. in NLP, word vector or output of lower RNN layer). The input of `chainer.links.LSTM` is the *incoming signal*.
 - **input array** The array which is linear transformed from *incoming signal* and the previous outgoing signal. The *input array* contains four sources, the sources of cell input, input gate, forget gate and output gate. The input of `chainer.functions.LSTM` is the *input array*.
-

chainer.functions.maxout

`chainer.functions.maxout(x, pool_size, axis=1)`

Maxout activation function.

It accepts an input tensor x , reshapes the axis dimension (say the size being $M * \text{pool_size}$) into two dimensions ($M, \text{pool_size}$), and takes maximum along the axis dimension.

Parameters

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable. A n -dimensional ($n \geq \text{axis}$) float array. In general, its first dimension is assumed to be the *minibatch dimension*. The other dimensions are treated as one concatenated dimension.
- **`pool_size`** (*int*) – The size used for downsampling of pooling layer.
- **`axis`** (*int*) – The axis dimension to be reshaped. The size of axis dimension should be $M * \text{pool_size}$.

Returns Output variable. The shape of the output is same as x except that axis dimension is transformed from $M * \text{pool_size}$ to M .

Return type *Variable*

See also:

Maxout

Example

Typically, `x` is the output of a linear layer or a convolution layer. The following is the example where we use `maxout()` in combination with a Linear link.

```
>>> in_size, out_size, pool_size = 10, 10, 10
>>> bias = np.arange(out_size * pool_size).astype(np.float32)
>>> l = L.Linear(in_size, out_size * pool_size, initial_bias=bias)
>>> x = np.zeros((1, in_size), np.float32) # prepare data
>>> x = l(x)
>>> y = F.maxout(x, pool_size)
>>> x.shape
(1, 100)
>>> y.shape
(1, 10)
>>> x.reshape((out_size, pool_size)).data
array([[ 0.,  1.,  2.,  3.,  4.,  5.,  6.,  7.,  8.,  9.],
       [10., 11., 12., 13., 14., 15., 16., 17., 18., 19.],
       [20., 21., 22., 23., 24., 25., 26., 27., 28., 29.],
       [30., 31., 32., 33., 34., 35., 36., 37., 38., 39.],
       [40., 41., 42., 43., 44., 45., 46., 47., 48., 49.],
       [50., 51., 52., 53., 54., 55., 56., 57., 58., 59.],
       [60., 61., 62., 63., 64., 65., 66., 67., 68., 69.],
       [70., 71., 72., 73., 74., 75., 76., 77., 78., 79.],
       [80., 81., 82., 83., 84., 85., 86., 87., 88., 89.],
       [90., 91., 92., 93., 94., 95., 96., 97., 98., 99.]], dtype=float32)
>>> y.data
array([[ 9., 19., 29., 39., 49., 59., 69., 79., 89., 99.]], dtype=float32)
```

chainer.functions.prelu

`chainer.functions.prelu(x, W)`

Parametric ReLU function.

It accepts two arguments: an input `x` and a weight array `W` and computes the output as $PReLU(x) = \max(x, W * x)$, where $*$ is an elementwise multiplication for each sample in the batch.

When the PReLU function is combined with two-dimensional convolution, the elements of parameter `W` are typically shared across the same filter of different pixels. In order to support such usage, this function supports the shape of parameter array that indicates leading dimensions of input arrays except the batch dimension.

For example, if `W` has the shape of $(2, 3, 4)$, `x` must have the shape of $(B, 2, 3, 4, S_1, \dots, S_N)$ where `B` is the batch size and the number of trailing `S`'s `N` is an arbitrary non-negative integer.

Parameters

- **x** (*Variable*) – Input variable. Its first argument is assumed to be the minibatch dimension.
- **W** (*Variable*) – Weight variable.

Returns Output variable

Return type *Variable*

See also:

PReLU

chainer.functions.relu

`chainer.functions.relu(x)`

Rectified Linear Unit function.

$$f(x) = \max(0, x).$$

Parameters **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable. A (s_1, s_2, \dots, s_N) -shaped float array.

Returns Output variable. A (s_1, s_2, \dots, s_N) -shaped float array.

Return type *Variable*

Example

```
>>> x = np.array([[ -1,  0], [ 2, -3], [-2,  1]], np.float32)
>>> np.any(x < 0)
True
>>> y = F.relu(x)
>>> np.any(y.data < 0)
False
>>> y.shape
(3, 2)
```

chainer.functions.selu

`chainer.functions.selu(x, alpha=1.6732632423543772, scale=1.0507009873554805)`

Scaled Exponential Linear Unit function.

For parameters α and λ , it is expressed as

$$f(x) = \lambda \begin{cases} x & \text{if } x \geq 0 \\ \alpha(\exp(x) - 1) & \text{if } x < 0, \end{cases}$$

See: <https://arxiv.org/abs/1706.02515>

Parameters

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable. A (s_1, s_2, \dots, s_N) -shaped float array.
- **alpha** (*float*) – Parameter α .
- **scale** (*float*) – Parameter λ .

Returns Output variable. A (s_1, s_2, \dots, s_N) -shaped float array.

Return type *Variable*

chainer.functions.sigmoid

`chainer.functions.sigmoid(x)`

Element-wise sigmoid logistic function.

$$f(x) = (1 + \exp(-x))^{-1}.$$

Parameters **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable. A (s_1, s_2, \dots, s_N) -shaped float array.

Returns Output variable. A (s_1, s_2, \dots, s_N) -shaped float array.

Return type *Variable*

Example

It maps the input values into the range of $[0, 1]$.

```
>>> x = np.arange(-2, 3, 2).astype(np.float32)
>>> x
array([-2.,  0.,  2.], dtype=float32)
>>> F.sigmoid(x)
variable([0.11920291, 0.5          , 0.8807971  ])
```

chainer.functions.slstm

`chainer.functions.slstm(c_prev1, c_prev2, x1, x2)`

S-LSTM units as an activation function.

This function implements S-LSTM unit. It is an extension of LSTM unit applied to tree structures. The function is applied to binary trees. Each node has two child nodes. It gets four arguments, previous cell states `c_prev1` and `c_prev2`, and input arrays `x1` and `x2`.

First both input arrays `x1` and `x2` are split into eight arrays a_1, i_1, f_1, o_1 , and a_2, i_2, f_2, o_2 . They have the same shape along the second axis. It means that `x1` and `x2` 's second axis must have 4 times the length of `c_prev1` and `c_prev2`.

The split input arrays are corresponding to:

- a_i : sources of cell input
- i_i : sources of input gate
- f_i : sources of forget gate
- o_i : sources of output gate

It computes the updated cell state `c` and the outgoing signal `h` as:

$$\begin{aligned} c &= \tanh(a_1 + a_2)\sigma(i_1 + i_2) + c_{\text{prev1}}\sigma(f_1) + c_{\text{prev2}}\sigma(f_2), \\ h &= \tanh(c)\sigma(o_1 + o_2), \end{aligned}$$

where σ is the elementwise sigmoid function. The function returns `c` and `h` as a tuple.

Parameters

- **c_prev1** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Variable that holds the previous cell state of the first child node. The cell state should be a zero array or the output of the previous call of LSTM.
- **c_prev2** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Variable that holds the previous cell state of the second child node.
- **x1** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Variable that holds the sources of cell input, input gate, forget gate and output gate from the first child node. It must have the second dimension whose size is four times of that of the cell state.

- **x2** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Variable that holds the input sources from the second child node.

Returns Two *Variable* objects *c* and *h*. *c* is the cell state. *h* indicates the outgoing signal.

Return type *tuple*

See detail in paper: [Long Short-Term Memory Over Tree Structures](#).

Example

Assuming *c1*, *c2* is the previous cell state of children, and *h1*, *h2* is the previous outgoing signal from children. Each of *c1*, *c2*, *h1* and *h2* has *n_units* channels. Most typical preparation of *x1*, *x2* is:

```
>>> n_units = 100
>>> h1 = chainer.Variable(np.zeros((1, n_units), np.float32))
>>> h2 = chainer.Variable(np.zeros((1, n_units), np.float32))
>>> c1 = chainer.Variable(np.zeros((1, n_units), np.float32))
>>> c2 = chainer.Variable(np.zeros((1, n_units), np.float32))
>>> model1 = chainer.Chain()
>>> with model1.init_scope():
...     model1.w = L.Linear(n_units, 4 * n_units)
...     model1.v = L.Linear(n_units, 4 * n_units)
>>> model2 = chainer.Chain()
>>> with model2.init_scope():
...     model2.w = L.Linear(n_units, 4 * n_units)
...     model2.v = L.Linear(n_units, 4 * n_units)
>>> x1 = model1.w(c1) + model1.v(h1)
>>> x2 = model2.w(c2) + model2.v(h2)
>>> c, h = F.slstm(c1, c2, x1, x2)
```

It corresponds to calculate the input array *x1*, or the input sources a_1, i_1, f_1, o_1 from the previous cell state of first child node *c1*, and the previous outgoing signal from first child node *h1*. Different parameters are used for different kind of input sources.

chainer.functions.softmax

`chainer.functions.softmax(x, axis=1)`

Softmax function.

This function computes its softmax along an axis. Let $c = (c_1, c_2, \dots, c_D)$ be the slice of *x* along with the axis. For each slice *c*, it computes the function $f(c)$ defined as $f(c) = \frac{\exp(c)}{\sum_d \exp(c_d)}$.

Parameters

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable. A *n*-dimensional ($n \geq 2$) float array.
- **axis** (*int*) – The axis along which the softmax is to be computed.

Returns Output variable. A *n*-dimensional ($n \geq 2$) float array, which is the same shape with *x*.

Return type *Variable*

Example

```
>>> x = np.array([[0, 1, 2], [0, 2, 4]], np.float32)
>>> x
array([[0., 1., 2.],
       [0., 2., 4.]], dtype=float32)
>>> y = F.softmax(x, axis=1)
>>> y.data
array([[0.09003057, 0.24472848, 0.66524094],
       [0.01587624, 0.11731043, 0.86681336]], dtype=float32)
>>> F.sum(y, axis=1).data
array([1., 1.], dtype=float32)
```

chainer.functions.softplus

chainer.functions.**softplus**(*x*, *beta*=1.0)

Element-wise softplus function.

The softplus function is the smooth approximation of ReLU.

$$f(x) = \frac{1}{\beta} \log(1 + \exp(\beta x)),$$

where β is a parameter. The function becomes curved and akin to ReLU as the β is increasing.

Parameters

- **x** (*Variable* or *numpy.ndarray* or *cupy.ndarray*) – Input variable. A (s_1, s_2, \dots, s_N) -shaped float array.
- **beta** (*float*) – Parameter β .

Returns Output variable. A (s_1, s_2, \dots, s_N) -shaped float array.

Return type *Variable*

Example

```
>>> x = np.arange(-2, 3, 2).astype(np.float32)
>>> x
array([-2.,  0.,  2.], dtype=float32)
>>> F.softplus(x, beta=1.0).data
array([0.126928,  0.6931472, 2.126928 ], dtype=float32)
```

chainer.functions.swish

chainer.functions.**swish**(*x*, *beta*)

Swish activation function.

$$f(x, \beta) = x \cdot \sigma(\beta x),$$

where $\sigma(\cdot)$ is the sigmoid function. It has the following properties:

$$f(x, 0) = \frac{x}{2},$$
$$\lim_{\beta \rightarrow \infty} f(x, \beta) = \max(0, x).$$

Parameters

- **x** (*Variable*) – Input variable of shape $(s_B, s_1, s_2, \dots, s_N)$, where s_B is assumed to be the *minibatch dimension*.
- **beta** (*Variable*) – Parameter variable β of shape (s_1, s_2, \dots, s_M) , where M is an arbitrary integer between $0 \leq M \leq N$. The number of dimensions of **beta** will be matched with **x** by reshaping it as $(1, s_1, \dots, s_M, 1, \dots, 1)$, then **beta** and **x** are multiplied together in an element-wise manner.

Returns Output variable of the same shape as **x**.

Return type *Variable*

See also:

`chainer.links.Swish`

chainer.functions.tanh

`chainer.functions.tanh(x)`

Elementwise hyperbolic tangent function.

$$f(x) = \tanh(x).$$

Parameters **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable. A (s_1, s_2, \dots, s_N) -shaped float array.

Returns Output variable. A (s_1, s_2, \dots, s_N) -shaped float array.

Return type *Variable*

Example

```
>>> x = np.arange(-1, 4, 2).astype(np.float32)
>>> x
array([-1.,  1.,  3.], dtype=float32)
>>> F.tanh(x).data
array([-0.7615942,  0.7615942,  0.9950548], dtype=float32)
```

chainer.functions.tree_lstm

`chainer.functions.tree_lstm(*inputs)`

TreeLSTM unit as an activation function.

This function implements TreeLSTM units both for N-ary TreeLSTM and Child-Sum TreeLSTM. Let the children cell states c_1, c_2, \dots, c_N , and the incoming signal x .

First, the incoming signal x is split into $(3 + N)$ arrays $a, i, o, f_1, f_2, \dots, f_N$ of the same shapes along the second axis. It means that x 's second axis must have $(3 + N)$ times of the length of each c_n .

The splitted input signals are corresponding to:

- a : sources of cell input
- i : sources of input gate

- o : sources of output gate
- f_n : sources of forget gate for n-th ary

Second, it computes outputs as:

$$\begin{aligned}c &= \tanh(a)\text{sigmoid}(i) \\ &+ c_1\text{sigmoid}(f_1), \\ &+ c_2\text{sigmoid}(f_2), \\ &+ \dots, \\ &+ c_N\text{sigmoid}(f_N), \\ h &= \tanh(c)\text{sigmoid}(o).\end{aligned}$$

These are returned as a tuple of $(N + 1)$ variables.

Parameters **inputs** (list of *Variable*) – Variable arguments which include all cell vectors from child-nodes, and an input vector. Each of the cell vectors and the input vector is *Variable*. The input vector must have the second dimension whose size is $(N + 3)$ times of that of each cell, where N denotes the total number of cells.

Returns Two *Variable* objects c and h . c is the updated cell state. h indicates the outgoing signal.

Return type tuple

See the papers for details: [Improved Semantic Representations From Tree-Structured Long Short-Term Memory Networks](#) and [A Fast Unified Model for Parsing and Sentence Understanding](#).

Tai et al.'s N-Ary TreeLSTM is little extended in Bowman et al., and this link is based on the variant by Bowman et al. Specifically, eq. 10 in Tai et al. only has one W matrix to be applied to x , consistently for all children. On the other hand, Bowman et al.'s model has multiple matrices, each of which affects the forget gate for each child's cell individually.

Example

Assuming y is the current input signal, c is the previous cell state, and h is the previous output signal from an `tree_lstm()` function. Each of y , c and h has `n_units` channels. Using 2-ary (binary) TreeLSTM, most typical preparation of x is:

```
>>> model = chainer.Chain()
>>> with model.init_scope():
...     model.w = L.Linear(10, 5 * 10)
...     model.v1 = L.Linear(10, 5 * 10)
...     model.v2 = L.Linear(10, 5 * 10)
>>> y = np.random.uniform(-1, 1, (4, 10)).astype(np.float32)
>>> h1 = np.random.uniform(-1, 1, (4, 10)).astype(np.float32)
>>> h2 = np.random.uniform(-1, 1, (4, 10)).astype(np.float32)
>>> c1 = np.random.uniform(-1, 1, (4, 10)).astype(np.float32)
>>> c2 = np.random.uniform(-1, 1, (4, 10)).astype(np.float32)
>>> x = model.w(y) + model.v1(h1) + model.v2(h2)
>>> c, h = F.tree_lstm(c1, c2, x)
```

It corresponds to calculate the input sources a, i, o, f_1, f_2 from the current input y and the children's outputs $h1$ and $h2$. Different parameters are used for different kind of input sources.

5.2.3 Array manipulations

<code>chainer.functions.broadcast</code>	Broadcast given variables.
<code>chainer.functions.broadcast_to</code>	Broadcast a given variable to a given shape.
<code>chainer.functions.cast</code>	Cast an input variable to a given type.
<code>chainer.functions.concat</code>	Concatenates given variables along an axis.
<code>chainer.functions.copy</code>	Copies the input variable onto the specified device.
<code>chainer.functions.depth2space</code>	Computes the depth2space transformation for subpixel calculations.
<code>chainer.functions.dstack</code>	Concatenate variables along third axis (depth wise).
<code>chainer.functions.expand_dims</code>	Expands dimensions of an input variable without copy.
<code>chainer.functions.flatten</code>	Flatten a given array into one dimension.
<code>chainer.functions.flip</code>	Flips an input variable in reverse order along the given axis.
<code>chainer.functions.fliplr</code>	Flip array in the left/right direction.
<code>chainer.functions.flipud</code>	Flip array in the up/down direction.
<code>chainer.functions.get_item</code>	Extract elements from array with specified shape, axes and offsets.
<code>chainer.functions.hstack</code>	Concatenate variables horizontally (column wise).
<code>chainer.functions.im2col</code>	Extract patches from an image based on the filter.
<code>chainer.functions.pad</code>	Pad an input variable.
<code>chainer.functions.pad_sequence</code>	Pad given arrays to make a matrix.
<code>chainer.functions.permutate</code>	Permutates a given variable along an axis.
<code>chainer.functions.repeat</code>	Construct an array by repeating a given array.
<code>chainer.functions.reshape</code>	Reshapes an input variable without copy.
<code>chainer.functions.resize_images</code>	Resize images to the given shape.
<code>chainer.functions.rollaxis</code>	Roll the axis backwards to the given position.
<code>chainer.functions.scatter_add</code>	Adds given values to specified elements of an array.
<code>chainer.functions.select_item</code>	Select elements stored in given indices.
<code>chainer.functions.separate</code>	Separates an array along a given axis.
<code>chainer.functions.space2depth</code>	Computes the space2depth transformation for subpixel calculations.
<code>chainer.functions.spatial_transformer_grid</code>	2D Spatial Transformer grid.
<code>chainer.functions.spatial_transformer_sampler</code>	2D Spatial Transformer sampler.
<code>chainer.functions.split_axis</code>	Splits given variables along an axis.
<code>chainer.functions.squeeze</code>	Remove demensions of size one from the shape of a ndarray.
<code>chainer.functions.stack</code>	Concatenate variables along a new axis.
<code>chainer.functions.swapaxes</code>	Swap two axes of a variable.
<code>chainer.functions.tile</code>	Construct an array by tiling a given array.
<code>chainer.functions.transpose</code>	Permute the dimensions of an input variable without copy.
<code>chainer.functions.transpose_sequence</code>	Transpose a list of Variables.
<code>chainer.functions.vstack</code>	Concatenate variables vertically (row wise).
<code>chainer.functions.where</code>	Choose elements depending on condition.

chainer.functions.broadcast

`chainer.functions.broadcast(*args)`
Broadcast given variables.

Parameters **args** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variables to be broadcasted. Each dimension of the shapes of the input variables must have the same size.

Returns *Variable* or tuple of *Variable* objects which are broadcasted from given arguments.

Return type *Variable*

Example

```
>>> x = np.random.uniform(0, 1, (3, 2)).astype(np.float32)
>>> y = F.broadcast(x)
>>> np.all(x == y.data)
True
>>> z = np.random.uniform(0, 1, (3, 2)).astype(np.float32)
>>> y, w = F.broadcast(x, z)
>>> np.all(x == y.data) & np.all(z == w.data)
True
```

chainer.functions.broadcast_to

`chainer.functions.broadcast_to(x, shape)`

Broadcast a given variable to a given shape.

Parameters

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable be broadcasted. A (s_1, s_2, \dots, s_N) -shaped float array.
- **shape** (*tuple*) – Tuple of `int` of the shape of the output variable.

Returns Output variable broadcasted to the given shape.

Return type *Variable*

Example

```
>>> x = np.arange(0, 3)
>>> x
array([0, 1, 2])
>>> y = F.broadcast_to(x, (3, 3))
>>> y.data
array([[0, 1, 2],
       [0, 1, 2],
       [0, 1, 2]])
```

chainer.functions.cast

`chainer.functions.cast(x, typ)`

Cast an input variable to a given type.

Parameters

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable to be casted. A (s_1, s_2, \dots, s_N) -shaped float array.

- **typ** (str of dtype or `numpy.dtype`) – Typecode or data type to cast.

Returns Variable holding a casted array.

Return type *Variable*

Example

```
>>> x = np.arange(0, 3, dtype=np.float64)
>>> x.dtype
dtype('float64')
>>> y = F.cast(x, np.float32)
>>> y.dtype
dtype('float32')
>>> y = F.cast(x, 'float16')
>>> y.dtype
dtype('float16')
```

chainer.functions.concat

`chainer.functions.concat(xs, axis=1)`

Concatenates given variables along an axis.

Parameters

- **xs** (tuple of *Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variables to be concatenated. The variables must have the same shape, except in the dimension corresponding to axis.
- **axis** (*int*) – The axis along which the arrays will be joined. Default is 1.

Returns The concatenated variable.

Return type *Variable*

Example

```
>>> x = np.arange(0, 12).reshape(3, 4)
>>> x
array([[ 0,  1,  2,  3],
       [ 4,  5,  6,  7],
       [ 8,  9, 10, 11]])
>>> y = np.arange(0, 3).reshape(3, 1)
>>> y
array([[0],
       [1],
       [2]])
>>> z = F.concat((x, y), axis=1)
>>> z.data
array([[ 0,  1,  2,  3,  0],
       [ 4,  5,  6,  7,  1],
       [ 8,  9, 10, 11,  2]])
```

chainer.functions.copy

`chainer.functions.copy(x, dst)`

Copies the input variable onto the specified device.

This function copies the array of input variable onto the device specified by `dst`. When `dst == -1`, it copies the array onto the host memory. This function supports copies from host to host, from host to device, from device to device and from device to host.

Parameters

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Variable to be copied.
- **dst** (*int*) – Target device specifier.

Returns Output variable.

Return type *Variable*

Example

```
>>> import chainer.backends.cuda as cuda
>>> x = np.random.uniform(-1, 1, (5, 10))
>>> cuda.get_device_from_array(x).id
-1
>>> y = F.copy(x, 0) # from host to device0
>>> cuda.get_device_from_array(y.data).id
0
>>> z = F.copy(y, -1) # from device0 to host
>>> cuda.get_device_from_array(z.data).id
-1
```

chainer.functions.depth2space

`chainer.functions.depth2space(X, r)`

Computes the depth2space transformation for subpixel calculations.

Parameters

- **X** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Variable holding a 4d array of shape (batch, channel * r * r, dim1, dim2).
- **r** (*int*) – the upscaling factor.

Returns A variable holding the upscaled array from interspersed depth layers. The shape is (batch, channel, dim1 * r, dim2 * r).

Return type *Variable*

Note: This can be used to compute super-resolution transformations. See <https://arxiv.org/abs/1609.05158> for details.

See also:

space2depth()

Example

```

>>> X = np.arange(24).reshape(1, 4, 2, 3).astype(np.float32)
>>> X.shape
(1, 4, 2, 3)
>>> X
array([[[[ 0.,  1.,  2.],
          [ 3.,  4.,  5.]],

        [[ 6.,  7.,  8.],
          [ 9., 10., 11.]],

        [[12., 13., 14.],
          [15., 16., 17.]],

        [[18., 19., 20.],
          [21., 22., 23.]]], dtype=float32)
>>> y = F.depth2space(X, 2)
>>> y.shape
(1, 1, 4, 6)
>>> y.data
array([[[[ 0.,  6.,  1.,  7.,  2.,  8.],
          [12., 18., 13., 19., 14., 20.],
          [ 3.,  9.,  4., 10.,  5., 11.],
          [15., 21., 16., 22., 17., 23.]]], dtype=float32)

```

chainer.functions.dstack

`chainer.functions.dstack(xs)`

Concatenate variables along third axis (depth wise).

Parameters *xs* (list of *Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variables to be concatenated. The variables must have the same *ndim*. When the variables have the third axis (i.e. $ndim \geq 3$), the variables must have the same shape along all but the third axis. When the variables do not have the third axis (i.e. $ndim < 3$), the variables must have the same shape.

Returns Output variable. When the input variables have the third axis (i.e. $ndim \geq 3$), the shapes of inputs and output are the same along all but the third axis. The length of third axis is the sum of the lengths of inputs' third axis. When the shape of variables are $(N1, N2)$ (i.e. $ndim = 2$), the shape of output is $(N1, N2, 2)$. When the shape of variables are $(N1,)$ (i.e. $ndim = 1$), the shape of output is $(1, N1, 2)$. When the shape of variables are $()$ (i.e. $ndim = 0$), the shape of output is $(1, 1, 2)$.

Return type *Variable*

Example

```

>>> x1 = np.array((1, 2, 3))
>>> x1.shape
(3,)
>>> x2 = np.array((2, 3, 4))
>>> x2.shape
(3,)
>>> y = F.dstack((x1, x2))
>>> y.shape
(1, 3, 2)

```

(continues on next page)

(continued from previous page)

```
>>> y.data
array([[1, 2],
       [2, 3],
       [3, 4]])
```

```
>>> x1 = np.arange(0, 6).reshape(3, 2)
>>> x1.shape
(3, 2)
>>> x1
array([[0, 1],
       [2, 3],
       [4, 5]])
>>> x2 = np.arange(6, 12).reshape(3, 2)
>>> x2.shape
(3, 2)
>>> x2
array([[ 6,  7],
       [ 8,  9],
       [10, 11]])
>>> y = F.dstack([x1, x2])
>>> y.shape
(3, 2, 2)
>>> y.data
array([[[ 0,  6],
        [ 1,  7]],

       [[ 2,  8],
        [ 3,  9]],

       [[ 4, 10],
        [ 5, 11]])])
```

```
>>> x1 = np.arange(0, 12).reshape(3, 2, 2)
>>> x2 = np.arange(12, 18).reshape(3, 2, 1)
>>> y = F.dstack([x1, x2])
>>> y.shape
(3, 2, 3)
>>> y.data
array([[[ 0,  1, 12],
        [ 2,  3, 13]],

       [[ 4,  5, 14],
        [ 6,  7, 15]],

       [[ 8,  9, 16],
        [10, 11, 17]])])
```

chainer.functions.expand_dims

chainer.functions.expand_dims(*x*, *axis*)

Expands dimensions of an input variable without copy.

Parameters

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable.

- **axis** (*int*) – Position where new axis is to be inserted. The `axis` parameter is acceptable when $-ndim - 1 \leq axis \leq ndim$. (`ndim` is the dimension of input variables). When $axis < 0$, the result is the same with $ndim + 1 - |axis|$.

Returns Variable that holds a expanded input. The `ndim` of output is one grater than that of `x`.

Return type *Variable*

Example

```
>>> x = np.array([1, 2, 3])
>>> x.shape
(3,)
>>> y = F.expand_dims(x, axis=0)
>>> y.shape
(1, 3)
>>> y.data
array([[1, 2, 3]])
>>> y = F.expand_dims(x, axis=1)
>>> y.shape
(3, 1)
>>> y.data
array([[1],
       [2],
       [3]])
>>> y = F.expand_dims(x, axis=-2)
>>> y.shape
(1, 3)
>>> y.data
array([[1, 2, 3]])
```

chainer.functions.flatten

`chainer.functions.flatten(x)`

Flatten a given array into one dimension.

Parameters `x` (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable.

Returns Output variable flatten to one dimension.

Return type *Variable*

Note: When you input a scalar array (i.e. the shape is `()`), you can also get the one dimension array whose shape is `(1,)`.

Example

```
>>> x = np.array([[1, 2], [3, 4]])
>>> x.shape
(2, 2)
>>> y = F.flatten(x)
>>> y.shape
(4,)
```

(continues on next page)

(continued from previous page)

```
>>> y.data
array([1, 2, 3, 4])
```

```
>>> x = np.arange(8).reshape(2, 2, 2)
>>> x.shape
(2, 2, 2)
>>> y = F.flatten(x)
>>> y.shape
(8,)
>>> y.data
array([0, 1, 2, 3, 4, 5, 6, 7])
```

chainer.functions.flip

`chainer.functions.flip(x, axis)`

Flips an input variable in reverse order along the given axis.

Parameters

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable.
- **axis** (*int*) – Axis along which the input variable is reversed.

Returns Output variable.

Return type *Variable*

chainer.functions.fliplr

`chainer.functions.fliplr(a)`

Flip array in the left/right direction.

Parameters **xs** (*Variable*) – Input variable.

Returns Output variable.

Return type *Variable*

chainer.functions.flipud

`chainer.functions.flipud(a)`

Flip array in the up/down direction.

Parameters **xs** (*Variable*) – Input variable.

Returns Output variable.

Return type *Variable*

chainer.functions.get_item

`chainer.functions.get_item(x, slices)`

Extract elements from array with specified shape, axes and offsets.

Parameters

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – A variable to be sliced.
- **slices** (*int*, *slice*, *Ellipsis*, *None*, *integer array-like*, *boolean array-like* or *tuple of them*) – An object to specify the selection of elements.

Returns A *Variable* object which contains sliced array of `x`.

Note: It only supports types that are supported by CUDA's `atomicAdd` when an integer array is included in `slices`. The supported types are `numpy.float32`, `numpy.int32`, `numpy.uint32`, `numpy.uint64` and `numpy.ulonglong`.

Note: It does not support `slices` that contains multiple boolean arrays.

Note: See NumPy document for details of [indexing](#).

Example

```
>>> x = np.arange(12).reshape((2, 2, 3))
>>> x
array([[[ 0,  1,  2],
        [ 3,  4,  5]],

       [[ 6,  7,  8],
        [ 9, 10, 11]]])
>>> F.get_item(x, 0)
variable([[0, 1, 2],
         [3, 4, 5]])
>>> F.get_item(x, (0, 0, slice(0, 2, 1))) # equals x[0, 0, 0:2:1]
variable([0, 1])
>>> F.get_item(x, (Ellipsis, 2)) # equals x[..., 2]
variable([[ 2,  5],
         [ 8, 11]])
>>> F.get_item(x, (1, np.newaxis, 1, 0)) # equals x[1, None, 1, 0]
variable([9])
```

chainer.functions.hstack

`chainer.functions.hstack(xs)`

Concatenate variables horizontally (column wise).

Parameters **xs** (list of *Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variables to be concatenated. The variables must have the same `ndim`. When the variables have the second axis (i.e. $ndim \geq 2$), the variables must have the same shape along all but the second axis. When the variables do not have the second axis (i.e. $ndim < 2$), the variables need not to have the same shape.

Returns Output variable. When the input variables have the second axis (i.e. $ndim \geq 2$), the shapes of inputs and output are the same along all but the second axis. The length of second axis is the

sum of the lengths of inputs' second axis. When the variables do not have the second axis (i.e. $ndim < 2$), the shape of output is $(N,)$ (N is the sum of the input variables' size).

Return type *Variable*

Example

```
>>> x1 = np.array((1, 2, 3))
>>> x1.shape
(3,)
>>> x2 = np.array((2, 3, 4))
>>> x2.shape
(3,)
>>> y = F.hstack((x1, x2))
>>> y.shape
(6,)
>>> y.data
array([1, 2, 3, 2, 3, 4])
>>> x1 = np.arange(0, 12).reshape(3, 4)
>>> x1.shape
(3, 4)
>>> x1
array([[ 0,  1,  2,  3],
       [ 4,  5,  6,  7],
       [ 8,  9, 10, 11]])
>>> x2 = np.arange(12, 18).reshape(3, 2)
>>> x2.shape
(3, 2)
>>> x2
array([[12, 13],
       [14, 15],
       [16, 17]])
>>> y = F.hstack([x1, x2])
>>> y.shape
(3, 6)
>>> y.data
array([[ 0,  1,  2,  3, 12, 13],
       [ 4,  5,  6,  7, 14, 15],
       [ 8,  9, 10, 11, 16, 17]])
```

chainer.functions.im2col

`chainer.functions.im2col(x, ksize, stride=1, pad=0, cover_all=False, dilate=1)`

Extract patches from an image based on the filter.

This function rearranges patches of an image and puts them in the channel dimension of the output.

Patches are extracted at positions shifted by multiples of `stride` from the first position `-pad` for each spatial axis. The right-most (or bottom-most) patches do not run over the padded spatial size.

Notation: here is a notation.

- n is the batch size.
- c is the number of the input channels.
- h and w are the height and width of the input image, respectively.

- k_H and k_W are the height and width of the filters, respectively.
- s_Y and s_X are the strides of the filter.
- p_H and p_W are the spatial padding sizes.
- d_Y and d_X are the dilation factors of filter application.

The output size (h_O, w_O) is determined by the following equations when `cover_all = False`:

$$h_O = (h + 2p_H - k_H - (k_H - 1) * (d_Y - 1)) / s_Y + 1,$$

$$w_O = (w + 2p_W - k_W - (k_W - 1) * (d_X - 1)) / s_X + 1.$$

When `cover_all = True`, the output size is determined by the following equations:

$$h_O = (h + 2p_H - k_H - (k_H - 1) * (d_Y - 1) + s_Y - 1) / s_Y + 1,$$

$$w_O = (w + 2p_W - k_W - (k_W - 1) * (d_X - 1) + s_X - 1) / s_X + 1.$$

Parameters

- **x** (*Variable*) – Input variable of shape (n, c, h, w) .
- **ksize** (*int or pair of ints*) – Size of filters (a.k.a. kernels). `ksize=k` and `ksize=(k, k)` are equivalent.
- **stride** (*int or pair of ints*) – Stride of filter applications. `stride=s` and `stride=(s, s)` are equivalent.
- **pad** (*int or pair of ints*) – Spatial padding width for input arrays. `pad=p` and `pad=(p, p)` are equivalent.
- **cover_all** (*bool*) – If `True`, all spatial locations are rearranged into some output pixels. It may make the output size larger.
- **dilate** (*int or pair of ints*) – Dilation factor of filter applications. `dilate=d` and `dilate=(d, d)` are equivalent.

Returns Output variable whose shape is $(n, c \cdot k_H \cdot k_W, h_O, w_O)$

Return type *Variable*

chainer.functions.pad

`chainer.functions.pad(x, pad_width, mode, **keywords)`

Pad an input variable.

Parameters

- **x** (*chainer.Variable or numpy.ndarray or cupy.ndarray*) – Input data.
- **pad_width** (*int or array-like*) – Number of values padded to the edges of each axis.
- **mode** (*str*) – Specifies how the function fills the periphery of the array. The mode is passed to `numpy.pad()` or `cupy.pad()`. If it is 'constant', the input is padded by a constant value specified by `constant_values`.
- **constant_values** (*int or array-like*) – Constant values to fill the periphery in the 'constant' mode.

Returns Output variable.

Return type *Variable*

chainer.functions.pad_sequence

`chainer.functions.pad_sequence(xs, length=None, padding=0)`

Pad given arrays to make a matrix.

Parameters

- **xs** (*list of ~chainer.Variable*) – Variables you want to concatenate.
- **length** (*None or int*) – Size of the first dimension of a padded array. If it is *None*, the longest size of the first dimension of *xs* is used.
- **padding** (*int or float*) – Value to fill.

Returns A padded matrix. Its shape is $(n, \text{length}, \dots)$, where $n == \text{len}(xs)$.

Return type *Variable*

chainer.functions.permutate

`chainer.functions.permutate(x, indices, axis=0, inv=False)`

Permutates a given variable along an axis.

This function permute *x* with given *indices*. That means $y[i] = x[\text{indices}[i]]$ for all *i*. Note that this result is same as $y = x.\text{take}(\text{indices})$. *indices* must be a permutation of $[0, 1, \dots, \text{len}(x) - 1]$.

When *inv* is *True*, *indices* is treated as its inverse. That means $y[\text{indices}[i]] = x[i]$.

Parameters

- **x** (*Variable or numpy.ndarray or cupy.ndarray*) – Variable to permute. A (s_1, s_2, \dots, s_N) -shaped float array.
- **indices** (*Variable or numpy.ndarray or cupy.ndarray*) – Indices to extract from the variable. A one-dimensional int array.
- **axis** (*int*) – Axis that the input array is permute along.
- **inv** (*bool*) – If *True*, *indices* is treated as its inverse.

Returns Output variable.

Return type *Variable*

Example

```
>>> x = np.arange(6).reshape((3, 2)).astype(np.float32)
>>> x
array([[0., 1.],
       [2., 3.],
       [4., 5.]], dtype=float32)
>>> indices = np.array([2, 0, 1], np.int32)
>>> y = F.permutate(x, indices)
>>> y.data
array([[4., 5.],
       [0., 1.],
       [2., 3.]], dtype=float32)
>>> y = F.permutate(x, indices, inv=True)
>>> y.data
array([[2., 3.],
```

(continues on next page)

(continued from previous page)

```

        [4., 5.],
        [0., 1.]], dtype=float32)
>>> indices = np.array([1, 0], np.int32)
>>> y = F.permutate(x, indices, axis=1)
>>> y.data
array([[1., 0.],
       [3., 2.],
       [5., 4.]], dtype=float32)

```

chainer.functions.repeat

`chainer.functions.repeat` (*x*, *repeats*, *axis=None*)

Construct an array by repeating a given array.

Parameters

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable.
- **repeats** (`int` or `tuple` of `int` s) – The number of times which each element of *x* is repeated.
- **axis** (`int`) – The axis along which to repeat values.

Returns The repeated output *Variable*.

Return type *Variable*

Example

```

>>> x = np.array([0, 1, 2])
>>> x.shape
(3,)
>>> y = F.repeat(x, 2)
>>> y.shape
(6,)
>>> y.data
array([0, 0, 1, 1, 2, 2])
>>> x = np.array([[1, 2], [3, 4]])
>>> x.shape
(2, 2)
>>> y = F.repeat(x, 3, axis=1)
>>> y.shape
(2, 6)
>>> y.data
array([[1, 1, 1, 2, 2, 2],
       [3, 3, 3, 4, 4, 4]])
>>> y = F.repeat(x, (1, 2), axis=0)
>>> y.shape
(3, 2)
>>> y.data
array([[1, 2],
       [3, 4],
       [3, 4]])

```

chainer.functions.reshape

`chainer.functions.reshape(x, shape)`

Reshapes an input variable without copy.

Parameters

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable.
- **shape** (tuple of `int`s) – Expected shape of the output array. The number of elements which the array of `shape` contains must be equal to that of input array. One shape dimension can be -1. In this case, the value is inferred from the length of the array and remaining dimensions.

Returns Variable that holds a reshaped version of the input variable.

Return type *Variable*

See also:

`numpy.reshape()`, `cupy.reshape()`

Example

```
>>> x = np.array([[1, 2, 3, 4], [5, 6, 7, 8]])
>>> y = F.reshape(x, (8,))
>>> y.shape
(8,)
>>> y.data
array([1, 2, 3, 4, 5, 6, 7, 8])
>>> y = F.reshape(x, (4, -1)) # the shape of output is inferred
>>> y.shape
(4, 2)
>>> y.data
array([[1, 2],
       [3, 4],
       [5, 6],
       [7, 8]])
>>> y = F.reshape(x, (4, 3)) # the shape of input and output are not consistent
Traceback (most recent call last):
...
chainer.utils.type_check.InvalidType:
Invalid operation is performed in: Reshape (Forward)

Expect: prod(in_types[0].shape) == prod((4, 3))
Actual: 8 != 12
```

chainer.functions.resize_images

`chainer.functions.resize_images(x, output_shape)`

Resize images to the given shape.

This function resizes 2D data to `output_shape`. Currently, only bilinear interpolation is supported as the sampling method.

Notation: here is a notation for dimensionalities.

- *n* is the batch size.

- c_I is the number of the input channels.
- h and w are the height and width of the input image, respectively.
- h_O and w_O are the height and width of the output image.

Parameters

- **x** (*Variable*) – Input variable of shape (n, c_I, h, w) .
- **output_shape** (*tuple*) – This is a tuple of length 2 whose values are (h_O, w_O) . Note that the order of height and width is opposite of the one in OpenCV.

Returns Resized image whose shape is (n, c_I, h_O, w_O) .

Return type *Variable*

chainer.functions.rollaxis

`chainer.functions.rollaxis(x, axis, start=0)`

Roll the axis backwards to the given position.

Parameters

- **x** (*Variable*) – Input variable.
- **axis** (*int*) – The axis to roll backwards.
- **start** (*int*) – The place to which the axis is moved.

Returns Variable whose axis is rolled.

Return type *Variable*

chainer.functions.scatter_add

`chainer.functions.scatter_add(a, slices, b)`

Adds given values to specified elements of an array.

This function adds b to the specified elements of the copy of a , and returns the copy. The value of the original a is not changed.

Parameters

- **a** (*Variable*) – A variable.
- **slices** (*int, slice, Ellipsis, None, integer array-like, boolean array-like or tuple of them*) – It is an integer, a slice, an ellipsis, a `numpy.newaxis`, an integer array-like, a boolean array-like or tuple of them.
- **b** (*Variable*) – A variable that is scatter added to a . Its shape has to equal $a[slices]$ because broadcasting of variables is not supported.

Returns A *Variable* object which is the result of scatter addition.

Note: It only supports types that are supported by CUDA's `atomicAdd` when an integer array is included in `slices`. The supported types are `numpy.float32`, `numpy.int32`, `numpy.uint32`, `numpy.uint64` and `numpy.ulonglong`.

Note: It does not support `slices` that contains multiple boolean arrays.

See also:

`numpy.add.at()` and `cupyx.scatter_add()`.

`chainer.functions.select_item`

`chainer.functions.select_item(x, t)`

Select elements stored in given indices.

This function returns `t.choose(x.T)`, that means `y[i] == x[i, t[i]]` for all `i`.

Parameters

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Variable storing arrays. A two-dimensional float array.
- **t** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Variable storing index numbers. A one-dimensional int array. Length of the `t` should be equal to `x.shape[0]`.

Returns Variable that holds `t`-th element of `x`.

Return type *Variable*

Example

```
>>> x = np.array([[0, 1, 2], [3, 4, 5]], np.float32)
>>> t = np.array([0, 2], np.int32)
>>> y = F.select_item(x, t)
>>> y.shape
(2,)
>>> y.data
array([0., 5.], dtype=float32)
```

`chainer.functions.separate`

`chainer.functions.separate(x, axis=0)`

Separates an array along a given axis.

This function separates an array along a given axis. For example, shape of an array is `(2, 3, 4)`. When it separates the array with `axis=1`, it returns three `(2, 4)` arrays.

This function is an inverse of `chainer.functions.stack()`.

Parameters

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Variable to be separated. A (s_1, s_2, \dots, s_N) -shaped float array.
- **axis** (*int*) – Axis along which variables are separated.

Returns Output variables.

Return type tuple of `chainer.Variable`

See also:`chainer.functions.stack()`**Example**

```

>>> x = np.arange(6).reshape((2, 3)).astype(np.float32)
>>> x
array([[0., 1., 2.],
       [3., 4., 5.]], dtype=float32)
>>> x.shape
(2, 3)
>>> y = F.separate(x) # split along axis=0
>>> isinstance(y, tuple)
True
>>> len(y)
2
>>> y[0].shape
(3,)
>>> y[0].data
array([0., 1., 2.], dtype=float32)
>>> y = F.separate(x, axis=1)
>>> len(y)
3
>>> y[0].shape
(2,)
>>> y[0].data
array([0., 3.], dtype=float32)

```

chainer.functions.space2depth`chainer.functions.space2depth(X, r)`

Computes the space2depth transformation for subpixel calculations.

Parameters

- **X** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Variable holding a 4d array of shape (batch, channel, dim1 * r, dim2 * r).
- **r** (*int*) – the downscaling factor.

Returns A variable holding the downscaled layer array from subpixel array sampling. The shape is (batch, channel * r * r, dim1, dim2).

Return type *Variable*

Note: This can be used to compute inverse super-resolution transformations. See <https://arxiv.org/abs/1609.05158> for details.

See also:`depth2space()`**Example**

```
>>> X = np.arange(24).reshape(1, 1, 4, 6).astype(np.float32)
>>> X.shape
(1, 1, 4, 6)
>>> X
array([[[[ 0.,  1.,  2.,  3.,  4.,  5.],
          [ 6.,  7.,  8.,  9., 10., 11.],
          [12., 13., 14., 15., 16., 17.],
          [18., 19., 20., 21., 22., 23.]]]], dtype=float32)
>>> y = F.space2depth(X, 2)
>>> y.shape
(1, 4, 2, 3)
>>> y.data
array([[[[ 0.,  2.,  4.],
          [12., 14., 16.]],

        [[ 1.,  3.,  5.],
          [13., 15., 17.]],

        [[ 6.,  8., 10.],
          [18., 20., 22.]],

        [[ 7.,  9., 11.],
          [19., 21., 23.]]]], dtype=float32)
```

chainer.functions.spatial_transformer_grid

`chainer.functions.spatial_transformer_grid(theta, output_shape, **kwargs)`
2D Spatial Transformer grid.

This function generates coordinates of the points sampled from an image to perform warping described in [Spatial Transformer Networks](#).

Given a coordinate in the warped image (x_i^t, y_i^t) , the point sampled from the source image (x_i^s, y_i^s) are calculated by the following equation.

Note: cuDNN supports SpatialTransformerGrid from version 5.0.0.

$$\begin{pmatrix} x_i^s \\ y_i^s \end{pmatrix} = \begin{pmatrix} \theta_{11} & \theta_{12} & \theta_{13} \\ \theta_{21} & \theta_{22} & \theta_{23} \end{pmatrix} \begin{pmatrix} x_i^t \\ y_i^t \\ 1 \end{pmatrix}$$

Notation: here is a notation for dimensionalities.

- n is the batch size.
- h_O and w_O are the height and the width of the output image.

Parameters

- **theta** (*Variable*) – An array of shape $(n, 2, 3)$. This is a batch of 2×3 matrix used for the warping described above.
- **output_shape** (*tuple*) – A tuple of 2 elements: h_O, w_O .

Returns A variable of shape $(n, 2, h_O, w_O)$. In the 2nd dimension, the first element is the coordinate along the x axis, and the second element is the coordinate along the y axis. All the coordinates in the image are scaled to fit range $[-1, 1]$. This means that the coordinate $(-1, -1)$ corresponds to the upper-left corner of the input image.

Return type *Variable*

chainer.functions.spatial_transformer_sampler

`chainer.functions.spatial_transformer_sampler(x, grid, **kwargs)`

2D Spatial Transformer sampler.

This is a differentiable image sampler. With a set of sampling points `grid` and an input feature map `x`, this produces a sampled output feature map.

This function currently only supports bilinear interpolation as a sampling kernel.

When coordinates in `grid` is outside range $[-1, 1]$, values are sampled from a zero padded input image.

Notation: here is a notation for dimensionalities.

- n is the batch size.
- c_I is the number of the input channels.
- h and w are the height and width of the input image, respectively.
- h_O and w_O are the height and width of the output image.

See detail in the following paper: [Spatial Transformer Networks](#).

Note: cuDNN supports SpatialTransformerSampler from version 5.0.0.

Parameters

- **x** (*Variable*) – Input variable of shape (n, c_I, h, w) .
- **grid** (*Variable*) – Coordinate variable of shape $(n, 2, h_O, w_O)$. Each coordinate defines the spatial location in the input where a sampling kernel is applied to get the value at a particular pixel in the output. `grid[idx, :, i, j]` corresponds to the coordinate that is used to sample the values for an output pixel at location (i, j) .

In the second dimension, the first coordinate corresponds to the location along the horizontal axis, and the second coordinate corresponds to the location along the vertical axis.

The coordinate $(-1, -1)$ corresponds to the upper-left corner of the input image.

Returns Output feature map of shape (n, c_I, h_O, w_O) .

Return type *Variable*

chainer.functions.split_axis

`chainer.functions.split_axis(x, indices_or_sections, axis, force_tuple=True)`

Splits given variables along an axis.

Parameters

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – A variable to be split.

- **indices_or_sections** (*int* or 1-D array) – If this argument is an integer, N , the array will be divided into N equal arrays along axis. If it is a 1-D array of sorted integers, it indicates the positions where the array is split.
- **axis** (*int*) – Axis that the input array is split along.
- **force_tuple** (*bool*) – If `True` (the default) this method returns a tuple even when the number of outputs is one. Otherwise, if `False` a `Variable` will be returned when the number of outputs is one.

Returns Tuple of `Variable` objects if the number of outputs is more than 1 or `Variable` otherwise. When `force_tuple` is `True`, returned value is always a tuple regardless of the number of outputs.

Return type tuple or `Variable`

Note: This function raises `ValueError` if at least one of the outputs is split to zero-size (i.e. `axis`-th value of its shape is zero).

chainer.functions.squeeze

`chainer.functions.squeeze(x, axis=None)`

Remove dimensions of size one from the shape of a ndarray.

Parameters

- **x** (`Variable` or `numpy.ndarray` or `cupy.ndarray`) – Input variable. A (s_1, s_2, \dots, s_N) -shaped float array.
- **axis** (`None` or *int* or tuple of ints) – A subset of the single-dimensional entries in the shape to remove. If `None` is supplied, all of them are removed. The dimension index starts at zero. If an axis with dimension greater than one is selected, an error is raised.

Returns Variable whose dimensions of size 1 are removed.

Return type `Variable`

Example

```
>>> x = np.array([[[[0, 1, 2]]], [[[3, 4, 5]]]], np.float32)
>>> x.shape
(2, 1, 1, 3)
>>> y = F.squeeze(x)
>>> y.shape
(2, 3)
>>> y.data
array([[0., 1., 2.],
       [3., 4., 5.]], dtype=float32)
>>> y = F.squeeze(x, axis=1)
>>> y.shape
(2, 1, 3)
>>> y.data
array([[[0., 1., 2.]],
       [[3., 4., 5.]]], dtype=float32)
>>> y = F.squeeze(x, axis=(1, 2))
>>> y.shape
```

(continues on next page)

(continued from previous page)

```
(2, 3)
>>> y.data
array([[0., 1., 2.],
       [3., 4., 5.]], dtype=float32)
```

chainer.functions.stack

`chainer.functions.stack(xs, axis=0)`

Concatenate variables along a new axis.

Parameters

- **xs** (list of *Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variables to be concatenated. The variables must have the same shape.
- **axis** (*int*) – The axis along which the arrays will be stacked. The `axis` parameter is acceptable when $-ndim - 1 \leq axis \leq ndim$. (`ndim` is the dimension of input variables). When `axis < 0`, the result is the same with $ndim + 1 - |axis|$.

Returns Output variable. Let `x_1`, `x_2`, ..., `x_n` and `y` be the input variables and the output variable, `y[:, ..., 0, ..., :]` is `x_1`, `y[:, ..., 1, ..., :]` is `x_2` and `y[:, ..., n-1, ..., :]` is `x_n` (The indexed axis indicates the `axis`).

Return type *Variable*

Example

```
>>> x1 = np.arange(0, 12).reshape(3, 4)
>>> x1.shape
(3, 4)
>>> x1
array([[ 0,  1,  2,  3],
       [ 4,  5,  6,  7],
       [ 8,  9, 10, 11]])
>>> x2 = np.arange(12, 24).reshape(3, 4)
>>> x2.shape
(3, 4)
>>> x2
array([[12, 13, 14, 15],
       [16, 17, 18, 19],
       [20, 21, 22, 23]])
>>> y = F.stack([x1, x2], axis=0)
>>> y.shape
(2, 3, 4)
>>> y.data
array([[[ 0,  1,  2,  3],
        [ 4,  5,  6,  7],
        [ 8,  9, 10, 11]],
       [[12, 13, 14, 15],
        [16, 17, 18, 19],
        [20, 21, 22, 23]])])
>>> y = F.stack([x1, x2], axis=1)
>>> y.shape
(3, 2, 4)
```

(continues on next page)

(continued from previous page)

```

>>> y.data
array([[ 0,  1,  2,  3],
       [12, 13, 14, 15]],

      [[ 4,  5,  6,  7],
       [16, 17, 18, 19]],

      [[ 8,  9, 10, 11],
       [20, 21, 22, 23]])
>>> y = F.stack([x1, x2], axis=2)
>>> y.shape
(3, 4, 2)
>>> y.data
array([[ 0, 12],
       [ 1, 13],
       [ 2, 14],
       [ 3, 15]],

      [[ 4, 16],
       [ 5, 17],
       [ 6, 18],
       [ 7, 19]],

      [[ 8, 20],
       [ 9, 21],
       [10, 22],
       [11, 23]])
>>> y = F.stack([x1, x2], axis=-1)
>>> y.shape
(3, 4, 2)

```

chainer.functions.swapaxes

`chainer.functions.swapaxes(x, axis1, axis2)`

Swap two axes of a variable.

Parameters

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable. A (s_1, s_2, \dots, s_N) -shaped float array.
- **axis1** (*int*) – The first axis to swap.
- **axis2** (*int*) – The second axis to swap.

Returns Variable whose axes are swapped.

Return type *Variable*

Example

```

>>> x = np.array([[[0, 1, 2], [3, 4, 5]]], np.float32)
>>> x.shape
(1, 2, 3)
>>> y = F.swapaxes(x, axis1=0, axis2=1)

```

(continues on next page)

(continued from previous page)

```
>>> y.shape
(2, 1, 3)
>>> y.data
array([[[0., 1., 2.],
        [3., 4., 5.]]], dtype=float32)
```

chainer.functions.tile

`chainer.functions.tile(x, reps)`

Construct an array by tiling a given array.

Parameters

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable. Let the length of `reps` be `d`. If `x.ndim < d`, `x` is treated as `d`-dimensional array by prepending new axes. For example, when the shape of `x` is `(2,)` and tiled with 2-dim repetitions, `x` is treated as the shape `(1, 2)`. If `x.ndim > d`, `reps` is treated as `x.ndim`-dimensional by pre-pending 1's. For example, when the shape of `x` is `(2, 3, 2, 3)`, the 2-dim `reps` of `(2, 2)` is treated as `(1, 1, 2, 2)`.
- **reps** (`int` or `tuple` of `int` s) – The number of times which `x` is replicated along each axis.

Returns The tiled output *Variable*. Let the length of `reps` be `d`, the output has the dimension of `max(d, x.ndim)`.

Return type *Variable*

Example

```
>>> x = np.array([0, 1, 2])
>>> x.shape
(3,)
>>> y = F.tile(x, 2)
>>> y.shape
(6,)
>>> y.data
array([0, 1, 2, 0, 1, 2])
>>> y = F.tile(x, (2, 2))
>>> y.shape
(2, 6)
>>> y.data
array([[0, 1, 2, 0, 1, 2],
       [0, 1, 2, 0, 1, 2]])
>>> y = F.tile(x, (2, 1, 2))
>>> y.shape
(2, 1, 6)
>>> y.data
array([[[0, 1, 2, 0, 1, 2]],
       [[0, 1, 2, 0, 1, 2]])])
```

```
>>> x = np.array([[1, 2], [3, 4]])
>>> x.shape
(2, 2)
>>> y = F.tile(x, 2)
>>> y.shape
(2, 4)
>>> y.data
array([[1, 2, 1, 2],
       [3, 4, 3, 4]])
>>> y = F.tile(x, (2, 2))
>>> y.shape
(4, 4)
>>> y.data
array([[1, 2, 1, 2],
       [3, 4, 3, 4],
       [1, 2, 1, 2],
       [3, 4, 3, 4]])
>>> y = F.tile(x, (2, 1, 2))
>>> y.shape
(2, 2, 4)
>>> y.data
array([[[1, 2, 1, 2],
        [3, 4, 3, 4]],
       [[1, 2, 1, 2],
        [3, 4, 3, 4]]])
```

chainer.functions.transpose

`chainer.functions.transpose(x, axes=None)`

Permute the dimensions of an input variable without copy.

Parameters

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable to be transposed. A (s_1, s_2, \dots, s_N) -shaped float array.
- **axes** (*tuple of ints*) – By default, reverse the dimensions, otherwise permute the axes according to the values given.

Returns Variable whose axes are permuted.

Return type *Variable*

Example

```
>>> x = np.array([[[0, 1, 2], [3, 4, 5]]], np.float32)
>>> x.shape
(1, 2, 3)
>>> y = F.transpose(x) # reverse the dimensions
>>> y.shape
(3, 2, 1)
>>> y.data
array([[[0.],
        [3.]],
       [[1.],
        [4.]],
       [[2.],
        [5.]])
```

(continues on next page)

(continued from previous page)

```

[[1.],
 [4.]],

[[2.],
 [5.]]], dtype=float32)
>>> y = F.transpose(x, axes=(1, 0, 2)) # swap 1st and 2nd axis
>>> y.shape
(2, 1, 3)
>>> y.data
array([[[0., 1., 2.]],

       [[3., 4., 5.]]], dtype=float32)

```

chainer.functions.transpose_sequence

`chainer.functions.transpose_sequence(xs)`

Transpose a list of Variables.

This function transposes a list of *Variables* and returns a list of Variables. For example a user gives `[(0, 1, 2, 3), (4, 5), (6)]`, the function returns `[(0, 4, 6), (1, 5), (2), (3)]`. Note that a given list needs to be sorted by each length of *Variable*.

Parameters *xs* (list of *Variable* or `numpy.ndarray` or `cupy.ndarray`) – Variables to transpose.

Returns Transposed list.

Return type tuple of *Variable*

Example

```

>>> lst = [chainer.Variable(np.array([1, 1, 1])),
...         chainer.Variable(np.array([2, 2])),
...         chainer.Variable(np.array([3]))]
>>> lst
[variable([1, 1, 1]), variable([2, 2]), variable([3])]
>>> transposed = F.transpose_sequence(lst)
>>> transposed
(variable([1, 2, 3]), variable([1, 2]), variable([1]))

```

chainer.functions.vstack

`chainer.functions.vstack(xs)`

Concatenate variables vertically (row wise).

Parameters *xs* (list of *Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variables to be concatenated. The variables must have the same *ndim*. When the variables have the second axis (i.e. $ndim \geq 2$), the variables must have the same shape along all but the first axis. When the variables do not have the second axis (i.e. $ndim < 2$), the variables must have the same shape.

Returns Output variable. When the input variables have the second axis (i.e. $ndim \geq 2$), the shapes of inputs and output are the same along all but the first axis. The length of first axis is the sum of

the lengths of inputs' first axis. When the variables do not have the second axis (i.e. `ndim < 2`), the shape of output is `(2, N)` (`N` is the size of the input variable).

Return type *Variable*

Example

```
>>> x1 = np.array((1, 2, 3))
>>> x1.shape
(3,)
>>> x2 = np.array((2, 3, 4))
>>> x2.shape
(3,)
>>> y = F.vstack((x1, x2))
>>> y.shape
(2, 3)
>>> y.data
array([[1, 2, 3],
       [2, 3, 4]])
>>> x1 = np.arange(0, 12).reshape(3, 4)
>>> x1.shape
(3, 4)
>>> x1
array([[ 0,  1,  2,  3],
       [ 4,  5,  6,  7],
       [ 8,  9, 10, 11]])
>>> x2 = np.arange(12, 20).reshape(2, 4)
>>> x2.shape
(2, 4)
>>> x2
array([[12, 13, 14, 15],
       [16, 17, 18, 19]])
>>> y = F.vstack([x1, x2])
>>> y.shape
(5, 4)
>>> y.data
array([[ 0,  1,  2,  3],
       [ 4,  5,  6,  7],
       [ 8,  9, 10, 11],
       [12, 13, 14, 15],
       [16, 17, 18, 19]])
```

chainer.functions.where

`chainer.functions.where(condition, x, y)`

Choose elements depending on condition.

This function choose values depending on a given `condition`. All `condition`, `x`, and `y` must have the same shape.

Parameters

- **condition** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable containing the condition. A (s_1, s_2, \dots, s_N) -shaped boolean array. Only boolean array is permitted.

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable chosen when condition is True. A (s_1, s_2, \dots, s_N) -shaped float array.
- **y** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable chosen when condition is False. A (s_1, s_2, \dots, s_N) -shaped float array.

Returns Variable containing chosen values.

Return type *Variable*

Example

```
>>> cond = np.array([[1, 0], [0, 1]], dtype=np.bool)
>>> cond
array([[ True, False],
       [False,  True]])
>>> x = np.array([[1, 2], [3, 4]], np.float32)
>>> y = np.zeros((2, 2), np.float32)
>>> F.where(cond, x, y).data
array([[1., 0.],
       [0., 4.]], dtype=float32)
```

5.2.4 Neural network connections

<code>chainer.functions.bilinear</code>	Applies a bilinear function based on given parameters.
<code>chainer.functions.convolution_2d</code>	Two-dimensional convolution function.
<code>chainer.functions.convolution_nd</code>	N-dimensional convolution function.
<code>chainer.functions.deconvolution_2d</code>	Two dimensional deconvolution function.
<code>chainer.functions.deconvolution_nd</code>	N-dimensional deconvolution function.
<code>chainer.functions.depthwise_convolution_2d</code>	Two-dimensional depthwise convolution function.
<code>chainer.functions.dilated_convolution_2d</code>	Two-dimensional dilated convolution function.
<code>chainer.functions.embed_id</code>	Efficient linear function for one-hot input.
<code>chainer.functions.linear</code>	Linear function, or affine transformation.
<code>chainer.functions.local_convolution_2d</code>	Two-dimensional local convolution function.
<code>chainer.functions.n_step_bigru</code>	Stacked Bi-directional Gated Recurrent Unit function.
<code>chainer.functions.n_step_bilstm</code>	Stacked Bi-directional Long Short-Term Memory function.
<code>chainer.functions.n_step_birnn</code>	Stacked Bi-directional RNN function for sequence inputs.
<code>chainer.functions.n_step_gru</code>	Stacked Uni-directional Gated Recurrent Unit function.
<code>chainer.functions.n_step_lstm</code>	Stacked Uni-directional Long Short-Term Memory function.
<code>chainer.functions.n_step_rnn</code>	Stacked Uni-directional RNN function for sequence inputs.
<code>chainer.functions.shift</code>	Shift function.

chainer.functions.bilinear

`chainer.functions.bilinear(e1, e2, W, V1=None, V2=None, b=None)`

Applies a bilinear function based on given parameters.

This is a building block of Neural Tensor Network (see the reference paper below). It takes two input variables and one or four parameters, and outputs one variable.

To be precise, denote six input arrays mathematically by $e^1 \in \mathbb{R}^{I \cdot J}$, $e^2 \in \mathbb{R}^{I \cdot K}$, $W \in \mathbb{R}^{J \cdot K \cdot L}$, $V^1 \in \mathbb{R}^{J \cdot L}$, $V^2 \in \mathbb{R}^{K \cdot L}$, and $b \in \mathbb{R}^L$, where I is mini-batch size. In this document, we call V^1 , V^2 , and b linear parameters.

The output of forward propagation is calculated as

$$y_{il} = \sum_{jk} e_{ij}^1 e_{ik}^2 W_{jkl} + \sum_j e_{ij}^1 V_{jl}^1 + \sum_k e_{ik}^2 V_{kl}^2 + b_l.$$

Note that V^1 , V^2 , b are optional. If these are not given, then this function omits the last three terms in the above equation.

Note: This function accepts an input variable `e1` or `e2` of a non-matrix array. In this case, the leading dimension is treated as the batch dimension, and the other dimensions are reduced to one dimension.

Note: In the original paper, J and K must be equal and the author denotes $[V^1 V^2]$ (concatenation of matrices) by V .

Parameters

- **e1** (*Variable*) – Left input variable.
- **e2** (*Variable*) – Right input variable.
- **w** (*Variable*) – Quadratic weight variable.
- **v1** (*Variable*) – Left coefficient variable.
- **v2** (*Variable*) – Right coefficient variable.
- **b** (*Variable*) – Bias variable.

Returns Output variable.

Return type *Variable*

See: [Reasoning With Neural Tensor Networks for Knowledge Base Completion](#) [Socher+, NIPS2013].

chainer.functions.convolution_2d

`chainer.functions.convolution_2d(x, W, b=None, stride=1, pad=0, cover_all=False, *, dilate=1, groups=1)`

Two-dimensional convolution function.

This is an implementation of two-dimensional convolution in ConvNets. It takes three variables: the input image x , the filter weight W , and the bias vector b .

Notation: here is a notation for dimensionalities.

- n is the batch size.

- c_I and c_O are the number of the input and output channels, respectively.
- h_I and w_I are the height and width of the input image, respectively.
- h_K and w_K are the height and width of the filters, respectively.
- h_P and w_P are the height and width of the spatial padding size, respectively.

Then the `Convolution2D` function computes correlations between filters and patches of size (h_K, w_K) in \mathbf{x} . Note that correlation here is equivalent to the inner product between expanded vectors. Patches are extracted at positions shifted by multiples of `stride` from the first position $(-h_P, -w_P)$ for each spatial axis. The right-most (or bottom-most) patches do not run over the padded spatial size.

Let (s_Y, s_X) be the stride of filter application. Then, the output size (h_O, w_O) is determined by the following equations:

$$\begin{aligned} h_O &= (h_I + 2h_P - h_K) / s_Y + 1, \\ w_O &= (w_I + 2w_P - w_K) / s_X + 1. \end{aligned}$$

If `cover_all` option is `True`, the filter will cover the all spatial locations. So, if the last stride of filter does not cover the end of spatial locations, an additional stride will be applied to the end part of spatial locations. In this case, the output size (h_O, w_O) is determined by the following equations:

$$\begin{aligned} h_O &= (h_I + 2h_P - h_K + s_Y - 1) / s_Y + 1, \\ w_O &= (w_I + 2w_P - w_K + s_X - 1) / s_X + 1. \end{aligned}$$

If the bias vector is given, then it is added to all spatial locations of the output of convolution.

The output of this function can be non-deterministic when it uses cuDNN. If `chainer.configuration.config.cudnn_deterministic` is `True` and cuDNN version is $\geq v3$, it forces cuDNN to use a deterministic algorithm.

Convolution links can use a feature of cuDNN called autotuning, which selects the most efficient CNN algorithm for images of fixed-size, can provide a significant performance boost for fixed neural nets. To enable, set `chainer.using_config('autotune', True)`

When the dilation factor is greater than one, cuDNN is not used unless the version is 6.0 or higher.

Warning: `deterministic` argument is not supported anymore since v2. Instead, use `chainer.using_config('cudnn_deterministic', value)` (value is either `True` or `False`). See `chainer.using_config()`.

Parameters

- **\mathbf{x}** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable of shape (n, c_I, h_I, w_I) .
- **\mathbf{W}** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Weight variable of shape (c_O, c_I, h_K, w_K) .
- **\mathbf{b}** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Bias variable of length c_O (optional).
- **`stride`** (`int` or pair of `int` s) – Stride of filter applications. `stride=s` and `stride=(s, s)` are equivalent.
- **`pad`** (`int` or pair of `int` s) – Spatial padding width for input arrays. `pad=p` and `pad=(p, p)` are equivalent.
- **`cover_all`** (`bool`) – If `True`, all spatial locations are convoluted into some output pixels.

- **dilate** (`int` or pair of `int` s) – Dilation factor of filter applications. `dilate=d` and `dilate=(d, d)` are equivalent.
- **groups** (`int`) – Number of groups of channels. If the number is greater than 1, input tensor W is divided into some blocks by this value. For each tensor blocks, convolution operation will be executed independently. Input channel size c_I and output channel size c_O must be exactly divisible by this value.

Returns Output variable of shape (n, c_O, h_O, w_O) .

Return type *Variable*

See also:

Convolution2D

Example

```
>>> n = 10
>>> c_i, c_o = 3, 1
>>> h_i, w_i = 30, 40
>>> h_k, w_k = 10, 10
>>> h_p, w_p = 5, 5
>>> x = np.random.uniform(0, 1, (n, c_i, h_i, w_i)).astype(np.float32)
>>> x.shape
(10, 3, 30, 40)
>>> W = np.random.uniform(0, 1, (c_o, c_i, h_k, w_k)).astype(np.float32)
>>> W.shape
(1, 3, 10, 10)
>>> b = np.random.uniform(0, 1, (c_o,)).astype(np.float32)
>>> b.shape
(1,)
>>> s_y, s_x = 5, 7
>>> y = F.convolution_2d(x, W, b, stride=(s_y, s_x), pad=(h_p, w_p))
>>> y.shape
(10, 1, 7, 6)
>>> h_o = int((h_i + 2 * h_p - h_k) / s_y + 1)
>>> w_o = int((w_i + 2 * w_p - w_k) / s_x + 1)
>>> y.shape == (n, c_o, h_o, w_o)
True
>>> y = F.convolution_2d(x, W, b, stride=(s_y, s_x), pad=(h_p, w_p), cover_
↪all=True)
>>> y.shape == (n, c_o, h_o, w_o + 1)
True
```

chainer.functions.convolution_nd

`chainer.functions.convolution_nd(x, W, b=None, stride=1, pad=0, cover_all=False)`

N-dimensional convolution function.

This is an implementation of N-dimensional convolution which is generalized two-dimensional convolution in ConvNets. It takes three variables: the input x , the filter weight W and the bias vector b .

Notation: here is a notation for dimensionalities.

- N is the number of spatial dimensions.
- n is the batch size.

- c_I and c_O are the number of the input and output channels, respectively.
- d_1, d_2, \dots, d_N are the size of each axis of the input's spatial dimensions, respectively.
- k_1, k_2, \dots, k_N are the size of each axis of the filters, respectively.
- l_1, l_2, \dots, l_N are the size of each axis of the output's spatial dimensions, respectively.
- p_1, p_2, \dots, p_N are the size of each axis of the spatial padding size, respectively.

Then the `convolution_nd` function computes correlations between filters and patches of size (k_1, k_2, \dots, k_N) in \mathbf{x} . Note that correlation here is equivalent to the inner product between expanded tensors. Patches are extracted at positions shifted by multiples of `stride` from the first position $(-p_1, -p_2, \dots, -p_N)$ for each spatial axis.

Let (s_1, s_2, \dots, s_N) be the stride of filter application. Then, the output size (l_1, l_2, \dots, l_N) is determined by the following equations:

$$l_n = (d_n + 2p_n - k_n) / s_n + 1 \quad (n = 1, \dots, N)$$

If `cover_all` option is `True`, the filter will cover the all spatial locations. So, if the last stride of filter does not cover the end of spatial locations, an additional stride will be applied to the end part of spatial locations. In this case, the output size is determined by the following equations:

$$l_n = (d_n + 2p_n - k_n + s_n - 1) / s_n + 1 \quad (n = 1, \dots, N)$$

The N-dimensional convolution function is defined as follows.

Parameters

- **\mathbf{x}** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable of shape $(n, c_I, d_1, d_2, \dots, d_N)$.
- **\mathbf{W}** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Weight variable of shape $(c_O, c_I, k_1, k_2, \dots, k_N)$.
- **\mathbf{b}** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – One-dimensional bias variable with length c_O (optional).
- **`stride`** (`int` or `tuple` of `int` s) – Stride of filter applications (s_1, s_2, \dots, s_N) . `stride=s` is equivalent to (s, s, \dots, s) .
- **`pad`** (`int` or `tuple` of `int` s) – Spatial padding width for input arrays (p_1, p_2, \dots, p_N) . `pad=p` is equivalent to (p, p, \dots, p) .
- **`cover_all`** (`bool`) – If `True`, all spatial locations are convoluted into some output pixels. It may make the output size larger. `cover_all` needs to be `False` if you want to use cuDNN.

Returns Output variable of shape $(n, c_O, l_1, l_2, \dots, l_N)$.

Return type *Variable*

Note: This function uses cuDNN implementation for its forward and backward computation if ALL of the following conditions are satisfied:

- `cuda.cudnn_enabled` is `True`
- `chainer.config.use_cudnn` is `'always'` or `'auto'`
- The number of spatial dimensions is more than one.
- `cover_all` is `False`
- The input's `dtype` is equal to the filter weight's.

- The dtype is FP16, FP32 or FP64. (FP16 is only available when cuDNN version \geq v3.)

Convolution links can use a feature of cuDNN called autotuning, which selects the most efficient CNN algorithm for images of fixed-size, can provide a significant performance boost for fixed neural nets. To enable, set `chainer.config('autotune', True)`

See also:

`ConvolutionND`, `convolution_2d()`

Example

```
>>> n = 10
>>> c_i, c_o = 3, 1
>>> d1, d2, d3 = 30, 40, 50
>>> k1, k2, k3 = 10, 10, 10
>>> p1, p2, p3 = 5, 5, 5
>>> x = np.random.uniform(0, 1, (n, c_i, d1, d2, d3)).astype(np.float32)
>>> x.shape
(10, 3, 30, 40, 50)
>>> W = np.random.uniform(0, 1, (c_o, c_i, k1, k2, k3)).astype(np.float32)
>>> W.shape
(1, 3, 10, 10, 10)
>>> b = np.random.uniform(0, 1, (c_o)).astype(np.float32)
>>> b.shape
(1,)
>>> s1, s2, s3 = 2, 4, 6
>>> y = F.convolution_nd(x, W, b, stride=(s1, s2, s3), pad=(p1, p2, p3))
>>> y.shape
(10, 1, 16, 11, 9)
>>> l1 = int((d1 + 2 * p1 - k1) / s1 + 1)
>>> l2 = int((d2 + 2 * p2 - k2) / s2 + 1)
>>> l3 = int((d3 + 2 * p3 - k3) / s3 + 1)
>>> y.shape == (n, c_o, l1, l2, l3)
True
>>> y = F.convolution_nd(x, W, b, stride=(s1, s2, s3), pad=(p1, p2, p3), cover_
↪all=True)
>>> y.shape == (n, c_o, l1, l2, l3 + 1)
True
```

chainer.functions.deconvolution_2d

`chainer.functions.deconvolution_2d(x, W, b=None, stride=1, pad=0, outsize=None, *, dilate=1, groups=1)`

Two dimensional deconvolution function.

This is an implementation of two-dimensional deconvolution. In most of deep learning frameworks and papers, this function is called **transposed convolution**. But because of historical reasons (e.g. paper by Ziller [Deconvolutional Networks](#)) and backward compatibility, this function is called **deconvolution** in Chainer.

It takes three variables: input image x , the filter weight W , and the bias vector b .

Notation: here is a notation for dimensionalities.

- n is the batch size.
- c_I and c_O are the number of the input and output channels, respectively.

- h_I and w_I are the height and width of the input image, respectively.
- h_K and w_K are the height and width of the filters, respectively.
- h_P and w_P are the height and width of the spatial padding size, respectively.

Let (s_Y, s_X) be the stride of filter application. Then, the output size (h_O, w_O) is estimated by the following equations:

$$\begin{aligned} h_O &= s_Y(h_I - 1) + h_K - 2h_P, \\ w_O &= s_X(w_I - 1) + w_K - 2w_P. \end{aligned}$$

The output of this function can be non-deterministic when it uses cuDNN. If `chainer.configuration.config.deterministic` is `True` and cuDNN version is $\geq v3$, it forces cuDNN to use a deterministic algorithm.

Deconvolution links can use a feature of cuDNN called autotuning, which selects the most efficient CNN algorithm for images of fixed-size, can provide a significant performance boost for fixed neural nets. To enable, set `chainer.using_config('autotune', True)`

Warning: `deterministic` argument is not supported anymore since v2. Instead, use `chainer.using_config('cudnn_deterministic', value)` (value is either `True` or `False`). See [chainer.using_config\(\)](#).

Parameters

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable of shape (n, c_I, h_I, w_I) .
- **w** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Weight variable of shape (c_I, c_O, h_K, w_K) .
- **b** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Bias variable of length c_O (optional).
- **stride** (`int` or pair of `int` s) – Stride of filter applications. `stride=s` and `stride=(s, s)` are equivalent.
- **pad** (`int` or pair of `int` s) – Spatial padding width for input arrays. `pad=p` and `pad=(p, p)` are equivalent.
- **outsize** (tuple of `int`) – Expected output size of deconvolutional operation. It should be pair of height and width (h_O, w_O) . Default value is `None` and the outsize is estimated by input size, stride and pad.
- **dilate** (`int` or pair of `int` s) – Dilation factor of filter applications. `dilate=d` and `dilate=(d, d)` are equivalent.
- **groups** (`int`) – The number of groups to use grouped deconvolution. The default is one, where grouped deconvolution is not used.

Returns Output variable of shape (n, c_O, h_O, w_O) .

Return type *Variable*

Example

```
>>> n = 10
>>> c_i, c_o = 1, 3
>>> h_i, w_i = 5, 10
>>> h_k, w_k = 10, 10
>>> h_p, w_p = 5, 5
>>> x = np.random.uniform(0, 1, (n, c_i, h_i, w_i)).astype(np.float32)
>>> x.shape
(10, 1, 5, 10)
>>> W = np.random.uniform(0, 1, (c_i, c_o, h_k, w_k)).astype(np.float32)
>>> W.shape
(1, 3, 10, 10)
>>> b = np.random.uniform(0, 1, c_o).astype(np.float32)
>>> b.shape
(3,)
>>> s_y, s_x = 5, 5
>>> y = F.deconvolution_2d(x, W, b, stride=(s_y, s_x), pad=(h_p, w_p))
>>> y.shape
(10, 3, 20, 45)
>>> h_o = s_y * (h_i - 1) + h_k - 2 * h_p
>>> w_o = s_x * (w_i - 1) + w_k - 2 * w_p
>>> y.shape == (n, c_o, h_o, w_o)
True
```

chainer.functions.deconvolution_nd

`chainer.functions.deconvolution_nd(x, W, b=None, stride=1, pad=0, outsize=None)`
N-dimensional deconvolution function.

This is an implementation of N-dimensional deconvolution which generalizes two-dimensional one. In most of deep learning frameworks and papers, this function is called **transposed convolution**. But because of historical reasons (e.g. paper by Ziller [Deconvolutional Networks](#)) and backward compatibility, this function is called **deconvolution** in Chainer.

It takes three variables: the input `x`, the filter weight `W`, and the bias vector `b`.

Notation: here is a notation for dimensionalities.

- N is the number of spatial dimensions.
- n is the batch size.
- c_I and c_O are the number of the input and output channels, respectively.
- d_1, d_2, \dots, d_N are the size of each axis of the input's spatial dimensions, respectively.
- k_1, k_2, \dots, k_N are the size of each axis of the filters, respectively.
- p_1, p_2, \dots, p_N are the size of each axis of the spatial padding size, respectively.
- s_1, s_2, \dots, s_N are the stride of each axis of filter application, respectively.

If `outsize` option is `None`, the output size (l_1, l_2, \dots, l_N) is determined by the following equations with the items in the above list:

$$l_n = s_n(d_n - 1) + k_n - 2p_n \quad (n = 1, \dots, N)$$

If `outsize` option is given, the output size is determined by `outsize`. In this case, the `outsize` (l_1, l_2, \dots, l_N) must satisfy the following equations:

$$d_n = \lfloor (l_n + 2p_n - k_n) / s_n \rfloor + 1 \quad (n = 1, \dots, N)$$

Deconvolution links can use a feature of cuDNN called autotuning, which selects the most efficient CNN algorithm for images of fixed-size, can provide a significant performance boost for fixed neural nets. To enable, set `chainer.config('autotune', True)`

Parameters

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable of shape $(n, c_I, d_1, d_2, \dots, d_N)$.
- **W** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Weight variable of shape $(c_I, c_O, k_1, k_2, \dots, k_N)$.
- **b** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – One-dimensional bias variable with length c_O (optional).
- **stride** (`int` or `tuple` of `int` s) – Stride of filter applications (s_1, s_2, \dots, s_N) . `stride=s` is equivalent to `(s, s, ..., s)`.
- **pad** (`int` or `tuple` of `int` s) – Spatial padding width for input arrays (p_1, p_2, \dots, p_N) . `pad=p` is equivalent to `(p, p, ..., p)`.
- **outsize** (`tuple` of `int` s) – Expected output size of deconvolutional operation. It should be a tuple of ints (l_1, l_2, \dots, l_N) . Default value is `None` and the outsize is estimated by input size, stride and pad.

Returns Output variable of shape $(n, c_O, l_1, l_2, \dots, l_N)$.

Return type *Variable*

See also:

`links.DeconvolutionND`, `deconvolution_2d()`

Example

Example1: the case when outsize is not given.

```
>>> n = 10
>>> c_i, c_o = 3, 1
>>> d1, d2, d3 = 5, 10, 15
>>> k1, k2, k3 = 10, 10, 10
>>> p1, p2, p3 = 5, 5, 5
>>> x = np.random.uniform(0, 1, (n, c_i, d1, d2, d3)).astype(np.float32)
>>> x.shape
(10, 3, 5, 10, 15)
>>> W = np.random.uniform(0, 1, (c_i, c_o, k1, k2, k3)).astype(np.float32)
>>> W.shape
(3, 1, 10, 10, 10)
>>> b = np.random.uniform(0, 1, (c_o)).astype(np.float32)
>>> b.shape
(1,)
>>> s1, s2, s3 = 2, 4, 6
>>> y = F.deconvolution_nd(x, W, b, stride=(s1, s2, s3), pad=(p1, p2, p3))
>>> y.shape
(10, 1, 8, 36, 84)
>>> l1 = s1 * (d1 - 1) + k1 - 2 * p1
>>> l2 = s2 * (d2 - 1) + k2 - 2 * p2
>>> l3 = s3 * (d3 - 1) + k3 - 2 * p3
>>> y.shape == (n, c_o, l1, l2, l3)
True
```

Example2: the case when outsize is given.

```

>>> n = 10
>>> c_i, c_o = 3, 1
>>> d1, d2, d3 = 5, 10, 15
>>> k1, k2, k3 = 10, 10, 10
>>> p1, p2, p3 = 5, 5, 5
>>> x = np.random.uniform(0, 1, (n, c_i, d1, d2, d3)).astype(np.float32)
>>> x.shape
(10, 3, 5, 10, 15)
>>> W = np.random.uniform(0, 1, (c_i, c_o, k1, k2, k3)).astype(np.float32)
>>> W.shape
(3, 1, 10, 10, 10)
>>> b = np.random.uniform(0, 1, (c_o)).astype(np.float32)
>>> b.shape
(1,)
>>> s1, s2, s3 = 2, 4, 6
>>> l1, l2, l3 = 9, 38, 87
>>> d1 == int((l1 + 2 * p1 - k1) / s1) + 1
True
>>> d2 == int((l2 + 2 * p2 - k2) / s2) + 1
True
>>> d3 == int((l3 + 2 * p3 - k3) / s3) + 1
True
>>> y = F.deconvolution_nd(x, W, b, stride=(s1, s2, s3), pad=(p1, p2, p3),
↪      outsize=(l1, l2, l3))
>>> y.shape
(10, 1, 9, 38, 87)
>>> y.shape == (n, c_o, l1, l2, l3)
True

```

chainer.functions.depthwise_convolution_2d

`chainer.functions.depthwise_convolution_2d(x, W, b=None, stride=1, pad=0)`

Two-dimensional depthwise convolution function.

This is an implementation of two-dimensional depthwise convolution. It takes two or three variables: the input image x , the filter weight W , and optionally, the bias vector b .

Notation: here is a notation for dimensionalities.

- n is the batch size.
- c_I is the number of the input.
- c_M is the channel multiplier.
- h and w are the height and width of the input image, respectively.
- h_O and w_O are the height and width of the output image, respectively.
- k_H and k_W are the height and width of the filters, respectively.

Parameters

- \mathbf{x} (`chainer.Variable` or `numpy.ndarray` or `cupy.ndarray`) – Input variable of shape (n, c_I, h, w) .
- \mathbf{W} (`Variable`) – Weight variable of shape (c_M, c_I, k_H, k_W) .
- \mathbf{b} (`Variable`) – Bias variable of length $c_M * c_I$ (optional).

- **stride** (*int* or *pair of ints*) – Stride of filter applications. `stride=s` and `stride=(s, s)` are equivalent.
- **pad** (*int* or *pair of ints*) – Spatial padding width for input arrays. `pad=p` and `pad=(p, p)` are equivalent.

Returns Output variable. Its shape is $(n, c_I * c_M, h_O, w_O)$.

Return type *Variable*

Like `Convolution2D`, `DepthwiseConvolution2D` function computes correlations between filters and patches of size (k_H, k_W) in x . But unlike `Convolution2D`, `DepthwiseConvolution2D` does not add up input channels of filters but concatenates them. For that reason, the shape of outputs of depthwise convolution are $(n, c_I * c_M, h_O, w_O)$, c_M is called `channel_multiplier`.

(h_O, w_O) is determined by the equivalent equation of `Convolution2D`.

If the bias vector is given, then it is added to all spatial locations of the output of convolution.

See: L. Sifre. Rigid-motion scattering for image classification

See also:

DepthwiseConvolution2D

Example

```
>>> x = np.random.uniform(0, 1, (2, 3, 4, 7))
>>> W = np.random.uniform(0, 1, (2, 3, 3, 3))
>>> b = np.random.uniform(0, 1, (6,))
>>> y = F.depthwise_convolution_2d(x, W, b)
>>> y.shape
(2, 6, 2, 5)
```

chainer.functions.dilated_convolution_2d

`chainer.functions.dilated_convolution_2d(x, W, b=None, stride=1, pad=0, dilate=1, cover_all=False)`

Two-dimensional dilated convolution function.

This is an implementation of two-dimensional dilated convolution in ConvNets. It takes three variables: the input image x , the filter weight W , and the bias vector b .

Note: You can also perform dilated convolution by passing `dilate` argument to `chainer.functions.convolution_2d`. The functionality is the same.

Notation: here is a notation for dimensionalities.

- n is the batch size.
- c_I and c_O are the number of the input and output, respectively.
- h and w are the height and width of the input image, respectively.
- k_H and k_W are the height and width of the filters, respectively.

Parameters

- **x** (*Variable*) – Input variable of shape (n, c_I, h, w) .
- **w** (*Variable*) – Weight variable of shape (c_O, c_I, k_H, k_W) .
- **b** (*Variable*) – Bias variable of length c_O (optional).
- **stride** (*int* or *pair of ints*) – Stride of filter applications. `stride=s` and `stride=(s, s)` are equivalent.
- **pad** (*int* or *pair of ints*) – Spatial padding width for input arrays. `pad=p` and `pad=(p, p)` are equivalent.
- **dilate** (*int* or *pair of ints*) – Dilation factor of filter applications. `dilate=d` and `dilate=(d, d)` are equivalent.
- **cover_all** (*bool*) – If `True`, all spatial locations are convoluted into some output pixels. It may make the output size larger.

Returns Output variable.

Return type *Variable*

The two-dimensional dilated convolution function is defined as follows. Then the `DilatedConvolution2D` function computes correlations between filters and patches of size (k_H, k_W) in **x**. Patches here are extracted at intervals of the dilation factor. Note that correlation here is equivalent to the inner product between expanded vectors. Patches are extracted at intervals of the dilation factor and at positions shifted by multiples of `stride` from the first position `-pad` for each spatial axis. The right-most (or bottom-most) patches do not run over the padded spatial size.

Let (s_Y, s_X) be the stride of filter application, (p_H, p_W) the spatial padding size, and (d_Y, d_X) the dilation factor of filter application. Then, the output size (h_O, w_O) is determined by the following equations:

$$\begin{aligned}h_O &= (h + 2p_H - k_H - (k_H - 1) * (d_Y - 1)) / s_Y + 1, \\w_O &= (w + 2p_W - k_W - (k_W - 1) * (d_X - 1)) / s_X + 1.\end{aligned}$$

If the bias vector is given, then it is added to all spatial locations of the output of convolution.

chainer.functions.embed_id

`chainer.functions.embed_id(x, W, ignore_label=None)`

Efficient linear function for one-hot input.

This function implements so called *word embeddings*. It takes two arguments: a set of IDs (words) **x** in B dimensional integer vector, and a set of all ID (word) embeddings **W** in $V \times d$ float32 matrix. It outputs $B \times d$ matrix whose *i*-th column is the **x**[*i*]-th column of **W**.

This function is only differentiable on the input **W**.

Parameters

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Batch vectors of IDs. Each element must be signed integer.
- **W** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Distributed representation of each ID (a.k.a. word embeddings).
- **ignore_label** (*int* or `None`) – If `ignore_label` is an int value, *i*-th column of return value is filled with 0.

Returns Output variable.

Return type *Variable*

See also:

[EmbedID](#)

Example

```
>>> x = np.array([2, 1]).astype(np.int32)
>>> x
array([2, 1], dtype=int32)
>>> W = np.array([[0, 0, 0],
...               [1, 1, 1],
...               [2, 2, 2]]).astype(np.float32)
>>> W
array([[0., 0., 0.],
       [1., 1., 1.],
       [2., 2., 2.]], dtype=float32)
>>> F.embed_id(x, W).data
array([[2., 2., 2.],
       [1., 1., 1.]], dtype=float32)
>>> F.embed_id(x, W, ignore_label=1).data
array([[2., 2., 2.],
       [0., 0., 0.]], dtype=float32)
```

chainer.functions.linear

`chainer.functions.linear(x, W, b=None)`

Linear function, or affine transformation.

It accepts two or three arguments: an input minibatch x , a weight matrix W , and optionally a bias vector b . It computes

$$Y = xW^T + b.$$

Parameters

- x ([Variable](#) or `numpy.ndarray` or `cupy.ndarray`) – Input variable, which is a $(s_B, s_1, s_2, \dots, s_n)$ -shaped float array. Its first dimension (s_B) is assumed to be the *mini-batch dimension*. The other dimensions are treated as concatenated one dimension whose size must be $(s_1 * \dots * s_n = N)$.
- W ([Variable](#) or `numpy.ndarray` or `cupy.ndarray`) – Weight variable of shape (M, N) , where $(N = s_1 * \dots * s_n)$.
- b ([Variable](#) or `numpy.ndarray` or `cupy.ndarray`) – Bias variable (optional) of shape $(M,)$.

Returns Output variable. A float array with shape of (s_B, M) .

Return type [Variable](#)

See also:

[Linear](#)

Example

```
>>> x = np.random.uniform(0, 1, (3, 4)).astype(np.float32)
>>> W = np.random.uniform(0, 1, (5, 4)).astype(np.float32)
>>> b = np.random.uniform(0, 1, (5,)).astype(np.float32)
>>> y = F.linear(x, W, b)
>>> y.shape
(3, 5)
```

chainer.functions.local_convolution_2d

`chainer.functions.local_convolution_2d(x, W, b=None, stride=1)`

Two-dimensional local convolution function.

Locally-connected function for 2D inputs. Works similarly to `convolution_2d`, except that weights are unshared, that is, a different set of filters is applied at each different patch of the input. It takes two or three variables: the input image x , the filter weight W , and optionally, the bias vector b .

Notation: here is a notation for dimensionalities.

- n is the batch size.
- c_I is the number of the input.
- c_O is the number of output channels.
- h and w are the height and width of the input image, respectively.
- h_O and w_O are the height and width of the output image, respectively.
- k_H and k_W are the height and width of the filters, respectively.

Parameters

- **x** (`chainer.Variable` or `numpy.ndarray` or `cupy.ndarray`) – Input variable of shape (n, c_I, h, w) .
- **W** (`Variable`) – Weight variable of shape $(c_O, h_O, w_O, c_I, k_H, k_W)$.
- **b** (`Variable`) – Bias variable of shape (c_O, h_O, w_O) (optional).
- **stride** (`int` or `pair of ints`) – Stride of filter applications. `stride=s` and `stride=(s, s)` are equivalent.

Returns Output variable. Its shape is $(n, c_I * c_O, h_O, w_O)$.

Return type `Variable`

Like `Convolution2D`, `LocalConvolution2D` function computes correlations between filters and patches of size (k_H, k_W) in x . But unlike `Convolution2D`, `LocalConvolution2D` has a separate filter for each patch of the input

(h_O, w_O) is determined by the equivalent equation of `Convolution2D`, without any padding

If the bias vector is given, then it is added to all spatial locations of the output of convolution.

See also:

`LocalConvolution2D`

Example


```

>>> x = np.random.uniform(0, 1, (2, 3, 7, 7))
>>> W = np.random.uniform(0, 1, (2, 5, 3, 3))
>>> b = np.random.uniform(0, 1, (2, 5, 5))
>>> y = F.local_convolution_2d(x, W, b)
>>> y.shape
(2, 2, 5, 5)

```

chainer.functions.n_step_bigru

`chainer.functions.n_step_bigru(n_layers, dropout_ratio, hx, ws, bs, xs)`

Stacked Bi-directional Gated Recurrent Unit function.

This function calculates stacked Bi-directional GRU with sequences. This function gets an initial hidden state h_0 , an input sequence x , weight matrices W , and bias vectors b . This function calculates hidden states h_t for each time t from input x_t .

$$\begin{aligned}
 r_t^f &= \sigma(W_0^f x_t + W_3^f h_{t-1} + b_0^f + b_3^f) \\
 z_t^f &= \sigma(W_1^f x_t + W_4^f h_{t-1} + b_1^f + b_4^f) \\
 h_t^{f'} &= \tanh(W_2^f x_t + b_2^f + r_t^f \cdot (W_5^f h_{t-1} + b_5^f)) \\
 h_t^f &= (1 - z_t^f) \cdot h_t^{f'} + z_t^f \cdot h_{t-1} \\
 r_t^b &= \sigma(W_0^b x_t + W_3^b h_{t-1} + b_0^b + b_3^b) \\
 z_t^b &= \sigma(W_1^b x_t + W_4^b h_{t-1} + b_1^b + b_4^b) \\
 h_t^{b'} &= \tanh(W_2^b x_t + b_2^b + r_t^b \cdot (W_5^b h_{t-1} + b_5^b)) \\
 h_t^b &= (1 - z_t^b) \cdot h_t^{b'} + z_t^b \cdot h_{t-1} \\
 h_t &= [h_t^f; h_t^b]
 \end{aligned}$$

where W^f is weight matrices for forward-GRU, W^b is weight matrices for backward-GRU.

As the function accepts a sequence, it calculates h_t for all t with one call. Six weight matrices and six bias vectors are required for each layers. So, when S layers exists, you need to prepare $6S$ weight matrices and $6S$ bias vectors.

If the number of layers `n_layers` is greater than 1, input of k -th layer is hidden state h_{t-1} of $k-1$ -th layer. Note that all input variables except first layer may have different shape from the first layer.

Warning: `train` and `use_cudnn` arguments are not supported anymore since v2. Instead, use `chainer.config('train', train)` and `chainer.config('use_cudnn', use_cudnn)` respectively. See [chainer.config\(\)](#).

Parameters

- **n_layers** (*int*) – Number of layers.
- **dropout_ratio** (*float*) – Dropout ratio.
- **hx** (*chainer.Variable*) – Variable holding stacked hidden states. Its shape is $(2S, B, N)$ where S is number of layers and is equal to `n_layers`, B is mini-batch size, and N is dimension of hidden units.

- **ws** (*list of list of chainer.Variable*) – Weight matrices. `ws[i]` represents weights for *i*-th layer. Each `ws[i]` is a list containing six matrices. `ws[i][j]` is corresponding with W_j in the equation. Only `ws[0][j]` where $0 \leq j < 3$ is (I, N) shape as they are multiplied with input variables. All other matrices has (N, N) shape.
- **bs** (*list of list of chainer.Variable*) – Bias vectors. `bs[i]` represents biases for *i*-th layer. Each `bs[i]` is a list containing six vectors. `bs[i][j]` is corresponding with b_j in the equation. Shape of each matrix is $(N,)$ where N is dimension of hidden units.
- **xs** (*list of chainer.Variable*) – A list of `Variable` holding input values. Each element `xs[t]` holds input value for time *t*. Its shape is (B_t, I) , where B_t is mini-batch size for time *t*, and I is size of input units. Note that this function supports variable length sequences. When sequences has different lengths, sort sequences in descending order by length, and transpose the sorted sequence. `transpose_sequence()` transpose a list of `Variable()` holding sequence. So `xs` needs to satisfy `xs[t].shape[0] >= xs[t + 1].shape[0]`.
- **use_bi_direction** (*bool*) – If `True`, this function uses Bi-direction GRU.

Returns

This function returns a tuple containing three elements, `hy` and `ys`.

- `hy` is an updated hidden states whose shape is same as `hx`.
- `ys` is a list of `Variable`. Each element `ys[t]` holds hidden states of the last layer corresponding to an input `xs[t]`. Its shape is (B_t, N) where B_t is mini-batch size for time *t*, and N is size of hidden units. Note that B_t is the same value as `xs[t]`.

Return type `tuple`

chainer.functions.n_step_bilstm

`chainer.functions.n_step_bilstm(n_layers, dropout_ratio, hx, cx, ws, bs, xs)`

Stacked Bi-directional Long Short-Term Memory function.

This function calculates stacked Bi-directional LSTM with sequences. This function gets an initial hidden state h_0 , an initial cell state c_0 , an input sequence x , weight matrices W , and bias vectors b . This function calculates

hidden states h_t and c_t for each time t from input x_t .

$$\begin{aligned}
 i_t^f &= \sigma(W_0^f x_t + W_4^f h_{t-1} + b_0^f + b_4^f), \\
 f_t^f &= \sigma(W_1^f x_t + W_5^f h_{t-1} + b_1^f + b_5^f), \\
 o_t^f &= \sigma(W_2^f x_t + W_6^f h_{t-1} + b_2^f + b_6^f), \\
 a_t^f &= \tanh(W_3^f x_t + W_7^f h_{t-1} + b_3^f + b_7^f), \\
 c_t^f &= f_t^f \cdot c_{t-1}^f + i_t^f \cdot a_t^f, \\
 h_t^f &= o_t^f \cdot \tanh(c_t^f), \\
 i_t^b &= \sigma(W_0^b x_t + W_4^b h_{t-1} + b_0^b + b_4^b), \\
 f_t^b &= \sigma(W_1^b x_t + W_5^b h_{t-1} + b_1^b + b_5^b), \\
 o_t^b &= \sigma(W_2^b x_t + W_6^b h_{t-1} + b_2^b + b_6^b), \\
 a_t^b &= \tanh(W_3^b x_t + W_7^b h_{t-1} + b_3^b + b_7^b), \\
 c_t^b &= f_t^b \cdot c_{t-1}^b + i_t^b \cdot a_t^b, \\
 h_t^b &= o_t^b \cdot \tanh(c_t^b), \\
 h_t &= [h_t^f; h_t^b]
 \end{aligned}$$

where W^f is the weight matrices for forward-LSTM, W^b is weight matrices for backward-LSTM.

As the function accepts a sequence, it calculates h_t for all t with one call. Eight weight matrices and eight bias vectors are required for each layer of each direction. So, when S layers exist, you need to prepare $16S$ weight matrices and $16S$ bias vectors.

If the number of layers `n_layers` is greater than 1, the input of the k -th layer is the hidden state `h_t` of the $k-1$ -th layer. Note that all input variables except the first layer may have different shape from the first layer.

Warning: `train` and `use_cudnn` arguments are not supported anymore since v2. Instead, use `chainer.config('train', train)` and `chainer.config('use_cudnn', use_cudnn)` respectively. See [chainer.config\(\)](#).

Parameters

- **n_layers** (*int*) – The number of layers.

- **dropout_ratio** (*float*) – Dropout ratio.
- **hx** (*Variable*) – Variable holding stacked hidden states. Its shape is $(2S, B, N)$ where S is the number of layers and is equal to `n_layers`, B is the mini-batch size, and N is the dimension of the hidden units. Because of bi-direction, the first dimension length is $2S$.
- **cx** (*Variable*) – Variable holding stacked cell states. It has the same shape as `hx`.
- **ws** (list of list of *Variable*) – Weight matrices. `ws[2 * l + m]` represents the weights for the l -th layer of the m -th direction. ($m == 0$ means the forward direction and $m == 1$ means the backward direction.) Each `ws[i]` is a list containing eight matrices. `ws[i][j]` corresponds to W_j in the equation. `ws[0][j]` and `ws[1][j]` where $0 \leq j < 4$ are (I, N) -shaped because they are multiplied with input variables, where I is the size of the input. `ws[i][j]` where $2 \leq i$ and $0 \leq j < 4$ are $(N, 2N)$ -shaped because they are multiplied with two hidden layers $h_t = [h_t^f; h_t^b]$. All other matrices are (N, N) -shaped.
- **bs** (list of list of *Variable*) – Bias vectors. `bs[2 * l + m]` represents the weights for the l -th layer of m -th direction. ($m == 0$ means the forward direction and $m == 1$ means the backward direction.) Each `bs[i]` is a list containing eight vectors. `bs[i][j]` corresponds to b_j in the equation. The shape of each matrix is $(N,)$.
- **xs** (list of *Variable*) – A list of *Variable* holding input values. Each element `xs[t]` holds input value for time t . Its shape is (B_t, I) , where B_t is the mini-batch size for time t . The sequences must be transposed. `transpose_sequence()` can be used to transpose a list of *Variables* each representing a sequence. When sequences has different lengths, they must be sorted in descending order of their lengths before transposing. So `xs` needs to satisfy `xs[t].shape[0] >= xs[t + 1].shape[0]`.

Returns

This function returns a tuple containing three elements, `hy`, `cy` and `ys`.

- `hy` is an updated hidden states whose shape is the same as `hx`.
- `cy` is an updated cell states whose shape is the same as `cx`.
- `ys` is a list of *Variable*. Each element `ys[t]` holds hidden states of the last layer corresponding to an input `xs[t]`. Its shape is $(B_t, 2N)$ where B_t is the mini-batch size for time t , and N is size of hidden units. Note that B_t is the same value as `xs[t]`.

Return type `tuple`

Example

```
>>> batches = [3, 2, 1] # support variable length sequences
>>> in_size, out_size, n_layers = 3, 2, 2
>>> dropout_ratio = 0.0
>>> xs = [np.ones((b, in_size)).astype(np.float32) for b in batches]
>>> [x.shape for x in xs]
[(3, 3), (2, 3), (1, 3)]
>>> h_shape = (n_layers * 2, batches[0], out_size)
>>> hx = np.ones(h_shape).astype(np.float32)
>>> cx = np.ones(h_shape).astype(np.float32)
>>> def w_in(i, j):
...     if i == 0 and j < 4:
...         return in_size
...     elif i > 0 and j < 4:
...         return out_size * 2
...     else:
```

(continues on next page)

(continued from previous page)

```

...         return out_size
...
>>> ws = []
>>> bs = []
>>> for n in range(n_layers):
...     for direction in (0, 1):
...         ws.append([np.ones((out_size, w_in(n, i))).astype(np.float32) for i_
↪in range(8)])
...         bs.append([np.ones((out_size,)).astype(np.float32) for _ in range(8)])
...
>>> ws[0][0].shape # ws[0:2][:4].shape are (out_size, in_size)
(2, 3)
>>> ws[2][0].shape # ws[2:][:4].shape are (out_size, 2 * out_size)
(2, 4)
>>> ws[0][4].shape # others are (out_size, out_size)
(2, 2)
>>> bs[0][0].shape
(2,)
>>> hy, cy, ys = F.n_step_bilstm(
...     n_layers, dropout_ratio, hx, cx, ws, bs, xs)
>>> hy.shape
(4, 3, 2)
>>> cy.shape
(4, 3, 2)
>>> [y.shape for y in ys]
[(3, 4), (2, 4), (1, 4)]

```

chainer.functions.n_step_birnn

`chainer.functions.n_step_birnn(n_layers, dropout_ratio, hx, ws, bs, xs, activation='tanh')`

Stacked Bi-directional RNN function for sequence inputs.

This function calculates stacked Bi-directional RNN with sequences. This function gets an initial hidden state h_0 , an initial cell state c_0 , an input sequence x , weight matrices W , and bias vectors b . This function calculates hidden states h_t and c_t for each time t from input x_t .

$$\begin{aligned}
 h_t^f &= \\
 &f(W_0^f x_t + W_1^f h_{t-1} + b_0^f + b_1^f), \\
 h_t^b &= \\
 &f(W_0^b x_t + W_1^b h_{t-1} + b_0^b + b_1^b), \\
 h_t &= \\
 &[h_t^f; h_t^b],
 \end{aligned}$$

where f is an activation function.

Weight matrices W contains two matrices W^f and W^b . W^f is weight matrices for forward directional RNN. W^b is weight matrices for backward directional RNN.

W^f contains W_0^f for an input sequence and W_1^f for a hidden state. W^b contains W_0^b for an input sequence and W_1^b for a hidden state.

Bias matrices b contains two matrices b^f and b^b . b^f contains b_0^f for an input sequence and b_1^f for a hidden state. b^b contains b_0^b for an input sequence and b_1^b for a hidden state.

As the function accepts a sequence, it calculates h_t for all t with one call. Two weight matrices and two bias vectors are required for each layer. So, when S layers exist, you need to prepare $2S$ weight matrices and $2S$ bias vectors.

If the number of layers `n_layers` is greater than 1, input of k -th layer is hidden state h_{t-1} of $k-1$ -th layer. Note that all input variables except first layer may have different shape from the first layer.

Warning: `train` and `use_cudnn` arguments are not supported anymore since v2. Instead, use `chainer.using_config('train', train)` and `chainer.using_config('use_cudnn', use_cudnn)` respectively. See `chainer.using_config()`.

Parameters

- **n_layers** (*int*) – Number of layers.
- **dropout_ratio** (*float*) – Dropout ratio.
- **hx** (*chainer.Variable*) – Variable holding stacked hidden states. Its shape is $(2S, B, N)$ where S is number of layers and is equal to `n_layers`, B is mini-batch size, and N is dimension of hidden units. Because of bi-direction, the first dimension length is $2S$.
- **ws** (*list of list of chainer.Variable*) – Weight matrices. `ws[i + di]` represents weights for i -th layer. Note that $di = 0$ for forward-RNN and $di = 1$ for backward-RNN. Each `ws[i + di]` is a list containing two matrices. `ws[i + di][j]` is corresponding with $W^{\{f\}}_j$ if $di = 0$ and corresponding with $W^{\{b\}}_j$ if $di = 1$ in the equation. Only `ws[0][j]` and `ws[1][j]` where $0 \leq j < 1$ are (I, N) shape as they are multiplied with input variables. All other matrices has (N, N) shape.
- **bs** (*list of list of chainer.Variable*) – Bias vectors. `bs[i + di]` represents biases for i -th layer. Note that $di = 0$ for forward-RNN and $di = 1$ for backward-RNN. Each `bs[i + di]` is a list containing two vectors. `bs[i + di][j]` is corresponding with $b^{\{f\}}_j$ if $di = 0$ and corresponding with $b^{\{b\}}_j$ if $di = 1$ in the equation. Shape of each matrix is $(N,)$ where N is dimension of hidden units.
- **xs** (*list of chainer.Variable*) – A list of *Variable* holding input values. Each element `xs[t]` holds input value for time t . Its shape is (B_t, I) , where B_t is mini-batch size for time t , and I is size of input units. Note that this function supports variable length sequences. When sequences has different lengths, sort sequences in descending order by length, and transpose the sorted sequence. `transpose_sequence()` transpose a list of *Variable()* holding sequence. So `xs` needs to satisfy `xs[t].shape[0] >= xs[t + 1].shape[0]`.
- **activation** (*str*) – Activation function name. Please select `tanh` or `relu`.

Returns

This function returns a tuple containing three elements, `hy` and `ys`.

- `hy` is an updated hidden states whose shape is same as `hx`.
- `ys` is a list of *Variable*. Each element `ys[t]` holds hidden states of the last layer corresponding to an input `xs[t]`. Its shape is (B_t, N) where B_t is mini-batch size for time t , and N is size of hidden units. Note that B_t is the same value as `xs[t]`.

Return type `tuple`

chainer.functions.n_step_gru

`chainer.functions.n_step_gru(n_layers, dropout_ratio, hx, ws, bs, xs)`

Stacked Uni-directional Gated Recurrent Unit function.

This function calculates stacked Uni-directional GRU with sequences. This function gets an initial hidden state h_0 , an input sequence x , weight matrices W , and bias vectors b . This function calculates hidden states h_t for each time t from input x_t .

$$\begin{aligned}
 r_t &= \sigma(W_0 x_t + W_3 h_{t-1} + b_0 + b_3) \\
 z_t &= \sigma(W_1 x_t + W_4 h_{t-1} + b_1 + b_4) \\
 h'_t &= \tanh(W_2 x_t + b_2 + r_t \cdot (W_5 h_{t-1} + b_5)) \\
 h_t &= (1 - z_t) \cdot h'_t + z_t \cdot h_{t-1}
 \end{aligned}$$

As the function accepts a sequence, it calculates h_t for all t with one call. Six weight matrices and six bias vectors are required for each layers. So, when S layers exists, you need to prepare $6S$ weight matrices and $6S$ bias vectors.

If the number of layers `n_layers` is greather than 1, input of k -th layer is hidden state h_{t-1} of $k-1$ -th layer. Note that all input variables except first layer may have different shape from the first layer.

Warning: `train` and `use_cudnn` arguments are not supported anymore since v2. Instead, use `chainer.using_config('train', train)` and `chainer.using_config('use_cudnn', use_cudnn)` respectively. See [chainer.using_config\(\)](#).

Parameters

- **n_layers** (*int*) – Number of layers.
- **dropout_ratio** (*float*) – Dropout ratio.
- **hx** (*chainer.Variable*) – Variable holding stacked hidden states. Its shape is (S, B, N) where S is number of layers and is equal to `n_layers`, B is mini-batch size, and N is dimension of hidden units.
- **ws** (*list of list of chainer.Variable*) – Weight matrices. `ws[i]` represents weights for i -th layer. Each `ws[i]` is a list containing six matrices. `ws[i][j]` is corresponding with W_j in the equation. Only `ws[0][j]` where $0 \leq j < 3$ is (I, N) shape as they are multiplied with input variables. All other matrices has (N, N) shape.
- **bs** (*list of list of chainer.Variable*) – Bias vectors. `bs[i]` represents biases for i -th layer. Each `bs[i]` is a list containing six vectors. `bs[i][j]` is corresponding with b_j in the equation. Shape of each matrix is $(N,)$ where N is dimension of hidden units.
- **xs** (*list of chainer.Variable*) – A list of *Variable* holding input values. Each element `xs[t]` holds input value for time t . Its shape is (B_t, I) , where B_t is mini-batch size for time t , and I is size of input units. Note that this function supports variable length sequences. When sequences has different lengths, sort sequences in descending order by length, and transpose the sorted sequence. `transpose_sequence()` transpose a list of *Variable()* holding sequence. So `xs` needs to satisfy `xs[t].shape[0] >= xs[t + 1].shape[0]`.

Returns

This function returns a tuple containing three elements, `hy` and `ys`.

- `hy` is an updated hidden states whose shape is same as `hx`.

- `ys` is a list of *Variable*. Each element `ys[t]` holds hidden states of the last layer corresponding to an input `xs[t]`. Its shape is (B_t, N) where B_t is mini-batch size for time t , and N is size of hidden units. Note that B_t is the same value as `xs[t]`.

Return type `tuple`

`chainer.functions.n_step_lstm`

`chainer.functions.n_step_lstm(n_layers, dropout_ratio, hx, cx, ws, bs, xs)`

Stacked Uni-directional Long Short-Term Memory function.

This function calculates stacked Uni-directional LSTM with sequences. This function gets an initial hidden state h_0 , an initial cell state c_0 , an input sequence x , weight matrices W , and bias vectors b . This function calculates hidden states h_t and c_t for each time t from input x_t .

$$\begin{aligned} i_t &= \sigma(W_0 x_t + W_4 h_{t-1} + b_0 + b_4) \\ f_t &= \sigma(W_1 x_t + W_5 h_{t-1} + b_1 + b_5) \\ o_t &= \sigma(W_2 x_t + W_6 h_{t-1} + b_2 + b_6) \\ a_t &= \tanh(W_3 x_t + W_7 h_{t-1} + b_3 + b_7) \\ c_t &= f_t \cdot c_{t-1} + i_t \cdot a_t \\ h_t &= o_t \cdot \tanh(c_t) \end{aligned}$$

As the function accepts a sequence, it calculates h_t for all t with one call. Eight weight matrices and eight bias vectors are required for each layer. So, when S layers exist, you need to prepare $8S$ weight matrices and $8S$ bias vectors.

If the number of layers `n_layers` is greater than 1, the input of the k -th layer is the hidden state h_t of the $k-1$ -th layer. Note that all input variables except the first layer may have different shape from the first layer.

Warning: `train` and `use_cudnn` arguments are not supported anymore since v2. Instead, use `chainer.config('train', train)` and `chainer.config('use_cudnn', use_cudnn)` respectively. See [chainer.config\(\)](#).

Parameters

- **`n_layers`** (*int*) – The number of layers.
- **`dropout_ratio`** (*float*) – Dropout ratio.
- **`hx`** (*Variable*) – Variable holding stacked hidden states. Its shape is (S, B, N) where S is the number of layers and is equal to `n_layers`, B is the mini-batch size, and N is the dimension of the hidden units.
- **`cx`** (*Variable*) – Variable holding stacked cell states. It has the same shape as `hx`.
- **`ws`** (list of list of *Variable*) – Weight matrices. `ws[i]` represents the weights for the i -th layer. Each `ws[i]` is a list containing eight matrices. `ws[i][j]` corresponds to W_j in the equation. Only `ws[0][j]` where $0 \leq j < 4$ are (I, N) -shaped as they are multiplied with input variables, where I is the size of the input and N is the dimension of the hidden units. All other matrices are (N, N) -shaped.
- **`bs`** (list of list of *Variable*) – Bias vectors. `bs[i]` represents the biases for the i -th layer. Each `bs[i]` is a list containing eight vectors. `bs[i][j]` corresponds to b_j in the equation. The shape of each matrix is $(N,)$ where N is the dimension of the hidden units.

- **xs** (list of *Variable*) – A list of *Variable* holding input values. Each element `xs[t]` holds input value for time `t`. Its shape is `(B_t, I)`, where `B_t` is the mini-batch size for time `t`. The sequences must be transposed. `transpose_sequence()` can be used to transpose a list of *Variables* each representing a sequence. When sequences has different lengths, they must be sorted in descending order of their lengths before transposing. So `xs` needs to satisfy `xs[t].shape[0] >= xs[t + 1].shape[0]`.

Returns

This function returns a tuple containing three elements, `hy`, `cy` and `ys`.

- `hy` is an updated hidden states whose shape is the same as `hx`.
- `cy` is an updated cell states whose shape is the same as `cx`.
- `ys` is a list of *Variable*. Each element `ys[t]` holds hidden states of the last layer corresponding to an input `xs[t]`. Its shape is `(B_t, N)` where `B_t` is the mini-batch size for time `t`, and `N` is size of hidden units. Note that `B_t` is the same value as `xs[t]`.

Return type `tuple`

Note: The dimension of hidden units is limited to only one size `N`. If you want to use variable dimension of hidden units, please use `chainer.functions.lstm`.

See also:

`chainer.functions.lstm()`

Example

```
>>> batches = [3, 2, 1] # support variable length sequences
>>> in_size, out_size, n_layers = 3, 2, 2
>>> dropout_ratio = 0.0
>>> xs = [np.ones((b, in_size)).astype(np.float32) for b in batches]
>>> [x.shape for x in xs]
[(3, 3), (2, 3), (1, 3)]
>>> h_shape = (n_layers, batches[0], out_size)
>>> hx = np.ones(h_shape).astype(np.float32)
>>> cx = np.ones(h_shape).astype(np.float32)
>>> w_in = lambda i, j: in_size if i == 0 and j < 4 else out_size
>>> ws = []
>>> bs = []
>>> for n in range(n_layers):
...     ws.append([np.ones((out_size, w_in(n, i))).astype(np.float32) for i in
... ↪range(8)])
...     bs.append([np.ones((out_size,)).astype(np.float32) for _ in range(8)])
...
>>> ws[0][0].shape # ws[0][:4].shape are (out_size, in_size)
(2, 3)
>>> ws[1][0].shape # others are (out_size, out_size)
(2, 2)
>>> bs[0][0].shape
(2,)
```

(continues on next page)

(continued from previous page)

```
>>> cy.shape
(2, 3, 2)
>>> [y.shape for y in ys]
[(3, 2), (2, 2), (1, 2)]
```

chainer.functions.n_step_rnn

`chainer.functions.n_step_rnn(n_layers, dropout_ratio, hx, ws, bs, xs, activation='tanh')`

Stacked Uni-directional RNN function for sequence inputs.

This function calculates stacked Uni-directional RNN with sequences. This function gets an initial hidden state h_0 , an initial cell state c_0 , an input sequence x , weight matrices W , and bias vectors b . This function calculates hidden states h_t and c_t for each time t from input x_t .

$$h_t = f(W_0 x_t + W_1 h_{t-1} + b_0 + b_1)$$

where f is an activation function.

Weight matrices W contains two matrices W_0 and W_1 . W_0 is a parameter for an input sequence. W_1 is a parameter for a hidden state. Bias matrices b contains two matrices b_0 and b_1 . b_0 is a parameter for an input sequence. b_1 is a parameter for a hidden state.

As the function accepts a sequence, it calculates h_t for all t with one call. Two weight matrices and two bias vectors are required for each layer. So, when S layers exist, you need to prepare $2S$ weight matrices and $2S$ bias vectors.

If the number of layers `n_layers` is greater than 1, input of k -th layer is hidden state h_t of $k-1$ -th layer. Note that all input variables except first layer may have different shape from the first layer.

Warning: `train` and `use_cudnn` arguments are not supported anymore since v2. Instead, use `chainer.config('train', train)` and `chainer.config('use_cudnn', use_cudnn)` respectively. See [chainer.config\(\)](#).

Parameters

- **n_layers** (*int*) – Number of layers.
- **dropout_ratio** (*float*) – Dropout ratio.
- **hx** (*chainer.Variable*) – Variable holding stacked hidden states. Its shape is (S, B, N) where S is number of layers and is equal to `n_layers`, B is mini-batch size, and N is dimension of hidden units.
- **ws** (*list of list of chainer.Variable*) – Weight matrices. `ws[i]` represents weights for i -th layer. Each `ws[i]` is a list containing two matrices. `ws[i][j]` is corresponding with W_j in the equation. Only `ws[0][j]` where $0 \leq j < 1$ is $(1, N)$ shape as they are multiplied with input variables. All other matrices has (N, N) shape.
- **bs** (*list of list of chainer.Variable*) – Bias vectors. `bs[i]` represents biases for i -th layer. Each `bs[i]` is a list containing two vectors. `bs[i][j]` is corresponding with b_j in the equation. Shape of each matrix is $(N,)$ where N is dimension of hidden units.

- **xs** (*list of chainer.Variable*) – A list of *Variable* holding input values. Each element `xs[t]` holds input value for time `t`. Its shape is (B_t, I) , where B_t is mini-batch size for time `t`, and I is size of input units. Note that this function supports variable length sequences. When sequences has different lengths, sort sequences in descending order by length, and transpose the sorted sequence. `transpose_sequence()` transpose a list of *Variable()* holding sequence. So `xs` needs to satisfy `xs[t].shape[0] >= xs[t + 1].shape[0]`.
- **activation** (*str*) – Activation function name. Please select `tanh` or `relu`.

Returns

This function returns a tuple containing three elements, `hy` and `ys`.

- `hy` is an updated hidden states whose shape is same as `hx`.
- `ys` is a list of *Variable*. Each element `ys[t]` holds hidden states of the last layer corresponding to an input `xs[t]`. Its shape is (B_t, N) where B_t is mini-batch size for time `t`, and N is size of hidden units. Note that B_t is the same value as `xs[t]`.

Return type *tuple*

chainer.functions.shift

`chainer.functions.shift(x, ksize=3, dilate=1)`

Shift function.

See: [Shift: A Zero FLOP, Zero Parameter Alternative to Spatial Convolutions](#)

Parameters

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable of shape (n, c, h, w) .
- **ksize** (*int* or *pair of ints*) – Size of filters (a.k.a. kernels). `ksize=k` and `ksize=(k, k)` are equivalent.
- **dilate** (*int* or *pair of ints*) – Dilation factor of filter applications. `dilate=d` and `dilate=(d, d)` are equivalent.

Returns Output variable of same shape as `x`.

Return type *Variable*

5.2.5 Evaluation functions

<code>chainer.functions.accuracy</code>	Computes multiclass classification accuracy of the minibatch.
<code>chainer.functions.binary_accuracy</code>	Computes binary classification accuracy of the mini-batch.
<code>chainer.functions.classification_summary</code>	Calculates Precision, Recall, F beta Score, and support.
<code>chainer.functions.f1_score</code>	
<code>chainer.functions.precision</code>	
<code>chainer.functions.r2_score</code>	Computes R^2 (coefficient of determination) regression score function.
<code>chainer.functions.recall</code>	

chainer.functions.accuracy

`chainer.functions.accuracy(y, t, ignore_label=None)`

Computes multiclass classification accuracy of the minibatch.

Parameters

- **y** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Array whose (i, j, k, ...) -th element indicates the score of the class j at the (i, k, ...) -th sample. The prediction label \hat{t} is calculated by the formula $\hat{t}(i, k, \dots) = \operatorname{argmax}_j y(i, j, k, \dots)$.
- **t** (*Variable* or `numpy.ndarray` or `cupy.ndarray` of signed integer) – Array of ground truth labels.
- **ignore_label** (*int* or *None*) – Skip calculating accuracy if the true label is `ignore_label`.

Returns A variable holding a scalar array of the accuracy.

Return type *Variable*

Note: This function is non-differentiable.

Example

We show the most common case, when `y` is the two dimensional array.

```
>>> y = np.array([[0.1, 0.7, 0.2], # prediction label is 1
...               [8.0, 1.0, 2.0], # prediction label is 0
...               [-8.0, 1.0, 2.0], # prediction label is 2
...               [-8.0, -1.0, -2.0]]) # prediction label is 1
>>> t = np.array([1, 0, 2, 1], np.int32)
>>> F.accuracy(y, t).data # 100% accuracy because all samples are correct
array(1.)
>>> t = np.array([1, 0, 0, 0], np.int32)
>>> F.accuracy(y, t).data # 50% accuracy because 1st and 2nd samples are correct.
array(0.5)
>>> F.accuracy(y, t, ignore_label=0).data # 100% accuracy because of ignoring the
↪2nd, 3rd and 4th samples.
array(1.)
```

chainer.functions.binary_accuracy

`chainer.functions.binary_accuracy(y, t)`

Computes binary classification accuracy of the minibatch.

Parameters

- **y** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Array whose i-th element indicates the score of positive at the i-th sample. The prediction label $\hat{t}[i]$ is 1 if `y[i] >= 0`, otherwise 0.
- **t** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Array holding a signed integer vector of ground truth labels. If `t[i] == 1`, it indicates that i-th sample is positive. If `t[i] == 0`, it indicates that i-th sample is negative. If `t[i] == -1`, corresponding `y[i]` is ignored. Accuracy is zero if all ground truth labels are -1.

Returns A variable holding a scalar array of the accuracy.

Return type *Variable*

Note: This function is non-differentiable.

Example

We show the most common case, when `y` is the two dimensional array.

```
>>> y = np.array([[ -2.0,  0.0], # prediction labels are [0, 1]
...               [ 3.0, -5.0]]) # prediction labels are [1, 0]
>>> t = np.array([[0, 1],
...               [1, 0]], np.int32)
>>> F.binary_accuracy(y, t).data # 100% accuracy because all samples are correct.
array(1.)
>>> t = np.array([[0, 0],
...               [1, 1]], np.int32)
>>> F.binary_accuracy(y, t).data # 50% accuracy because y[0][0] and y[1][0] are
↳correct.
array(0.5)
>>> t = np.array([[0, -1],
...               [1, -1]], np.int32)
>>> F.binary_accuracy(y, t).data # 100% accuracy because of ignoring y[0][1] and
↳y[1][1].
array(1.)
```

chainer.functions.classification_summary

`chainer.functions.classification_summary(y, t, label_num=None, beta=1.0, ignore_label=-1)`

Calculates Precision, Recall, F beta Score, and support.

This function calculates the following quantities for each class.

- Precision: $\frac{tp}{tp+fp}$
- Recall: $\frac{tp}{tp+tn}$
- F beta Score: The weighted harmonic average of Precision and Recall.
- Support: The number of instances of each ground truth label.

Here, `tp`, `fp`, and `tn` stand for the number of true positives, false positives, and true negative, respectively.

`label_num` specifies the number of classes, that is, each value in `t` must be an integer in the range of `[0, label_num)`. If `label_num` is `None`, this function regards `label_num` as a maximum of in `t` plus one.

`ignore_label` determines which instances should be ignored. Specifically, instances with the given label are not taken into account for calculating the above quantities. By default, it is set to `-1` so that all instances are taken into consideration, as labels are supposed to be non-negative integers. Setting `ignore_label` to a non-negative integer less than `label_num` is illegal and yields undefined behavior. In the current implementation, it arises `RuntimeWarning` and `ignore_label`-th entries in output arrays do not contain correct quantities.

Parameters

- `y` (*Variable*) – Variable holding a vector of scores.

- **t** (*Variable*) – Variable holding a vector of ground truth labels.
- **label_num** (*int*) – The number of classes.
- **beta** (*float*) – The parameter which determines the weight of precision in the F-beta score.
- **ignore_label** (*int*) – Instances with this label are ignored.

Returns 4-tuple of ~chainer.Variable of size (label_num,). Each element represents precision, recall, F beta score, and support of this minibatch.

chainer.functions.f1_score

`chainer.functions.f1_score(y, t, label_num=None, ignore_label=-1)`

chainer.functions.precision

`chainer.functions.precision(y, t, label_num=None, ignore_label=-1)`

chainer.functions.r2_score

`chainer.functions.r2_score(pred, true, sample_weight=None, multioutput='uniform_average')`

Computes R²(coefficient of determination) regression score function.

Parameters

- **pred** (*Variable*) – Variable holding a vector, matrix or tensor of estimated target values.
- **true** (*Variable*) – Variable holding a vector, matrix or tensor of correct target values.
- **sample_weight** – This argument is for compatibility with scikit-learn’s implementation of r2_score. Current implementation admits None only.
- **multioutput** (*string*) – [‘uniform_average’, ‘raw_values’]. If ‘uniform_average’, this function returns an average of R² score of multiple output. If ‘raw_average’, this function return a set of R² score of multiple output.

Returns A Variable holding a scalar array of the R² score if ‘multioutput’ is ‘uniform_average’ or a vector of R² scores if ‘multioutput’ is ‘raw_values’.

Return type *Variable*

Note: This function is non-differentiable.

chainer.functions.recall

`chainer.functions.recall(y, t, label_num=None, ignore_label=-1)`

5.2.6 Loss functions

chainer.functions.absolute_error

Element-wise absolute error function.

Continued on next page

Table 7 – continued from previous page

<code>chainer.functions.bernoulli_nll</code>	Computes the negative log-likelihood of a Bernoulli distribution.
<code>chainer.functions.black_out</code>	BlackOut loss function.
<code>chainer.functions.connectionist_temporal_classification</code>	Connectionist Temporal Classification loss function.
<code>chainer.functions.contrastive</code>	Computes contrastive loss.
<code>chainer.functions.crfld</code>	Calculates negative log-likelihood of linear-chain CRF.
<code>chainer.functions.argmax_crfld</code>	Computes a state that maximizes a joint probability of the given CRF.
<code>chainer.functions.cross_covariance</code>	Computes the sum-squared cross-covariance penalty between y and z .
<code>chainer.functions.decov</code>	Computes the DeCov loss of h .
<code>chainer.functions.gaussian_kl_divergence</code>	Computes the KL-divergence of Gaussian variables from the standard one.
<code>chainer.functions.gaussian_nll</code>	Computes the negative log-likelihood of a Gaussian distribution.
<code>chainer.functions.hinge</code>	Computes the hinge loss for a one-of-many classification task.
<code>chainer.functions.huber_loss</code>	Computes the Huber loss.
<code>chainer.functions.mean_absolute_error</code>	Mean absolute error function.
<code>chainer.functions.mean_squared_error</code>	Mean squared error function.
<code>chainer.functions.negative_sampling</code>	Negative sampling loss function.
<code>chainer.functions.sigmoid_cross_entropy</code>	Computes cross entropy loss for pre-sigmoid activations.
<code>chainer.functions.softmax_cross_entropy</code>	Computes cross entropy loss for pre-softmax activations.
<code>chainer.functions.squared_error</code>	Squared error function.
<code>chainer.functions.triplet</code>	Computes triplet loss.

chainer.functions.absolute_error

`chainer.functions.absolute_error(x0, x1)`

Element-wise absolute error function.

Computes the element-wise absolute error L between two inputs x_0 and x_1 defined as follows.

$$L = |x_0 - x_1|$$

Parameters

- **x0** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – First input variable.
- **x1** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Second input variable.

Returns An array representing the element-wise absolute error between the two inputs.

Return type *Variable*

chainer.functions.bernoulli_nll

`chainer.functions.bernoulli_nll(x, y, reduce='sum')`

Computes the negative log-likelihood of a Bernoulli distribution.

This function calculates the negative log-likelihood of a Bernoulli distribution.

$$-\log B(x; p) = -\sum_i \{x_i \log(p_i) + (1 - x_i) \log(1 - p_i)\},$$

where $p = \sigma(y)$, $\sigma(\cdot)$ is a sigmoid function, and $B(x; p)$ is a Bernoulli distribution.

The output is a variable whose value depends on the value of the option `reduce`. If it is `'no'`, it holds the elementwise loss values. If it is `'sum'`, loss values are summed up.

Note: As this function uses a sigmoid function, you can pass a result of fully-connected layer (that means `Linear`) to this function directly.

Parameters

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable.
- **y** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – A variable representing the parameter of Bernoulli distribution.
- **reduce** (*str*) – Reduction option. Its value must be either `'sum'` or `'no'`. Otherwise, `ValueError` is raised.

Returns A variable representing the negative log-likelihood. If `reduce` is `'no'`, the output variable holds array whose shape is same as one of (hence both of) input variables. If it is `'sum'`, the output variable holds a scalar value.

Return type *Variable*

chainer.functions.black_out

`chainer.functions.black_out(x, t, W, samples, reduce='mean')`

BlackOut loss function.

BlackOut loss function is defined as

$$-\log(p(t)) - \sum_{s \in S} \log(1 - p(s)),$$

where t is the correct label, S is a set of negative examples and $p(\cdot)$ is likelihood of a given label. And, p is defined as

$$p(y) = \frac{\exp(W_y^\top x)}{\sum_{s \in \text{samples}} \exp(W_s^\top x)}.$$

The output is a variable whose value depends on the value of the option `reduce`. If it is `'no'`, it holds the no loss values. If it is `'mean'`, this function takes a mean of loss values.

Parameters

- **x** (*Variable*) – Batch of input vectors. Its shape should be (N, D) .
- **t** (*Variable*) – Vector of ground truth labels. Its shape should be $(N,)$. Each elements v should satisfy $0 \leq v \leq V$ or -1 where V is the number of label types.
- **W** (*Variable*) – Weight matrix. Its shape should be (V, D)
- **samples** (*Variable*) – Negative samples. Its shape should be (N, S) where S is the number of negative samples.

- **reduce** (*str*) – Reduction option. Its value must be either 'no' or 'mean'. Otherwise, `ValueError` is raised.

Returns A variable object holding loss value(s). If `reduce` is 'no', the output variable holds an array whose shape is $(N,)$. If it is 'mean', it holds a scalar.

Return type *Variable*

See: [BlackOut: Speeding up Recurrent Neural Network Language Models With Very Large Vocabularies](#)

See also:

[BlackOut](#).

chainer.functions.connectionist_temporal_classification

```
chainer.functions.connectionist_temporal_classification(x, t, blank_symbol,
                                                         input_length=None,
                                                         label_length=None, reduce='mean')
```

Connectionist Temporal Classification loss function.

Connectionist Temporal Classification(CTC) [[Graves2006](#)] is a loss function of sequence labeling where the alignment between the inputs and target is unknown. See also [[Graves2012](#)]

The output is a variable whose value depends on the value of the option `reduce`. If it is 'no', it holds the samplewise loss values. If it is 'mean', it takes the mean of loss values.

Parameters

- **x** (list or tuple of *Variable*) – A list of unnormalized probabilities for labels. Each element of `x`, `x[i]` is a *Variable* object, which has shape (B, V) , where B is the batch size and V is the number of labels. The softmax of `x[i]` represents the probabilities of the labels at time i .
- **t** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – A matrix including expected label sequences. Its shape is (B, M) , where B is the batch size and M is the maximum length of the label sequences. All elements in `t` must be less than V , the number of labels.
- **blank_symbol** (*int*) – Index of blank_symbol. This value must be non-negative.
- **input_length** (*Variable* or `numpy.ndarray` or `cupy.ndarray` or `None`) – Length of sequence for each of mini batch `x` (optional). Its shape must be $(B,)$. If the `input_length` is omitted or `None`, it assumes that all of `x` is valid input.
- **label_length** (*Variable* or `numpy.ndarray` or `cupy.ndarray` or `None`) – Length of sequence for each of mini batch `t` (optional). Its shape must be $(B,)$. If the `label_length` is omitted or `None`, it assumes that all of `t` is valid input.
- **reduce** (*str*) – Reduction option. Its value must be either 'mean' or 'no'. Otherwise, `ValueError` is raised.

Returns A variable holding a scalar value of the CTC loss. If `reduce` is 'no', the output variable holds array whose shape is $(B,)$ where B is the number of samples. If it is 'mean', it holds a scalar.

Return type *Variable*

Note: You need to input `x` without applying to activation functions(e.g. softmax function), because this function applies softmax functions to `x` before calculating CTC loss to avoid numerical limitations. You also

need to apply softmax function to forwarded values before you decode it.

Note: This function is differentiable only by `x`.

Note: This function supports (batch, sequence, 1-dimensional input)-data.

chainer.functions.contrastive

`chainer.functions.contrastive(x0, x1, y, margin=1, reduce='mean')`

Computes contrastive loss.

It takes a pair of samples and a label as inputs. The label is 1 when those samples are similar, or 0 when they are dissimilar.

Let N and K denote mini-batch size and the dimension of input variables, respectively. The shape of both input variables `x0` and `x1` should be (N, K) . The loss value of the n -th sample pair L_n is

$$L_n = \frac{1}{2} (y_n d_n^2 + (1 - y_n) \max(\text{margin} - d_n, 0)^2)$$

where $d_n = \|\mathbf{x}_{0n} - \mathbf{x}_{1n}\|_2$, \mathbf{x}_{0n} and \mathbf{x}_{1n} are n -th K -dimensional vectors of `x0` and `x1`.

The output is a variable whose value depends on the value of the option `reduce`. If it is `'no'`, it holds the elementwise loss values. If it is `'mean'`, this function takes a mean of loss values.

Parameters

- **x0** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – The first input variable. The shape should be (N, K) , where N denotes the mini-batch size, and K denotes the dimension of `x0`.
- **x1** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – The second input variable. The shape should be the same as `x0`.
- **y** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Labels. All values should be 0 or 1. The shape should be $(N,)$, where N denotes the mini-batch size.
- **margin** (*float*) – A parameter for contrastive loss. It should be positive value.
- **reduce** (*str*) – Reduction option. Its value must be either `'mean'` or `'no'`. Otherwise, `ValueError` is raised.

Returns A variable holding the loss value(s) calculated by the above equation. If `reduce` is `'no'`, the output variable holds array whose shape is same as one of (hence both of) input variables. If it is `'mean'`, the output variable holds a scalar value.

Return type *Variable*

Note: This cost can be used to train siamese networks. See [Learning a Similarity Metric Discriminatively, with Application to Face Verification](#) for details.

Example

```

>>> x0 = np.array([[-2.0, 3.0, 0.5], [5.0, 2.0, -0.5]]).astype(np.float32)
>>> x1 = np.array([[-1.0, 3.0, 1.0], [3.5, 0.5, -2.0]]).astype(np.float32)
>>> y = np.array([1, 0]).astype(np.int32)
>>> F.contrastive(x0, x1, y)
variable(0.3125)
>>> F.contrastive(x0, x1, y, margin=3.0) # harder penalty
variable(0.3528857)
>>> z = F.contrastive(x0, x1, y, reduce='no')
>>> z.shape
(2,)
>>> z.data
array([0.625, 0.   ], dtype=float32)

```

chainer.functions.crf1d

`chainer.functions.crf1d(cost, xs, ys, reduce='mean')`

Calculates negative log-likelihood of linear-chain CRF.

It takes a transition cost matrix, a sequence of costs, and a sequence of labels. Let c_{st} be a transition cost from a label s to a label t , x_{it} be a cost of a label t at position i , and y_i be an expected label at position i . The negative log-likelihood of linear-chain CRF is defined as

$$L = - \left(\sum_{i=1}^l x_{iy_i} + \sum_{i=1}^{l-1} c_{y_i y_{i+1}} - \log(Z) \right),$$

where l is the length of the input sequence and Z is the normalizing constant called partition function.

Note: When you want to calculate the negative log-likelihood of sequences which have different lengths, sort the sequences in descending order of lengths and transpose the sequences. For example, you have three input sequences:

```

>>> a1 = a2 = a3 = a4 = np.random.uniform(-1, 1, 3).astype(np.float32)
>>> b1 = b2 = b3 = np.random.uniform(-1, 1, 3).astype(np.float32)
>>> c1 = c2 = np.random.uniform(-1, 1, 3).astype(np.float32)

```

```

>>> a = [a1, a2, a3, a4]
>>> b = [b1, b2, b3]
>>> c = [c1, c2]

```

where `a1` and all other variables are arrays with $(K,)$ shape. Make a transpose of the sequences:

```

>>> x1 = np.stack([a1, b1, c1])
>>> x2 = np.stack([a2, b2, c2])
>>> x3 = np.stack([a3, b3])
>>> x4 = np.stack([a4])

```

and make a list of the arrays:

```

>>> xs = [x1, x2, x3, x4]

```

You need to make label sequences in the same fashion. And then, call the function:

```
>>> cost = chainer.Variable(
...     np.random.uniform(-1, 1, (3, 3)).astype(np.float32))
>>> ys = [np.zeros(x.shape[0:1], dtype=np.int32) for x in xs]
>>> loss = F.crflld(cost, xs, ys)
```

It calculates mean of the negative log-likelihood of the three sequences.

The output is a variable whose value depends on the value of the option `reduce`. If it is `'no'`, it holds the elementwise loss values. If it is `'mean'`, it holds mean of the loss values.

Parameters

- **cost** (*Variable*) – A $K \times K$ matrix which holds transition cost between two labels, where K is the number of labels.
- **xs** (*list of Variable*) – Input vector for each label. `len(xs)` denotes the length of the sequence, and each *Variable* holds a $B \times K$ matrix, where B is mini-batch size, K is the number of labels. Note that B s in all the variables are not necessary the same, i.e., it accepts the input sequences with different lengths.
- **ys** (*list of Variable*) – Expected output labels. It needs to have the same length as `xs`. Each *Variable* holds a B integer vector. When `x` in `xs` has the different B , corresponding `y` has the same B . In other words, `ys` must satisfy `ys[i].shape == xs[i].shape[0:1]` for all `i`.
- **reduce** (*str*) – Reduction option. Its value must be either `'mean'` or `'no'`. Otherwise, `ValueError` is raised.

Returns A variable holding the average negative log-likelihood of the input sequences.

Return type *Variable*

Note: See detail in the original paper: [Conditional Random Fields: Probabilistic Models for Segmenting and Labeling Sequence Data](#).

chainer.functions.argmax_crflld

`chainer.functions.argmax_crflld(cost, xs)`

Computes a state that maximizes a joint probability of the given CRF.

Parameters

- **cost** (*Variable*) – A $K \times K$ matrix which holds transition cost between two labels, where K is the number of labels.
- **xs** (*list of Variable*) – Input vector for each label. `len(xs)` denotes the length of the sequence, and each *Variable* holds a $B \times K$ matrix, where B is mini-batch size, K is the number of labels. Note that B s in all the variables are not necessary the same, i.e., it accepts the input sequences with different lengths.

Returns A tuple of *Variable* object `s` and a *list* `ps`. The shape of `s` is $(B,)$, where B is the mini-batch size. `i`-th element of `s`, `s[i]`, represents log-likelihood of `i`-th data. `ps` is a list of `numpy.ndarray` or `cupy.ndarray`, and denotes the state that maximizes the point probability. `len(ps)` is equal to `len(xs)`, and shape of each `ps[i]` is the mini-batch size of the corresponding `xs[i]`. That means, `ps[i].shape == xs[i].shape[0:1]`.

Return type `tuple`

`chainer.functions.cross_covariance`

`chainer.functions.cross_covariance(y, z, reduce='half_squared_sum')`

Computes the sum-squared cross-covariance penalty between `y` and `z`

The output is a variable whose value depends on the value of the option `reduce`. If it is `'no'`, it holds the covariant matrix that has as many rows (resp. columns) as the dimension of `y` (resp. `z`). If it is `'half_squared_sum'`, it holds the half of the Frobenius norm (i.e. L2 norm of a matrix flattened to a vector) of the covariant matrix.

Parameters

- **y** (`Variable`) – Variable holding a matrix where the first dimension corresponds to the batches.
- **z** (`Variable`) – Variable holding a matrix where the first dimension corresponds to the batches.
- **reduce** (`str`) – Reduction option. Its value must be either `'half_squared_sum'` or `'no'`. Otherwise, `ValueError` is raised.

Returns A variable holding the cross covariance loss. If `reduce` is `'no'`, the output variable holds 2-dimensional array matrix of shape (M, N) where M (resp. N) is the number of columns of `y` (resp. `z`). If it is `'half_squared_sum'`, the output variable holds a scalar value.

Return type `Variable`

Note: This cost can be used to disentangle variables. See <https://arxiv.org/abs/1412.6583v3> for details.

`chainer.functions.decov`

`chainer.functions.decov(h, reduce='half_squared_sum')`

Computes the DeCov loss of `h`

The output is a variable whose value depends on the value of the option `reduce`. If it is `'no'`, it holds a matrix whose size is same as the number of columns of `y`. If it is `'half_squared_sum'`, it holds the half of the squared Frobenius norm (i.e. squared of the L2 norm of a matrix flattened to a vector) of the matrix.

Parameters

- **h** (`Variable`) – Variable holding a matrix where the first dimension corresponds to the batches.
- **recude** (`str`) – Reduction option. Its value must be either `'half_squared_sum'` or `'no'`. Otherwise, `ValueError` is raised.

Returns A variable holding a scalar of the DeCov loss. If `reduce` is `'no'`, the output variable holds 2-dimensional array matrix of shape (N, N) where N is the number of columns of `y`. If it is `'half_squared_sum'`, the output variable holds a scalar value.

Return type `Variable`

Note: See <https://arxiv.org/abs/1511.06068> for details.

chainer.functions.gaussian_kl_divergence

`chainer.functions.gaussian_kl_divergence(mean, ln_var, reduce='sum')`

Computes the KL-divergence of Gaussian variables from the standard one.

Given two variable `mean` representing μ and `ln_var` representing $\log(\sigma^2)$, this function calculates the KL-divergence in elementwise manner between the given multi-dimensional Gaussian $N(\mu, S)$ and the standard Gaussian $N(0, I)$

$$D_{\text{KL}}(N(\mu, S) \| N(0, I)),$$

where S is a diagonal matrix such that $S_{ii} = \sigma_i^2$ and I is an identity matrix.

The output is a variable whose value depends on the value of the option `reduce`. If it is `'no'`, it holds the elementwise loss values. If it is `'sum'`, loss values are summed up.

Parameters

- **mean** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – A variable representing mean of given gaussian distribution, μ .
- **ln_var** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – A variable representing logarithm of variance of given gaussian distribution, $\log(\sigma^2)$.
- **reduce** (*str*) – Reduction option. Its value must be either `'sum'` or `'no'`. Otherwise, `ValueError` is raised.

Returns A variable representing KL-divergence between given gaussian distribution and the standard gaussian. If `reduce` is `'no'`, the output variable holds array whose shape is same as one of (hence both of) input variables. If it is `'sum'`, the output variable holds a scalar value.

Return type *Variable*

chainer.functions.gaussian_nll

`chainer.functions.gaussian_nll(x, mean, ln_var, reduce='sum')`

Computes the negative log-likelihood of a Gaussian distribution.

Given two variable `mean` representing μ and `ln_var` representing $\log(\sigma^2)$, this function computes in elementwise manner the negative log-likelihood of x on a Gaussian distribution $N(\mu, S)$,

$$-\log N(x; \mu, \sigma^2) = \log \left(\sqrt{(2\pi)^D |S|} \right) + \frac{1}{2} (x - \mu)^\top S^{-1} (x - \mu),$$

where D is a dimension of x and S is a diagonal matrix where $S_{ii} = \sigma_i^2$.

The output is a variable whose value depends on the value of the option `reduce`. If it is `'no'`, it holds the elementwise loss values. If it is `'sum'`, loss values are summed up.

Parameters

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable.
- **mean** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – A variable representing mean of a Gaussian distribution, μ .
- **ln_var** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – A variable representing logarithm of variance of a Gaussian distribution, $\log(\sigma^2)$.
- **reduce** (*str*) – Reduction option. Its value must be either `'sum'` or `'no'`. Otherwise, `ValueError` is raised.

Returns A variable representing the negative log-likelihood. If `reduce` is `'no'`, the output variable holds array whose shape is same as one of (hence both of) input variables. If it is `'sum'`, the output variable holds a scalar value.

Return type *Variable*

chainer.functions.hinge

`chainer.functions.hinge(x, t, norm='L1', reduce='mean')`

Computes the hinge loss for a one-of-many classification task.

$$L = \frac{1}{N} \sum_{n=1}^N \sum_{k=1}^K [\max(0, 1 - \delta\{t_n = k\}x_{nk})]^p$$

where N denotes the batch size and K is the number of classes of interest,

$$\delta\{\text{condition}\} = \begin{cases} 1 & \text{if condition is true} \\ -1 & \text{otherwise,} \end{cases}$$

and

$$p = \begin{cases} 1 & \text{if norm = L1} \\ 2 & \text{if norm = L2.} \end{cases}$$

Let the hinge loss function $l(x, \delta)$ be $[\max(0, 1 - \delta x)]^p$. When x and δ have the same sign (meaning x predicts the proper score for classification) and $|x| \geq 1$, the hinge loss $l(x, \delta) = 0$, but when they have opposite sign, $l(x, \delta)$ increases linearly with x .

The output is a variable whose value depends on the value of the option `reduce`. If it is `'no'`, it holds the elementwise loss values. If it is `'mean'`, it takes the mean of loss values.

Parameters

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray` of `numpy.float`) – Input variable. The shape of `x` should be (N, K) .
- **t** (*Variable* or `numpy.ndarray` or `cupy.ndarray` of signed integer) – The N -dimensional label vector with values $t_n \in \{0, 1, 2, \dots, K - 1\}$. The shape of `t` should be $(N,)$.
- **norm** (*string*) – Specifies norm type. Either `'L1'` or `'L2'` is acceptable.
- **reduce** (*str*) – Reduction option. Its value must be either `'mean'` or `'no'`. Otherwise, `ValueError` is raised.

Returns A variable object holding a scalar array of the hinge loss L . If `reduce` is `'no'`, the output variable holds array whose shape is same as one of (hence both of) input variables. If it is `'mean'`, the output variable holds a scalar value.

Return type *Variable*

Example

In this case, the batch size N is 2 and the number of classes K is 3.

```
>>> x = np.array([[-2.0, 3.0, 0.5],
...               [5.0, 2.0, -0.5]]).astype(np.float32)
>>> x
array([[ -2. ,  3. ,  0.5],
       [ 5. ,  2. , -0.5]], dtype=float32)
>>> t = np.array([1, 0]).astype(np.int32)
>>> t
array([1, 0], dtype=int32)
>>> F.hinge(x, t)
variable(2.5)
>>> F.hinge(x, t, reduce='no')
variable([[0. , 0. , 1.5],
          [0. , 3. , 0.5]])
>>> F.hinge(x, t, norm='L2')
variable(5.75)
```

chainer.functions.huber_loss

`chainer.functions.huber_loss(x, t, delta, reduce='sum_along_second_axis')`

Computes the Huber loss.

The Huber loss is similar to the `mean_squared_error()` but is less sensitive to outliers in the data. It is defined as

$$L_{\delta}(a) = \begin{cases} \frac{1}{2}a^2 & \text{if } |a| \leq \delta \\ \delta(|a| - \frac{1}{2}\delta) & \text{otherwise,} \end{cases}$$

where $a = x - t$ is the difference between the input x and the target t .

The loss is a variable whose value depends on the value of the option `reduce`. If it is `'no'`, it holds the elementwise loss values. If it is `'sum_along_second_axis'`, loss values are summed up along the second axis (i.e. `axis=1`).

See: [Huber loss - Wikipedia](#).

Parameters

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable. The shape of `x` should be (N, K) .
- **t** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Target variable for regression. The shape of `t` should be (N, K) .
- **delta** (*float*) – Constant variable for Huber loss function as used in definition.
- **reduce** (*str*) – Reduction option. Its value must be either `'sum_along_second_axis'` or `'no'`. Otherwise, `ValueError` is raised.

Returns A variable object holding a scalar array of the Huber loss L_{δ} . If `reduce` is `'no'`, the output variable holds array whose shape is same as one of (hence both of) input variables. If it is `'sum_along_second_axis'`, the shape of the array is same as the input variables, except the second axis is removed.

Return type *Variable*

Example

Example without reduction, in which case the output `y` will have the same shape as the inputs `x` and `t`.


```
>>> import numpy as np
>>> from chainer import functions as F
>>> x = np.array([[ -2.0,  3.0,  0.5], [ 5.0,  2.0, -0.5]]).astype(np.float32)
>>> x.shape
(2, 3)
>>> t = np.array([[ -2.0,  3.0,  0.0], [10.0,  2.0, -0.5]]).astype(np.float32)
>>> t.shape
(2, 3)
>>> y = F.huber_loss(x, t, delta=1.0, reduce='no')
>>> y.shape
(2, 3)
>>> y
variable([[0.    , 0.    , 0.125],
          [4.5   , 0.    , 0.    ]])
```

Example with reduction along the second axis.

```
>>> y = F.huber_loss(x, t, delta=1.0, reduce='sum_along_second_axis')
>>> y.shape
(2,)
>>> y
variable([0.125, 4.5  ])
```

chainer.functions.mean_absolute_error

`chainer.functions.mean_absolute_error(x0, x1)`

Mean absolute error function.

This function computes mean absolute error between two variables. The mean is taken over the minibatch.

chainer.functions.mean_squared_error

`chainer.functions.mean_squared_error(x0, x1)`

Mean squared error function.

This function computes mean squared error between two variables. The mean is taken over the minibatch. Note that the error is not scaled by 1/2.

Parameters

- **x0** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable.
- **x1** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable.

Returns A variable holding an array representing the mean squared error of two inputs.

Return type *Variable*

chainer.functions.negative_sampling

`chainer.functions.negative_sampling(x, t, W, sampler, sample_size, reduce='sum')`

Negative sampling loss function.

In natural language processing, especially language modeling, the number of words in a vocabulary can be very large. Therefore, you need to spend a lot of time calculating the gradient of the embedding matrix.

By using the negative sampling trick you only need to calculate the gradient for a few sampled negative examples.

The loss is defined as follows.

$$f(x, p) = -\log \sigma(x^\top w_p) - k E_{i \sim P(i)} [\log \sigma(-x^\top w_i)]$$

where $\sigma(\cdot)$ is a sigmoid function, w_i is the weight vector for the word i , and p is a positive example. It is approximated with k examples N sampled from probability $P(i)$.

$$f(x, p) \approx -\log \sigma(x^\top w_p) - \sum_{n \in N} \log \sigma(-x^\top w_n)$$

Each sample of N is drawn from the word distribution $P(w) = \frac{1}{Z} c(w)^\alpha$, where $c(w)$ is the unigram count of the word w , α is a hyper-parameter, and Z is the normalization constant.

Parameters

- **x** (*Variable*) – Batch of input vectors.
- **t** (*Variable*) – Vector of ground truth labels.
- **W** (*Variable*) – Weight matrix.
- **sampler** (*FunctionType*) – Sampling function. It takes a shape and returns an integer array of the shape. Each element of this array is a sample from the word distribution. A *WalkerAlias* object built with the power distribution of word frequency is recommended.
- **sample_size** (*int*) – Number of samples.
- **reduce** (*str*) – Reduction option. Its value must be either 'sum' or 'no'. Otherwise, *ValueError* is raised.

Returns A variable holding the loss value(s) calculated by the above equation. If `reduce` is 'no', the output variable holds array whose shape is same as one of (hence both of) input variables. If it is 'sum', the output variable holds a scalar value.

Return type *Variable*

See: [Distributed Representations of Words and Phrases and their Compositionality](#)

See also:

NegativeSampling.

chainer.functions.sigmoid_cross_entropy

`chainer.functions.sigmoid_cross_entropy(x, t, normalize=True, reduce='mean')`

Computes cross entropy loss for pre-sigmoid activations.

Parameters

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – A variable object holding a matrix whose (i, j)-th element indicates the unnormalized log probability of the j-th unit at the i-th example.
- **t** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – A variable object holding a matrix whose (i, j)-th element indicates a signed integer vector of ground truth labels 0 or 1. If `t[i, j] == -1`, corresponding `x[i, j]` is ignored. Loss is zero if all ground truth labels are -1.

- **normalize** (*bool*) – Variable holding a boolean value which determines the normalization constant. If true, this function normalizes the cross entropy loss across all instances. If else, it only normalizes along a batch size.
- **reduce** (*str*) – Variable holding a str which determines whether to reduce the shape of the input. If it is 'mean', it computes the sum of cross entropy and normalize it according to `normalize` option. If it is 'no', this function computes cross entropy for each instance and does not normalize it (`normalize` option is ignored). In this case, the loss value of the ignored instance, which has -1 as its target value, is set to 0.

Returns A variable object holding an array of the cross entropy. If `reduce` is 'mean', it is a scalar array. If `reduce` is 'no', the shape is same as `x`.

Return type *Variable*

Note: This function is differentiable only by `x`.

Example

```
>>> x = np.array([[-2.0, 3.0, 0.5], [5.0, 2.0, -0.5]]).astype(np.float32)
>>> x
array([[ -2. ,  3. ,  0.5],
       [ 5. ,  2. , -0.5]], dtype=float32)
>>> t = np.array([[0, 1, 0], [1, 1, -1]]).astype(np.int32)
>>> t
array([[ 0,  1,  0],
       [ 1,  1, -1]], dtype=int32)
>>> F.sigmoid_cross_entropy(x, t)
variable(0.25664714)
>>> F.sigmoid_cross_entropy(x, t, normalize=False)
variable(0.64161783)
>>> y = F.sigmoid_cross_entropy(x, t, reduce='no')
>>> y.shape
(2, 3)
>>> y.data
array([[ 0.126928 ,  0.04858735,  0.974077 ],
       [ 0.00671535,  0.126928 , -0.          ]], dtype=float32)
```

chainer.functions.softmax_cross_entropy

`chainer.functions.softmax_cross_entropy`(*x*, *t*, *normalize=True*, *cache_score=True*, *class_weight=None*, *ignore_label=-1*, *reduce='mean'*, *enable_double_backprop=False*)

Computes cross entropy loss for pre-softmax activations.

Parameters

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Variable holding a multidimensional array whose element indicates unnormalized log probability: the first axis of the variable represents the number of samples, and the second axis represents the number of classes. While this function computes a usual softmax cross entropy if the number of dimensions is equal to 2, it computes a cross entropy of the replicated softmax if the number of dimensions is greater than 2.

- **t** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Variable holding a signed integer vector of ground truth labels. If `t[i] == ignore_label`, corresponding `x[i]` is ignored.
- **normalize** (*bool*) – If `True`, this function normalizes the cross entropy loss across all instances. If `False`, it only normalizes along a batch size.
- **cache_score** (*bool*) – When it is `True`, the function stores result of forward computation to use it on backward computation. It reduces computational cost though consumes more memory. If `enable_double_backprop` option is `True`, this option is forcibly turned off and the function does not cache the intermediate value.
- **class_weight** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – An array that contains constant weights that will be multiplied with the loss values along with the second dimension. The shape of this array should be `(x.shape[1],)`. If this is not `None`, each class weight `class_weight[i]` is actually multiplied to `y[:, i]` that is the corresponding log-softmax output of `x` and has the same shape as `x` before calculating the actual loss value.
- **ignore_label** (*int*) – Label value you want to ignore. Its default value is `-1`. See description of the argument *t*.
- **reduce** (*str*) – A string that determines whether to reduce the loss values. If it is `'mean'`, it computes the sum of the individual cross entropy and normalize it according to `normalize` option. If it is `'no'`, this function computes cross entropy for each instance and does not normalize it (`normalize` option is ignored). In this case, the loss value of the ignored instance, which has `ignore_label` as its target value, is set to 0.
- **enable_double_backprop** (*bool*) – If `True`, this function uses implementation that supports higher order differentiation. If `False`, it uses single-backprop implementation. This function use the single-backprop version because we expect it is faster. So, if you need second or higher derivatives, you need to turn it on explicitly.

Returns A variable holding a scalar array of the cross entropy loss. If `reduce` is `'mean'`, it is a scalar array. If `reduce` is `'no'`, the shape is same as that of `x`.

Return type *Variable*

Note: This function is differentiable only by `x`.

Example

```
>>> x = np.array([[ -1,  0,  1,  2], [ 2,  0,  1, -1]]).astype(np.float32)
>>> x
array([[ -1.,  0.,  1.,  2.],
       [  2.,  0.,  1., -1.]], dtype=float32)
>>> t = np.array([3, 0]).astype(np.int32)
>>> t
array([3, 0], dtype=int32)
>>> y = F.softmax_cross_entropy(x, t)
>>> y
variable(0.44018972)
>>> log_softmax = -F.log_softmax(x)
>>> expected_loss = np.mean([log_softmax[row, column].data for row, column in
↪ enumerate(t)])
>>> y.array == expected_loss
True
```

chainer.functions.squared_error

`chainer.functions.squared_error(x0, x1)`

Squared error function.

This function computes the squared error between two variables:

$$(x_0 - x_1)^2$$

where operation is done in elementwise manner. Note that the error is not scaled by 1/2:

Parameters

- **x0** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable.
- **x1** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable.

Returns A variable holding an array representing the squared error of two inputs.

Return type *Variable*

chainer.functions.triplet

`chainer.functions.triplet(anchor, positive, negative, margin=0.2, reduce='mean')`

Computes triplet loss.

It takes a triplet of variables as inputs, *a*, *p* and *n*: anchor, positive example and negative example respectively. The triplet defines a relative similarity between samples. Let *N* and *K* denote mini-batch size and the dimension of input variables, respectively. The shape of all input variables should be (*N*, *K*).

$$L(a, p, n) = \frac{1}{N} \left(\sum_{i=1}^N \max\{d(a_i, p_i) - d(a_i, n_i) + \text{margin}, 0\} \right)$$

where $d(x_i, y_i) = \|\mathbf{x}_i - \mathbf{y}_i\|_2^2$.

The output is a variable whose value depends on the value of the option `reduce`. If it is 'no', it holds the elementwise loss values. If it is 'mean', this function takes a mean of loss values.

Parameters

- **anchor** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – The anchor example variable. The shape should be (*N*, *K*), where *N* denotes the minibatch size, and *K* denotes the dimension of the anchor.
- **positive** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – The positive example variable. The shape should be the same as anchor.
- **negative** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – The negative example variable. The shape should be the same as anchor.
- **margin** (*float*) – A parameter for triplet loss. It should be a positive value.
- **reduce** (*str*) – Reduction option. Its value must be either 'mean' or 'no'. Otherwise, `ValueError` is raised.

Returns A variable holding a scalar that is the loss value calculated by the above equation. If `reduce` is 'no', the output variable holds array whose shape is same as one of (hence both of) input variables. If it is 'mean', the output variable holds a scalar value.

Return type *Variable*

Note: This cost can be used to train triplet networks. See [Learning Fine-grained Image Similarity with Deep Ranking](#) for details.

Example

```
>>> anchor = np.array([[-2.0, 3.0, 0.5], [5.0, 2.0, -0.5]]).astype(np.float32)
>>> pos = np.array([[-2.1, 2.8, 0.5], [4.9, 2.0, -0.4]]).astype(np.float32)
>>> neg = np.array([[-2.1, 2.7, 0.7], [4.9, 2.0, -0.7]]).astype(np.float32)
>>> F.triplet(anchor, pos, neg)
variable(0.14000003)
>>> y = F.triplet(anchor, pos, neg, reduce='no')
>>> y.shape
(2,)
>>> y.data
array([0.11000005, 0.17      ], dtype=float32)
>>> F.triplet(anchor, pos, neg, margin=0.5) # harder penalty
variable(0.44000003)
```

5.2.7 Mathematical functions

<code>chainer.functions.absolute</code>	Element-wise absolute.
<code>chainer.functions.arccos</code>	Elementwise arccosine function.
<code>chainer.functions.arcsin</code>	Elementwise arcsine function.
<code>chainer.functions.arctan</code>	Elementwise arctangent function.
<code>chainer.functions.arctan2</code>	Elementwise arctangent function with two arguments.
<code>chainer.functions.argmax</code>	Returns index which holds maximum of array elements over a given axis.
<code>chainer.functions.argmin</code>	Returns index which holds minimum of array elements over a given axis.
<code>chainer.functions.average</code>	Calculate weighted average of array elements over a given axis.
<code>chainer.functions.batch_inv</code>	Computes the inverse of a batch of square matrices.
<code>chainer.functions.batch_l2_norm_squared</code>	L2 norm (a.k.a. Euclidean norm) squared.
<code>chainer.functions.batch_matmul</code>	Computes the batch matrix multiplications of two sets of arrays.
<code>chainer.functions.bias</code>	Elementwise summation with broadcasting.
<code>chainer.functions.ceil</code>	Elementwise ceil function.
<code>chainer.functions.clip</code>	Clips (limits) elements of input variable.
<code>chainer.functions.cos</code>	Elementwise cos function.
<code>chainer.functions.cosh</code>	Elementwise hyperbolic cosine function.
<code>chainer.functions.cumsum</code>	Cumulative sum of array elements over a given axis.
<code>chainer.functions.det</code>	Computes the determinant of a single square matrix.
<code>chainer.functions.batch_det</code>	Computes the determinant of a batch of square matrices.
<code>chainer.functions.erf</code>	Elementwise error function.
<code>chainer.functions.erfc</code>	Elementwise complementary error function.

Continued on next page

Table 8 – continued from previous page

<code>chainer.functions.exp</code>	Elementwise exponential function.
<code>chainer.functions.expm1</code>	Elementwise exponential minus one function.
<code>chainer.functions.fft</code>	Fast Fourier transform.
<code>chainer.functions.fix</code>	Elementwise fix function.
<code>chainer.functions.fmod</code>	Elementwise mod function.
<code>chainer.functions.floor</code>	Elementwise floor function.
<code>chainer.functions.identity</code>	Just returns input variables.
<code>chainer.functions.ifft</code>	Inverse fast Fourier transform.
<code>chainer.functions.inv</code>	Computes the inverse of square matrix.
<code>chainer.functions.linear_interpolate</code>	Elementwise linear-interpolation function.
<code>chainer.functions.log</code>	Elementwise natural logarithm function.
<code>chainer.functions.log10</code>	Elementwise logarithm function to the base 10.
<code>chainer.functions.log1p</code>	Elementwise natural logarithm plus one function.
<code>chainer.functions.log2</code>	Elementwise logarithm function to the base 2.
<code>chainer.functions.logsumexp</code>	Log-sum-exp of array elements over a given axis.
<code>chainer.functions.matmul</code>	Computes the matrix multiplication of two arrays.
<code>chainer.functions.max</code>	Maximum of array elements over a given axis.
<code>chainer.functions.maximum</code>	Element-wise maximum of input variables.
<code>chainer.functions.mean</code>	Calculate weighted average of array elements over a given axis.
<code>chainer.functions.min</code>	Minimum of array elements over a given axis.
<code>chainer.functions.minimum</code>	Element-wise minimum of input variables.
<code>chainer.functions.prod</code>	Product of array elements over a given axis.
<code>chainer.functions.rsqrt</code>	Computes elementwise reciprocal of square root of input x_i .
<code>chainer.functions.scale</code>	Elementwise product with broadcasting.
<code>chainer.functions.sin</code>	Elementwise sin function.
<code>chainer.functions.sinh</code>	Elementwise hyperbolic sine function.
<code>chainer.functions.sign</code>	Elementwise sign function.
<code>chainer.functions.sqrt</code>	Elementwise square root function.
<code>chainer.functions.square</code>	Elementwise square function.
<code>chainer.functions.squared_difference</code>	Squared difference of input variables.
<code>chainer.functions.sum</code>	Sum of array elements over a given axis.
<code>chainer.functions.tanh</code>	Elementwise hyperbolic tangent function.
<code>chainer.functions.tan</code>	Elementwise tan function.
<code>chainer.functions.tensordot</code>	Returns the tensor dot product of two arrays along specified axes.

chainer.functions.absolute

`chainer.functions.absolute` (*self*)

Element-wise absolute.

Returns Output variable.

Return type *Variable*

chainer.functions.arccos

`chainer.functions.arccos(x)`
Elementwise arccosine function.

$$y_i = \arccos x_i.$$

Parameters **x** (*Variable*) – Input variable.

Returns Output variable.

Return type *Variable*

chainer.functions.arcsin

`chainer.functions.arcsin(x)`
Elementwise arcsine function.

$$y_i = \arcsin x_i.$$

Parameters **x** (*Variable*) – Input variable.

Returns Output variable.

Return type *Variable*

chainer.functions.arctan

`chainer.functions.arctan(x)`
Elementwise arctangent function.

$$y_i = \arctan x_i.$$

Parameters **x** (*Variable*) – Input variable.

Returns Output variable.

Return type *Variable*

chainer.functions.arctan2

`chainer.functions.arctan2(x1, x2)`
Elementwise arctangent function with two arguments.

Parameters

- **x1** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Y-coordinates.
- **x2** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – X-coordinates.

Returns Angles in radians, in the range $[-\pi, \pi]$.

Return type *Variable*

chainer.functions.argmax

`chainer.functions.argmax(x, axis=None)`

Returns index which holds maximum of array elements over a given axis.

Parameters

- **x** (*Variable*) – Array to find maximum elements.
- **axis** (*None* or *int*) – Axis over which a max is performed. The default (axis = None) is perform a max over all the dimensions of the input array.

Returns Output variable.

Return type *Variable*

chainer.functions.argmin

`chainer.functions.argmin(x, axis=None)`

Returns index which holds minimum of array elements over a given axis.

Parameters

- **x** (*Variable*) – Array to find minimum elements.
- **axis** (*None* or *int*) – Axis over which a min is performed. The default (axis = None) is perform a min over all the dimensions of the input array.

Returns Output variable.

Return type *Variable*

chainer.functions.average

`chainer.functions.average(x, axis=None, weights=None, keepdims=False)`

Calculate weighted average of array elements over a given axis.

Parameters

- **x** (*Variable*) – Elements to sum.
- **axis** (*None* or *int* or *tuple of int*) – Axis which the method is performed. With the default (axis = None) it performs a mean over all the dimensions of the input array.
- **weights** (*None* or *chainer.Variable*) – An array holding weights to calculate weighted average. If it is None, all weights are assumed to be one. When axis is None, weights must have the same shape of x. And when axis is int, it must be 1-D array satisfying `weights.shape == (x.shape[axis],)`.
- **keepdims** (*bool*) – If True, the specified axes are remained as axes of length one.

Returns Output variable.

Return type *Variable*

chainer.functions.batch_inv

`chainer.functions.batch_inv(a)`

Computes the inverse of a batch of square matrices.

Parameters **a** (*Variable*) – Input array to compute the inverse for. Shape of the array should be (m, n, n) where m is the number of matrices in the batch, and n is the dimensionality of a square matrix.

Returns Inverse of every matrix in the batch of matrices.

Return type *Variable*

`chainer.functions.batch_l2_norm_squared`

`chainer.functions.batch_l2_norm_squared(x)`

L2 norm (a.k.a. Euclidean norm) squared.

This function implements the square of L2 norm on a vector. No reduction along batch axis is done.

Parameters **x** (*Variable*) – Input variable. The first dimension is assumed to be the *minibatch dimension*. If x has more than two dimensions all but the first dimension are flattened to one dimension.

Returns Two dimensional output variable.

Return type *Variable*

`chainer.functions.batch_matmul`

`chainer.functions.batch_matmul(a, b, transa=False, transb=False)`

Computes the batch matrix multiplications of two sets of arrays.

Parameters

- **a** (*Variable*) – The left operand of the batch matrix multiplications. A 2-D array of shape (B, N) is considered as $B \times N \times 1$ matrices. A 3-D array of shape (B, M, N) is considered as $B \times M \times N$ matrices.
- **b** (*Variable*) – The right operand of the batch matrix multiplications. Its array is treated as matrices in the same way as a 's array.
- **transa** (*bool*) – If `True`, transpose each matrix in a .
- **transb** (*bool*) – If `True`, transpose each matrix in b .

Returns The result of the batch matrix multiplications as a 3-D array.

Return type *Variable*

Deprecated since version v3.0.0: `batch_matmul` is deprecated. Use `matmul` instead.

`chainer.functions.bias`

`chainer.functions.bias(x, y, axis=1)`

Elementwise summation with broadcasting.

Computes a elementwise summation of two input variables, with the shape of the latter variable broadcasted to match the shape of the former. `axis` is the first axis of the first variable along which the second variable is applied.

The term “broadcasting” here comes from Caffe’s bias layer so the “broadcasting” with the following arguments:

```
x : 100 x 3 x 40 x 5 x 6
y : 3 x 40
axis : 1
```

is equivalent to the following numpy broadcasting:

```
x : 100 x 3 x 40 x 5 x 6
y : (1 x) 3 x 40 x 1 x 1
```

Note that the axis of `x` to which we apply `y` is specified by the argument `axis`, whose meaning is different from numpy's `axis`.

Parameters

- **x** (*Variable*) – Input variable to be summed.
- **y** (*Variable*) – Input variable to sum, broadcasted.
- **axis** (*int*) – The first axis of `x` along which `y` is applied.

Returns Output variable.

Return type *Variable*

chainer.functions.ceil

`chainer.functions.ceil(x)`
Elementwise ceil function.

$$y_i = \lceil x_i \rceil$$

Parameters **x** (*Variable*) – Input variable.

Returns Output variable.

Return type *Variable*

chainer.functions.clip

`chainer.functions.clip(x, x_min, x_max)`
Clips (limits) elements of input variable.

Given an interval `[x_min, xmax]`, elements outside the interval are clipped to the interval edges.

Parameters

- **x** (*Variable*) – Input variable to be clipped.
- **x_min** (*float*) – Minimum value.
- **x_max** (*float*) – Maximum value.

Returns Output variable.

Return type *Variable*

chainer.functions.cos

`chainer.functions.cos(x)`
Elementwise cos function.

chainer.functions.cosh

`chainer.functions.cosh(x)`
Elementwise hyperbolic cosine function.

$$y_i = \cosh x_i.$$

Parameters *x* (*Variable*) – Input variable.

Returns Output variable.

Return type *Variable*

chainer.functions.cumsum

`chainer.functions.cumsum(x, axis=None)`
Cumulative sum of array elements over a given axis.

Parameters

- *x* (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Elements to calculate the cumulative sum.
- *axis* (*int* or *None*) – Axis along which the cumulative sum is taken. If it is not specified, the input is flattened.

Returns Output variable.

Return type *Variable*

chainer.functions.det

`chainer.functions.det(a)`
Computes the determinant of a single square matrix.

Parameters *a* (*Variable*) – Input array to compute the determinant for.

Returns Scalar determinant of the matrix *a*.

Return type *Variable*

chainer.functions.batch_det

`chainer.functions.batch_det(a)`
Computes the determinant of a batch of square matrices.

Parameters

- *a* (*Variable*) – Input array to compute the determinant for.
- **first dimension should iterate over each matrix and be** (*The*) –
- **the batchsize.** (*of*) –

Returns vector of determinants for every matrix in the batch.

Return type *Variable*

chainer.functions.erf

`chainer.functions.erf(x)`
Elementwise error function.

Note: Forward computation in CPU can be slow if SciPy is not available.

Parameters **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable.

Returns Output variable.

Return type *Variable*

chainer.functions.erfc

`chainer.functions.erfc(x)`
Elementwise complementary error function.

Note: Forward computation in CPU can be slow if SciPy is not available.

Parameters **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable.

Returns Output variable.

Return type *Variable*

chainer.functions.exp

`chainer.functions.exp(x)`
Elementwise exponential function.

chainer.functions.expm1

`chainer.functions.expm1(x)`
Elementwise exponential minus one function.

chainer.functions.fft

`chainer.functions.fft(x)`
Fast Fourier transform.

Parameters **x** (*tuple*) – (real, imag) where real is a *Variable* storing the real part and imag is a *Variable* storing the imaginary part.

Returns Returns (ry, ri) where ry is the real part of the result and ri is the imaginary part of the result.

Return type *tuple*

Note: Currently this function supports a tuple as input. It will support a complex numbers directly in the future.

chainer.functions.fix

`chainer.functions.fix(x)`
Elementwise fix function.

$$y_i = x_i$$

Parameters `x` (*Variable*) – Input variable.

Returns Output variable.

Return type *Variable*

chainer.functions.fmod

`chainer.functions.fmod(x, divisor)`
Elementwise mod function.

$$y_i = x_i \bmod \text{divisor}.$$

Parameters

- `x` (*Variable*) – Input variable.
- `divisor` (*Variable*) – Input divisor.

Returns Output variable.

Return type *Variable*

chainer.functions.floor

`chainer.functions.floor(x)`
Elementwise floor function.

$$y_i = \lfloor x_i \rfloor$$

Parameters `x` (*Variable*) – Input variable.

Returns Output variable.

Return type *Variable*

chainer.functions.identity

`chainer.functions.identity(*inputs)`
Just returns input variables.

chainer.functions.iff

`chainer.functions.iff(x)`
Inverse fast Fourier transform.

Parameters `x` (*tuple*) – (real, imag) where real is a *Variable* storing the real part and imag is a *Variable* storing the imaginary part.

Returns Returns `(ry, ri)` where `ry` is the real part of the result and `ri` is the imaginary part of the result.

Return type `tuple`

Note: Currently this function supports a tuple as input. It will support a complex numbers directly in the future.

chainer.functions.inv

`chainer.functions.inv(a)`

Computes the inverse of square matrix.

Parameters `a` (`Variable`) – Input array to compute the inverse for. Shape of the array should be `(n, n)` where `n` is the dimensionality of a square matrix.

Returns Matrix inverse of `a`.

Return type `Variable`

chainer.functions.linear_interpolate

`chainer.functions.linear_interpolate(p, x, y)`

Elementwise linear-interpolation function.

This function is defined as

$$f(p, x, y) = px + (1 - p)y.$$

Parameters

- `p` (`Variable`) – Input variable.
- `x` (`Variable`) – Input variable.
- `y` (`Variable`) – Input variable.

Returns Output variable.

Return type `Variable`

chainer.functions.log

`chainer.functions.log(x)`

Elementwise natural logarithm function.

chainer.functions.log10

`chainer.functions.log10(x)`

Elementwise logarithm function to the base 10.

$$y_i = \log_{10} x_i.$$

Parameters `x` (`Variable`) – Input variable.

Returns Output variable.

Return type `Variable`

chainer.functions.log1p

`chainer.functions.log1p(x)`

Elementwise natural logarithm plus one function.

chainer.functions.log2

`chainer.functions.log2(x)`

Elementwise logarithm function to the base 2.

$$y_i = \log_2 x_i.$$

Parameters **x** (*Variable*) – Input variable.

Returns Output variable.

Return type *Variable*

chainer.functions.logsumexp

`chainer.functions.logsumexp(x, axis=None)`

Log-sum-exp of array elements over a given axis.

This function calculates logarithm of sum of exponential of array elements.

$$y_i = \log \left(\sum_j \exp(x_{ij}) \right)$$

Parameters

- **x** (*Variable*) – Elements to log-sum-exp.
- **axis** (*None, int, or tuple of int*) – Axis which a sum is performed. The default (axis = None) is perform a sum over all the dimensions of the input array.

Returns Output variable.

Return type *Variable*

chainer.functions.matmul

`chainer.functions.matmul(a, b, transa=False, transb=False)`

Computes the matrix multiplication of two arrays.

Parameters

- **a** (*Variable*) – The left operand of the matrix multiplication. If **a** and **b** are both 1-D arrays, `matmul` returns a dot product of vector *a* and vector *b*. If 2-D arrays, `matmul` returns matrix product of **a** and **b**. If arrays' dimension is larger than 2, they are treated as a stack of matrices residing in the last two indexes. `matmul` returns a stack of each two arrays. **a** and **b** must have the same dimension.
- **b** (*Variable*) – The right operand of the matrix multiplication. Its array is treated as a matrix in the same way as **a**'s array.
- **transa** (*bool*) – If `True`, each matrices in **a** will be transposed. If `a.ndim == 1`, do nothing.

- **transb** (*bool*) – If True, each matrices in *b* will be transposed. If *b.ndim == 1*, do nothing.

Returns The result of the matrix multiplication.

Return type *Variable*

Example

```
>>> a = np.array([[1, 0], [0, 1]], np.float32)
>>> b = np.array([[4, 1], [2, 2]], np.float32)
>>> F.matmul(a, b).data
array([[4., 1.],
       [2., 2.]], dtype=float32)
```

chainer.functions.max

`chainer.functions.max(x, axis=None, keepdims=False)`

Maximum of array elements over a given axis.

Parameters

- **x** (*Variable*) – Array to be maximized.
- **axis** (*None, int, or tuple of int*) – Axis over which a max is performed. The default (*axis = None*) is perform a max over all the dimensions of the input array.

Returns Output variable.

Return type *Variable*

chainer.functions.maximum

`chainer.functions.maximum(x1, x2)`

Element-wise maximum of input variables.

Parameters

- **x1** (*Variable*) – Input variables to be compared.
- **x2** (*Variable*) – Input variables to be compared.

Returns Output variable.

Return type *Variable*

chainer.functions.mean

`chainer.functions.mean(x, axis=None, weights=None, keepdims=False)`

Calculate weighted average of array elements over a given axis.

Parameters

- **x** (*Variable*) – Elements to sum.
- **axis** (*None or int or tuple of int*) – Axis which the method is performed. With the default (*axis = None*) it performs a mean over all the dimensions of the input array.

- **weights** (*None* or *chainer.Variable*) – An array holding weights to calculate weighted average. If it is *None*, all weights are assumed to be one. When *axis* is *None*, *weights* must have the same shape of *x*. And when *axis* is *int*, it must be 1-D array satisfying *weights.shape == (x.shape[axis],)*.
- **keepdims** (*bool*) – If *True*, the specified axes are remained as axes of length one.

Returns Output variable.

Return type *Variable*

chainer.functions.min

`chainer.functions.min(x, axis=None, keepdims=False)`

Minimum of array elements over a given axis.

Parameters

- **x** (*Variable*) – Array to be minimized.
- **axis** (*None*, *int*, or *tuple of int*) – Axis over which a min is performed. The default (*axis = None*) is perform a min over all the dimensions of the input array.

Returns Output variable.

Return type *Variable*

chainer.functions.minimum

`chainer.functions.minimum(x1, x2)`

Element-wise minimum of input variables.

Parameters

- **x1** (*Variable*) – Input variables to be compared.
- **x2** (*Variable*) – Input variables to be compared.

Returns Output variable.

Return type *Variable*

chainer.functions.prod

`chainer.functions.prod(x, axis=None, keepdims=False)`

Product of array elements over a given axis.

Parameters

- **x** (*Variable* or *numpy.ndarray* or *cupy.ndarray*) – Elements to calculate the product.
- **axis** (*None*, *int*, or *tuple of int*) – Axis which a product is performed. The default (*axis = None*) is perform a product over all the dimensions of the input array.
- **keepdims** (*bool*) – If *True*, the specified axes are remained as axes of length one.

Returns Output variable.

Return type *Variable*

chainer.functions.rsqrt

`chainer.functions.rsqrt(x)`

Computes elementwise reciprocal of square root of input x_i .

$$y_i = \frac{1}{\sqrt{x_i}}.$$

Parameters **x** (*Variable*) – Input variable.

Returns Output variable.

Return type *Variable*

See also:

`sqrt()`

chainer.functions.scale

`chainer.functions.scale(x, y, axis=1)`

Elementwise product with broadcasting.

Computes a elementwise product of two input variables, with the shape of the latter variable broadcasted to match the shape of the former. `axis` is the first axis of the first variable along which the second variable is applied.

The term “broadcasting” here comes from Caffe’s scale layer so the “broadcasting” with the following arguments:

```
x : 100 x 3 x 40 x 5 x 6
y : 3 x 40
axis : 1
```

is equivalent to the following numpy broadcasting:

```
x : 100 x 3 x 40 x 5 x 6
y : (1 x) 3 x 40 x 1 x 1
```

Note that the axis of `x` to which we apply `y` is specified by the argument `axis`, whose meaning is different from numpy’s `axis`.

Parameters

- **x** (*Variable*) – Input variable to be scaled.
- **y** (*Variable*) – Input variable to scale, broadcasted.
- **axis** (*int*) – The first axis of `x` along which `y` is applied.

Returns Output variable.

Return type *Variable*

chainer.functions.sin

`chainer.functions.sin(x)`

Elementwise sin function.

chainer.functions.sinh

`chainer.functions.sinh(x)`

Elementwise hyperbolic sine function.

$$y_i = \sinh x_i.$$

Parameters **x** (`Variable`) – Input variable.

Returns Output variable.

Return type `Variable`

chainer.functions.sign

`chainer.functions.sign(x)`

Elementwise sign function.

For a given input x , this function returns $sgn(x)$ defined as

$$sgn(x) = \begin{cases} -1 & \text{if } x < 0 \\ 0 & \text{if } x = 0 \\ 1 & \text{if } x > 0 \end{cases}$$

Note: The gradient of this function is `None` everywhere and therefore unchains the computational graph.

Parameters **x** (`Variable`) – Input variable for which the sign is computed.

Returns Output variable.

Return type `Variable`

chainer.functions.sqrt

`chainer.functions.sqrt(x)`

Elementwise square root function.

$$y_i = \sqrt{x_i}.$$

If the value of x_i is negative, it returns `Nan` for y_i respect to underlying numpy and cupy specification.

Parameters **x** (`Variable`) – Input variable.

Returns Output variable.

Return type `Variable`

chainer.functions.square

`chainer.functions.square(x)`

Elementwise square function.

$$y_i = x_i^2.$$

Parameters **x** (`chainer.Variable` or `numpy.ndarray` or `cupy.ndarray`) – Input variable.

Returns Output variable.

Return type `Variable`

chainer.functions.squared_difference

`chainer.functions.squared_difference(x1, x2)`
 Squared difference of input variables.

Parameters

- **x1** (*Variable*) – Input variables to be compared.
- **x2** (*Variable*) – Input variables to be compared.

Returns $(x1 - x2) ** 2$ element-wise.

Return type *Variable*

chainer.functions.sum

`chainer.functions.sum(x, axis=None, keepdims=False)`
 Sum of array elements over a given axis.

Parameters

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Elements to sum. A (s_1, s_2, \dots, s_N) -shaped float array.
- **axis** (*None*, *int*, or *tuple of int*) – Axis along which a sum is performed. The default (axis = None) is perform a sum over all the dimensions of the input array.
- **keepdims** (*bool*) – If `True`, the specified axes are remained as axes of length one.

Returns Output variable.

Return type *Variable*

Example

```
>>> x = np.arange(6).reshape(2,3).astype(np.float32)
>>> x
array([[0., 1., 2.],
       [3., 4., 5.]], dtype=float32)
>>> y = F.sum(x)
>>> y.shape
()
>>> y.data
array(15., dtype=float32)
>>> y = F.sum(x, axis=1)
>>> y.shape
(2,)
>>> y.data
array([ 3., 12.], dtype=float32)
>>> y = F.sum(x, keepdims=True)
>>> y.shape
(1, 1)
>>> y.data
array([[15.]], dtype=float32)
```

chainer.functions.tan

`chainer.functions.tan(x)`
Elementwise tan function.

chainer.functions.tensordot

`chainer.functions.tensordot(a, b, axes=2)`
Returns the tensor dot product of two arrays along specified axes.

This is equivalent to compute dot product along the specified axes which are treated as one axis by reshaping.

Parameters

- **a** (*Variable*) – The first argument.
- **b** (*Variable*) – The second argument.
- **axes** –
 - If it is an integer, then `axes` axes at the last of `a` and the first of `b` are used.
 - If it is a pair of sequences of integers, then these two sequences specify the list of axes for `a` and `b`. The corresponding axes are paired for sum-product.

Returns The tensor dot product of `a` and `b` along the axes specified by `axes`.

Return type *Variable*

Example

```
>>> a = np.random.rand(5, 3, 2)
>>> b = np.random.rand(3, 2, 4)
>>> c = F.tensordot(a, b, axes=2)
>>> c.shape
(5, 4)
```

See also:

`numpy.tensordot()`

5.2.8 Noise injections

<code>chainer.functions.dropout</code>	Drops elements of input variable randomly.
<code>chainer.functions.gaussian</code>	Gaussian sampling function.
<code>chainer.functions.gumbel_softmax</code>	Gumbel-Softmax sampling function.
<code>chainer.functions.simplified_dropconnect</code>	Linear unit regularized by simplified dropconnect.
<code>chainer.functions.zoneout</code>	Drops elements of input variable and sets to previous variable randomly.

chainer.functions.dropout

`chainer.functions.dropout(x, ratio=.5, *, mask=None, return_mask=False)`
Drops elements of input variable randomly.

This function drops input elements randomly with probability `ratio` and scales the remaining elements by factor $1 / (1 - \text{ratio})$. In testing mode (i.e., `chainer.config.train` is set to `False`), it does nothing and just returns `x`.

Warning: `train` argument is not supported anymore since v2. Instead, use `chainer.using_config('train', boolean)`. See `chainer.using_config()`.

Parameters

- **x** (*Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input variable. A (s_1, s_2, \dots, s_N) -shaped float array.
- **ratio** (*float*) – Dropout ratio. The ratio must be $0.0 \leq \text{ratio} < 1.0$.
- **mask** (*ndarray* or `None`) – The mask to be used for dropout. You do not have to specify this value, unless you need to make results deterministic. If `mask` is not specified or set to `None`, a mask will be generated randomly according to the given `ratio`. If `mask` is specified, `ratio` will be ignored. The shape and dtype must be the same as `x` and should be on the same device. Note that `iDeep` will not be used for this function if `mask` is specified, as `iDeep` does not support it.
- **return_mask** (*bool*) – If `True`, the mask used for dropout is returned together with the output variable. The returned mask can later be reused by passing it to `mask` argument.

Returns When `return_mask` is `False` (default), returns the output variable. When `True`, returns the tuple of the output variable and mask (*ndarray*). The mask will be on the same device as the input. The mask will become `None` when `chainer.config.train` is set to `False`.

Return type *Variable* or *tuple*

See the paper by G. Hinton: [Improving neural networks by preventing co-adaptation of feature detectors](#).

Example

```
>>> x = np.array([[ -1,  0], [ 2, -3], [-2,  1]], np.float32)
>>> with chainer.using_config('train', True):
...     y = F.dropout(x)
>>> y.data
array([[ -2.,  0.],
       [  4., -6.],
       [ -0.,  2.]], dtype=float32)
>>> with chainer.using_config('train', True):
...     y = F.dropout(x, ratio=0.0) # dropout returns original input if ratio=0.0
>>> (x == y.data).all()
True
>>> with chainer.using_config('train', False):
...     y = F.dropout(x) # dropout in test mode returns original input
>>> (x == y.data).all()
True
```

chainer.functions.gaussian

`chainer.functions.gaussian` (*mean*, *ln_var*, *, *eps*=`None`, *return_eps*=`False`)

Gaussian sampling function.

This function takes a mean μ and the logarithm of a variance $\log(\sigma^2)$ as inputs and outputs a sample drawn from a Gaussian distribution $N(\mu, \sigma)$.

The inputs must have the same shape.

Parameters

- **mean** (*Variable*) – Input variable representing the mean μ .
- **ln_var** (*Variable*) – Input variable representing the logarithm of a variance $\log(\sigma^2)$.
- **eps** (*ndarray* or *None*) – The eps value to be used. You do not have to specify this value, unless you need to make results deterministic. If `eps` is not specified or set to `None`, an eps value will be generated randomly. The shape and dtype must be the same as `ln_var` and should be on the same device.
- **return_eps** (*bool*) – If `True`, the eps value used in this function is returned together with the output variable. The returned eps can later be reused by passing it to the `eps` argument.

Returns When `return_eps` is `False` (default), returns the output variable with the shape of `mean` and/or `ln_var`. When `True`, returns the tuple of the output variable and `eps` (*ndarray*). The eps will be on the same device as the input (`ln_var`).

Return type *Variable* or *tuple*

chainer.functions.gumbel_softmax

`chainer.functions.gumbel_softmax(log_pi, tau=0.1, axis=1)`
Gumbel-Softmax sampling function.

This function draws samples y_i from Gumbel-Softmax distribution,

$$y_i = \frac{\exp((g_i + \log \pi_i)/\tau)}{\sum_j \exp((g_j + \log \pi_j)/\tau)},$$

where τ is a temperature parameter and g_i s are samples drawn from Gumbel distribution $Gumbel(0, 1)$

See [Categorical Reparameterization with Gumbel-Softmax](#).

Parameters

- **log_pi** (*Variable* or *numpy.ndarray* or *cupy.ndarray*) – Input variable representing pre-normalized log-probability $\log \pi$.
- **tau** (*float* or *Variable*) – Input variable representing temperature τ .

Returns Output variable.

Return type *Variable*

chainer.functions.simplified_dropconnect

`chainer.functions.simplified_dropconnect(x, W, b=None, ratio=0.5, train=True, mask=None, use_batchwise_mask=True)`

Linear unit regularized by simplified dropconnect.

Simplified dropconnect drops weight matrix elements randomly with probability `ratio` and scales the remaining elements by factor $1 / (1 - \text{ratio})$. It accepts two or three arguments: an input minibatch `x`, a weight matrix `W`, and optionally a bias vector `b`. It computes $Y = xW^\top + b$.

In testing mode, zero will be used as simplified dropconnect ratio instead of `ratio`.

Notice: This implementation cannot be used for reproduction of the paper. There is a difference between the current implementation and the original one. The original version uses sampling with gaussian distribution before passing activation function, whereas the current implementation averages before activation.

Parameters

- **x** (`chainer.Variable` or `numpy.ndarray` or `cupy.ndarray`) – Input variable. Its first dimension `n` is assumed to be the *minibatch dimension*. The other dimensions are treated as concatenated one dimension whose size must be `N`.
- **W** (`Variable`) – Weight variable of shape `(M, N)`.
- **b** (`Variable`) – Bias variable (optional) of shape `(M,)`.
- **ratio** (`float`) – Dropconnect ratio.
- **train** (`bool`) – If `True`, executes simplified dropconnect. Otherwise, simplified dropconnect function works as a linear function.
- **mask** (`None` or `chainer.Variable` or `numpy.ndarray` or `cupy.ndarray`) – If `None`, randomized dropconnect mask is generated. Otherwise, The mask must be `(n, M, N)` or `(M, N)` shaped array, and `use_batchwise_mask` is ignored. Main purpose of this option is debugging. `mask` array will be used as a dropconnect mask.
- **use_batchwise_mask** (`bool`) – If `True`, dropped connections depend on each sample in mini-batch.

Returns Output variable.

Return type `Variable`

See also:

`Dropconnect`

See also:

Li, W., Matthew Z., Sixin Z., Yann L., Rob F. (2013). Regularization of Neural Network using DropConnect. International Conference on Machine Learning. [URL](#)

chainer.functions.zoneout

`chainer.functions.zoneout(h, x, ratio=.5)`

Drops elements of input variable and sets to previous variable randomly.

This function drops input elements randomly with probability `ratio` and instead sets dropping element to their previous variable. In testing mode, it does nothing and just returns `x`.

Warning: `train` argument is not supported anymore since v2. Instead, use `chainer.using_config('train', train)`. See `chainer.using_config()`.

Parameters

- **h** (`Variable`) – Previous variable.
- **x** (`Variable`) – Input variable.
- **ratio** (`float`) – Zoneout ratio.

Returns Output variable.

Return type `Variable`

See the paper: [Zoneout: Regularizing RNNs by Randomly Preserving Hidden Activations](#).

5.2.9 Normalization functions

<code>chainer.functions.batch_normalization</code>	Batch normalization function.
<code>chainer.functions.batch_renormalization</code>	Batch renormalization function.
<code>chainer.functions.fixed_batch_normalization</code>	Batch normalization function with fixed statistics.
<code>chainer.functions.fixed_batch_renormalization</code>	
<code>chainer.functions.layer_normalization</code>	Layer normalization.
<code>chainer.functions.local_response_normalization</code>	Local response normalization across neighboring channels.
<code>chainer.functions.normalize</code>	L2 norm squared (a.k.a. Euclidean norm).

chainer.functions.batch_normalization

`chainer.functions.batch_normalization` (*x*, *gamma*, *beta*, *eps*=2e-5, *running_mean*=None, *running_var*=None, *decay*=0.9)

Batch normalization function.

It takes the input variable *x* and two parameter variables *gamma* and *beta*. The parameter variables must both have the same dimensionality, which is referred to as the channel shape. This channel shape corresponds to the dimensions in the input which are not averaged over. Since the first dimension of the input corresponds to the batch size, the second dimension of *x* will correspond to the first dimension of the channel shape, the third dimension of *x* will correspond to the second channel dimension (if it exists) and so on. Therefore, the dimensionality of the input must be at least one plus the number of channel dimensions. The total effective “batch size” will then be considered to be the product of all dimensions in *x* except for the channel dimensions.

As an example, if the input is four dimensional and the parameter variables are one dimensional, then it is assumed that the first dimension of the input is the batch size, the second dimension is the channel size, and the remaining two dimensions are considered to be spatial dimensions that will be averaged over along with the batch size in the batch normalization computations. That is, the total batch size will be considered to be the product of all input dimensions except the second dimension.

Warning: `train` argument is not supported anymore since v2. Instead, use `chainer.using_config('train', train)`. See `chainer.using_config()`.

Parameters

- ***x*** (*Variable*) – Input variable.
- ***gamma*** (*Variable*) – Scaling parameter of normalized data.
- ***beta*** (*Variable*) – Shifting parameter of scaled normalized data.
- ***eps*** (*float*) – Epsilon value for numerical stability.
- ***running_mean*** (*numpy.ndarray* or *cupy.ndarray*) – Running average of the mean. This is a running average of the mean over several mini-batches using the decay parameter. The function takes a previous running average, and updates the array in-place by the new running average. If None, the running average is not computed. If this is None, then *running_var* must also be None.

- **running_var** (*numpy.ndarray* or *cupy.ndarray*) – Running average of the variance. This is a running average of the variance over several mini-batches using the decay parameter. The function takes a previous running average, and updates the array in-place by the new running average. If *None*, the running average is not computed. If this is *None*, then *running_mean* must also be *None*.
- **decay** (*float*) – Decay rate of moving average. It is used during training.

See: [Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift](#)

See also:

`BatchNormalization`

chainer.functions.batch_renormalization

`chainer.functions.batch_renormalization(x, gamma, beta, rmax, dmax, eps=2e-05, running_mean=None, running_var=None, decay=0.9)`

Batch renormalization function.

This is an extension of batch normalization, which ensures that the training and inference models generate the same outputs that depend on individual examples rather than the entire minibatch.

Note: This function does not perform in-place update to *running_mean* and *running_var*, contrary to `batch_normalization()`. If the function is called, it will not be possible to access the updated running mean and variance statistics, because they are members of the function object, which cannot be accessed by the caller. If it is desired to access the updated running statistics, it is necessary to get a new instance of the function object, call the object, and then access the *running_mean* and/or *running_var* attributes. See the corresponding Link class for an example of how to do this.

See: [Batch Renormalization: Towards Reducing Minibatch Dependence in Batch-Normalized Models](#)

See also:

`links.BatchRenormalization`

See also:

`functions.BatchNormalization()`

chainer.functions.fixed_batch_normalization

`chainer.functions.fixed_batch_normalization(x, gamma, beta, mean, var, eps=2e-05)`

Batch normalization function with fixed statistics.

This is a variant of batch normalization, where the mean and variance statistics are given by the caller as fixed variables. This is used on testing mode of the batch normalization layer, where batch statistics cannot be used for prediction consistency.

Parameters

- **x** (*Variable*) – Input variable.
- **gamma** (*Variable*) – Scaling parameter of normalized data.
- **beta** (*Variable*) – Shifting parameter of scaled normalized data.
- **mean** (*Variable*) – Shifting parameter of input.
- **var** (*Variable*) – Square of scaling parameter of input.

- **eps** (*float*) – Epsilon value for numerical stability.

See also:

`batch_normalization()`, `BatchNormalization`

`chainer.functions.fixed_batch_renormalization`

`chainer.functions.fixed_batch_renormalization(x, gamma, beta, mean, var, eps=2e-05)`

`chainer.functions.layer_normalization`

`chainer.functions.layer_normalization(x, gamma, beta, eps=1e-05)`

Layer normalization.

This function implements a “layer normalization” which normalizes the input units by statistics that are computed along the second axis, scales and shifts them.

Parameters

- **x** (*Variable*) – Batch vectors. Shape of this value must be $(batch_size, unit_size)$, e.g., the output of `linear()`.
- **gamma** (*Variable*) – Scaling vectors.
- **beta** (*Variable*) – Shifting vectors.

Returns The output variable which has the same shape as x .

Return type *Variable*

See: [Layer Normalization](#)

`chainer.functions.local_response_normalization`

`chainer.functions.local_response_normalization(x, n=5, k=2, alpha=0.0001, beta=0.75)`

Local response normalization across neighboring channels.

This function implements normalization across channels. Let x an input image with N channels. Then, this function computes an output image y by following formula:

$$y_i = \frac{x_i}{\left(k + \alpha \sum_{j=\max(1, i-n/2)}^{\min(N, i+n/2)} x_j^2\right)^\beta}.$$

Parameters

- **x** (*Variable*) – Input variable.
- **n** (*int*) – Normalization window width.
- **k** (*float*) – Smoothing parameter.
- **alpha** (*float*) – Normalizer scaling parameter.
- **beta** (*float*) – Normalizer power parameter.

Returns Output variable.

Return type *Variable*

See: [Section 3.3 of ImageNet Classification with Deep Convolutional Neural Networks](#)

chainer.functions.normalize

`chainer.functions.normalize(x, eps=1e-05, axis=1)`
 L2 norm squared (a.k.a. Euclidean norm).

This function implements L2 normalization on a vector along the given axis. No reduction is done along the normalization axis.

In the case when `axis=1` and x is a vector of dimension (N, K) , where N and K denote mini-batch size and the dimension of the input variable, this function computes an output vector y by the following equation:

$$y_i = \frac{x_i}{\|x_i\|_2 + \epsilon}$$

`eps` is used to avoid division by zero when norm of x along the given axis is zero.

The default value of `axis` is determined for backward compatibility.

Parameters

- **x** (*Variable*) – Two dimensional output variable. The first dimension is assumed to be the mini-batch dimension.
- **eps** (*float*) – Epsilon value for numerical stability.
- **axis** (*int*) – Axis along which to normalize.

Returns The output variable which has the same shape as x .

Return type *Variable*

5.2.10 Spatial pooling

<code>chainer.functions.average_pooling_2d</code>	Spatial average pooling function.
<code>chainer.functions.average_pooling_nd</code>	N-dimensionally spatial average pooling function.
<code>chainer.functions.max_pooling_2d</code>	Spatial max pooling function.
<code>chainer.functions.max_pooling_nd</code>	N-dimensionally spatial max pooling function.
<code>chainer.functions.roi_pooling_2d</code>	Spatial Region of Interest (ROI) pooling function.
<code>chainer.functions.spatial_pyramid_pooling_2d</code>	Spatial pyramid pooling function.
<code>chainer.functions.unpooling_2d</code>	Inverse operation of pooling for 2d array.
<code>chainer.functions.unpooling_nd</code>	Inverse operation of N-dimensional spatial pooling.
<code>chainer.functions.upsampling_2d</code>	Upsampling using pooling indices.

chainer.functions.average_pooling_2d

`chainer.functions.average_pooling_2d(x, ksize, stride=None, pad=0)`
 Spatial average pooling function.

This function acts similarly to `Convolution2D`, but it computes the average of input spatial patch for each channel without any parameter instead of computing the inner products.

Parameters

- **x** (*Variable*) – Input variable.
- **ksize** (*int or pair of ints*) – Size of pooling window. `ksize=k` and `ksize=(k, k)` are equivalent.

- **stride** (*int or pair of ints or None*) – Stride of pooling applications. `stride=s` and `stride=(s, s)` are equivalent. If `None` is specified, then it uses same stride as the pooling window size.
- **pad** (*int or pair of ints*) – Spatial padding width for the input array. `pad=p` and `pad=(p, p)` are equivalent.

Returns Output variable.

Return type *Variable*

Note: This function currently does not support `cover_all` mode as `max_pooling_2d()`. Average pooling runs in non-cover-all mode.

chainer.functions.average_pooling_nd

`chainer.functions.average_pooling_nd(x, ksize, stride=None, pad=0)`
N-dimensionally spatial average pooling function.

Warning: This feature is experimental. The interface can change in the future.

This function provides a N-dimensionally generalized version of `average_pooling_2d()`. This acts similarly to `ConvolutionND`, but it computes the average of input spatial patch for each channel without any parameter instead of computing the inner products.

Parameters

- **x** (*Variable*) – Input variable.
- **ksize** (*int or tuple of ints*) – Size of pooling window. `ksize=k` and `ksize=(k, k, ..., k)` are equivalent.
- **stride** (*int or tuple of ints or None*) – Stride of pooling applications. `stride=s` and `stride=(s, s, ..., s)` are equivalent. If `None` is specified, then it uses same stride as the pooling window size.
- **pad** (*int or tuple of ints*) – Spatial padding width for the input array. `pad=p` and `pad=(p, p, ..., p)` are equivalent.

Returns Output variable.

Return type *Variable*

Note: This function currently does not support `cover_all` mode as `max_pooling_nd()`. Average pooling runs in non-cover-all mode.

chainer.functions.max_pooling_2d

`chainer.functions.max_pooling_2d(x, ksize, stride=None, pad=0, cover_all=True, return_indices=False)`
Spatial max pooling function.

This function acts similarly to `Convolution2D`, but it computes the maximum of input spatial patch for each channel without any parameter instead of computing the inner products.

Parameters

- **x** (*Variable*) – Input variable.
- **ksize** (*int or pair of ints*) – Size of pooling window. `ksize=k` and `ksize=(k, k)` are equivalent.
- **stride** (*int or pair of ints or None*) – Stride of pooling applications. `stride=s` and `stride=(s, s)` are equivalent. If `None` is specified, then it uses same stride as the pooling window size.
- **pad** (*int or pair of ints*) – Spatial padding width for the input array. `pad=p` and `pad=(p, p)` are equivalent.
- **cover_all** (*bool*) – If `True`, all spatial locations are pooled into some output pixels. It may make the output size larger.
- **return_indices** (*bool*) – If `True`, pooling indices array is returned together with the output variable. The returned indices are expected for use by `chainer.functions.upsampling_2d()`. Note that cuDNN will not be used for this function if `return_indices` is set to `True`, as cuDNN does not return indices information.

Returns When `return_indices` is `False` (default), returns the output variable. When `True`, returns the tuple of the output variable and pooling indices (*ndarray*). Pooling indices will be on the same device as the input.

Return type *Variable* or *tuple*

chainer.functions.max_pooling_nd

`chainer.functions.max_pooling_nd(x, ksize, stride=None, pad=0, cover_all=True, return_indices=False)`
 N-dimensionally spatial max pooling function.

Warning: This feature is experimental. The interface can change in the future.

This function provides a N-dimensionally generalized version of `max_pooling_2d()`. This acts similarly to `ConvolutionND`, but it computes the maximum of input spatial patch for each channel without any parameter instead of computing the inner products.

Parameters

- **x** (*Variable*) – Input variable.
- **ksize** (*int or tuple of ints*) – Size of pooling window. `ksize=k` and `ksize=(k, k, ..., k)` are equivalent.
- **stride** (*int or tuple of ints or None*) – Stride of pooling applications. `stride=s` and `stride=(s, s, ..., s)` are equivalent. If `None` is specified, then it uses same stride as the pooling window size.
- **pad** (*int or tuple of ints*) – Spatial padding width for the input array. `pad=p` and `pad=(p, p, ..., p)` are equivalent.
- **cover_all** (*bool*) – If `True`, all spatial locations are pooled into some output pixels. It may make the output size larger.

- **return_indices** (*bool*) – If `True`, pooling indices array is returned together with the output variable. The returned indices are expected for use by `chainer.functions.upsampling_nd()`. Note that cuDNN will not be used for this function if `return_indices` is set to `True`, as cuDNN does not return indices information.

Returns When `return_indices` is `False` (default), returns the output variable. When `True`, returns the tuple of the output variable and pooling indices (*ndarray*). Pooling indices will be on the same device as the input.

Return type *Variable* or *tuple*

`chainer.functions.roi_pooling_2d`

`chainer.functions.roi_pooling_2d(x, rois, outh, outw, spatial_scale)`
Spatial Region of Interest (ROI) pooling function.

This function acts similarly to `MaxPooling2D`, but it computes the maximum of input spatial patch for each channel with the region of interest.

Parameters

- **x** (*Variable*) – Input variable. The shape is expected to be 4 dimensional: (n: batch, c: channel, h: height, w: width).
- **rois** (*Variable*) – Input roi variable. The shape is expected to be (n: data size, 5), and each datum is set as below: (batch_index, x_min, y_min, x_max, y_max).
- **outh** (*int*) – Height of output image after pooled.
- **outw** (*int*) – Width of output image after pooled.
- **spatial_scale** (*float*) – Scale of the roi is resized.

Returns Output variable.

Return type *Variable*

See the original paper proposing ROI Pooling: [Fast R-CNN](#).

`chainer.functions.spatial_pyramid_pooling_2d`

`chainer.functions.spatial_pyramid_pooling_2d(x, pyramid_height, pooling_class=None, pooling=None)`
Spatial pyramid pooling function.

Spatial pyramid pooling function.

It outputs a fixed-length vector regardless of input feature map size.

It performs pooling operation to the input 4D-array `x` with different kernel sizes and padding sizes, and then flattens all dimensions except first dimension of all pooling results, and finally concatenates them along second dimension.

At i -th pyramid level, the kernel size $(k_h^{(i)}, k_w^{(i)})$ and padding size $(p_h^{(i)}, p_w^{(i)})$ of pooling operation are calculated as below:

$$\begin{aligned}k_h^{(i)} &= \lceil b_h / 2^i \rceil, \\k_w^{(i)} &= \lceil b_w / 2^i \rceil, \\p_h^{(i)} &= (2^i k_h^{(i)} - b_h) / 2, \\p_w^{(i)} &= (2^i k_w^{(i)} - b_w) / 2,\end{aligned}$$

where $\lceil \cdot \rceil$ denotes the ceiling function, and b_h, b_w are height and width of input variable x , respectively. Note that index of pyramid level i is zero-based.

See detail in paper: [Spatial Pyramid Pooling in Deep Convolutional Networks for Visual Recognition](#).

Parameters

- **x** (*Variable*) – Input variable. The shape of x should be (batchsize, # of channels, height, width).
- **pyramid_height** (*int*) – Number of pyramid levels
- **pooling_class** (*MaxPooling2D*) – (*deprecated since v4.0.0*) Only MaxPooling2D is supported. Please use the `pooling` argument instead since this argument is deprecated.
- **pooling** (*str*) – Currently, only `max` is supported, which performs a 2d max pooling operation. Replaces the `pooling_class` argument.

Returns Output variable. The shape of the output variable will be $(batchsize, c \sum_{h=0}^{H-1} 2^{2h}, 1, 1)$, where c is the number of channels of input variable x and H is the number of pyramid levels.

Return type *Variable*

Note: This function uses some pooling classes as components to perform spatial pyramid pooling. Currently, it only supports MaxPooling2D as elemental pooling operator so far.

chainer.functions.unpooling_2d

`chainer.functions.unpooling_2d(x, ksize, stride=None, pad=0, outsize=None, cover_all=True)`

Inverse operation of pooling for 2d array.

This function acts similarly to `Deconvolution2D`, but it spreads input 2d array's value without any parameter instead of computing the inner products.

Parameters

- **x** (*Variable*) – Input variable.
- **ksize** (*int or pair of ints*) – Size of pooling window. `ksize=k` and `ksize=(k, k)` are equivalent.
- **stride** (*int, pair of ints or None*) – Stride of pooling applications. `stride=s` and `stride=(s, s)` are equivalent. If `None` is specified, then it uses same stride as the pooling window size.
- **pad** (*int or pair of ints*) – Spatial padding width for the input array. `pad=p` and `pad=(p, p)` are equivalent.
- **outsize** (*None or pair of ints*) – Expected output size (height, width) of array after the operation. If `None`, the size (height or width) is estimated from the size of input array in first batch with `get_deconv_outsize()`. If `outsize` is not `None`, the result of `outsize` applied to `get_conv_outsize()` must be equal to the shape of the 2d array in the input batch x .
- **cover_all** (*bool*) – If `True`, the output size may be smaller than the size if `cover_all` is `False`. This flag serves to align behavior to the pooling functions which can cover all input locations, see `max_pooling_2d()` and `convolution_2d()`.

Returns Output variable.

Return type *Variable*

chainer.functions.unpooling_nd

`chainer.functions.unpooling_nd(x, ksize, stride=None, pad=0, outsize=None, cover_all=True)`
Inverse operation of N-dimensional spatial pooling.

Warning: This feature is experimental. The interface can change in the future.

This function acts similarly to `DeconvolutionND`, but it spreads input N-dimensional array's value without any parameter instead of computing the inner products.

Parameters

- **x** (*Variable*) – Input variable.
- **ksize** (*int or pair of ints*) – Size of pooling window (k_1, k_2, \dots, k_N) . `ksize=k` is equivalent to (k, k, \dots, k) .
- **stride** (*int, pair of ints or None*) – Stride of pooling applications (s_1, s_2, \dots, s_N) . `stride=s` is equivalent to (s, s, \dots, s) . If `None` is specified, then it uses same stride as the pooling window size.
- **pad** (*int or pair of ints*) – Spatial padding width for the input array (p_1, p_2, \dots, p_N) . `pad=p` is equivalent to (p, p, \dots, p) .
- **outsize** (*None or pair of ints*) – Expected output size of unpooling operation $(out_1, out_2, \dots, out_N)$. If `None`, the size is estimated from input size, stride and padding.
- **cover_all** (*bool*) – If `True`, the pooling window is assumed to cover all of the output array, eventually the output size may be smaller than that in the case `cover_all` is `False`.

Returns Output variable.

Return type *Variable*

chainer.functions.upsampling_2d

`chainer.functions.upsampling_2d(x, indexes, ksize, stride=None, pad=0, outsize=None, cover_all=True)`
Upsampling using pooling indices.

This function produces an upsampled image using pooling indices.

Example

```
>>> x = np.arange(1, 37).reshape(1, 1, 6, 6).astype(np.float32)
>>> x = chainer.Variable(x)
>>> x.array
array([[[[ 1.,  2.,  3.,  4.,  5.,  6.],
          [ 7.,  8.,  9., 10., 11., 12.],
          [13., 14., 15., 16., 17., 18.],
          [19., 20., 21., 22., 23., 24.],
          [25., 26., 27., 28., 29., 30.],
          [31., 32., 33., 34., 35., 36.]]]], dtype=float32)
```

This is the original `x` before max pooling.

```
>>> pooled_x, indexes = F.max_pooling_2d(
...     x, ksize=2, stride=2, return_indices=True)
>>> pooled_x.array
array([[[[ 8., 10., 12.],
          [20., 22., 24.],
          [32., 34., 36.]]]], dtype=float32)
>>> indexes
array([[[[3, 3, 3],
          [3, 3, 3],
          [3, 3, 3]]]])
```

These are the outputs from the max pooling operation including the resulting indices that will be used to upsample `pooled_x`. Note that the indices all point to the largest, in the case the last, elements in each window.

```
>>> upsampled_x = F.upsampling_2d(
...     pooled_x, indexes, ksize=2, stride=2, outsize=x.shape[2:])
>>> upsampled_x.shape
(1, 1, 6, 6)
>>> upsampled_x.data
array([[[[ 0., 0., 0., 0., 0., 0.],
          [ 0., 8., 0., 10., 0., 12.],
          [ 0., 0., 0., 0., 0., 0.],
          [ 0., 20., 0., 22., 0., 24.],
          [ 0., 0., 0., 0., 0., 0.],
          [ 0., 32., 0., 34., 0., 36.]]]], dtype=float32)
```

Parameters

- **x** (*Variable*) – Input variable.
- **indexes** (*ndarray or ndarray*) – Index array returned from preceding call to `max_pooling_2d()`.
- **ksize** (*int or pair of ints*) – Size of pooling window. `ksize=k` and `ksize=(k, k)` are equivalent.
- **stride** (*int or pair of ints or None*) – Stride of pooling applications. `stride=s` and `stride=(s, s)` are equivalent. If `None` is specified, then it uses same stride as the pooling window size.
- **pad** (*int or pair of ints*) – Spatial padding width for the input array. `pad=p` and `pad=(p, p)` are equivalent.
- **outsize** (*(int, int)*) – Expected output size (height, width).
- **cover_all** (*bool*) – Should be set to `True` if all spatial locations were pooled into some output pixels during the preceding pooling operation. `False` otherwise. See `max_pooling_2d()`.

Returns Output variable.

Return type *Variable*

5.2.11 Utility functions

*chainer.functions.forget*Calls a function without storing intermediate results.

chainer.functions.forget

`chainer.functions.forget(func, *xs)`

Calls a function without storing intermediate results.

On a forward propagation, Chainer normally stores all intermediate results of *VariableNodes* on a computational graph as they are required on backward propagation. Sometimes these results consume too much memory. `F.forget` forgets such intermediate results on forward propagation, and still supports backpropagation with recalculation.

On a forward propagation, `F.forget` calls a given function with given variables without creating a computational graph. That means, no intermediate results are stored. On a backward propagation, `F.forget` calls the given function again to create a computational graph for backpropagation.

`F.forget` reduces internal memory usage, whereas it requires more calculation time as it calls the function twice.

Example

Let `f` be a function defined as:

```
>>> def f(a, b):  
...     return a + b + a
```

and, `x` and `y` be *Variables*:

```
>>> x = chainer.Variable(np.random.uniform(-1, 1, 5).astype(np.float32))  
>>> y = chainer.Variable(np.random.uniform(-1, 1, 5).astype(np.float32))
```

When `z` is calculated as `z = f(x, y)`, its intermediate result `x + y` is stored in memory. Instead, if you call `f` with `F.forget`:

```
>>> z = F.forget(f, x, y)
```

intermediate `x + y` is forgotten.

Note: `F.forget` does not support functions which behave differently in multiple calls with the same inputs, such as `F.dropout()` and `F.negative_sampling()`.

Note: In case input argument variables are of class `numpy.ndarray` or `cupy.ndarray` objects, arguments will automatically be converted to *Variables*. This conversion takes place to ensure that this function is included in the computational graph to enable backward computations.

Parameters

- **func** (*callable*) – A function to call. It needs to be called with *Variable* object(s) and to return a *Variable* object or a tuple of *Variable* objects.
- **xs** (*Variable*) – Argument variables of the function.

Returns A variable `func` returns. If it returns a tuple, the method returns a tuple too.

Return type *Variable*

5.2.12 Function base

<code>chainer.Function</code>	Old-style interface of a differentiable function.
<code>chainer.FunctionAdapter</code>	Adapter class to wrap Function with FunctionNode.
<code>chainer.FunctionNode</code>	Function node of the computational graph.
<code>chainer.force_backprop_mode</code>	Make a context manager which enables back-propagation.
<code>chainer.no_backprop_mode</code>	Make a context manager which disables back-propagation.
<code>chainer.grad</code>	Computes the gradient of output variables w.r.t. the input variables.

chainer.Function

class `chainer.Function`

Old-style interface of a differentiable function.

This class provides an interface to implement an old-style differentiable function (i.e., the function application is recorded to the computational graph). The subclass of *Function* that implement *forward()* and *backward()* can be used to run the forward computation and automatically induce the backpropagation procedure.

There is another way to implement such a function: subclassing *FunctionNode*. There are mainly two differences between them.

1. The *differentiable backward* is available for *FunctionNode*, while it is not for *Function* because the *backward()* of the latter directly operates on the arrays instead of *Variable* objects so that it cannot record the history of the computation.
2. The information passed to *backward()* is different. In *FunctionNode*, which inputs the function node has to compute the gradients w.r.t. is passed so that it can omit unnecessary computations, while *Function* always has to compute gradients w.r.t. all the input nodes. The *FunctionNode* also accepts the current gradient values of the input nodes so that the accumulation work can be merged with the gradient computation if an efficient kernel is available.

This class uses *FunctionAdapter* to convert the interface to that of *FunctionNode* and adds the *FunctionNode* object to the computational graph.

See *FunctionNode* for the details of building the computational graph in Chainer.

Methods

__call__ (*inputs)

Applies forward propagation with chaining backward references.

This method creates a new *FunctionAdapter* object and runs the forward propagation using it.

See *FunctionNode* for the detailed behavior of building the computational graph.

Parameters *inputs* – Tuple of input *Variable*, `numpy.ndarray` or `cupy.ndarray` objects. If the input is an `numpy.ndarray` or a `cupy.ndarray`, it is automatically wrapped with *Variable*.

Returns One *Variable* object or a tuple of multiple *Variable* objects.

add_hook (*hook*, *name=None*)

Registers a function hook.

See `FunctionNode.add_hook()` for the detail.

Parameters

- **hook** (`FunctionHook`) – Function hook to be registered.
- **name** (`str`) – Name of the function hook. `name` must be unique among function hooks registered to the function. If `None`, default name of the function hook is used.

backward (*inputs*, *grad_outputs*)

Applies backprop to output gradient arrays.

It delegates the procedure to `backward_cpu()` or `backward_gpu()` by default. Which it selects is determined by the type of input arrays and output gradient arrays. Implementations of `Function` must implement either CPU/GPU methods or this method, if the function is intended to be backprop-ed.

Parameters

- **inputs** – Tuple of input arrays.
- **grad_outputs** – Tuple of output gradient arrays.

Returns Tuple of input gradient arrays. Some or all of them can be `None`, if the function is not differentiable on inputs.

Return type `tuple`

Warning: Implementations of `Function` must take care that the return value must be a tuple even if it returns only one array.

backward_cpu (*inputs*, *grad_outputs*)

Applies backprop to output gradient arrays on CPU.

Parameters

- **inputs** – Tuple of input `numpy.ndarray` object(s).
- **grad_outputs** – Tuple of output gradient `numpy.ndarray` object(s).

Returns Tuple of input gradient `numpy.ndarray` object(s). Some or all of them can be `None`, if the function is not differentiable on corresponding inputs.

Return type `tuple`

Warning: Implementations of `Function` must take care that the return value must be a tuple even if it returns only one array.

backward_gpu (*inputs*, *grad_outputs*)

Applies backprop to output gradient arrays on GPU.

Parameters

- **inputs** – Tuple of input `cupy.ndarray` object(s).
- **grad_outputs** – Tuple of output gradient `cupy.ndarray` object(s).

Returns Tuple of input gradient `cupy.ndarray` object(s). Some or all of them can be `None`, if the function is not differentiable on corresponding inputs.

Return type `tuple`

Warning: Implementations of *Function* must take care that the return value must be a tuple even if it returns only one array.

check_type_forward (*in_types*)

Checks types of input data before forward propagation.

Before *forward()* is called, this function is called. You need to validate types of input data in this function using *the type checking utilities*.

Parameters *in_types* (`TypeInfoTuple`) – The type information of input data for *forward()*.

delete_hook (*name*)

Unregisters the specified function hook.

Parameters *name* (`str`) – the name of the function hook to be unregistered.

forward (*inputs*)

Applies forward propagation to input arrays.

It delegates the procedure to *forward_cpu()* or *forward_gpu()* by default. Which it selects is determined by the type of input arrays. Implementations of *Function* must implement either CPU/GPU methods or this method.

Parameters *inputs* – Tuple of input array(s).

Returns Tuple of output array(s).

Warning: Implementations of *Function* must take care that the return value must be a tuple even if it returns only one array.

forward_cpu (*inputs*)

Applies forward propagation to input arrays on CPU.

Parameters *inputs* – Tuple of `numpy.ndarray` object(s).

Returns Tuple of `numpy.ndarray` object(s).

Return type `tuple`

Warning: Implementations of *Function* must take care that the return value must be a tuple even if it returns only one array.

forward_gpu (*inputs*)

Applies forward propagation to input arrays on GPU.

Parameters *inputs* – Tuple of `cupy.ndarray` object(s).

Returns Tuple of `cupy.ndarray` object(s).

Return type `tuple`

Warning: Implementations of `Function` must take care that the return value must be a tuple even if it returns only one array.

retain_inputs (*indexes*)

Lets specified input variable nodes keep data arrays.

By calling this method from `forward()`, the function can specify which inputs are required for backprop.

If this method is not called, the function keeps all input arrays. If you want to release all input arrays, call this method by passing an empty sequence. *Note that this behavior is different from that of `FunctionNode.retain_inputs()`.*

Note that **this method must not be called from the outside of `forward()`**.

Parameters *indexes* (*iterable of int*) – Indexes of input variables that the function will require for backprop.

retain_outputs (*indexes, retain_after_backward=False*)

Lets specified output variable nodes keep data arrays.

By calling this method from `forward()`, the function can specify which outputs are required for backprop. If this method is not called, any output variables are not marked to keep the data array at the point of returning from `__call__()`. The retained arrays are stored to `output_data`.

Note: It is STRONGLY RECOMMENDED to use this method if the function requires some or all output arrays in backprop. The function can also use output arrays just by keeping references to them directly, whereas it might influence on the performance of later function applications to the output variables.

Note that **this method must not be called from the outside of `forward()`**.

Parameters

- **indexes** (*iterable of int*) – Indexes of input variables that the function will require for backprop.
- **retain_after_backward** (*bool*) – This option has no effect. It is left only for the backward compatibility.

unchain()

Purges in/out nodes and this function itself from the graph.

See `FunctionNode.unchain()` for the detail.

Attributes**inputs**

The input nodes of the function.

label

Short text that represents the function.

The default implementation returns its type name. Each function should override it to give more information.

local_function_hooks

Ordered Dictionary of registered function hooks.

See `FunctionNode.local_function_hooks` for the detail.

node

The *FunctionAdapter* object that wraps this Function.

If the Function does not have a node object, this property automatically creates a new one.

output_data

A tuple of the retained output arrays.

It has the same length as the *outputs*. Elements that are not retained are set to *None*.

outputs

Weak references to the output nodes of the function.

rank

The topological ordinal of the corresponding function node.

stack**chainer.FunctionAdapter**

class `chainer.FunctionAdapter` (*function*)

Adapter class to wrap Function with FunctionNode.

While *FunctionNode* provides the interface of new-style differentiable functions, the old-style *Function* can still be used for the backward compatibility. This class provides an adapter of there interface; it adds *FunctionNode* interface to any *Function* object by delegation.

Note: The ownership of *FunctionAdapter* and *Function* is a bit tricky. At the initialization, *FunctionAdapter* is owned by the *Function* object. Once the function is applied to variables, the ownership is reversed; the adapter becomes the owner of the *Function* object and the *Function* object changes the reference to a weak one.

Parameters *function* (*Function*) – The function object to wrap.

New in version 3.0.0.

Methods

__call__ (**args, **kwargs*)

Call self as a function.

add_hook (*hook, name=None*)

Registers a function hook.

Parameters

- **hook** (*FunctionHook*) – Function hook to be registered.
- **name** (*str*) – Name of the function hook. The name must be unique among function hooks registered to this function. If *None*, the default name of the function hook is used.

apply (*inputs*)

Computes output variables and grows the computational graph.

Basic behavior is expressed in the documentation of *FunctionNode*.

Note: If the `data` attribute of input variables exist on a GPU device, that device is made current before calling `forward()`, so implementors do not need to take care of device selection in most cases.

Parameters **inputs** – Tuple of input variables. Each element can be either `Variable`, `numpy.ndarray`, or `cupy.ndarray`. If the element is an ndarray, it is automatically wrapped with `Variable`.

Returns A tuple of output `Variable` objects.

backward (*target_input_indexes*, *grad_outputs*)

Computes gradients w.r.t. specified inputs given output gradients.

This method is used to compute one step of the backpropagation corresponding to the forward computation of this function node. Given the gradients w.r.t. output variables, this method computes the gradients w.r.t. specified input variables. Note that this method does not need to compute any input gradients not specified by `target_input_indexes`.

Unlike `Function.backward()`, gradients are given as `Variable` objects and this method itself has to return input gradients as `Variable` objects. It enables the function node to return the input gradients with the full computational history, in which case it supports *differentiable backpropagation* or *higher-order differentiation*.

The default implementation returns `None`s, which means the function is not differentiable.

Parameters

- **target_input_indexes** (*tuple of int*) – Indices of the input variables w.r.t. which the gradients are required. It is guaranteed that this tuple contains at least one element.
- **grad_outputs** (*tuple of Variables*) – Gradients w.r.t. the output variables. If the gradient w.r.t. an output variable is not given, the corresponding element is `None`.

Returns Tuple of variables that represent the gradients w.r.t. specified input variables. The length of the tuple can be same as either `len(target_input_indexes)` or the number of inputs. In the latter case, the elements not specified by `target_input_indexes` will be discarded.

See also:

`backward_accumulate()` provides an alternative interface that allows you to implement the backward computation fused with the gradient accumulation.

backward_accumulate (*target_input_indexes*, *grad_outputs*, *grad_inputs*)

Computes gradients w.r.t. specified inputs and accumulates them.

This method provides a way to fuse the backward computation and the gradient accumulations in the case that the multiple functions are applied to the same variable.

Users have to override either of this method or `backward()`. It is often simpler to implement `backward()` and is recommended if you do not need to provide efficient gradient accumulation.

Parameters

- **target_input_indexes** (*tuple of int*) – Indices of the input variables w.r.t. which the gradients are required. It is guaranteed that this tuple contains at least one element.
- **grad_outputs** (*tuple of Variable*) – Gradients w.r.t. the output variables. If the gradient w.r.t. an output variable is not given, the corresponding element is `None`.

- **grad_inputs** (*tuple of Variable*) – Gradients w.r.t. the input variables specified by `target_input_indexes`. These values are computed by other computation paths. If there is no gradient value existing for the variable, the corresponding element is `None`. See also the note below.

Returns Tuple of variables that represent the gradients w.r.t. specified input variables. Unlike `backward()`, the length of the tuple **must** be same as that of `target_input_indices`.

Note: When the same variable is passed to the multiple input arguments of a function, only the first position of `grad_inputs` corresponding to these input arguments may contain the gradient variable corresponding to that input variable, and other entries are set to `None`. This is an implementation-detail convention to avoid the complication of correctly accumulating gradients in such a case. This behavior might be changed in a future version.

check_type_forward (*in_types*)

Checks types of input data before forward propagation.

This method is called before `forward()` and validates the types of input variables using *the type checking utilities*.

Parameters **in_types** (*TypeInfoTuple*) – The type information of input variables for `forward()`.

delete_hook (*name*)

Unregisters the function hook.

Parameters **name** (*str*) – The name of the function hook to be unregistered.

forward (*inputs*)

Computes the output arrays from the input arrays.

It delegates the procedure to `forward_cpu()` or `forward_gpu()` by default. Which of them this method selects is determined by the type of input arrays. Implementations of *FunctionNode* must implement either CPU/GPU methods or this method.

Parameters **inputs** – Tuple of input array(s).

Returns Tuple of output array(s).

Warning: Implementations of *FunctionNode* must take care that the return value must be a tuple even if it returns only one array.

forward_cpu (*inputs*)

Computes the output arrays from the input NumPy arrays.

Parameters **inputs** – Tuple of input `numpy.ndarray` objects.

Returns Tuple of output arrays. Each element can be NumPy or CuPy arrays.

Warning: Implementation of *FunctionNode* must take care that the return value must be a tuple even if it returns only one array.

forward_gpu (*inputs*)

Computes the output arrays from the input CuPy arrays.

Parameters `inputs` – Tuple of input `cupy.ndarray` objects.

Returns Tuple of output arrays. Each element can be NumPy or CuPy arrays.

Warning: Implementation of `FunctionNode` must take care that the return value must be a tuple even if it returns only one array.

`get_retained_inputs()`

Returns a tuple of retained input variables.

This method is used to retrieve the input variables retained in `forward()`.

Returns A tuple of retained input variables.

`get_retained_outputs()`

Returns a tuple of retained output variables.

This method is used to retrieve the output variables retained in `forward()`.

Returns A tuple of retained output variables.

Note: This method does a tricky thing to support the case of an output node garbage-collected before this method is called; in this case, this method creates a fresh variable node that acts as an output node of the function node.

`retain_inputs(indexes)`

Lets specified input variable nodes keep data arrays.

By calling this method from `forward()`, the function node can specify which inputs are required for backprop. The input variables with retained arrays can then be obtained by calling `get_retained_inputs()` from inside `backward()`.

Unlike `Function`, the function node **DOES NOT** keep input arrays by default. If you want to keep some or all input arrays, do not forget to call this method.

Note that **this method must not be called from the outside of `forward()`**.

Parameters `indexes` (*iterable of int*) – Indexes of input variables that the function will require for backprop.

`retain_outputs(indexes)`

Lets specified output variable nodes keep data arrays.

By calling this method from `forward()`, the function node can specify which outputs are required for backprop. If this method is not called, no output variables will be marked to keep their data array at the point of returning from `apply()`. The output variables with retained arrays can then be obtained by calling `get_retained_outputs()` from inside `backward()`.

Note: It is recommended to use this method if the function requires some or all output arrays in backprop. The function can also use output arrays just by keeping references to them directly, although it might affect the performance of later function applications on the output variables.

Note that **this method must not be called from the outside of `forward()`**.

Parameters `indexes` (*iterable of int*) – Indexes of output variables that the function will require for backprop.

unchain()

Purges in/out nodes and this function node itself from the graph.

Attributes

function

The *Function* object that this adapter is wrapping.

inputs = None

label

Short text that represents the function.

The default implementation returns its type name. Each function should override it to give more information.

lazy_grad_sum = False

local_function_hooks

Ordered dictionary of registered function hooks.

Contrary to `chainer.thread_local.function_hooks`, which registers its elements to all functions, Function hooks in this property is specific to this function.

output_data

A tuple of the retained output arrays.

This property is mainly used by *Function*. Users basically do not have to use this property; use `get_retained_outputs()` instead.

outputs = None

rank = 0

stack = None

chainer.FunctionNode

class chainer.FunctionNode

Function node of the computational graph.

FunctionNode is a class representing a node in a computational graph. The node corresponds to an application of a differentiable function to input variables.

When a differentiable function is applied to *Variable* objects, it creates an instance of FunctionNode implementation and calls its `apply()` method. The `apply()` method basically does the following three things.

1. Adding an edge from the function node to the variable node corresponding to each input. The node of each input is extracted by `Variable.node`.
2. Computing the output arrays of the function.
3. Creating a *Variable* object for each output array and adding an edge from the node of the variable to the function node.

The output variables are then returned.

Example

Let `x` be an instance of *Variable* and `f` be an instance of *FunctionNode* taking only one argument. Then the following code

```
>>> import numpy, chainer, chainer.functions as F
>>> x = chainer.Variable(numpy.zeros(10))
>>> f = F.Identity()
>>> y = f.apply((x,))[0]
```

computes a new variable `y` and creates backward references. The backward references are actually set as per the following diagram:

```
x.node <--- f <--- y.node
```

If an application of another function `g` occurs as

```
>>> g = F.Identity()
>>> z = g.apply((x,))[0]
```

then the graph grows with a branch:

```
      |--- f <--- y.node
x.node <-+
      |--- g <--- z.node
```

Note that the branching is correctly managed on backward computation, i.e. the gradients from `f` and `g` are accumulated to the gradient of `x`.

Every function-node implementation should provide `forward()` and `backward()`. Instead of overriding `forward()`, one can also implement `forward_cpu()` and `forward_gpu()` when the implementations for CPU and GPU arrays are totally different.

Note that the input and output variables are inaccessible from `backward()` by default. If it needs accesses to these variables, the `forward()` method (or its CPU/GPU variants) has to call `retain_inputs()` and `retain_outputs()` appropriately. The retained input/output variables can be accessed from `backward()` by calling `get_retained_inputs()` and `get_retained_outputs()`.

Note: There are two types of differentiable functions in Chainer (since v3). The first type is of a function using a subclass of `Function`, which is called *old-style differentiable function*. The second type is of a function using a subclass of `FunctionNode`, which is called **new-style differentiable function**. There are several advantages on using the new-style differentiable function.

- The new-style differentiable function supports *differentiable backpropagation*. The backpropagated gradients computed through the new-style differentiable functions themselves support further backpropagations so that the automatic higher-order differentiation is available.
- The backpropagation of the new-style differentiable function can be more computationally efficient because the interface allows an implementation to omit the computation of unneeded input gradients.

Note that the new-style differentiable function is the standard way of defining a function node of the computational graph in Chainer; old-style differentiable functions are implemented as wrappers of the new-style differentiable functions.

Variables

- **inputs** – A tuple of the input `VariableNode` objects.
- **outputs** – A tuple of weak references to the output `VariableNode` objects.
- **rank** (`int`) – An ordinal following the topological order of the computational graph.

- **stack** – Stack trace retrieved at the forward computation. The stack trace is available only in the debug mode.

New in version 3.0.0.

Methods

__call__ (*args, **kwargs)
Call self as a function.

add_hook (hook, name=None)
Registers a function hook.

Parameters

- **hook** ([FunctionHook](#)) – Function hook to be registered.
- **name** (*str*) – Name of the function hook. The name must be unique among function hooks registered to this function. If `None`, the default name of the function hook is used.

apply (inputs)
Computes output variables and grows the computational graph.
Basic behavior is expressed in the documentation of [FunctionNode](#).

Note: If the `data` attribute of input variables exist on a GPU device, that device is made current before calling `forward()`, so implementors do not need to take care of device selection in most cases.

Parameters inputs – Tuple of input variables. Each element can be either [Variable](#), `numpy.ndarray`, or `cupy.ndarray`. If the element is an ndarray, it is automatically wrapped with [Variable](#).

Returns A tuple of output [Variable](#) objects.

backward (target_input_indexes, grad_outputs)
Computes gradients w.r.t. specified inputs given output gradients.

This method is used to compute one step of the backpropagation corresponding to the forward computation of this function node. Given the gradients w.r.t. output variables, this method computes the gradients w.r.t. specified input variables. Note that this method does not need to compute any input gradients not specified by `target_input_indexes`.

Unlike `Function.backward()`, gradients are given as [Variable](#) objects and this method itself has to return input gradients as [Variable](#) objects. It enables the function node to return the input gradients with the full computational history, in which case it supports *differentiable backpropagation* or *higher-order differentiation*.

The default implementation returns `None` s, which means the function is not differentiable.

Parameters

- **target_input_indexes** (*tuple of int*) – Indices of the input variables w.r.t. which the gradients are required. It is guaranteed that this tuple contains at least one element.
- **grad_outputs** (*tuple of Variables*) – Gradients w.r.t. the output variables. If the gradient w.r.t. an output variable is not given, the corresponding element is `None`.

Returns Tuple of variables that represent the gradients w.r.t. specified input variables. The length of the tuple can be same as either `len(target_input_indexes)` or the number of inputs. In the latter case, the elements not specified by `target_input_indexes` will be discarded.

See also:

`backward_accumulate()` provides an alternative interface that allows you to implement the backward computation fused with the gradient accumulation.

backward_accumulate (*target_input_indexes*, *grad_outputs*, *grad_inputs*)

Computes gradients w.r.t. specified inputs and accumulates them.

This method provides a way to fuse the backward computation and the gradient accumulations in the case that the multiple functions are applied to the same variable.

Users have to override either of this method or `backward()`. It is often simpler to implement `backward()` and is recommended if you do not need to provide efficient gradient accumulation.

Parameters

- **target_input_indexes** (*tuple of int*) – Indices of the input variables w.r.t. which the gradients are required. It is guaranteed that this tuple contains at least one element.
- **grad_outputs** (*tuple of Variable*) – Gradients w.r.t. the output variables. If the gradient w.r.t. an output variable is not given, the corresponding element is `None`.
- **grad_inputs** (*tuple of Variable*) – Gradients w.r.t. the input variables specified by `target_input_indexes`. These values are computed by other computation paths. If there is no gradient value existing for the variable, the corresponding element is `None`. See also the note below.

Returns Tuple of variables that represent the gradients w.r.t. specified input variables. Unlike `backward()`, the length of the tuple **must** be same as that of `target_input_indexes`.

Note: When the same variable is passed to the multiple input arguments of a function, only the first position of `grad_inputs` corresponding to these input arguments may contain the gradient variable corresponding to that input variable, and other entries are set to `None`. This is an implementation-detail convention to avoid the complication of correctly accumulating gradients in such a case. This behavior might be changed in a future version.

check_type_forward (*in_types*)

Checks types of input data before forward propagation.

This method is called before `forward()` and validates the types of input variables using *the type checking utilities*.

Parameters **in_types** (*TypeInfoTuple*) – The type information of input variables for `forward()`.

delete_hook (*name*)

Unregisters the function hook.

Parameters **name** (*str*) – The name of the function hook to be unregistered.

forward (*inputs*)

Computes the output arrays from the input arrays.

It delegates the procedure to `forward_cpu()` or `forward_gpu()` by default. Which of them this method selects is determined by the type of input arrays. Implementations of `FunctionNode` must implement either CPU/GPU methods or this method.

Parameters `inputs` – Tuple of input array(s).

Returns Tuple of output array(s).

Warning: Implementations of `FunctionNode` must take care that the return value must be a tuple even if it returns only one array.

`forward_cpu(inputs)`

Computes the output arrays from the input NumPy arrays.

Parameters `inputs` – Tuple of input `numpy.ndarray` objects.

Returns Tuple of output arrays. Each element can be NumPy or CuPy arrays.

Warning: Implementation of `FunctionNode` must take care that the return value must be a tuple even if it returns only one array.

`forward_gpu(inputs)`

Computes the output arrays from the input CuPy arrays.

Parameters `inputs` – Tuple of input `cupy.ndarray` objects.

Returns Tuple of output arrays. Each element can be NumPy or CuPy arrays.

Warning: Implementation of `FunctionNode` must take care that the return value must be a tuple even if it returns only one array.

`get_retained_inputs()`

Returns a tuple of retained input variables.

This method is used to retrieve the input variables retained in `forward()`.

Returns A tuple of retained input variables.

`get_retained_outputs()`

Returns a tuple of retained output variables.

This method is used to retrieve the output variables retained in `forward()`.

Returns A tuple of retained output variables.

Note: This method does a tricky thing to support the case of an output node garbage-collected before this method is called; in this case, this method creates a fresh variable node that acts as an output node of the function node.

`retain_inputs(indexes)`

Lets specified input variable nodes keep data arrays.

By calling this method from `forward()`, the function node can specify which inputs are required for backprop. The input variables with retained arrays can then be obtained by calling `get_retained_inputs()` from inside `backward()`.

Unlike *Function*, the function node **DOES NOT** keep input arrays by default. If you want to keep some or all input arrays, do not forget to call this method.

Note that **this method must not be called from the outside of `forward()`**.

Parameters `indexes` (*iterable of int*) – Indexes of input variables that the function will require for backprop.

`retain_outputs` (*indexes*)

Lets specified output variable nodes keep data arrays.

By calling this method from `forward()`, the function node can specify which outputs are required for backprop. If this method is not called, no output variables will be marked to keep their data array at the point of returning from `apply()`. The output variables with retained arrays can then be obtained by calling `get_retained_outputs()` from inside `backward()`.

Note: It is recommended to use this method if the function requires some or all output arrays in backprop. The function can also use output arrays just by keeping references to them directly, although it might affect the performance of later function applications on the output variables.

Note that **this method must not be called from the outside of `forward()`**.

Parameters `indexes` (*iterable of int*) – Indexes of output variables that the function will require for backprop.

`unchain()`

Purges in/out nodes and this function node itself from the graph.

Attributes

`inputs` = `None`

`label`

Short text that represents the function.

The default implementation returns its type name. Each function should override it to give more information.

`lazy_grad_sum` = `False`

`local_function_hooks`

Ordered dictionary of registered function hooks.

Contrary to `chainer.thread_local.function_hooks`, which registers its elements to all functions, Function hooks in this property is specific to this function.

`output_data`

A tuple of the retained output arrays.

This property is mainly used by *Function*. Users basically do not have to use this property; use `get_retained_outputs()` instead.

`outputs` = `None`

`rank` = `0`

`stack` = `None`

chainer.force_backprop_mode

`chainer.force_backprop_mode()`

Make a context manager which enables back-propagation.

When you want to enable back-propagation in `no_backprop_mode()`, call this method. A `Variable` created in this context always has a computational graph unless overridden by deeper contexts. If you call this method outside of `no_backprop_mode()` context, it changes nothing.

In the following example, `y` has a computational graph and calling `backward()` on `y` will compute and accumulate the gradients of the variables in the graph, in this case only `x`.

```

>>> x = chainer.Variable(np.array([1,], np.float32))
>>> with chainer.no_backprop_mode():
...     with chainer.force_backprop_mode():
...         y = x + 1
>>> y.backward()
>>> x.grad
array([1.], dtype=float32)

```

See also:

See `no_backprop_mode()` for details on disabled back-propagation mode.

chainer.no_backprop_mode

`chainer.no_backprop_mode()`

Make a context manager which disables back-propagation.

In this context, Chainer does not make a computational graph. It has the benefit of reducing memory consumption. However, a `Variable` created in this context does not hold a reference to the `FunctionNode` that created itself so no gradients are accumulated by `backward()`.

In the following example, `y` is created in this context, which means that calling `backward()` on `y` has no effect on the gradients of `x`.

```

>>> x = chainer.Variable(np.array([1,], np.float32))
>>> with chainer.no_backprop_mode():
...     y = x + 1
>>> y.backward()
>>> x.grad is None
True

```

See also:

See `force_backprop_mode()` for details on how to override this context.

chainer.grad

`chainer.grad(outputs, inputs, grad_outputs=None, grad_inputs=None, set_grad=False, retain_grad=False, enable_double_backprop=False, loss_scale=None)`

Computes the gradient of output variables w.r.t. the input variables.

This function implements the backpropagation algorithm. While `Variable.backward()` also implements backprop, this function selects the smallest paths in the computational graph needed to compute the gradients w.r.t. inputs. The error is backpropagated only through these selected paths, which may reduce the overall computational cost.

This function also differs from `Variable.backward()` in the way to return the gradients; it directly returns the gradient variables as a list instead of setting gradients to the `Variable.grad_var` attribute of the original variable. It means users do not need to clear the gradient w.r.t. each variable before computing the gradient using this function. If `set_grad` option is set to `True`, the computed gradient is also stored in the `Variable.grad_var` attribute of each variable, in which case any original value of `Variable.grad_var` will be updated even if it had already been set.

Parameters

- **outputs** (tuple or list of `Variable`) – A sequence of output variables from which back-prop starts.
- **inputs** (tuple or list of `Variable`) – A sequence of input variables each of which this function computes the gradient w.r.t.
- **grad_outputs** (tuple or list of `Variable` or `None`) – A sequence of variables that gives the initial value of each output gradient. If an element is set to `None`, an array filled with 1 is used. If this argument itself is `None`, it is treated as a sequence of `Nones`.
- **grad_inputs** (tuple or list of `Variable` or `None`) – A sequence of variables that gives the initial value of each input gradient. The gradients computed by the backprop algorithm are accumulated to them (not in-place). If an element is set to `None`, the gradient is not accumulated to this value. If this argument itself is `None`, it is treated as a sequence of `Nones`.
- **set_grad** (`bool`) – If it is `True`, the `Variable.grad_var` attribute of each input variable is set to the corresponding computed gradient variable.
- **retain_grad** (`bool`) – If it is `True`, the gradients w.r.t. all the intermediate variables are stored in the `Variable.grad_var` attribute. In this case, the `set_grad` option is ignored.
- **enable_double_backprop** (`bool`) – If it is `True`, the computed gradients can be further backpropagated. Enabling it may increase the memory consumption (and possibly the computational time) to remember the intermediate gradient values for the second back-propagation.
- **loss_scale** (`float`) – Loss scaling factor. Loss scaling is a useful technique to mitigate vanishing gradient issue that tends to happen when low precision data type like `float16` is used during training. If you set loss scaling factor, gradients of loss values are to be multiplied by the factor before backprop starts. The factor is propagated to whole gradients in a computational graph along the backprop. The gradients of parameters are divided by the factor just before the parameters are to be updated.

Returns A list of gradient variables w.r.t. the inputs.

5.2.13 Function hooks

Chainer provides a function-hook mechanism that enriches the behavior of forward and backward propagation of `FunctionNode` and `Function`.

<code>chainer.function_hooks.CUDAProfileHook</code>	
<code>chainer.function_hooks.CupyMemoryProfileHook</code>	Function hook for measuring memory usage of functions in cupy memory pool.
<code>chainer.function_hooks.PrintHook</code>	Function hook that prints debug information.
<code>chainer.function_hooks.TimerHook</code>	Function hook for measuring elapsed time of functions.

chainer.function_hooks.CUDAProfileHook

```
class chainer.function_hooks.CUDAProfileHook
```

Methods

```
__enter__()
```

```
__exit__(*_)
```

```
added(function=None)
```

Callback function invoked when a function hook is added

Parameters **function** (`FunctionNode`) – Function object to which the function hook is added.

```
backward_postprocess(function, in_data, out_grad)
```

Callback function invoked after backward propagation.

Parameters

- **function** (`FunctionNode`) – Function object to which the function hook is registered.
- **in_data** (*tuple of numpy.ndarray or tuple of cupy.ndarray*) – Input of forward propagation.
- **out_grad** (*tuple of numpy.ndarray or tuple of cupy.ndarray*) – Gradient data of backward propagation.

```
backward_preprocess(function, in_data, out_grad)
```

Callback function invoked before backward propagation.

Parameters

- **function** (`FunctionNode`) – Function object to which the function hook is registered.
- **in_data** (*tuple of numpy.ndarray or tuple of cupy.ndarray*) – Input data of forward propagation.
- **out_grad** (*tuple of numpy.ndarray or tuple of cupy.ndarray*) – Gradient data of backward propagation.

```
deleted(function=None)
```

Callback function invoked when a function hook is deleted

Parameters **function** (`FunctionNode`) – Function object to which the function hook is deleted.

```
forward_postprocess(function, in_data)
```

Callback function invoked after forward propagation.

Parameters

- **function** (`FunctionNode`) – Function object to which the function hook is registered.
- **in_data** (*tuple of numpy.ndarray or tuple of cupy.ndarray*) – Input data of forward propagation.

```
forward_preprocess(function, in_data)
```

Callback function invoked before forward propagation.

Parameters

- **function** (`FunctionNode`) – Function object to which the function hook is registered.
- **in_data** (*tuple of `numpy.ndarray` or `tuple of cupy.ndarray`*) – Input data of forward propagation.

Attributes

`name` = 'CUDAProfileHook'

chainer.function_hooks.CupyMemoryProfileHook

class `chainer.function_hooks.CupyMemoryProfileHook`

Function hook for measuring memory usage of functions in cupy memory pool.

Example

Code example:

```
from chainer.function_hooks import CupyMemoryProfileHook
hook = CupyMemoryProfileHook()
with hook:
    trainer.run()
hook.print_report()
```

Output example:

FunctionName	UsedBytes	AcquiredBytes	Occurrence
LinearFunction	5.16GB	179.98MB	3900
ReLU	991.82MB	458.97MB	2600
SoftmaxCrossEntropy	7.71MB	5.08MB	1300
Accuracy	617.97KB	351.00KB	700

where *FunctionName* is the name of function that calls the hook, and *UsedBytes* is the memory bytes the function used from cupy memory pool, and *AcquiredBytes* is the actual memory bytes the cupy memory pool acquired from GPU device on the function call, and *Occurrence* is the number of calls.

Variables `call_history` – List of measurement results. It consists of the name of the function that calls this hook, the memory bytes the function used from cupy memory pool, and the memory bytes the cupy memory pool acquired from GPU device on the function call.

Methods

`__enter__()`

`__exit__(*_)`

`added` (*function=None*)

Callback function invoked when a function hook is added

Parameters **function** (`FunctionNode`) – Function object to which the function hook is added.

backward_postprocess (*function, in_data, out_grad*)
 Callback function invoked after backward propagation.

Parameters

- **function** (`FunctionNode`) – Function object to which the function hook is registered.
- **in_data** (*tuple of numpy.ndarray or tuple of cupy.ndarray*) – Input of forward propagation.
- **out_grad** (*tuple of numpy.ndarray or tuple of cupy.ndarray*) – Gradient data of backward propagation.

backward_preprocess (*function, in_data, out_grad*)
 Callback function invoked before backward propagation.

Parameters

- **function** (`FunctionNode`) – Function object to which the function hook is registered.
- **in_data** (*tuple of numpy.ndarray or tuple of cupy.ndarray*) – Input data of forward propagation.
- **out_grad** (*tuple of numpy.ndarray or tuple of cupy.ndarray*) – Gradient data of backward propagation.

deleted (*function=None*)
 Callback function invoked when a function hook is deleted

Parameters **function** (`FunctionNode`) – Function object to which the function hook is deleted.

forward_postprocess (*function, in_data*)
 Callback function invoked after forward propagation.

Parameters

- **function** (`FunctionNode`) – Function object to which the function hook is registered.
- **in_data** (*tuple of numpy.ndarray or tuple of cupy.ndarray*) – Input data of forward propagation.

forward_preprocess (*function, in_data*)
 Callback function invoked before forward propagation.

Parameters

- **function** (`FunctionNode`) – Function object to which the function hook is registered.
- **in_data** (*tuple of numpy.ndarray or tuple of cupy.ndarray*) – Input data of forward propagation.

print_report (*file=<_io.TextIOWrapper name='<stdout>' mode='w' encoding='UTF-8'>*)
 Prints a summary report of memory profiling in functions.

summary ()
 Returns a summary of memory profiling in functions.

Returns A summarized dictionary whose keys are function names and values are dictionaries of `used_bytes`, `acquired_bytes`, and `occurrence`.

total_acquired_bytes()

Returns total bytes that cupy memory pool acquired from GPU.

total_used_bytes()

Returns total bytes that functions used from cupy memory pool.

Attributes

name = 'CupyMemoryProfileHook'

chainer.function_hooks.PrintHook

```
class chainer.function_hooks.PrintHook(sep=None, end='n', file=<_io.TextIOWrapper
                                         name='<stdout>' mode='w' encoding='UTF-8'>,
                                         flush=True)
```

Function hook that prints debug information.

This function hook outputs the debug information of input arguments of `forward` and `backward` methods involved in the hooked functions at preprocessing time (that is, just before each method is called).

Unlike simple “debug print” technique, where users insert print functions at every function to be inspected, we can show the information of all functions involved with single `with` statement.

Further, this hook enables us to show the information of `backward` methods without inserting print functions into Chainer’s library code.

Parameters

- **sep** – (*deprecated since v4.0.0*) Ignored.
- **end** – Character to be added at the end of print function.
- **file** – Output file_like object that that redirect to.
- **flush** – If `True`, this hook forcibly flushes the text stream at the end of preprocessing.

Example

The basic usage is to use it with `with` statement.

```
>>> from chainer import function_hooks
>>> l = L.Linear(10, 10)
>>> x = chainer.Variable(np.zeros((1, 10), np.float32))
>>> with chainer.function_hooks.PrintHook():
...     y = l(x)
...     z = F.sum(y)
...     z.backward()
```

In this example, `PrintHook` shows the debug information of forward propagation of `LinearFunction` (which is implicitly called by `l`) and `Sum` (called by `F.sum`) and backward propagation of `z` and `y`.

Methods

__enter__()

__exit__(*_)

added (*function=None*)

Callback function invoked when a function hook is added

Parameters **function** (`FunctionNode`) – Function object to which the function hook is added.

backward_postprocess (*function, in_data, out_grad*)

Callback function invoked after backward propagation.

Parameters

- **function** (`FunctionNode`) – Function object to which the function hook is registered.
- **in_data** (*tuple of numpy.ndarray or tuple of cupy.ndarray*) – Input of forward propagation.
- **out_grad** (*tuple of numpy.ndarray or tuple of cupy.ndarray*) – Gradient data of backward propagation.

backward_preprocess (*function, in_data, out_grad*)

Callback function invoked before backward propagation.

Parameters

- **function** (`FunctionNode`) – Function object to which the function hook is registered.
- **in_data** (*tuple of numpy.ndarray or tuple of cupy.ndarray*) – Input data of forward propagation.
- **out_grad** (*tuple of numpy.ndarray or tuple of cupy.ndarray*) – Gradient data of backward propagation.

deleted (*function=None*)

Callback function invoked when a function hook is deleted

Parameters **function** (`FunctionNode`) – Function object to which the function hook is deleted.

forward_postprocess (*function, in_data*)

Callback function invoked after forward propagation.

Parameters

- **function** (`FunctionNode`) – Function object to which the function hook is registered.
- **in_data** (*tuple of numpy.ndarray or tuple of cupy.ndarray*) – Input data of forward propagation.

forward_preprocess (*function, in_data*)

Callback function invoked before forward propagation.

Parameters

- **function** (`FunctionNode`) – Function object to which the function hook is registered.
- **in_data** (*tuple of numpy.ndarray or tuple of cupy.ndarray*) – Input data of forward propagation.

Attributes

`name = 'PrintHook'`

chainer.function_hooks.TimerHook

class `chainer.function_hooks.TimerHook`
Function hook for measuring elapsed time of functions.

Example

Code example:

```
from chainer.function_hooks import TimerHook
hook = TimerHook()
with hook:
    trainer.run()
hook.print_report()
```

Output example:

FunctionName	ElapsedTime	Occurrence
LinearFunction	1.24sec	3900
ReLU	593.05ms	2600
SoftmaxCrossEntropy	824.11ms	1300
Accuracy	176.54ms	700

where *FunctionName* is the name of function that calls the hook, and *ElapsedTime* is the elapsed time the function consumed, and *Occurrence* is the number of calls.

Variables `call_history` – List of measurement results. It consists of pairs of the name of the function that calls this hook and the elapsed time the function consumes.

Methods

`__enter__()`

`__exit__(*_)`

`added(function=None)`

Callback function invoked when a function hook is added

Parameters `function` (`FunctionNode`) – Function object to which the function hook is added.

`backward_postprocess(function, in_data, out_grad)`

Callback function invoked after backward propagation.

Parameters

- **function** (`FunctionNode`) – Function object to which the function hook is registered.
- **in_data** (*tuple of numpy.ndarray or tuple of cupy.ndarray*) – Input of forward propagation.

- **out_grad** (*tuple of numpy.ndarray or tuple of cupy.ndarray*) – Gradient data of backward propagation.

backward_preprocess (*function, in_data, out_grad*)
 Callback function invoked before backward propagation.

Parameters

- **function** (`FunctionNode`) – Function object to which the function hook is registered.
- **in_data** (*tuple of numpy.ndarray or tuple of cupy.ndarray*) – Input data of forward propagation.
- **out_grad** (*tuple of numpy.ndarray or tuple of cupy.ndarray*) – Gradient data of backward propagation.

deleted (*function=None*)
 Callback function invoked when a function hook is deleted

Parameters **function** (`FunctionNode`) – Function object to which the function hook is deleted.

forward_postprocess (*function, in_data*)
 Callback function invoked after forward propagation.

Parameters

- **function** (`FunctionNode`) – Function object to which the function hook is registered.
- **in_data** (*tuple of numpy.ndarray or tuple of cupy.ndarray*) – Input data of forward propagation.

forward_preprocess (*function, in_data*)
 Callback function invoked before forward propagation.

Parameters

- **function** (`FunctionNode`) – Function object to which the function hook is registered.
- **in_data** (*tuple of numpy.ndarray or tuple of cupy.ndarray*) – Input data of forward propagation.

print_report (*file=<_io.TextIOWrapper name='<stdout>' mode='w' encoding='UTF-8'>*)
 Prints a summary report of time profiling in functions.

summary ()
 Returns a summary of time profiling in functions.

Returns A summarized dictionary whose keys are function names and values are dictionaries of *elapsed_time* and *occurrence*.

total_time ()
 Returns total elapsed time in seconds.

Attributes

name = 'TimerHook'

You can also implement your own function-hook to inject arbitrary code before/after the forward/backward propagation.

*chainer.FunctionHook*Base class of hooks for Functions.

chainer.FunctionHook

class `chainer.FunctionHook`

Base class of hooks for Functions.

FunctionHook is a callback object that is registered to *FunctionNode*. Registered function hooks are invoked before and after forward and backward operations of each function.

Function hooks that derive *FunctionHook* are required to implement four methods: *forward_preprocess()*, *forward_postprocess()*, *backward_preprocess()*, and *backward_postprocess()*. By default, these methods do nothing.

Specifically, when `__call__()` method of some function is invoked, *forward_preprocess()* (resp. *forward_postprocess()*) of all function hooks registered to this function are called before (resp. after) forward propagation.

Likewise, when *backward()* of some *Variable* is invoked, *backward_preprocess()* (resp. *backward_postprocess()*) of all function hooks registered to the function which holds this variable as a gradient are called before (resp. after) backward propagation.

There are two ways to register *FunctionHook* objects to *FunctionNode* objects.

First one is to use `with` statement. Function hooks hooked in this way are registered to all functions within `with` statement and are unregistered at the end of `with` statement.

Example

The following code is a simple example in which we measure the elapsed time of a part of forward propagation procedure with *TimerHook*, which is a subclass of *FunctionHook*.

```
>>> from chainer import function_hooks
>>> class Model(chainer.Chain):
...     def __init__(self):
...         super(Model, self).__init__()
...         with self.init_scope():
...             self.l = L.Linear(10, 10)
...     def __call__(self, x1):
...         return F.exp(self.l(x1))
>>> model1 = Model()
>>> model2 = Model()
>>> x = chainer.Variable(np.zeros((1, 10), np.float32))
>>> with chainer.function_hooks.TimerHook() as m:
...     _ = model1(x)
...     y = model2(x)
...     print("Total time : " + str(m.total_time()))
...     model3 = Model()
...     z = model3(y)
Total time : ...
```

In this example, we measure the elapsed times for each forward propagation of all functions in `model1` and `model2` (specifically, `LinearFunction` and `Exp` of `model1` and `model2`). Note that `model3` is not a

target of measurement as `TimerHook` is unregistered before forward propagation of `model3`.

Note: Chainer stores the dictionary of registered function hooks as a thread local object. So, function hooks registered are different depending on threads.

The other one is to register directly to `FunctionNode` object with `add_hook()` method. Function hooks registered in this way can be removed by `delete_hook()` method. Contrary to former registration method, function hooks are registered only to the function which `add_hook()` is called.

Parameters `name` (*str*) – Name of this function hook.

Methods

`__enter__()`

`__exit__(*_)`

added (*function=None*)

Callback function invoked when a function hook is added

Parameters `function` (`FunctionNode`) – Function object to which the function hook is added.

backward_postprocess (*function, in_data, out_grad*)

Callback function invoked after backward propagation.

Parameters

- **function** (`FunctionNode`) – Function object to which the function hook is registered.
- **in_data** (*tuple of numpy.ndarray or tuple of cupy.ndarray*) – Input of forward propagation.
- **out_grad** (*tuple of numpy.ndarray or tuple of cupy.ndarray*) – Gradient data of backward propagation.

backward_preprocess (*function, in_data, out_grad*)

Callback function invoked before backward propagation.

Parameters

- **function** (`FunctionNode`) – Function object to which the function hook is registered.
- **in_data** (*tuple of numpy.ndarray or tuple of cupy.ndarray*) – Input data of forward propagation.
- **out_grad** (*tuple of numpy.ndarray or tuple of cupy.ndarray*) – Gradient data of backward propagation.

deleted (*function=None*)

Callback function invoked when a function hook is deleted

Parameters `function` (`FunctionNode`) – Function object to which the function hook is deleted.

forward_postprocess (*function, in_data*)

Callback function invoked after forward propagation.

Parameters

- **function** (`FunctionNode`) – Function object to which the function hook is registered.
- **in_data** (*tuple of `numpy.ndarray` or `tuple of cupy.ndarray`*) – Input data of forward propagation.

forward_preprocess (*function, in_data*)

Callback function invoked before forward propagation.

Parameters

- **function** (`FunctionNode`) – Function object to which the function hook is registered.
- **in_data** (*tuple of `numpy.ndarray` or `tuple of cupy.ndarray`*) – Input data of forward propagation.

Attributes

name = `'FunctionHook'`

5.3 Link and Chains

Chainer provides many *Link* implementations in the `chainer.links` package.

Note: Some of the links are originally defined in the `chainer.functions` namespace. They are still left in the namespace for backward compatibility, though it is strongly recommended to use them via the `chainer.links` package.

5.3.1 Learnable connections

<code>chainer.links.Bias</code>	Broadcasted elementwise summation with learnable parameters.
<code>chainer.links.Bilinear</code>	Bilinear layer that performs tensor multiplication.
<code>chainer.links.ChildSumTreeLSTM</code>	Child-Sum TreeLSTM unit.
<code>chainer.links.Convolution2D</code>	Two-dimensional convolutional layer.
<code>chainer.links.ConvolutionND</code>	N-dimensional convolution layer.
<code>chainer.links.Deconvolution2D</code>	Two dimensional deconvolution function.
<code>chainer.links.DeconvolutionND</code>	N-dimensional deconvolution function.
<code>chainer.links.DepthwiseConvolution2D</code>	Two-dimensional depthwise convolutional layer.
<code>chainer.links.DilatedConvolution2D</code>	Two-dimensional dilated convolutional layer.
<code>chainer.links.EmbedID</code>	Efficient linear layer for one-hot input.
<code>chainer.links.GRU</code>	Stateful Gated Recurrent Unit function (GRU)
<code>chainer.links.Highway</code>	Highway module.
<code>chainer.links.Inception</code>	Inception module of GoogLeNet.
<code>chainer.links.InceptionBN</code>	Inception module of the new GoogLeNet with Batch-Normalization.
<code>chainer.links.Linear</code>	Linear layer (a.k.a. fully-connected layer).

Continued on next page

Table 16 – continued from previous page

<code>chainer.links.LocalConvolution2D</code>	Two-dimensional local convolutional layer.
<code>chainer.links.LSTM</code>	Fully-connected LSTM layer.
<code>chainer.links.MLPConvolution2D</code>	Two-dimensional MLP convolution layer of Network in Network.
<code>chainer.links.NaryTreeLSTM</code>	N-ary TreeLSTM unit.
<code>chainer.links.NStepBiGRU</code>	Stacked Bi-directional GRU for sequences.
<code>chainer.links.NStepBiLSTM</code>	Stacked Bi-directional LSTM for sequences.
<code>chainer.links.NStepBiRNNReLU</code>	Stacked Bi-directional RNN for sequences.
<code>chainer.links.NStepBiRNNTanh</code>	Stacked Bi-directional RNN for sequences.
<code>chainer.links.NStepGRU</code>	Stacked Uni-directional GRU for sequences.
<code>chainer.links.NStepLSTM</code>	Stacked Uni-directional LSTM for sequences.
<code>chainer.links.NStepRNNReLU</code>	Stacked Uni-directional RNN for sequences.
<code>chainer.links.NStepRNNTanh</code>	Stacked Uni-directional RNN for sequences.
<code>chainer.links.Parameter</code>	Link that just holds a parameter and returns it.
<code>chainer.links.Scale</code>	Broadcasted elementwise product with learnable parameters.
<code>chainer.links.StatefulGRU</code>	Stateful Gated Recurrent Unit function (GRU).
<code>chainer.links.StatelessGRU</code>	Stateless Gated Recurrent Unit function (GRU).
<code>chainer.links.StatefulMGU</code>	
<code>chainer.links.StatelessMGU</code>	
<code>chainer.links.StatefulPeepholeLSTM</code>	Fully-connected LSTM layer with peephole connections.
<code>chainer.links.StatefulZoneoutLSTM</code>	
<code>chainer.links.StatelessLSTM</code>	Stateless LSTM layer.

chainer.links.Bias

class `chainer.links.Bias` (*axis=1, shape=None*)

Broadcasted elementwise summation with learnable parameters.

Computes a elementwise summation as `bias()` function does except that its second input is a learnable bias parameter *b* the link has.

Parameters

- **axis** (*int*) – The first axis of the first input of `bias()` function along which its second input is applied.
- **shape** (*tuple of ints*) – Shape of the learnable bias parameter. If `None`, this link does not have learnable parameters so an explicit bias needs to be given to its `__call__` method's second input.

See also:

See `bias()` for details.

Variables **b** (*Variable*) – Bias parameter if `shape` is given. Otherwise, no attributes.

Methods

`__call__` (**xs*)

Applies broadcasted elementwise summation.

Parameters **xs** (*list of Variables*) – Input variables whose length should be one if the link has a learnable bias parameter, otherwise should be two.

add_param (*name*, *shape=None*, *dtype=<class 'numpy.float32'>*, *initializer=None*)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a `Parameter` object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name*, *value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** (`Link`) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode='share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters `mode` (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type *Link*

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters `link` (*Link*) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope ()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a *Parameter* object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters **name** (*str*) – Name of the attribute to be registered.

repeat (*n_repeat, mode='init'*)

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The `mode` argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))
```

(continues on next page)

(continued from previous page)

```
net = ConvBNReLU().repeat(16, mode='init')
```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the mode was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

`update_enabled`

True if at least one parameter has an update rule enabled.

`within_init_scope`

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

`xp`

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

`chainer.links.Bilinear`

```
class chainer.links.Bilinear(left_size, right_size, out_size, nobias=False, initialW=None, initial_bias=None)
```

Bilinear layer that performs tensor multiplication.

Bilinear is a primitive link that wraps the `bilinear()` functions. It holds parameters W , $V1$, $V2$, and b corresponding to the arguments of `bilinear()`.

Parameters

- **left_size** (*int*) – Dimension of input vector e^1 (J)
- **right_size** (*int*) – Dimension of input vector e^2 (K)
- **out_size** (*int*) – Dimension of output vector y (L)
- **nobias** (*bool*) – If True, parameters $V1$, $V2$, and b are omitted.
- **initialW** (*initializer*) – Initializer to initialize the weight. When it is `numpy.ndarray`, its `ndim` should be 3.
- **initial_bias** (tuple of *initializer*) – Initial values of V^1 , V^2 and b . The length of this argument must be 3. Each element of this tuple must have the shapes of `(left_size, out_size)`, `(right_size, out_size)`, and `(out_size,)`, respectively if they are `numpy.ndarray`. If None, V^1 and V^2 are initialized by the default initializer and b is set to 0.

See also:

See `chainer.functions.bilinear()` for details.

Variables

- **W** (*Variable*) – Bilinear weight parameter.
- **V1** (*Variable*) – Linear weight parameter for the first argument.
- **V2** (*Variable*) – Linear weight parameter for the second argument.
- **b** (*Variable*) – Bias parameter.

Methods

`__call__(e1, e2)`

Applies the bilinear function to inputs and the internal parameters.

Parameters

- **e1** (*Variable*) – Left input.
- **e2** (*Variable*) – Right input.

Returns Output variable.

Return type *Variable*

add_param (*name*, *shape=None*, *dtype=<class 'numpy.float32'>*, *initializer=None*)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a *Parameter* object directly to an attribute within *init_scope()* instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not None, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, *dtype* argument is ignored.

add_persistent (*name*, *value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** (*Link*) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode*='share')

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument *mode* below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters *mode* (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameter variables under the returned link object is re-initialized by calling their *initialize()* method, so that all the parameters may have different initial values from the original link. *copy* means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. *share* means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default *mode* is *share*.

Returns Copied link object.

Return type *Link*

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters *link* (*Link*) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the *enabled* flag of the update rule of each parameter variable to *False*.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the *enabled* flag of the update rule of each parameter variable to *True*.

init_scope ()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the *init_scope* method, we can simply assign a *Parameter* object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters **name** (*str*) – Name of the attribute to be registered.

repeat (*n_repeat, mode='init'*)

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The `mode` argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
```

(continues on next page)

(continued from previous page)

```

        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')

```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the mode was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: `self`

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: `self`

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zero_grads ()

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

`update_enabled`

True if at least one parameter has an update rule enabled.

`within_init_scope`

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

`xp`

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

`chainer.links.ChildSumTreeLSTM`

class `chainer.links.ChildSumTreeLSTM(in_size, out_size)`

Child-Sum TreeLSTM unit.

Warning: This feature is experimental. The interface can change in the future.

This is a Child-Sum TreeLSTM unit as a chain. This link is a variable arguments function, which compounds the states of all children nodes into the new states of a current (parent) node. *states* denotes the cell state, *c*, and the output, *h*, which are produced by this link. This link doesn't keep cell and hidden states internally.

For example, this link is called such as `func(c1, c2, h1, h2, x)` if the number of children nodes is 2, while `func(c1, c2, c3, h1, h2, h3, x)` if that is 3. This function is *independent* from an order of children nodes. Thus, the returns of `func(c1, c2, h1, h2, x)` equal to those of `func(c2, c1, h2, h1, x)`.

Parameters

- **`in_size`** (*int*) – Dimension of input vectors.
- **`out_size`** (*int*) – Dimensionality of cell and output vectors.

Variables

- **`W_x`** (`chainer.links.Linear`) – Linear layer of connections from input vectors.
- **`W_h_aio`** (`chainer.links.Linear`) – Linear layer of connections between (*a*, *i*, *o*) and summation of children's output vectors. *a*, *i* and *o* denotes input compound, input gate and output gate, respectively. *a*, input compound, equals to *u* in the paper by Tai et al.
- **`W_h_f`** (`chainer.links.Linear`) – Linear layer of connections between forget gate *f* and the output of each child.

See the paper for details: [Improved Semantic Representations From Tree-Structured Long Short-Term Memory Networks](#).

Methods

`__call__` (**cshsx*)

Returns new cell state and output of Child-Sum TreeLSTM.

Parameters *cshsx* (list of *Variable*) – Variable arguments which include all cell vectors and all output vectors of variable children, and an input vector.

Returns Returns (*c_{new}*, *h_{new}*), where *c_{new}* represents new cell state vector, and *h_{new}* is new output vector.

Return type tuple of ~chainer.Variable

`__getitem__` (*name*)

Equivalent to `getattr`.

`add_link` (*name*, *link*)

Registers a child link to this chain.

Deprecated since version v2.0.0: Assign the child link directly to an attribute within `init_scope()` instead. For example, the following code

```
chain.add_link('l1', L.Linear(3, 5))
```

can be replaced by the following line.

```
with chain.init_scope():
    chain.l1 = L.Linear(3, 5)
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the child link. This name is also used as the attribute name.
- **link** (*Link*) – The link object to be registered.

`add_param` (*name*, *shape=None*, *dtype=<class 'numpy.float32'>*, *initializer=None*)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a *Parameter* object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.

- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name*, *value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** (*Link*) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode*=*'share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type *Link*

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters **link** (*Link*) – Source link object.

count_params()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a *Parameter* object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (skipself=False)

Returns a generator of all links under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (skipself=False)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (include_uninit=True)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If *name* has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters *name* (*str*) – Name of the attribute to be registered.

repeat (*n_repeat, mode='init'*)

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The *mode* argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The *net* object contains 16 blocks, each of which is *ConvBNReLU*. And the *mode* was *init*, so each block is re-initialized with different parameters. If you give *copy* to this argument, each block has same values for its parameters but its object ID is different from others. If it is *share*, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. *copy* means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. *share* means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.Convolution2D

```
class chainer.links.Convolution2D(self, in_channels, out_channels, ksize=None, stride=1,  
                                pad=0, nobias=False, initialW=None, initial_bias=None, *,  
                                dilate=1, groups=1)
```

Two-dimensional convolutional layer.

This link wraps the `convolution_2d()` function and holds the filter weight and bias vector as parameters.

The output of this function can be non-deterministic when it uses cuDNN. If `chainer.configuration.config.deterministic` is True and cuDNN version is \geq v3, it forces cuDNN to use a deterministic algorithm.

Convolution links can use a feature of cuDNN called autotuning, which selects the most efficient CNN algorithm for images of fixed-size, can provide a significant performance boost for fixed neural nets. To enable, set `chainer.using_config('autotune', True)`

Warning: `deterministic` argument is not supported anymore since v2. Instead, use `chainer.using_config('cudnn_deterministic', value)` (value is either `True` or `False`). See `chainer.using_config()`.

Parameters

- **in_channels** (*int or None*) – Number of channels of input arrays. If `None`, parameter initialization will be deferred until the first forward data pass at which time the size will be determined.
- **out_channels** (*int*) – Number of channels of output arrays.
- **ksize** (*int or pair of ints*) – Size of filters (a.k.a. kernels). `ksize=k` and `ksize=(k, k)` are equivalent.
- **stride** (*int or pair of ints*) – Stride of filter applications. `stride=s` and `stride=(s, s)` are equivalent.
- **pad** (*int or pair of ints*) – Spatial padding width for input arrays. `pad=p` and `pad=(p, p)` are equivalent.
- **nobias** (*bool*) – If `True`, then this link does not use the bias term.
- **initialW** (*initializer*) – Initializer to initialize the weight. When it is `numpy.ndarray`, its `ndim` should be 4.
- **initial_bias** (*initializer*) – Initializer to initialize the bias. If `None`, the bias will be initialized to zero. When it is `numpy.ndarray`, its `ndim` should be 1.
- **dilate** (*int or pair of ints*) – Dilation factor of filter applications. `dilate=d` and `dilate=(d, d)` are equivalent.
- **groups** (*int*) – Number of groups of channels. If the number is greater than 1, input tensor W is divided into some blocks by this value channel-wise. For each tensor blocks, convolution operation will be executed independently. Input channel size `in_channels` and output channel size `out_channels` must be exactly divisible by this value.

See also:

See `chainer.functions.convolution_2d()` for the definition of two-dimensional convolution.

Variables

- **W** (*Variable*) – Weight parameter.
- **b** (*Variable*) – Bias parameter.

Example

There are several ways to make a Convolution2D link.

Let an input vector x be:

```
>>> x = np.arange(1 * 3 * 10 * 10, dtype=np.float32).reshape(1, 3, 10, 10)
```

1. Give the first three arguments explicitly:

```
>>> l = L.Convolution2D(3, 7, 5)
>>> y = l(x)
>>> y.shape
(1, 7, 6, 6)
```

2. Omit `in_channels` or fill it with `None`:

The below two cases are the same.

```
>>> l = L.Convolution2D(7, 5)
>>> y = l(x)
>>> y.shape
(1, 7, 6, 6)
```

```
>>> l = L.Convolution2D(None, 7, 5)
>>> y = l(x)
>>> y.shape
(1, 7, 6, 6)
```

When you omit the first argument, you need to specify the other subsequent arguments from `stride` as keyword arguments. So the below two cases are the same.

```
>>> l = L.Convolution2D(7, 5, stride=1, pad=0)
>>> y = l(x)
>>> y.shape
(1, 7, 6, 6)
```

```
>>> l = L.Convolution2D(None, 7, 5, 1, 0)
>>> y = l(x)
>>> y.shape
(1, 7, 6, 6)
```

Methods

`__call__(x)`

Applies the convolution layer.

Parameters **x** (*Variable*) – Input image.

Returns Output of the convolution.

Return type *Variable*

add_param (*name*, *shape=None*, *dtype=<class 'numpy.float32'>*, *initializer=None*)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a *Parameter* object directly to an attribute within *init_scope()* instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.


```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name, value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** (*Link*) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode='share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays

are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type *Link*

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters **link** (*Link*) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope ()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a *Parameter* object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters **name** (*str*) – Name of the attribute to be registered.

repeat (*n_repeat, mode='init'*)

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The `mode` argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The `net` object contains 16 blocks, each of which is *ConvBNReLU*. And the `mode` was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.ConvolutionND

```
class chainer.links.ConvolutionND(ndim, in_channels, out_channels, ksize, stride=1,
                                   pad=0, nobias=False, initialW=None, initial_bias=None,
                                   cover_all=False)
```

N-dimensional convolution layer.

This link wraps the `convolution_nd()` function and holds the filter weight and bias vector as parameters.

Convolution links can use a feature of cuDNN called autotuning, which selects the most efficient CNN algorithm for images of fixed-size, can provide a significant performance boost for fixed neural nets. To enable, set `chainer.config('autotune', True)`

Parameters

- **ndim** (*int*) – Number of spatial dimensions.
- **in_channels** (*int*) – Number of channels of input arrays.
- **out_channels** (*int*) – Number of channels of output arrays.
- **ksize** (*int or tuple of ints*) – Size of filters (a.k.a. kernels). `ksize=k` and `ksize=(k, k, ..., k)` are equivalent.
- **stride** (*int or tuple of ints*) – Stride of filter application. `stride=s` and `stride=(s, s, ..., s)` are equivalent.
- **pad** (*int or tuple of ints*) – Spatial padding width for input arrays. `pad=p` and `pad=(p, p, ..., p)` are equivalent.
- **nobias** (*bool*) – If True, then this function does not use the bias.
- **initialW** (*initializer*) – Initializer to initialize the weight. When it is `numpy.ndarray`, its `ndim` should be $n + 2$ where n is the number of spatial dimensions.
- **initial_bias** (*initializer*) – Initializer to initialize the bias. If None, the bias will be initialized to zero. When it is `numpy.ndarray`, its `ndim` should 1.
- **cover_all** (*bool*) – If True, all spatial locations are convoluted into some output pixels. It may make the output size larger. `cover_all` needs to be False if you want to use cuDNN.

See also:

See `convolution_nd()` for the definition of N-dimensional convolution. See `convolution_2d()` for the definition of two-dimensional convolution.

Variables

- **W** (*Variable*) – Weight parameter.
- **b** (*Variable*) – Bias parameter. If `initial_bias` is None, set to None.

Methods

`__call__(x)`

Applies N-dimensional convolution layer.

Parameters **x** (*Variable*) – Input image.

Returns Output of convolution.

Return type *Variable*

add_param (*name*, *shape=None*, *dtype=<class 'numpy.float32'>*, *initializer=None*)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a `Parameter` object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name*, *value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** (`Link`) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode='share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters `mode` (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type *Link*

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters `link` (*Link*) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope ()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a *Parameter* object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters **name** (*str*) – Name of the attribute to be registered.

repeat (*n_repeat, mode='init'*)

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The `mode` argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))
```

(continues on next page)

(continued from previous page)

```
net = ConvBNReLU().repeat(16, mode='init')
```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the mode was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

`update_enabled`

True if at least one parameter has an update rule enabled.

`within_init_scope`

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

`xp`

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

`chainer.links.Deconvolution2D`

```
class chainer.links.Deconvolution2D(self, in_channels, out_channels, ksize=None, stride=1,
                                     pad=0, nobias=False, outsize=None, initialW=None,
                                     initial_bias=None, *, groups=1)
```

Two dimensional deconvolution function.

This link wraps the `deconvolution_2d()` function and holds the filter weight and bias vector as parameters.

Deconvolution links can use a feature of cuDNN called autotuning, which selects the most efficient CNN algorithm for images of fixed-size, can provide a significant performance boost for fixed neural nets. To enable, set `chainer.config('autotune', True)`

Warning: deterministic argument is not supported anymore since v2. Instead, use `chainer.config('cudnn_deterministic', value)` (value is either True or False). See `chainer.config()`.

Parameters

- **`in_channels`** (*int* or *None*) – Number of channels of input arrays. If *None*, parameter initialization will be deferred until the first forward data pass at which time the size will be determined.
- **`out_channels`** (*int*) – Number of channels of output arrays.
- **`ksize`** (*int* or *pair of ints*) – Size of filters (a.k.a. kernels). `ksize=k` and `ksize=(k, k)` are equivalent.
- **`stride`** (*int* or *pair of ints*) – Stride of filter applications. `stride=s` and `stride=(s, s)` are equivalent.
- **`pad`** (*int* or *pair of ints*) – Spatial padding width for input arrays. `pad=p` and `pad=(p, p)` are equivalent.
- **`nobias`** (*bool*) – If *True*, then this function does not use the bias term.
- **`outsize`** (*tuple*) – Expected output size of deconvolutional operation. It should be pair of height and width (out_H, out_W). Default value is *None* and the outsize is estimated by input size, stride and pad.
- **`initialW`** (*initializer*) – Initializer to initialize the weight. When it is `numpy.ndarray`, its `ndim` should be 4.

- **initial_bias** (*initializer*) – Initializer to initialize the bias. If `None`, the bias will be initialized to zero. When it is `numpy.ndarray`, its `ndim` should be 1.
- **groups** (*int*) – The number of groups to use grouped deconvolution. The default is one, where grouped deconvolution is not used.

The filter weight has four dimensions (c_I, c_O, k_H, k_W) which indicate the number of input channels, output channels, height and width of the kernels, respectively. The filter weight is initialized with i.i.d. Gaussian random samples, each of which has zero mean and deviation $\sqrt{1/(c_I k_H k_W)}$ by default.

The bias vector is of size c_O . Its elements are initialized by `bias` argument. If `nobias` argument is set to `True`, then this function does not hold the bias parameter.

The output of this function can be non-deterministic when it uses cuDNN. If `chainer.configuration.config.cudnn_deterministic` is `True` and cuDNN version is $\geq v3$, it forces cuDNN to use a deterministic algorithm.

See also:

See `chainer.functions.deconvolution_2d()` for the definition of two-dimensional convolution.

See also:

See `chainer.links.Convolution2D()` for the examples of ways to give arguments to this link.

Example

There are several ways to make a Deconvolution2D link.

Let an input vector `x` be:

```
>>> x = np.arange(1 * 3 * 10 * 10, dtype=np.float32).reshape(1, 3, 10, 10)
```

1. Give the first three arguments explicitly:

In this case, all the other arguments are set to the default values.

```
>>> l = L.Deconvolution2D(3, 7, 4)
>>> y = l(x)
>>> y.shape
(1, 7, 13, 13)
```

2. Omit `in_channels` or fill it with `None`:

The below two cases are the same.

```
>>> l = L.Deconvolution2D(7, 4)
>>> y = l(x)
>>> y.shape
(1, 7, 13, 13)
```

```
>>> l = L.Deconvolution2D(None, 7, 4)
>>> y = l(x)
>>> y.shape
(1, 7, 13, 13)
```

When you omit the first argument, you need to specify the other subsequent arguments from `stride` as keyword arguments. So the below two cases are the same.

```
>>> l = L.Deconvolution2D(None, 7, 4, 2, 1)
>>> y = l(x)
>>> y.shape
(1, 7, 20, 20)
```

```
>>> l = L.Deconvolution2D(7, 4, stride=2, pad=1)
>>> y = l(x)
>>> y.shape
(1, 7, 20, 20)
```

Methods

__call__(x)

Call self as a function.

add_param(name, shape=None, dtype=<class 'numpy.float32'>, initializer=None)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a [Parameter](#) object directly to an attribute within [init_scope\(\)](#) instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent(name, value)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads(link)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters `link` (`Link`) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (`mode='share'`)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters `mode` (`str`) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type `Link`

copyparams (`link`)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters `link` (`Link`) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the `Parameters` held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a *Parameter* object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (skipself=False)

Returns a generator of all links under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (skipself=False)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (include_uninit=True)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (include_uninit=True)

Returns a generator of all parameters under the link hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (name)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If *name* has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters *name* (*str*) – Name of the attribute to be registered.

repeat (n_repeat, mode='init')

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The *mode* argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The *net* object contains 16 blocks, each of which is *ConvBNReLU*. And the *mode* was *init*, so each block is re-initialized with different parameters. If you give *copy* to this argument, each block has same values for its parameters but its object ID is different from others. If it is *share*, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. *copy* means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. *share* means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters *serializer* (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters *device* – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64()

Copies parameter variables and persistent values to CPU.

zerograds()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.DeconvolutionND

```
class chainer.links.DeconvolutionND(ndim, in_channels, out_channels, ksize, stride=1,  
                                     pad=0, nobias=False, outsize=None, initialW=None,  
                                     initial_bias=None)
```

N-dimensional deconvolution function.

This link wraps `deconvolution_nd()` function and holds the filter weight and bias vector as its parameters.

Deconvolution links can use a feature of cuDNN called autotuning, which selects the most efficient CNN algorithm for images of fixed-size, can provide a significant performance boost for fixed neural nets. To enable, set `chainer.config.autotune = True`

Parameters

- **ndim** (*int*) – Number of spatial dimensions.
- **in_channels** (*int*) – Number of channels of input arrays.
- **out_channels** (*int*) – Number of channels of output arrays.
- **ksize** (*int* or *tuple of ints*) – Size of filters (a.k.a. kernels). `ksize=k` and `ksize=(k, k, ..., k)` are equivalent.
- **stride** (*int* or *tuple of ints*) – Stride of filter application. `stride=s` and `stride=(s, s, ..., s)` are equivalent.
- **pad** (*int* or *tuple of ints*) – Spatial padding width for input arrays. `pad=p` and `pad=(p, p, ..., p)` are equivalent.
- **nobias** (*bool*) – If True, then this function does not use the bias.

- **outsize** (*tuple of ints*) – Expected output size of deconvolutional operation. It should be a tuple of ints that represents the output size of each dimension. Default value is `None` and the outsize is estimated with input size, stride and pad.
- **initialW** (*initializer*) – Initializer to initialize the weight. When it is `numpy.ndarray`, its `ndim` should be $n + 2$ where n is the number of spatial dimensions.
- **initial_bias** (*initializer*) – Initializer to initialize the bias. If `None`, the bias will be initialized to zero. When it is `numpy.ndarray`, its `ndim` should be 1.

See also:

`deconvolution_nd()`

Variables

- **W** (*Variable*) – Weight parameter.
- **b** (*Variable*) – Bias parameter. If `initial_bias` is `None`, set to `None`.

Methods

`__call__(x)`

Call self as a function.

`add_param(name, shape=None, dtype=<class 'numpy.float32'>, initializer=None)`

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a `Parameter` object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

`add_persistent(name, value)`

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.

- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** ([Link](#)) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode='share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument *mode* below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameter variables under the returned link object is re-initialized by calling their *initialize()* method, so that all the parameters may have different initial values from the original link. *copy* means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. *share* means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is *share*.

Returns Copied link object.

Return type [Link](#)

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters **link** ([Link](#)) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the [Parameters](#) held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the *enabled* flag of the update rule of each parameter variable to *False*.

enable_update()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a *Parameter* object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (skipself=False)

Returns a generator of all links under the hierarchy.

Parameters `skipself (bool)` – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (skipself=False)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters `skipself (bool)` – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (include_uninit=True)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters `include_uninit (bool)` – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (include_uninit=True)

Returns a generator of all parameters under the link hierarchy.

Parameters `include_uninit (bool)` – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (name)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters **name** (*str*) – Name of the attribute to be registered.

repeat (*n_repeat*, *mode*='init')

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The *mode* argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The *net* object contains 16 blocks, each of which is *ConvBNReLU*. And the *mode* was *init*, so each block is re-initialized with different parameters. If you give *copy* to this argument, each block has same values for its parameters but its object ID is different from others. If it is *share*, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. *copy* means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. *share* means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters `device` – Target device specifier. If omitted, the current device is used.

Returns: self

`to_intel64()`

Copies parameter variables and persistent values to CPU.

`zerograds()`

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

`update_enabled`

True if at least one parameter has an update rule enabled.

`within_init_scope`

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

`xp`

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.DepthwiseConvolution2D

```
class chainer.links.DepthwiseConvolution2D(in_channels, channel_multiplier, ksize,
                                           stride=1, pad=0, nobias=False, initialW=None, initial_bias=None)
```

Two-dimensional depthwise convolutional layer.

This link wraps the `depthwise_convolution_2d()` function and holds the filter weight and bias vector as parameters.

Parameters

- **`in_channels`** (*int*) – Number of channels of input arrays. If `None`, parameter initialization will be deferred until the first forward data pass at which time the size will be determined.
- **`channel_multiplier`** (*int*) – Channel multiplier number. Number of output arrays equal `in_channels * channel_multiplier`.
- **`ksize`** (*int or pair of ints*) – Size of filters (a.k.a. kernels). `ksize=k` and `ksize=(k, k)` are equivalent.
- **`stride`** (*int or pair of ints*) – Stride of filter applications. `stride=s` and `stride=(s, s)` are equivalent.
- **`pad`** (*int or pair of ints*) – Spatial padding width for input arrays. `pad=p` and `pad=(p, p)` are equivalent.

- **nobias** (*bool*) – If True, then this link does not use the bias term.
- **initialW** (*initializer*) – Initializer to initialize the weight. When it is `numpy.ndarray`, its `ndim` should be 4.
- **initial_bias** (*initializer*) – Initializer to initialize the bias. If None, the bias will be initialized to zero. When it is `numpy.ndarray`, its `ndim` should be 1.

See also:

See `chainer.functions.depthwise_convolution_2d()`.

Variables

- **W** (*Variable*) – Weight parameter.
- **b** (*Variable*) – Bias parameter.

Methods

`__call__(x)`

Applies the depthwise convolution layer.

Parameters **x** (*chainer.Variable* or `numpy.ndarray` or `cupy.ndarray`) – Input image.

Returns Output of the depthwise convolution.

Return type *Variable*

add_param (*name*, *shape=None*, *dtype=<class 'numpy.float32'>*, *initializer=None*)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a *Parameter* object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int* or *tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not None, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name*, *value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** ([Link](#)) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode='share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument *mode* below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameter variables under the returned link object is re-initialized by calling their *initialize()* method, so that all the parameters may have different initial values from the original link. *copy* means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. *share* means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is *share*.

Returns Copied link object.

Return type [Link](#)

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters **link** ([Link](#)) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the [Parameters](#) held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for [Chain](#)) by an assignment. A [Parameter](#) object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a [Parameter](#) object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters `skipself` (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters `skipself` (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters `include_uninit` (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters `include_uninit` (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters `name` (*str*) – Name of the attribute to be registered.

repeat (*n_repeat*, *mode*='init')

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The *mode* argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the *mode* was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.DilatedConvolution2D

```
class chainer.links.DilatedConvolution2D (in_channels, out_channels, ksize=None,  
                                           stride=1, pad=0, dilate=1, nobias=False,  
                                           initialW=None, initial_bias=None)
```

Two-dimensional dilated convolutional layer.

This link wraps the `dilated_convolution_2d()` function and holds the filter weight and bias vector as parameters.

Note: You can also define a dilated convolutional layer by passing `dilate` argument to `chainer.links.Convolution2D`. The functionality is the same.

Parameters

- **in_channels** (*int* or *None*) – Number of channels of input arrays. If *None*, parameter initialization will be deferred until the first forward data pass at which time the size will be determined.
- **out_channels** (*int*) – Number of channels of output arrays.

- **ksize** (*int or pair of ints*) – Size of filters (a.k.a. kernels). `ksize=k` and `ksize=(k, k)` are equivalent.
- **stride** (*int or pair of ints*) – Stride of filter applications. `stride=s` and `stride=(s, s)` are equivalent.
- **pad** (*int or pair of ints*) – Spatial padding width for input arrays. `pad=p` and `pad=(p, p)` are equivalent.
- **dilate** (*int or pair of ints*) – Dilation factor of filter applications. `dilate=d` and `dilate=(d, d)` are equivalent.
- **nobias** (*bool*) – If True, then this link does not use the bias term.
- **initialW** (*initializer*) – Initializer to initialize the weight. When it is `numpy.ndarray`, its `ndim` should be 4.
- **initial_bias** (*initializer*) – Initializer to initialize the bias. If None, the bias will be initialized to zero. When it is `numpy.ndarray`, its `ndim` should be 1.

See also:

See `chainer.functions.dilated_convolution_2d()` for the definition of two-dimensional dilated convolution.

Variables

- **W** (*Variable*) – Weight parameter.
- **b** (*Variable*) – Bias parameter.

Example

There are several ways to make a DilatedConvolution2D link.

Let an input vector `x` be:

```
>>> x = np.arange(1 * 3 * 10 * 10, dtype=np.float32).reshape(1, 3, 10, 10)
```

1. Give the first three arguments explicitly:

```
>>> l = L.DilatedConvolution2D(3, 7, 5)
>>> y = l(x)
>>> y.shape
(1, 7, 6, 6)
```

2. Omit `in_channels` or fill it with None:

The below two cases are the same.

```
>>> l = L.DilatedConvolution2D(7, 5)
>>> y = l(x)
>>> y.shape
(1, 7, 6, 6)
```

```
>>> l = L.DilatedConvolution2D(None, 7, 5)
>>> y = l(x)
>>> y.shape
(1, 7, 6, 6)
```

When you omit the first argument, you need to specify the other subsequent arguments from `stride` as keyword arguments. So the below two cases are the same.

```
>>> l = L.DilatedConvolution2D(None, 7, 5, 1, 0, 2)
>>> y = l(x)
>>> y.shape
(1, 7, 2, 2)
```

```
>>> l = L.DilatedConvolution2D(7, 5, stride=1, pad=0, dilate=2)
>>> y = l(x)
>>> y.shape
(1, 7, 2, 2)
```

Methods

`__call__(x)`

Applies the convolution layer.

Parameters `x` (*Variable*) – Input image.

Returns Output of the convolution.

Return type *Variable*

`add_param(name, shape=None, dtype=<class 'numpy.float32'>, initializer=None)`

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a *Parameter* object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

`add_persistent(name, value)`

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** ([Link](#)) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode*=*'share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument *mode* below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameter variables under the returned link object is re-initialized by calling their *initialize()* method, so that all the parameters may have different initial values from the original link. *copy* means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. *share* means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is *share*.

Returns Copied link object.

Return type [Link](#)

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters **link** ([Link](#)) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the [Parameters](#) held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a *Parameter* object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters `name` (*str*) – Name of the attribute to be registered.

repeat (*n_repeat*, *mode*='init')

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The *mode* argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the *mode* was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.EmbedID

class `chainer.links.EmbedID` (*in_size, out_size, initialW=None, ignore_label=None*)

Efficient linear layer for one-hot input.

This is a link that wraps the `embed_id()` function. This link holds the ID (word) embedding matrix *W* as a parameter.

Parameters

- **in_size** (*int*) – Number of different identifiers (a.k.a. vocabulary size).
- **out_size** (*int*) – Size of embedding vector.
- **initialW** (*initializer*) – Initializer to initialize the weight. When it is `numpy.ndarray`, its `ndim` should be 2.
- **ignore_label** (*int* or *None*) – If `ignore_label` is an `int` value, *i*-th column of return value is filled with 0.

See also:

`embed_id()`

Variables `W` (*Variable*) – Embedding parameter matrix.

Example

```
>>> W = np.array([[0, 0, 0],
...               [1, 1, 1],
...               [2, 2, 2]]).astype(np.float32)
>>> W
array([[0., 0., 0.],
       [1., 1., 1.],
       [2., 2., 2.]], dtype=float32)
>>> l = L.EmbedID(W.shape[0], W.shape[1], initialW=W)
>>> x = np.array([2, 1]).astype(np.int32)
>>> x
array([2, 1], dtype=int32)
>>> y = l(x)
>>> y.data
array([[2., 2., 2.],
       [1., 1., 1.]], dtype=float32)
```

Methods

`__call__(x)`

Extracts the word embedding of given IDs.

Parameters `x` (*Variable*) – Batch vectors of IDs.

Returns Batch of corresponding embeddings.

Return type *Variable*

`add_param(name, shape=None, dtype=<class 'numpy.float32'>, initializer=None)`

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a *Parameter* object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name*, *value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** ([Link](#)) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode*=*'share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument *mode* below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameter variables under the returned link object is re-initialized by calling their *initialize()* method, so that all the parameters may have different initial values from the original link. *copy* means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. *share* means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is *share*.

Returns Copied link object.

Return type [Link](#)

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters **link** ([Link](#)) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the [Parameters](#) held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for [Chain](#)) by an assignment. A [Parameter](#) object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a [Parameter](#) object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If *name* has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters *name* (*str*) – Name of the attribute to be registered.

repeat (*n_repeat*, *mode*='init')

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The *mode* argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The *net* object contains 16 blocks, each of which is *ConvBNReLU*. And the *mode* was *init*, so each block is re-initialized with different parameters. If you give *copy* to this argument, each block has same values for its parameters but its object ID is different from others. If it is *share*, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. *copy* means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. *share* means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters *serializer* (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

ignore_label = None

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.GRU

class `chainer.links.GRU` (*in_size, out_size, init=None, inner_init=None, bias_init=0*)

Stateful Gated Recurrent Unit function (GRU)

This is an alias of `StatefulGRU`.

Warning: In Chainer v1, GRU was *stateless*, as opposed to the current implementation. To align with LSTM links, we have changed the naming convention from Chainer v2 so that the shorthand name points the stateful links. You can use `StatelessGRU` for stateless version, whose implementation is identical to GRU in v1.

See issue #2537 for details.

Methods

__call__ (*self*, *x*)
Does forward propagation.

__getitem__ (*name*)
Equivalent to `getattr`.

add_link (*name*, *link*)
Registers a child link to this chain.

Deprecated since version v2.0.0: Assign the child link directly to an attribute within `init_scope()` instead. For example, the following code

```
chain.add_link('l1', L.Linear(3, 5))
```

can be replaced by the following line.

```
with chain.init_scope():  
    chain.l1 = L.Linear(3, 5)
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the child link. This name is also used as the attribute name.
- **link** (*Link*) – The link object to be registered.

add_param (*name*, *shape=None*, *dtype=<class 'numpy.float32'>*, *initializer=None*)
Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a `Parameter` object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():  
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name*, *value*)
Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** (*Link*) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode*=*'share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument *mode* below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameter variables under the returned link object is re-initialized by calling their *initialize()* method, so that all the parameters may have different initial values from the original link. *copy* means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. *share* means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is *share*.

Returns Copied link object.

Return type *Link*

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters **link** (*Link*) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for [Chain](#)) by an assignment. A [Parameter](#) object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a [Parameter](#) object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (skipself=False)

Returns a generator of all links under the hierarchy.

Parameters `skipself (bool)` – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (skipself=False)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters `skipself (bool)` – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (include_uninit=True)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters `include_uninit (bool)` – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (include_uninit=True)

Returns a generator of all parameters under the link hierarchy.

Parameters `include_uninit (bool)` – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (name)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters `name` (*str*) – Name of the attribute to be registered.

repeat (*n_repeat*, *mode*='init')

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The *mode* argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the *mode* was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

reset_state ()

serialize (*serializer*)

Serializes the link object.

Parameters `serializer` (*AbstractSerializer*) – Serializer object.

set_state (*h*)

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.Highway

class `chainer.links.Highway` (*in_out_size*, *nobias=False*, *activate=<function relu>*,
init_Wh=None, *init_Wt=None*, *init_bh=None*, *init_bt=-1*)

Highway module.

In highway network, two gates are added to the ordinal non-linear transformation ($H(x) = \text{activate}(W_h x + b_h)$). One gate is the transform gate $T(x) = \sigma(W_t x + b_t)$, and the other is the carry gate $C(x)$. For simplicity, the author defined $C = 1 - T$. Highway module returns y defined as

$$y = \text{activate}(W_h x + b_h) \odot \sigma(W_t x + b_t) + x \odot (1 - \sigma(W_t x + b_t))$$

The output array has the same spatial size as the input. In order to satisfy this, W_h and W_t must be square matrices.

Parameters

- **in_out_size** (*int*) – Dimension of input and output vectors.
- **nobias** (*bool*) – If True, then this function does not use the bias.

- **activate** – Activation function of plain array. *tanh* is also available.
- **init_Wh** (*initializer*) – Initializer to initialize the weight. When it is `numpy.ndarray`, its `ndim` should be 2.
- **init_bh** (*initializer*) – Initializer to initialize the bias. If `None`, the bias will be initialized to zero. When it is `numpy.ndarray`, its `ndim` should be 1.
- **init_Wt** (*initializer*) – Initializer to initialize the weight. When it is `numpy.ndarray`, its `ndim` should be 2.
- **init_bt** (*initializer*) – Initializer to initialize the bias. If `None`, the bias will be initialized to zero. When it is `numpy.ndarray`, its `ndim` should be 1. Negative value is recommended by the author of the paper. (e.g. -1, -3, ...).

See: [Highway Networks](#).

Methods

__call__ (*x*)

Computes the output of the Highway module.

Parameters *x* (*Variable*) – Input variable.

Returns Output variable. Its array has the same spatial size and the same minibatch size as the input array.

Return type *Variable*

__getitem__ (*name*)

Equivalent to `getattr`.

add_link (*name*, *link*)

Registers a child link to this chain.

Deprecated since version v2.0.0: Assign the child link directly to an attribute within `init_scope()` instead. For example, the following code

```
chain.add_link('l1', L.Linear(3, 5))
```

can be replaced by the following line.

```
with chain.init_scope():
    chain.l1 = L.Linear(3, 5)
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the child link. This name is also used as the attribute name.
- **link** (*Link*) – The link object to be registered.

add_param (*name*, *shape=None*, *dtype=<class 'numpy.float32'>*, *initializer=None*)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a *Parameter* object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name, value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** (*Link*) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode='share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays

are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type *Link*

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters **link** (*Link*) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope ()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a *Parameter* object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters **name** (*str*) – Name of the attribute to be registered.

repeat (*n_repeat, mode='init'*)

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The `mode` argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the `mode` was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.Inception

```
class chainer.links.Inception(in_channels, out1, proj3, out3, proj5, out5, proj_pool,
                             conv_init=None, bias_init=None)
```

Inception module of GoogLeNet.

It applies four different functions to the input array and concatenates their outputs along the channel dimension. Three of them are 2D convolutions of sizes 1x1, 3x3 and 5x5. Convolution paths of 3x3 and 5x5 sizes have 1x1 convolutions (called projections) ahead of them. The other path consists of 1x1 convolution (projection) and 3x3 max pooling.

The output array has the same spatial size as the input. In order to satisfy this, Inception module uses appropriate padding for each convolution and pooling.

See: [Going Deeper with Convolutions](#).

Parameters

- **in_channels** (*int* or *None*) – Number of channels of input arrays.
- **out1** (*int*) – Output size of 1x1 convolution path.
- **proj3** (*int*) – Projection size of 3x3 convolution path.
- **out3** (*int*) – Output size of 3x3 convolution path.
- **proj5** (*int*) – Projection size of 5x5 convolution path.
- **out5** (*int*) – Output size of 5x5 convolution path.
- **proj_pool** (*int*) – Projection size of max pooling path.
- **conv_init** (*initializer*) – Initializer to initialize the convolution matrix weights. When it is `numpy.ndarray`, its `ndim` should be 4.
- **bias_init** (*initializer*) – Initializer to initialize the convolution matrix weights. When it is `numpy.ndarray`, its `ndim` should be 1.

Methods

```
__call__(x)
```

Computes the output of the Inception module.

Parameters **x** (*Variable*) – Input variable.

Returns Output variable. Its array has the same spatial size and the same minibatch size as the input array. The channel dimension has size `out1 + out3 + out5 + proj_pool`.

Return type *Variable*

```
__getitem__(name)
```

Equivalent to `getattr`.

```
add_link(name, link)
```

Registers a child link to this chain.

Deprecated since version v2.0.0: Assign the child link directly to an attribute within `init_scope()` instead. For example, the following code

```
chain.add_link('l1', L.Linear(3, 5))
```

can be replaced by the following line.


```
with chain.init_scope():
    chain.ll = L.Linear(3, 5)
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the child link. This name is also used as the attribute name.
- **link** (*Link*) – The link object to be registered.

add_param (*name*, *shape=None*, *dtype=<class 'numpy.float32'>*, *initializer=None*)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a *Parameter* object directly to an attribute within *init_scope()* instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not *None*, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, *dtype* argument is ignored.

add_persistent (*name*, *value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** (*Link*) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (mode='share')

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters `mode (str)` – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type *Link*

copyparams (link)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters `link (Link)` – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope ()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a `Parameter` object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If *name* has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters *name* (*str*) – Name of the attribute to be registered.

repeat (*n_repeat*, *mode='init'*)

Repeats this link multiple times to make a `Sequential`.

This method returns a `Sequential` object which has the same `Link` multiple times repeatedly. The *mode* argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer `Sequential` block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the mode was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes**update_enabled**

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.InceptionBN

```
class chainer.links.InceptionBN(in_channels, out1, proj3, out3, proj33, out33, pooltype,
                                proj_pool=None, stride=1, conv_init=None, dtype=<class
                                'numpy.float32'>)
```

Inception module of the new GoogLeNet with BatchNormalization.

This chain acts like `Inception`, while InceptionBN uses the `BatchNormalization` on top of each convolution, the 5x5 convolution path is replaced by two consecutive 3x3 convolution applications, and the pooling method is configurable.

See: [Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift](#).

Parameters

- **in_channels** (*int* or *None*) – Number of channels of input arrays.
- **out1** (*int*) – Output size of the 1x1 convolution path.
- **proj3** (*int*) – Projection size of the single 3x3 convolution path.
- **out3** (*int*) – Output size of the single 3x3 convolution path.
- **proj33** (*int*) – Projection size of the double 3x3 convolutions path.
- **out33** (*int*) – Output size of the double 3x3 convolutions path.
- **pooltype** (*str*) – Pooling type. It must be either 'max' or 'avg'.
- **proj_pool** (*int* or *None*) – Projection size in the pooling path. If *None*, no projection is done.
- **stride** (*int*) – Stride parameter of the last convolution of each path.
- **conv_init** (*initializer*) – Initializer to initialize the convolution matrix weights. When it is `numpy.ndarray`, its `ndim` should be 4.
- **dtype** (*numpy.dtype*) – Type to use in `BatchNormalization`.

See also:

Inception

Methods

`__call__(x)`

Call self as a function.

`__getitem__(name)`

Equivalent to `getattr`.

`add_link(name, link)`

Registers a child link to this chain.

Deprecated since version v2.0.0: Assign the child link directly to an attribute within `init_scope()` instead. For example, the following code

```
chain.add_link('l1', L.Linear(3, 5))
```

can be replaced by the following line.

```
with chain.init_scope():
    chain.l1 = L.Linear(3, 5)
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the child link. This name is also used as the attribute name.
- **link** (*Link*) – The link object to be registered.

`add_param(name, shape=None, dtype=<class 'numpy.float32'>, initializer=None)`

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a *Parameter* object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name*, *value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** ([Link](#)) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode*=*'share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument *mode* below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameter variables under the returned link object is re-initialized by calling their *initialize()* method, so that all the parameters may have different initial values from the original link. *copy* means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. *share* means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is *share*.

Returns Copied link object.

Return type [Link](#)

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters **link** ([Link](#)) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the [Parameters](#) held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for [Chain](#)) by an assignment. A [Parameter](#) object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a [Parameter](#) object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If *name* has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters *name* (*str*) – Name of the attribute to be registered.

repeat (*n_repeat*, *mode*='init')

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The *mode* argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The *net* object contains 16 blocks, each of which is *ConvBNReLU*. And the *mode* was *init*, so each block is re-initialized with different parameters. If you give *copy* to this argument, each block has same values for its parameters but its object ID is different from others. If it is *share*, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. *copy* means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. *share* means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters *serializer* (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.Linear

class `chainer.links.Linear` (*in_size*, *out_size=None*, *nobias=False*, *initialW=None*, *initial_bias=None*)

Linear layer (a.k.a. fully-connected layer).

This is a link that wraps the `linear()` function, and holds a weight matrix `W` and optionally a bias vector `b` as parameters.

If `initialW` is left to the default value of `None`, the weight matrix `W` is initialized with i.i.d. Gaussian samples, each of which has zero mean and deviation $\sqrt{1/}$

Parameters

- **in_size** (*int* or *None*) – Dimension of input vectors. If unspecified or `None`, parameter initialization will be deferred until the first forward data pass at which time the size will be determined.
- **out_size** (*int*) – Dimension of output vectors. If only one value is passed for `in_size` and `out_size`, that value will be used for the `out_size` dimension.
- **nobias** (*bool*) – If `True`, then this function does not use the bias.

- **initialW** (*initializer*) – Initializer to initialize the weight. When it is `numpy.ndarray`, its `ndim` should be 2. If `initialW` is `None`, then the weights are initialized with i.i.d. Gaussian samples, each of which has zero mean and deviation $\sqrt{1/}$

See also:

`linear()`

. Variables

- **W** (*Variable*) – Weight parameter.
- **b** (*Variable*) – Bias parameter.

Example

There are several ways to make a Linear link.

Define an input vector `x` as:

```
>>> x = np.array([[0, 1, 2, 3, 4]], np.float32)
```

1. Give the first two arguments explicitly:

Those numbers are considered as the input size and the output size.

```
>>> l = L.Linear(5, 10)
>>> y = l(x)
>>> y.shape
(1, 10)
```

2. Omit `in_size` (give the output size only as the first argument) or fill it with `None`:

In this case, the size of second axis of `x` is used as the input size. So the below two cases are the same.

```
>>> l = L.Linear(10)
>>> y = l(x)
>>> y.shape
(1, 10)
```

```
>>> l = L.Linear(None, 10)
>>> y = l(x)
>>> y.shape
(1, 10)
```

When you omit the first argument, you need to specify the other subsequent arguments from `nobias` as keyword arguments. So the below two cases are the same.

```
>>> l = L.Linear(None, 10, False, None, 0)
>>> y = l(x)
>>> y.shape
(1, 10)
```

```
>>> l = L.Linear(10, nobias=False, initialW=None, initial_bias=0)
>>> y = l(x)
>>> y.shape
(1, 10)
```

Methods

__call__ (*x*)

Applies the linear layer.

Parameters **x** (*Variable*) – Batch of input vectors.

Returns Output of the linear layer.

Return type *Variable*

add_param (*name*, *shape=None*, *dtype=<class 'numpy.float32'>*, *initializer=None*)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a *Parameter* object directly to an attribute within *init_scope()* instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not *None*, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, *dtype* argument is ignored.

add_persistent (*name*, *value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** (*Link*) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (mode='share')

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters `mode` (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default `mode` is `share`.

Returns Copied link object.

Return type *Link*

copyparams (link)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters `link` (*Link*) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope ()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a `Parameter` object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If *name* has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters *name* (*str*) – Name of the attribute to be registered.

repeat (*n_repeat, mode='init'*)

Repeats this link multiple times to make a `Sequential`.

This method returns a `Sequential` object which has the same `Link` multiple times repeatedly. The *mode* argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer `Sequential` block like this:

```

class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')

```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the `mode` was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **`n_repeat`** (*int*) – Number of times to repeat.
- **`mode`** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

`serialize` (*serializer*)

Serializes the link object.

Parameters **`serializer`** (*AbstractSerializer*) – Serializer object.

`to_cpu` ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

`to_gpu` (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **`device`** – Target device specifier. If omitted, the current device is used.

Returns: self

`to_intel64` ()

Copies parameter variables and persistent values to CPU.

zerograds()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes**update_enabled**

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.LocalConvolution2D

```
class chainer.links.LocalConvolution2D(in_channels, out_channels, in_size=None,  
                                       ksize=None, stride=1, nobias=False, initialW=None, initial_bias=None, **kwargs)
```

Two-dimensional local convolutional layer.

This link wraps the `local_convolution_2d()` function and holds the filter weight and bias array as parameters.

Parameters

- **in_channels** (*int*) – Number of channels of input arrays. If either `in_channels` or `in_size` is `None`, parameter initialization will be deferred until the first forward data pass at which time the size will be determined.
- **out_channels** (*int*) – Number of channels of output arrays
- **in_size** (*int or pair of ints*) – Size of each image channel `in_size=k` and `in_size=(k, k)` are equivalent. If either `in_channels` or `in_size` is `None`, parameter initialization will be deferred until the first forward data pass when the size will be determined.
- **ksize** (*int or pair of ints*) – Size of filters (a.k.a. kernels). `ksize=k` and `ksize=(k, k)` are equivalent.
- **stride** (*int or pair of ints*) – Stride of filter applications. `stride=s` and `stride=(s, s)` are equivalent.
- **nobias** (*bool*) – If `True`, then this link does not use the bias term.
- **initialW** (*initializer*) – Initializer to initialize the weight. When it is `numpy.ndarray`, its `ndim` should be 6.
- **initial_bias** (*initializer*) – Initializer to initialize the bias. If `None`, the bias will be initialized to zero. When it is `numpy.ndarray`, its `ndim` should be 3.

See also:

See `chainer.functions.local_convolution_2d()`.

Variables

- **W** (`Variable`) – Weight parameter.
- **b** (`Variable`) – Bias parameter.

Methods

`__call__(x)`

Applies the local convolution layer.

Parameters **x** (`Variable`) – Input image.

Returns Output of the convolution.

Return type `Variable`

add_param (*name*, *shape=None*, *dtype=<class 'numpy.float32'>*, *initializer=None*)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a `Parameter` object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (`str`) – Name of the parameter. This name is also used as the attribute name.
- **shape** (`int` or `tuple of ints`) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name*, *value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (`str`) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** ([Link](#)) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode='share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument *mode* below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameter variables under the returned link object is re-initialized by calling their *initialize()* method, so that all the parameters may have different initial values from the original link. *copy* means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. *share* means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is *share*.

Returns Copied link object.

Return type [Link](#)

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters **link** ([Link](#)) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the [Parameters](#) held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the *enabled* flag of the update rule of each parameter variable to *False*.

enable_update()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for [Chain](#)) by an assignment. A [Parameter](#) object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a [Parameter](#) object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (skipself=False)

Returns a generator of all links under the hierarchy.

Parameters `skipself (bool)` – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (skipself=False)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters `skipself (bool)` – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (include_uninit=True)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters `include_uninit (bool)` – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (include_uninit=True)

Returns a generator of all parameters under the link hierarchy.

Parameters `include_uninit (bool)` – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (name)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters **name** (*str*) – Name of the attribute to be registered.

repeat (*n_repeat*, *mode*='init')

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The *mode* argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The *net* object contains 16 blocks, each of which is *ConvBNReLU*. And the *mode* was *init*, so each block is re-initialized with different parameters. If you give *copy* to this argument, each block has same values for its parameters but its object ID is different from others. If it is *share*, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. *copy* means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. *share* means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters `device` – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64()

Copies parameter variables and persistent values to CPU.

zerograds()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.LSTM

```
class chainer.links.LSTM(in_size, out_size=None, lateral_init=None, upward_init=None,
                        bias_init=None, forget_bias_init=None)
```

Fully-connected LSTM layer.

This is a fully-connected LSTM layer as a chain. Unlike the `lstm()` function, which is defined as a stateless activation function, this chain holds upward and lateral connections as child links.

It also maintains *states*, including the cell state and the output at the previous time step. Therefore, it can be used as a *stateful LSTM*.

This link supports variable length inputs. The mini-batch size of the current input must be equal to or smaller than that of the previous one. The mini-batch size of `c` and `h` is determined as that of the first input `x`. When mini-batch size of *i*-th input is smaller than that of the previous input, this link only updates `c[0:len(x)]` and `h[0:len(x)]` and doesn't change the rest of `c` and `h`. So, please sort input sequences in descending order of lengths before applying the function.

Parameters

- **in_size** (*int*) – Dimension of input vectors. If it is `None` or omitted, parameter initialization will be deferred until the first forward data pass at which time the size will be determined.
- **out_size** (*int*) – Dimensionality of output vectors.

- **lateral_init** – A callable that takes `numpy.ndarray` or `cupy.ndarray` and edits its value. It is used for initialization of the lateral connections. May be `None` to use default initialization.
- **upward_init** – A callable that takes `numpy.ndarray` or `cupy.ndarray` and edits its value. It is used for initialization of the upward connections. May be `None` to use default initialization.
- **bias_init** – A callable that takes `numpy.ndarray` or `cupy.ndarray` and edits its value. It is used for initialization of the biases of cell input, input gate and output gate and gates of the upward connection. May be a scalar, in that case, the bias is initialized by this value. If it is `None`, the cell-input bias is initialized to zero.
- **forget_bias_init** – A callable that takes `numpy.ndarray` or `cupy.ndarray` and edits its value. It is used for initialization of the biases of the forget gate of the upward connection. May be a scalar, in that case, the bias is initialized by this value. If it is `None`, the forget bias is initialized to one.

Variables

- **upward** (`Linear`) – Linear layer of upward connections.
- **lateral** (`Linear`) – Linear layer of lateral connections.
- **c** (`Variable`) – Cell states of LSTM units.
- **h** (`Variable`) – Output at the previous time step.

Example

There are several ways to make a LSTM link.

Let a two-dimensional input array x be:

```
>>> x = np.zeros((1, 10), dtype=np.float32)
```

1. Give both `in_size` and `out_size` arguments:

```
>>> l = L.LSTM(10, 20)
>>> h_new = l(x)
>>> h_new.shape
(1, 20)
```

2. Omit `in_size` argument or fill it with `None`:

The below two cases are the same.

```
>>> l = L.LSTM(20)
>>> h_new = l(x)
>>> h_new.shape
(1, 20)
```

```
>>> l = L.LSTM(None, 20)
>>> h_new = l(x)
>>> h_new.shape
(1, 20)
```

Methods

__call__ (*x*)

Updates the internal state and returns the LSTM outputs.

Parameters *x* (*Variable*) – A new batch from the input sequence.

Returns Outputs of updated LSTM units.

Return type *Variable*

__getitem__ (*name*)

Equivalent to `getattr`.

add_link (*name*, *link*)

Registers a child link to this chain.

Deprecated since version v2.0.0: Assign the child link directly to an attribute within `init_scope()` instead. For example, the following code

```
chain.add_link('l1', L.Linear(3, 5))
```

can be replaced by the following line.

```
with chain.init_scope():
    chain.l1 = L.Linear(3, 5)
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the child link. This name is also used as the attribute name.
- **link** (*Link*) – The link object to be registered.

add_param (*name*, *shape=None*, *dtype=<class 'numpy.float32'>*, *initializer=None*)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a *Parameter* object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name*, *value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** ([Link](#)) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode*=*'share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument *mode* below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameter variables under the returned link object is re-initialized by calling their *initialize()* method, so that all the parameters may have different initial values from the original link. *copy* means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. *share* means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is *share*.

Returns Copied link object.

Return type [Link](#)

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters **link** ([Link](#)) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the [Parameters](#) held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for [Chain](#)) by an assignment. A [Parameter](#) object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a [Parameter](#) object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If *name* has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters *name* (*str*) – Name of the attribute to be registered.

repeat (*n_repeat*, *mode*='init')

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The *mode* argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The *net* object contains 16 blocks, each of which is *ConvBNReLU*. And the *mode* was *init*, so each block is re-initialized with different parameters. If you give *copy* to this argument, each block has same values for its parameters but its object ID is different from others. If it is *share*, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. *copy* means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. *share* means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

reset_state ()

Resets the internal state.

It sets *None* to the *c* and *h* attributes.

serialize (*serializer*)

Serializes the link object.

Parameters `serializer` (`AbstractSerializer`) – Serializer object.

set_state (`c`, `h`)

Sets the internal state.

It sets the `c` and `h` attributes.

Parameters

- `c` (`Variable`) – A new cell states of LSTM units.
- `h` (`Variable`) – A new output at the previous time step.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (`device=None`)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters `device` – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.MLPConvolution2D

```
class chainer.links.MLPConvolution2D (self, in_channels, out_channels, ksize=None, stride=1,
                                     pad=0, activation=relu.relu, conv_init=None,
                                     bias_init=None)
```

Two-dimensional MLP convolution layer of Network in Network.

This is an “mlpconv” layer from the Network in Network paper. This layer is a two-dimensional convolution layer followed by 1x1 convolution layers and interleaved activation functions.

Note that it does not apply the activation function to the output of the last 1x1 convolution layer.

Parameters

- **in_channels** (*int or None*) – Number of channels of input arrays. If it is `None` or omitted, parameter initialization will be deferred until the first forward data pass at which time the size will be determined.
- **out_channels** (*tuple of ints*) – Tuple of number of channels. The *i*-th integer indicates the number of filters of the *i*-th convolution.
- **ksize** (*int or pair of ints*) – Size of filters (a.k.a. kernels) of the first convolution layer. `ksize=k` and `ksize=(k, k)` are equivalent.
- **stride** (*int or pair of ints*) – Stride of filter applications at the first convolution layer. `stride=s` and `stride=(s, s)` are equivalent.
- **pad** (*int or pair of ints*) – Spatial padding width for input arrays at the first convolution layer. `pad=p` and `pad=(p, p)` are equivalent.
- **activation** (*callable*) – Activation function for internal hidden units. You can specify one of activation functions from *built-in activation functions* or your own function. It should not be an *activation functions with parameters* (i.e., *Link* instance). The function must accept one argument (the output from each child link), and return a value. Returned value must be a *Variable* derived from the input *Variable* to perform backpropagation on the variable. Note that this function is not applied to the output of this link.
- **conv_init** – An initializer of weight matrices passed to the convolution layers. This option must be specified as a keyword argument.
- **bias_init** – An initializer of bias vectors passed to the convolution layers. This option must be specified as a keyword argument.

See: [Network in Network](#).

Variables **activation** (*callable*) – Activation function. See the description in the arguments for details.

Methods

__call__ (*x*)

Computes the output of the mlpconv layer.

Parameters **x** (*Variable*) – Input image.

Returns Output of the mlpconv layer.

Return type *Variable*

__getitem__ (*index*)

Returns the child at given index.

Parameters **index** (*int*) – Index of the child in the list.

Returns The *index*-th child link.

Return type *Link*

__len__ ()

Returns the number of children.

`__iter__()`

`add_link(link)`

Registers a child link and adds it to the tail of the list.

Parameters `link` (`Link`) – The link object to be registered.

`add_param(name, shape=None, dtype=<class 'numpy.float32'>, initializer=None)`

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a `Parameter` object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (`str`) – Name of the parameter. This name is also used as the attribute name.
- **shape** (`int` or `tuple of ints`) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

`add_persistent(name, value)`

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (`str`) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

`addgrads(link)`

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters `link` (`Link`) – Source link object.

`append(link)`

Registers a child link and adds it to the tail of the list.

This is equivalent to `add_link()`. This method has been added to emulate the `list` interface.

Parameters `link` (`Link`) – The link object to be registered.

`children()`

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (mode='share')

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters `mode (str)` – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type *Link*

copyparams (link)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters `link (Link)` – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope ()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a `Parameter` object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If *name* has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters *name* (*str*) – Name of the attribute to be registered.

repeat (*n_repeat*, *mode='init'*)

Repeats this link multiple times to make a `Sequential`.

This method returns a `Sequential` object which has the same `Link` multiple times repeatedly. The *mode* argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer `Sequential` block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the mode was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes**update_enabled**

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.NaryTreeLSTM

class `chainer.links.NaryTreeLSTM(in_size, out_size, n_ary=2)`

N-ary TreeLSTM unit.

Warning: This feature is experimental. The interface can change in the future.

This is a N-ary TreeLSTM unit as a chain. This link is a fixed-length arguments function, which compounds the states of all children nodes into the new states of a current (parent) node. *states* denotes the cell state, *c*, and the output, *h*, which are produced by this link. This link doesn't keep cell and hidden states internally.

For example, this link is called such as `func(c1, c2, h1, h2, x)` if the number of children nodes was set 2 (`n_ary = 2`), while `func(c1, c2, c3, h1, h2, h3, x)` if that was 3 (`n_ary = 3`). This function is *dependent* from an order of children nodes unlike Child-Sum TreeLSTM. Thus, the returns of `func(c1, c2, h1, h2, x)` are different from those of `func(c2, c1, h2, h1, x)`.

Parameters

- **in_size** (*int*) – Dimension of input vectors.
- **out_size** (*int*) – Dimensionality of cell and output vectors.
- **n_ary** (*int*) – The number of children nodes in a tree structure.

Variables

- **W_x** (`chainer.links.Linear`) – Linear layer of connections from input vectors.
- **W_h** (`chainer.links.Linear`) – Linear layer of connections between (*a*, *i*, *o*, all *f*) and the output of each child. *a*, *i*, *o* and *f* denotes input compound, input gate, output gate and forget gate, respectively. *a*, input compound, equals to *u* in the paper by Tai et al.

See the papers for details: [Improved Semantic Representations From Tree-Structured Long Short-Term Memory Networks](#), and [A Fast Unified Model for Parsing and Sentence Understanding](#).

Tai et al.'s N-Ary TreeLSTM is little extended in Bowman et al., and this link is based on the variant by Bowman et al. Specifically, eq. 10 in Tai et al. has only one W matrix to be applied to x , consistently for all children. On the other hand, Bowman et al.'s model has multiple matrices, each of which affects the forget gate for each child's cell individually.

Methods

`__call__` (**cshsx*)

Returns new cell state and output of N-ary TreeLSTM.

Parameters *cshsx* (list of *Variable*) – Arguments which include all cell vectors and all output vectors of fixed-length children, and an input vector. The number of arguments must be same as `n_ary * 2 + 1`.

Returns Returns (c_{new}, h_{new}) , where c_{new} represents new cell state vector, and h_{new} is new output vector.

Return type tuple of ~chainer.Variable

`__getitem__` (*name*)

Equivalent to `getattr`.

`add_link` (*name*, *link*)

Registers a child link to this chain.

Deprecated since version v2.0.0: Assign the child link directly to an attribute within `init_scope()` instead. For example, the following code

```
chain.add_link('l1', L.Linear(3, 5))
```

can be replaced by the following line.

```
with chain.init_scope():
    chain.l1 = L.Linear(3, 5)
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the child link. This name is also used as the attribute name.
- **link** (*Link*) – The link object to be registered.

`add_param` (*name*, *shape=None*, *dtype=<class 'numpy.float32'>*, *initializer=None*)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a *Parameter* object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.

- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name, value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** (*Link*) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode='share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type *Link*

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters `link` (`Link`) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the `Parameters` held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope ()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for `Chain`) by an assignment. A `Parameter` object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a `Parameter` object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (`skipself=False`)

Returns a generator of all links under the hierarchy.

Parameters `skipself` (`bool`) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (`skipself=False`)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters `skipself` (`bool`) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (`include_uninit=True`)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters `include_uninit` (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters `include_uninit` (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters `name` (*str*) – Name of the attribute to be registered.

repeat (*n_repeat, mode='init'*)

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The `mode` argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The `net` object contains 16 blocks, each of which is *ConvBNReLU*. And the `mode` was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial

parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.NStepBiGRU

class `chainer.links.NStepBiGRU` (*self, n_layers, in_size, out_size, dropout*)

Stacked Bi-directional GRU for sequences.

This link is stacked version of Bi-directional GRU for sequences. It calculates hidden and cell states of all layer at end-of-string, and all hidden states of the last layer for each time.

Unlike `chainer.functions.n_step_bigru()`, this function automatically sort inputs in descending order by length, and transpose the sequence. Users just need to call the link with a list of `chainer.Variable` holding sequences.

Warning: `use_cudnn` argument is not supported anymore since v2. Instead, use `chainer.using_config('use_cudnn', use_cudnn)`. See `chainer.using_config()`.

Parameters

- **n_layers** (*int*) – Number of layers.
- **in_size** (*int*) – Dimensionality of input vectors.
- **out_size** (*int*) – Dimensionality of hidden states and output vectors.
- **dropout** (*float*) – Dropout ratio.

See also:

`chainer.functions.n_step_bigru()`

Methods

`__call__(self, hx, xs)`

Calculate all hidden states and cell states.

Warning: `train` argument is not supported anymore since v2. Instead, use `chainer.using_config('train', train)`. See `chainer.using_config()`.

Parameters

- **hx** (*Variable* or *None*) – Initial hidden states. If *None* is specified zero-vector is used. Its shape is (S, B, N) for uni-directional RNN and $(2S, B, N)$ for bi-directional RNN where S is the number of layers and is equal to `n_layers`, B is the mini-batch size, and N is the dimension of the hidden units.
- **xs** (*list of ~chainer.Variable*) – List of input sequences. Each element `xs[i]` is a `chainer.Variable` holding a sequence. Its shape is (L_t, I) , where L_t is the length of a sequence for time t , and I is the size of the input and is equal to `in_size`.

Returns

This function returns a tuple containing three elements, `hy` and `ys`.

- `hy` is an updated hidden states whose shape is same as `hx`.
- `ys` is a list of *Variable*. Each element `ys[t]` holds hidden states of the last layer corresponding to an input `xs[t]`. Its shape is (L_t, N) for uni-directional RNN and $(L_t, 2N)$ for bi-directional RNN where L_t is the length of a sequence for time t , and N is size of hidden units.

Return type `tuple`

`__getitem__(index)`

Returns the child at given index.

Parameters `index` (*int*) – Index of the child in the list.

Returns The `index`-th child link.

Return type *Link*

`__len__()`

Returns the number of children.

`__iter__()`

`add_link(link)`

Registers a child link and adds it to the tail of the list.

Parameters `link` (*Link*) – The link object to be registered.

`add_param(name, shape=None, dtype=<class 'numpy.float32'>, initializer=None)`

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a *Parameter* object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

`add_persistent(name, value)`

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

`addgrads(link)`

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters `link` (*Link*) – Source link object.

append (*link*)

Registers a child link and adds it to the tail of the list.

This is equivalent to `add_link()`. This method has been added to emulate the `list` interface.

Parameters **link** (*Link*) – The link object to be registered.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode='share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type *Link*

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters **link** (*Link*) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_hx (*xs*)

init_scope ()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for [Chain](#)) by an assignment. A [Parameter](#) object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a [Parameter](#) object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters **name** (*str*) – Name of the attribute to be registered.

repeat (*n_repeat*, *mode*='init')

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The *mode* argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The *net* object contains 16 blocks, each of which is *ConvBNReLU*. And the *mode* was *init*, so each block is re-initialized with different parameters. If you give *copy* to this argument, each block has same values for its parameters but its object ID is different from others. If it is *share*, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. *copy* means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. *share* means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

rnn (**args*)

Calls RNN function.

This function must be implemented in a child class.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

n_cells

Returns the number of cells.

This function must be implemented in a child class.

n_weights = 6

update_enabled

True if at least one parameter has an update rule enabled.

use_bi_direction = True

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.NStepBiLSTM

class `chainer.links.NStepBiLSTM` (*self, n_layers, in_size, out_size, dropout*)

Stacked Bi-directional LSTM for sequences.

This link is stacked version of Bi-directional LSTM for sequences. It calculates hidden and cell states of all layer at end-of-string, and all hidden states of the last layer for each time.

Unlike `chainer.functions.n_step_bilstm()`, this function automatically sort inputs in descending order by length, and transpose the sequence. Users just need to call the link with a list of `chainer.Variable` holding sequences.

Warning: `use_cudnn` argument is not supported anymore since v2. Instead, use `chainer.using_config('use_cudnn', use_cudnn)`. See `chainer.using_config()`.

Parameters

- **n_layers** (*int*) – Number of layers.
- **in_size** (*int*) – Dimensionality of input vectors.
- **out_size** (*int*) – Dimensionality of hidden states and output vectors.
- **dropout** (*float*) – Dropout ratio.

See also:

`chainer.functions.n_step_bilstm()`

Methods

__call__ (*self, hx, cx, xs*)

Calculate all hidden states and cell states.

Warning: `train` argument is not supported anymore since v2. Instead, use `chainer.using_config('train', train)`. See `chainer.using_config()`.

Parameters

- **hx** (*Variable* or *None*) – Initial hidden states. If *None* is specified zero-vector is used. Its shape is (S, B, N) for uni-directional LSTM and $(2S, B, N)$ for bi-directional LSTM where S is the number of layers and is equal to `n_layers`, B is the mini-batch size, and N is the dimension of the hidden units.
- **cx** (*Variable* or *None*) – Initial cell states. If *None* is specified zero-vector is used. It has the same shape as `hx`.
- **xs** (*list of ~chainer.Variable*) – List of input sequences. Each element `xs[i]` is a *chainer.Variable* holding a sequence. Its shape is (L_t, I) , where L_t is the length of a sequence for time t , and I is the size of the input and is equal to `in_size`.

Returns

This function returns a tuple containing three elements, `hy`, `cy` and `ys`.

- `hy` is an updated hidden states whose shape is the same as `hx`.
- `cy` is an updated cell states whose shape is the same as `cx`.
- `ys` is a list of *Variable*. Each element `ys[t]` holds hidden states of the last layer corresponding to an input `xs[t]`. Its shape is (L_t, N) for uni-directional LSTM and $(L_t, 2N)$ for bi-directional LSTM where L_t is the length of a sequence for time t , and N is size of hidden units.

Return type *tuple*

__getitem__ (*index*)

Returns the child at given index.

Parameters **index** (*int*) – Index of the child in the list.

Returns The `index`-th child link.

Return type *Link*

`__len__()`

Returns the number of children.

`__iter__()`

`add_link(link)`

Registers a child link and adds it to the tail of the list.

Parameters `link` (`Link`) – The link object to be registered.

`add_param(name, shape=None, dtype=<class 'numpy.float32'>, initializer=None)`

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a `Parameter` object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (`str`) – Name of the parameter. This name is also used as the attribute name.
- **shape** (`int` or `tuple of ints`) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

`add_persistent(name, value)`

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (`str`) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

`addgrads(link)`

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters `link` (`Link`) – Source link object.

`append(link)`

Registers a child link and adds it to the tail of the list.

This is equivalent to `add_link()`. This method has been added to emulate the `list` interface.

Parameters `link` (`Link`) – The link object to be registered.

children()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy(mode='share')

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters mode (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type *Link*

copyparams(link)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters link (*Link*) – Source link object.

count_params()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_hx(xs)

init_scope()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a *Parameter* object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If *name* has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters *name* (*str*) – Name of the attribute to be registered.

repeat (*n_repeat*, *mode='init'*)

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The *mode* argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```

class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')

```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the mode was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

rnn (*args)

Calls RNN function.

This function must be implemented in a child class.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters `device` – Target device specifier. If omitted, the current device is used.

Returns: self

`to_intel64()`

Copies parameter variables and persistent values to CPU.

`zerograds()`

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

`n_cells`

Returns the number of cells.

This function must be implemented in a child class.

`n_weights = 8`

`update_enabled`

True if at least one parameter has an update rule enabled.

`use_bi_direction = True`

`within_init_scope`

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

`xp`

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.NStepBiRNNReLU

class `chainer.links.NStepBiRNNReLU` (*self*, *n_layers*, *in_size*, *out_size*, *dropout*)

Stacked Bi-directional RNN for sequences.

This link is stacked version of Bi-directional RNN for sequences. Note that the activation function is `relu`. It calculates hidden and cell states of all layer at end-of-string, and all hidden states of the last layer for each time.

Unlike `chainer.functions.n_step_birnn()`, this function automatically sort inputs in descending order by length, and transpose the sequence. Users just need to call the link with a list of `chainer.Variable` holding sequences.

Warning: `use_cudnn` argument is not supported anymore since v2. Instead, use `chainer.using_config('use_cudnn', use_cudnn)`. See `chainer.using_config()`.

Parameters

- `n_layers` (*int*) – Number of layers.
- `in_size` (*int*) – Dimensionality of input vectors.

- **out_size** (*int*) – Dimensionality of hidden states and output vectors.
- **dropout** (*float*) – Dropout ratio.

See also:

`chainer.functions.n_step_birnn()`

Methods

__call__ (*self*, *hx*, *xs*)

Calculate all hidden states and cell states.

Warning: `train` argument is not supported anymore since v2. Instead, use `chainer.using_config('train', train)`. See `chainer.using_config()`.

Parameters

- **hx** (*Variable* or *None*) – Initial hidden states. If *None* is specified zero-vector is used. Its shape is (*S*, *B*, *N*) for uni-directional RNN and (*2S*, *B*, *N*) for bi-directional RNN where *S* is the number of layers and is equal to `n_layers`, *B* is the mini-batch size, and *N* is the dimension of the hidden units.
- **xs** (*list of ~chainer.Variable*) – List of input sequences. Each element `xs[i]` is a *chainer.Variable* holding a sequence. Its shape is (*L_t*, *I*), where *L_t* is the length of a sequence for time *t*, and *I* is the size of the input and is equal to `in_size`.

Returns

This function returns a tuple containing three elements, *hy* and *ys*.

- *hy* is an updated hidden states whose shape is same as *hx*.
- *ys* is a list of *Variable*. Each element `ys[t]` holds hidden states of the last layer corresponding to an input `xs[t]`. Its shape is (*L_t*, *N*) for uni-directional RNN and (*L_t*, *2N*) for bi-directional RNN where *L_t* is the length of a sequence for time *t*, and *N* is size of hidden units.

Return type tuple

__getitem__ (*index*)

Returns the child at given index.

Parameters **index** (*int*) – Index of the child in the list.

Returns The *index*-th child link.

Return type *Link*

__len__ ()

Returns the number of children.

__iter__ ()

add_link (*link*)

Registers a child link and adds it to the tail of the list.

Parameters **link** (*Link*) – The link object to be registered.

add_param (*name*, *shape=None*, *dtype=<class 'numpy.float32'>*, *initializer=None*)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a `Parameter` object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int* or *tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name*, *value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** (`Link`) – Source link object.

append (*link*)

Registers a child link and adds it to the tail of the list.

This is equivalent to `add_link()`. This method has been added to emulate the `list` interface.

Parameters **link** (`Link`) – The link object to be registered.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode*='share')

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument *mode* below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters *mode* (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameter variables under the returned link object is re-initialized by calling their *initialize()* method, so that all the parameters may have different initial values from the original link. *copy* means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. *share* means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default *mode* is *share*.

Returns Copied link object.

Return type *Link*

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters *link* (*Link*) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the *enabled* flag of the update rule of each parameter variable to *False*.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the *enabled* flag of the update rule of each parameter variable to *True*.

init_hx (*xs*)

init_scope ()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the *init_scope* method, we can simply assign a *Parameter* object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters **name** (*str*) – Name of the attribute to be registered.

repeat (*n_repeat, mode='init'*)

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The `mode` argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
```

(continues on next page)

(continued from previous page)

```

        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

        def __call__(self, x):
            return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')

```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the mode was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

rnn (*args)

Calls RNN function.

This function must be implemented in a child class.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes**n_cells**

Returns the number of cells.

This function must be implemented in a child class.

n_weights = 2**update_enabled**

True if at least one parameter has an update rule enabled.

use_bi_direction = True**within_init_scope**

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.NStepBiRNNTanh

class `chainer.links.NStepBiRNNTanh` (*self*, *n_layers*, *in_size*, *out_size*, *dropout*)

Stacked Bi-directional RNN for sequences.

This link is stacked version of Bi-directional RNN for sequences. Note that the activation function is `tanh`. It calculates hidden and cell states of all layer at end-of-string, and all hidden states of the last layer for each time.

Unlike `chainer.functions.n_step_birnn()`, this function automatically sort inputs in descending order by length, and transpose the sequence. Users just need to call the link with a list of `chainer.Variable` holding sequences.

Warning: `use_cudnn` argument is not supported anymore since v2. Instead, use `chainer.using_config('use_cudnn', use_cudnn)`. See `chainer.using_config()`.

Parameters

- **n_layers** (*int*) – Number of layers.
- **in_size** (*int*) – Dimensionality of input vectors.
- **out_size** (*int*) – Dimensionality of hidden states and output vectors.
- **dropout** (*float*) – Dropout ratio.

See also:

`chainer.functions.n_step_birnn()`

Methods

__call__ (*self*, *hx*, *xs*)

Calculate all hidden states and cell states.

Warning: `train` argument is not supported anymore since v2. Instead, use `chainer.using_config('train', train)`. See `chainer.using_config()`.

Parameters

- **hx** (*Variable* or *None*) – Initial hidden states. If *None* is specified zero-vector is used. Its shape is (*S*, *B*, *N*) for uni-directional RNN and (*2S*, *B*, *N*) for bi-directional RNN where *S* is the number of layers and is equal to `n_layers`, *B* is the mini-batch size, and *N* is the dimension of the hidden units.
- **xs** (*list of ~chainer.Variable*) – List of input sequences. Each element `xs[i]` is a *chainer.Variable* holding a sequence. Its shape is (*L_t*, *I*), where *L_t* is the length of a sequence for time *t*, and *I* is the size of the input and is equal to `in_size`.

Returns

This function returns a tuple containing three elements, *hy* and *ys*.

- *hy* is an updated hidden states whose shape is same as *hx*.
- *ys* is a list of *Variable*. Each element `ys[t]` holds hidden states of the last layer corresponding to an input `xs[t]`. Its shape is (*L_t*, *N*) for uni-directional RNN and (*L_t*, *2N*) for bi-directional RNN where *L_t* is the length of a sequence for time *t*, and *N* is size of hidden units.

Return type *tuple*

__getitem__ (*index*)

Returns the child at given index.

Parameters **index** (*int*) – Index of the child in the list.

Returns The *index*-th child link.

Return type *Link*

__len__ ()

Returns the number of children.

__iter__ ()

add_link (*link*)

Registers a child link and adds it to the tail of the list.

Parameters **link** (*Link*) – The link object to be registered.

add_param (*name*, *shape=None*, *dtype=<class 'numpy.float32'>*, *initializer=None*)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a `Parameter` object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name, value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** (`Link`) – Source link object.

append (*link*)

Registers a child link and adds it to the tail of the list.

This is equivalent to `add_link()`. This method has been added to emulate the `list` interface.

Parameters **link** (`Link`) – The link object to be registered.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode='share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters `mode` (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type *Link*

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters `link` (*Link*) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_hx (*xs*)

init_scope ()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a *Parameter* object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters **name** (*str*) – Name of the attribute to be registered.

repeat (*n_repeat, mode='init'*)

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The `mode` argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
```

(continues on next page)

(continued from previous page)

```

        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

        def __call__(self, x):
            return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')

```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the mode was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

rnn (*args)

Calls RNN function.

This function must be implemented in a child class.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes**n_cells**

Returns the number of cells.

This function must be implemented in a child class.

n_weights = 2**update_enabled**

True if at least one parameter has an update rule enabled.

use_bi_direction = True**within_init_scope**

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.NStepGRU

class `chainer.links.NStepGRU` (*self*, *n_layers*, *in_size*, *out_size*, *dropout*)

Stacked Uni-directional GRU for sequences.

This link is stacked version of Uni-directional GRU for sequences. It calculates hidden and cell states of all layer at end-of-string, and all hidden states of the last layer for each time.

Unlike `chainer.functions.n_step_gru()`, this function automatically sort inputs in descending order by length, and transpose the sequence. Users just need to call the link with a list of `chainer.Variable` holding sequences.

Warning: `use_cudnn` argument is not supported anymore since v2. Instead, use `chainer.using_config('use_cudnn', use_cudnn)`. See `chainer.using_config()`.

Parameters

- **n_layers** (*int*) – Number of layers.
- **in_size** (*int*) – Dimensionality of input vectors.
- **out_size** (*int*) – Dimensionality of hidden states and output vectors.
- **dropout** (*float*) – Dropout ratio.

See also:

`chainer.functions.n_step_gru()`

Methods

__call__ (*self*, *hx*, *xs*)

Calculate all hidden states and cell states.

Warning: `train` argument is not supported anymore since v2. Instead, use `chainer.using_config('train', train)`. See `chainer.using_config()`.

Parameters

- **hx** (*Variable* or *None*) – Initial hidden states. If *None* is specified zero-vector is used. Its shape is (*S*, *B*, *N*) for uni-directional RNN and (*2S*, *B*, *N*) for bi-directional RNN where *S* is the number of layers and is equal to `n_layers`, *B* is the mini-batch size, and *N* is the dimension of the hidden units.
- **xs** (*list of ~chainer.Variable*) – List of input sequences. Each element `xs[i]` is a *chainer.Variable* holding a sequence. Its shape is (*L_t*, *I*), where *L_t* is the length of a sequence for time *t*, and *I* is the size of the input and is equal to `in_size`.

Returns

This function returns a tuple containing three elements, *hy* and *ys*.

- *hy* is an updated hidden states whose shape is same as *hx*.
- *ys* is a list of *Variable*. Each element `ys[t]` holds hidden states of the last layer corresponding to an input `xs[t]`. Its shape is (*L_t*, *N*) for uni-directional RNN and (*L_t*, *2N*) for bi-directional RNN where *L_t* is the length of a sequence for time *t*, and *N* is size of hidden units.

Return type *tuple*

__getitem__ (*index*)

Returns the child at given index.

Parameters **index** (*int*) – Index of the child in the list.

Returns The *index*-th child link.

Return type *Link*

__len__ ()

Returns the number of children.

__iter__ ()

add_link (*link*)

Registers a child link and adds it to the tail of the list.

Parameters **link** (*Link*) – The link object to be registered.

add_param (*name*, *shape=None*, *dtype=<class 'numpy.float32'>*, *initializer=None*)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a `Parameter` object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name, value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** (`Link`) – Source link object.

append (*link*)

Registers a child link and adds it to the tail of the list.

This is equivalent to `add_link()`. This method has been added to emulate the `list` interface.

Parameters **link** (`Link`) – The link object to be registered.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode='share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters `mode` (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type *Link*

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters `link` (*Link*) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_hx (*xs*)

init_scope ()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a *Parameter* object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters **name** (*str*) – Name of the attribute to be registered.

repeat (*n_repeat, mode='init'*)

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The `mode` argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
```

(continues on next page)

(continued from previous page)

```

        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')

```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the mode was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

rnn (*args)

Calls RNN function.

This function must be implemented in a child class.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes**n_cells**

Returns the number of cells.

This function must be implemented in a child class.

n_weights = 6**update_enabled**

True if at least one parameter has an update rule enabled.

use_bi_direction = False**within_init_scope**

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.NStepLSTM

class `chainer.links.NStepLSTM`(*self*, *n_layers*, *in_size*, *out_size*, *dropout*)

Stacked Uni-directional LSTM for sequences.

This link is stacked version of Uni-directional LSTM for sequences. It calculates hidden and cell states of all layer at end-of-string, and all hidden states of the last layer for each time.

Unlike `chainer.functions.n_step_lstm()`, this function automatically sort inputs in descending order by length, and transpose the sequence. Users just need to call the link with a list of `chainer.Variable` holding sequences.

Warning: `use_cudnn` argument is not supported anymore since v2. Instead, use `chainer.using_config('use_cudnn', use_cudnn)`. See `chainer.using_config()`.

Parameters

- **n_layers** (*int*) – Number of layers.
- **in_size** (*int*) – Dimensionality of input vectors.
- **out_size** (*int*) – Dimensionality of hidden states and output vectors.
- **dropout** (*float*) – Dropout ratio.

See also:

`chainer.functions.n_step_lstm()`

Methods

__call__ (*self*, *hx*, *cx*, *xs*)

Calculate all hidden states and cell states.

Warning: `train` argument is not supported anymore since v2. Instead, use `chainer.using_config('train', train)`. See `chainer.using_config()`.

Parameters

- **hx** (*Variable* or *None*) – Initial hidden states. If *None* is specified zero-vector is used. Its shape is (S, B, N) for uni-directional LSTM and $(2S, B, N)$ for bi-directional LSTM where S is the number of layers and is equal to `n_layers`, B is the mini-batch size, and N is the dimension of the hidden units.
- **cx** (*Variable* or *None*) – Initial cell states. If *None* is specified zero-vector is used. It has the same shape as *hx*.
- **xs** (*list of ~chainer.Variable*) – List of input sequences. Each element `xs[i]` is a *chainer.Variable* holding a sequence. Its shape is (L_t, I) , where L_t is the length of a sequence for time t , and I is the size of the input and is equal to `in_size`.

Returns

This function returns a tuple containing three elements, *hy*, *cy* and *ys*.

- *hy* is an updated hidden states whose shape is the same as *hx*.
- *cy* is an updated cell states whose shape is the same as *cx*.
- *ys* is a list of *Variable*. Each element `ys[t]` holds hidden states of the last layer corresponding to an input `xs[t]`. Its shape is (L_t, N) for uni-directional LSTM and $(L_t, 2N)$ for bi-directional LSTM where L_t is the length of a sequence for time t , and N is size of hidden units.

Return type *tuple*

__getitem__ (*index*)

Returns the child at given index.

Parameters **index** (*int*) – Index of the child in the list.

Returns The *index*-th child link.

Return type *Link*

__len__ ()

Returns the number of children.

__iter__ ()

add_link (*link*)

Registers a child link and adds it to the tail of the list.

Parameters **link** (*Link*) – The link object to be registered.

add_param (*name*, *shape=None*, *dtype=<class 'numpy.float32'>*, *initializer=None*)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a `Parameter` object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int* or *tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name*, *value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** (`Link`) – Source link object.

append (*link*)

Registers a child link and adds it to the tail of the list.

This is equivalent to `add_link()`. This method has been added to emulate the `list` interface.

Parameters **link** (`Link`) – The link object to be registered.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode*='share')

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument *mode* below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters *mode* (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameter variables under the returned link object is re-initialized by calling their *initialize()* method, so that all the parameters may have different initial values from the original link. *copy* means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. *share* means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default *mode* is *share*.

Returns Copied link object.

Return type *Link*

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters *link* (*Link*) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the *enabled* flag of the update rule of each parameter variable to *False*.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the *enabled* flag of the update rule of each parameter variable to *True*.

init_hx (*xs*)

init_scope ()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the *init_scope* method, we can simply assign a *Parameter* object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters **name** (*str*) – Name of the attribute to be registered.

repeat (*n_repeat, mode='init'*)

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The `mode` argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
```

(continues on next page)

(continued from previous page)

```

        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

        def __call__(self, x):
            return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')

```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the mode was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

rnn (*args)

Calls RNN function.

This function must be implemented in a child class.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes**n_cells**

Returns the number of cells.

This function must be implemented in a child class.

n_weights = 8**update_enabled**

True if at least one parameter has an update rule enabled.

use_bi_direction = False**within_init_scope**

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.NStepRNNReLU

class `chainer.links.NStepRNNReLU` (*self*, *n_layers*, *in_size*, *out_size*, *dropout*)

Stacked Uni-directional RNN for sequences.

This link is stacked version of Uni-directional RNN for sequences. Note that the activation function is `relu`. It calculates hidden and cell states of all layer at end-of-string, and all hidden states of the last layer for each time.

Unlike `chainer.functions.n_step_rnn()`, this function automatically sort inputs in descending order by length, and transpose the sequence. Users just need to call the link with a list of `chainer.Variable` holding sequences.

Warning: `use_cudnn` argument is not supported anymore since v2. Instead, use `chainer.using_config('use_cudnn', use_cudnn)`. See `chainer.using_config()`.

Parameters

- **n_layers** (*int*) – Number of layers.
- **in_size** (*int*) – Dimensionality of input vectors.
- **out_size** (*int*) – Dimensionality of hidden states and output vectors.
- **dropout** (*float*) – Dropout ratio.

See also:

`chainer.functions.n_step_rnn()`

Methods

__call__ (*self*, *hx*, *xs*)

Calculate all hidden states and cell states.

Warning: `train` argument is not supported anymore since v2. Instead, use `chainer.using_config('train', train)`. See `chainer.using_config()`.

Parameters

- **hx** (*Variable* or *None*) – Initial hidden states. If *None* is specified zero-vector is used. Its shape is (*S*, *B*, *N*) for uni-directional RNN and (*2S*, *B*, *N*) for bi-directional RNN where *S* is the number of layers and is equal to `n_layers`, *B* is the mini-batch size, and *N* is the dimension of the hidden units.
- **xs** (*list of ~chainer.Variable*) – List of input sequences. Each element `xs[i]` is a *chainer.Variable* holding a sequence. Its shape is (*L_t*, *I*), where *L_t* is the length of a sequence for time *t*, and *I* is the size of the input and is equal to `in_size`.

Returns

This function returns a tuple containing three elements, *hy* and *ys*.

- *hy* is an updated hidden states whose shape is same as *hx*.
- *ys* is a list of *Variable*. Each element `ys[t]` holds hidden states of the last layer corresponding to an input `xs[t]`. Its shape is (*L_t*, *N*) for uni-directional RNN and (*L_t*, *2N*) for bi-directional RNN where *L_t* is the length of a sequence for time *t*, and *N* is size of hidden units.

Return type *tuple*

__getitem__ (*index*)

Returns the child at given index.

Parameters **index** (*int*) – Index of the child in the list.

Returns The *index*-th child link.

Return type *Link*

__len__ ()

Returns the number of children.

__iter__ ()

add_link (*link*)

Registers a child link and adds it to the tail of the list.

Parameters **link** (*Link*) – The link object to be registered.

add_param (*name*, *shape=None*, *dtype=<class 'numpy.float32'>*, *initializer=None*)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a `Parameter` object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name, value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** (`Link`) – Source link object.

append (*link*)

Registers a child link and adds it to the tail of the list.

This is equivalent to `add_link()`. This method has been added to emulate the `list` interface.

Parameters **link** (`Link`) – The link object to be registered.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode='share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters `mode` (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type *Link*

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters `link` (*Link*) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_hx (*xs*)

init_scope ()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a *Parameter* object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters **name** (*str*) – Name of the attribute to be registered.

repeat (*n_repeat, mode='init'*)

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The `mode` argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
```

(continues on next page)

(continued from previous page)

```

        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

        def __call__(self, x):
            return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')

```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the mode was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

rnn (*args)

Calls RNN function.

This function must be implemented in a child class.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes**n_cells**

Returns the number of cells.

This function must be implemented in a child class.

n_weights = 2**update_enabled**

True if at least one parameter has an update rule enabled.

use_bi_direction = False**within_init_scope**

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.NStepRNNTanh

class `chainer.links.NStepRNNTanh` (*self*, *n_layers*, *in_size*, *out_size*, *dropout*)

Stacked Uni-directional RNN for sequences.

This link is stacked version of Uni-directional RNN for sequences. Note that the activation function is `tanh`. It calculates hidden and cell states of all layer at end-of-string, and all hidden states of the last layer for each time.

Unlike `chainer.functions.n_step_rnn()`, this function automatically sort inputs in descending order by length, and transpose the sequence. Users just need to call the link with a list of `chainer.Variable` holding sequences.

Warning: `use_cudnn` argument is not supported anymore since v2. Instead, use `chainer.using_config('use_cudnn', use_cudnn)`. See `chainer.using_config()`.

Parameters

- **n_layers** (*int*) – Number of layers.
- **in_size** (*int*) – Dimensionality of input vectors.
- **out_size** (*int*) – Dimensionality of hidden states and output vectors.
- **dropout** (*float*) – Dropout ratio.

See also:

`chainer.functions.n_step_rnn()`

Methods

__call__ (*self*, *hx*, *xs*)

Calculate all hidden states and cell states.

Warning: `train` argument is not supported anymore since v2. Instead, use `chainer.using_config('train', train)`. See `chainer.using_config()`.

Parameters

- **hx** (*Variable* or *None*) – Initial hidden states. If *None* is specified zero-vector is used. Its shape is (*S*, *B*, *N*) for uni-directional RNN and (*2S*, *B*, *N*) for bi-directional RNN where *S* is the number of layers and is equal to `n_layers`, *B* is the mini-batch size, and *N* is the dimension of the hidden units.
- **xs** (*list of ~chainer.Variable*) – List of input sequences. Each element `xs[i]` is a *chainer.Variable* holding a sequence. Its shape is (*L_t*, *I*), where *L_t* is the length of a sequence for time *t*, and *I* is the size of the input and is equal to `in_size`.

Returns

This function returns a tuple containing three elements, *hy* and *ys*.

- *hy* is an updated hidden states whose shape is same as *hx*.
- *ys* is a list of *Variable*. Each element `ys[t]` holds hidden states of the last layer corresponding to an input `xs[t]`. Its shape is (*L_t*, *N*) for uni-directional RNN and (*L_t*, *2N*) for bi-directional RNN where *L_t* is the length of a sequence for time *t*, and *N* is size of hidden units.

Return type *tuple*

__getitem__ (*index*)

Returns the child at given index.

Parameters **index** (*int*) – Index of the child in the list.

Returns The *index*-th child link.

Return type *Link*

__len__ ()

Returns the number of children.

__iter__ ()

add_link (*link*)

Registers a child link and adds it to the tail of the list.

Parameters **link** (*Link*) – The link object to be registered.

add_param (*name*, *shape=None*, *dtype=<class 'numpy.float32'>*, *initializer=None*)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a `Parameter` object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name, value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** (`Link`) – Source link object.

append (*link*)

Registers a child link and adds it to the tail of the list.

This is equivalent to `add_link()`. This method has been added to emulate the `list` interface.

Parameters **link** (`Link`) – The link object to be registered.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode='share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters `mode` (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type *Link*

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters `link` (*Link*) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_hx (*xs*)

init_scope ()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a *Parameter* object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters **name** (*str*) – Name of the attribute to be registered.

repeat (*n_repeat, mode='init'*)

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The `mode` argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
```

(continues on next page)

(continued from previous page)

```

        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

        def __call__(self, x):
            return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')

```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the mode was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

rnn (*args)

Calls RNN function.

This function must be implemented in a child class.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes**n_cells**

Returns the number of cells.

This function must be implemented in a child class.

n_weights = 2**update_enabled**

True if at least one parameter has an update rule enabled.

use_bi_direction = False**within_init_scope**

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.Parameter

class `chainer.links.Parameter(array)`

Link that just holds a parameter and returns it.

Deprecated since version v1.5: The parameters are stored as variables as of v1.5. Use them directly instead.

Parameters **array** – Initial parameter array.

Variables **w**(*Variable*) – Parameter variable.

Methods

__call__(*volatile='off'*)

Returns the parameter variable.

Parameters **volatile**(*Flag*) – The volatility of the returned variable.

Returns A copy of the parameter variable with given volatility.

Return type *Variable*

add_param(*name, shape=None, dtype=<class 'numpy.float32'>, initializer=None*)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a `Parameter` object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name, value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** (*Link*) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode='share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link.

`copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type *Link*

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters **link** (*Link*) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope ()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a *Parameter* object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters **name** (*str*) – Name of the attribute to be registered.

repeat (*n_repeat, mode='init'*)

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The `mode` argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the `mode` was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to

others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.Scale

class `chainer.links.Scale` (*axis=1, W_shape=None, bias_term=False, bias_shape=None*)

Broadcasted elementwise product with learnable parameters.

Computes a elementwise product as `scale()` function does except that its second input is a learnable weight parameter *W* the link has.

Parameters

- **axis** (*int*) – The first axis of the first input of `scale()` function along which its second input is applied.
- **W_shape** (*tuple of ints*) – Shape of learnable weight parameter. If `None`, this link does not have learnable weight parameter so an explicit weight needs to be given to its `__call__` method's second input.
- **bias_term** (*bool*) – Whether to also learn a bias (equivalent to Scale link + Bias link).
- **bias_shape** (*tuple of ints*) – Shape of learnable bias. If `W_shape` is `None`, this should be given to determine the shape. Otherwise, the bias has the same shape `W_shape` with the weight parameter and `bias_shape` is ignored.

See also:

See `scale()` for details.

Variables

- **W** (*Parameter*) – Weight parameter if `W_shape` is given. Otherwise, no `W` attribute.
- **bias** (*Bias*) – Bias term if `bias_term` is `True`. Otherwise, no `bias` attribute.

Methods

__call__ (**xs*)

Applies broadcasted elementwise product.

Parameters *xs* (*list of Variables*) – Input variables whose length should be one if the link has a learnable weight parameter, otherwise should be two.**__getitem__** (*name*)Equivalent to `getattr`.**add_link** (*name, link*)

Registers a child link to this chain.

Deprecated since version v2.0.0: Assign the child link directly to an attribute within `init_scope()` instead. For example, the following code

```
chain.add_link('l1', L.Linear(3, 5))
```

can be replaced by the following line.

```
with chain.init_scope():
    chain.ll = L.Linear(3, 5)
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the child link. This name is also used as the attribute name.
- **link** (*Link*) – The link object to be registered.

add_param (*name*, *shape=None*, *dtype=<class 'numpy.float32'>*, *initializer=None*)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a *Parameter* object directly to an attribute within *init_scope()* instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not *None*, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, *dtype* argument is ignored.

add_persistent (*name*, *value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** (*Link*) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (mode='share')

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters `mode` (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type *Link*

copyparams (link)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters `link` (*Link*) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope ()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a `Parameter` object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If *name* has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters *name* (*str*) – Name of the attribute to be registered.

repeat (*n_repeat*, *mode='init'*)

Repeats this link multiple times to make a `Sequential`.

This method returns a `Sequential` object which has the same `Link` multiple times repeatedly. The *mode* argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer `Sequential` block like this:

```

class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')

```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the mode was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes**update_enabled**

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.StatefulGRU

class `chainer.links.StatefulGRU` (*in_size*, *out_size*, *init=None*, *inner_init=None*, *bias_init=0*)

Stateful Gated Recurrent Unit function (GRU).

Stateful GRU function has six parameters W_r , W_z , W , U_r , U_z , and U . The three parameters W_r , W_z , and W are $n \times m$ matrices, and the others U_r , U_z , and U are $n \times n$ matrices, where m is the length of input vectors and n is the length of hidden vectors.

Given input vector x , Stateful GRU returns the next hidden vector h' defined as

$$\begin{aligned} r &= \sigma(W_r x + U_r h), \\ z &= \sigma(W_z x + U_z h), \\ \tilde{h} &= \tanh(W x + U(r \odot h)), \\ h' &= (1 - z) \odot h + z \odot \tilde{h}, \end{aligned}$$

where h is current hidden vector.

As the name indicates, `StatefulGRU` is *stateful*, meaning that it also holds the next hidden vector h' as a state. For a *stateless* GRU, use `StatelessGRU`.

Parameters

- **in_size** (*int*) – Dimension of input vector x .
- **out_size** (*int*) – Dimension of hidden vector h .
- **init** – Initializer for GRU's input units (W). It is a callable that takes `numpy.ndarray` or `cupy.ndarray` and edits its value. If it is `None`, the default initializer is used.

- **inner_init** – Initializer for the GRU’s inner recurrent units (U). It is a callable that takes `numpy.ndarray` or `cupy.ndarray` and edits its value. If it is `None`, the default initializer is used.
- **bias_init** – Bias initializer. It is a callable that takes `numpy.ndarray` or `cupy.ndarray` and edits its value. If `None`, the bias is set to zero.

Variables **h** (*Variable*) – Hidden vector that indicates the state of *StatefulGRU*.

See also:

- *StatelessGRU*
- *GRU*: an alias of *StatefulGRU*

Example

There are several ways to make a *StatefulGRU* link. Let `x` be a two-dimensional input array:

```
>>> in_size = 10
>>> out_size = 20
>>> x = np.zeros((1, in_size), dtype=np.float32)
```

1. Give only `in_size` and `out_size` arguments:

```
>>> l = L.StatefulGRU(in_size, out_size)
>>> h_new = l(x)
>>> h_new.shape
(1, 20)
```

2. Give all optional arguments:

```
>>> init = np.zeros((out_size, in_size), dtype=np.float32)
>>> inner_init = np.zeros((out_size, out_size), dtype=np.float32)
>>> bias = np.zeros((1, out_size), dtype=np.float32)
>>> l = L.StatefulGRU(in_size, out_size, init=init,
...                  inner_init=inner_init, bias_init=bias)
>>> h_new = l(x)
>>> h_new.shape
(1, 20)
```

Methods

__call__(*x*)

Call self as a function.

__getitem__(*name*)

Equivalent to `getattr`.

add_link(*name*, *link*)

Registers a child link to this chain.

Deprecated since version v2.0.0: Assign the child link directly to an attribute within `init_scope()` instead. For example, the following code

```
chain.add_link('l1', L.Linear(3, 5))
```

can be replaced by the following line.

```
with chain.init_scope():
    chain.l1 = L.Linear(3, 5)
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the child link. This name is also used as the attribute name.
- **link** (*Link*) – The link object to be registered.

add_param (*name, shape=None, dtype=<class 'numpy.float32'>, initializer=None*)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a *Parameter* object directly to an attribute within *init_scope()* instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not *None*, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, *dtype* argument is ignored.

add_persistent (*name, value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** (*Link*) – Source link object.

children()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy(mode='share')

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters mode (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type *Link*

copyparams(link)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters link (*Link*) – Source link object.

count_params()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a *Parameter* object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If *name* has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters *name* (*str*) – Name of the attribute to be registered.

repeat (*n_repeat*, *mode='init'*)

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The *mode* argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```

class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')

```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the mode was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

reset_state()

serialize(*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

set_state(*h*)

to_cpu()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: `self`

to_gpu(*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters `device` – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64()

Copies parameter variables and persistent values to CPU.

zerograds()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.StatelessGRU

class `chainer.links.StatelessGRU` (*in_size*, *out_size*, *init=None*, *inner_init=None*, *bias_init=None*)

Stateless Gated Recurrent Unit function (GRU).

GRU function has six parameters W_r , W_z , W , U_r , U_z , and U . The three parameters W_r , W_z , and W are $n \times m$ matrices, and the others U_r , U_z , and U are $n \times n$ matrices, where m is the length of input vectors and n is the length of hidden vectors.

Given two inputs a previous hidden vector h and an input vector x , GRU returns the next hidden vector h' defined as

$$\begin{aligned} r &= \sigma(W_r x + U_r h), \\ z &= \sigma(W_z x + U_z h), \\ \bar{h} &= \tanh(W x + U(r \odot h)), \\ h' &= (1 - z) \odot h + z \odot \bar{h}, \end{aligned}$$

where σ is the sigmoid function, and \odot is the element-wise product.

As the name indicates, `StatelessGRU` is *stateless*, meaning that it does not hold the value of hidden vector h . For a *stateful* GRU, use `StatefulGRU`.

Parameters

- **in_size** (*int*) – Dimension of input vector x . If `None`, parameter initialization will be deferred until the first forward data pass at which time the size will be determined.
- **out_size** (*int*) – Dimension of hidden vector h , \bar{h} and h' .

See:

- [On the Properties of Neural Machine Translation: Encoder-Decoder Approaches](#) [Cho+, SSST2014].
- [Empirical Evaluation of Gated Recurrent Neural Networks on Sequence Modeling](#) [Chung+NIPS2014 DLWorkshop].

See also:

StatefulGRU

Example

There are several ways to make a `StatelessGRU` link. Let `x` be a two-dimensional input array:

```
>>> in_size = 10
>>> out_size = 20
>>> x = np.zeros((1, in_size), dtype=np.float32)
>>> h = np.zeros((1, out_size), dtype=np.float32)
```

1. Give both `in_size` and `out_size` arguments:

```
>>> l = L.StatelessGRU(in_size, out_size)
>>> h_new = l(h, x)
>>> h_new.shape
(1, 20)
```

2. Omit `in_size` argument or fill it with `None`:

```
>>> l = L.StatelessGRU(None, out_size)
>>> h_new = l(h, x)
>>> h_new.shape
(1, 20)
```

Methods

__call__ (*h, x*)

Call self as a function.

__getitem__ (*name*)

Equivalent to `getattr`.

add_link (*name, link*)

Registers a child link to this chain.

Deprecated since version v2.0.0: Assign the child link directly to an attribute within `init_scope()` instead. For example, the following code

```
chain.add_link('l1', L.Linear(3, 5))
```

can be replaced by the following line.

```
with chain.init_scope():
    chain.ll = L.Linear(3, 5)
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the child link. This name is also used as the attribute name.
- **link** (*Link*) – The link object to be registered.

add_param (*name*, *shape=None*, *dtype=<class 'numpy.float32'>*, *initializer=None*)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a *Parameter* object directly to an attribute within *init_scope()* instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not *None*, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, *dtype* argument is ignored.

add_persistent (*name*, *value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** (*Link*) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (mode='share')

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters `mode` (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type *Link*

copyparams (link)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters `link` (*Link*) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope ()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a *Parameter* object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If *name* has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters *name* (*str*) – Name of the attribute to be registered.

repeat (*n_repeat*, *mode='init'*)

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The *mode* argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```

class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')

```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the mode was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes**update_enabled**

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.StatefulMGU

```
class chainer.links.StatefulMGU(in_size, out_size)
```

Methods**__call__(x)**

Call self as a function.

__getitem__(name)

Equivalent to `getattr`.

add_link(name, link)

Registers a child link to this chain.

Deprecated since version v2.0.0: Assign the child link directly to an attribute within `init_scope()` instead. For example, the following code

```
chain.add_link('l1', L.Linear(3, 5))
```

can be replaced by the following line.

```
with chain.init_scope():
    chain.l1 = L.Linear(3, 5)
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the child link. This name is also used as the attribute name.
- **link** (*Link*) – The link object to be registered.

add_param (*name*, *shape=None*, *dtype=<class 'numpy.float32'>*, *initializer=None*)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a `Parameter` object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name*, *value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** (`Link`) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode='share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters `mode` (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type *Link*

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters `link` (*Link*) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope ()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a *Parameter* object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters **name** (*str*) – Name of the attribute to be registered.

repeat (*n_repeat, mode='init'*)

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The `mode` argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))
```

(continues on next page)

(continued from previous page)

```
net = ConvBNReLU().repeat(16, mode='init')
```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the mode was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

reset_state ()

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

set_state (*h*)

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: `self`

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: `self`

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

`update_enabled`

True if at least one parameter has an update rule enabled.

`within_init_scope`

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

`xp`

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

`chainer.links.StatelessMGU`

```
class chainer.links.StatelessMGU(n_inputs, n_units)
```

Methods

`__call__(h, x)`

`__getitem__(name)`

Equivalent to `getattr`.

`add_link(name, link)`

Registers a child link to this chain.

Deprecated since version v2.0.0: Assign the child link directly to an attribute within `init_scope()` instead. For example, the following code

```
chain.add_link('l1', L.Linear(3, 5))
```

can be replaced by the following line.

```
with chain.init_scope():
    chain.l1 = L.Linear(3, 5)
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the child link. This name is also used as the attribute name.
- **link** (*Link*) – The link object to be registered.

`add_param(name, shape=None, dtype=<class 'numpy.float32'>, initializer=None)`

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a `Parameter` object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name, value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** (*Link*) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode='share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type *Link*

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters **link** (*Link*) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope ()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a *Parameter* object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters **name** (*str*) – Name of the attribute to be registered.

repeat (*n_repeat, mode='init'*)

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The `mode` argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the `mode` was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.

- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.StatefulPeepholeLSTM

class chainer.links.StatefulPeepholeLSTM(*in_size*, *out_size*)

Fully-connected LSTM layer with peephole connections.

This is a fully-connected LSTM layer with peephole connections as a chain. Unlike the [LSTM](#) link, this chain holds `peep_i`, `peep_f` and `peep_o` as child links besides `upward` and `lateral`.

Given a input vector x , Peephole returns the next hidden vector h' defined as

$$\begin{aligned} a &= \tanh(\text{upward}x + \text{lateral}h), \\ i &= \sigma(\text{upward}x + \text{lateral}h + \text{peep}_i c), \\ f &= \sigma(\text{upward}x + \text{lateral}h + \text{peep}_f c), \\ c' &= a \odot i + f \odot c, \\ o &= \sigma(\text{upward}x + \text{lateral}h + \text{peep}_o c'), \\ h' &= o \tanh(c'), \end{aligned}$$

where σ is the sigmoid function, \odot is the element-wise product, c is the current cell state, c' is the next cell state and h is the current hidden vector.

Parameters

- **in_size** (*int*) – Dimension of the input vector x .
- **out_size** (*int*) – Dimension of the hidden vector h .

Variables

- **upward** ([Linear](#)) – Linear layer of upward connections.
- **lateral** ([Linear](#)) – Linear layer of lateral connections.
- **peep_i** ([Linear](#)) – Linear layer of peephole connections to the input gate.
- **peep_f** ([Linear](#)) – Linear layer of peephole connections to the forget gate.
- **peep_o** ([Linear](#)) – Linear layer of peephole connections to the output gate.
- **c** ([Variable](#)) – Cell states of LSTM units.
- **h** ([Variable](#)) – Output at the current time step.

Methods

__call__ (x)

Updates the internal state and returns the LSTM outputs.

Parameters \mathbf{x} ([Variable](#)) – A new batch from the input sequence.

Returns Outputs of updated LSTM units.

Return type [Variable](#)

`__getitem__` (*name*)
Equivalent to `getattr`.

`add_link` (*name*, *link*)
Registers a child link to this chain.

Deprecated since version v2.0.0: Assign the child link directly to an attribute within `init_scope()` instead. For example, the following code

```
chain.add_link('l1', L.Linear(3, 5))
```

can be replaced by the following line.

```
with chain.init_scope():
    chain.l1 = L.Linear(3, 5)
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the child link. This name is also used as the attribute name.
- **link** (*Link*) – The link object to be registered.

`add_param` (*name*, *shape*=None, *dtype*=<class 'numpy.float32'>, *initializer*=None)
Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a `Parameter` object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not None, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

`add_persistent` (*name*, *value*)
Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** ([Link](#)) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode='share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument *mode* below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameter variables under the returned link object is re-initialized by calling their *initialize()* method, so that all the parameters may have different initial values from the original link. *copy* means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. *share* means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is *share*.

Returns Copied link object.

Return type [Link](#)

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters **link** ([Link](#)) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the [Parameters](#) held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the *enabled* flag of the update rule of each parameter variable to *False*.

enable_update()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a *Parameter* object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (skipself=False)

Returns a generator of all links under the hierarchy.

Parameters `skipself (bool)` – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (skipself=False)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters `skipself (bool)` – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (include_uninit=True)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters `include_uninit (bool)` – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (include_uninit=True)

Returns a generator of all parameters under the link hierarchy.

Parameters `include_uninit (bool)` – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (name)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters **name** (*str*) – Name of the attribute to be registered.

repeat (*n_repeat*, *mode*='init')

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The *mode* argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The *net* object contains 16 blocks, each of which is *ConvBNReLU*. And the *mode* was *init*, so each block is re-initialized with different parameters. If you give *copy* to this argument, each block has same values for its parameters but its object ID is different from others. If it is *share*, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. *copy* means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. *share* means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

reset_state ()

Resets the internal states.

It sets *None* to the *c* and *h* attributes.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.StatefulZoneoutLSTM

```
class chainer.links.StatefulZoneoutLSTM(in_size, out_size, c_ratio=0.5, h_ratio=0.5,
                                         **kwargs)
```

Methods

__call__ (*x*)

Updates the internal state and returns the LSTM outputs.

Parameters **x** (*Variable*) – A new batch from the input sequence.

Returns Outputs of updated LSTM units.

Return type *Variable*

__getitem__ (*name*)

Equivalent to `getattr`.

add_link (*name*, *link*)

Registers a child link to this chain.

Deprecated since version v2.0.0: Assign the child link directly to an attribute within `init_scope()` instead. For example, the following code

```
chain.add_link('l1', L.Linear(3, 5))
```

can be replaced by the following line.

```
with chain.init_scope():
    chain.l1 = L.Linear(3, 5)
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the child link. This name is also used as the attribute name.
- **link** (*Link*) – The link object to be registered.

add_param (*name*, *shape*=None, *dtype*=<class 'numpy.float32'>, *initializer*=None)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a `Parameter` object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not None, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name*, *value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters `link` (`Link`) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (`mode='share'`)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters `mode` (`str`) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type `Link`

copyparams (`link`)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters `link` (`Link`) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the `Parameters` held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a *Parameter* object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (skipself=False)

Returns a generator of all links under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (skipself=False)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (include_uninit=True)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (include_uninit=True)

Returns a generator of all parameters under the link hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (name)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If *name* has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters *name* (*str*) – Name of the attribute to be registered.

repeat (n_repeat, mode='init')

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The *mode* argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The *net* object contains 16 blocks, each of which is *ConvBNReLU*. And the *mode* was *init*, so each block is re-initialized with different parameters. If you give *copy* to this argument, each block has same values for its parameters but its object ID is different from others. If it is *share*, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. *copy* means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. *share* means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

reset_state()

Resets the internal state.

It sets *None* to the *c* and *h* attributes.

serialize(*serializer*)

Serializes the link object.

Parameters *serializer* (*AbstractSerializer*) – Serializer object.

set_state(*c*, *h*)

Sets the internal state.

It sets the *c* and *h* attributes.

Parameters

- **c** (*Variable*) – A new cell states of LSTM units.
- **h** (*Variable*) – A new output at the previous time step.

to_cpu()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu(device=None)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64()

Copies parameter variables and persistent values to CPU.

zerograds()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.StatelessLSTM

```
class chainer.links.StatelessLSTM(in_size, out_size=None, lateral_init=None, upward_init=None, bias_init=None, forget_bias_init=None)
```

Stateless LSTM layer.

This is a fully-connected LSTM layer as a chain. Unlike the `lstm()` function, this chain holds upward and lateral connections as child links. This link doesn't keep cell and hidden states.

Parameters

- **in_size** (*int* or *None*) – Dimension of input vectors. If *None*, parameter initialization will be deferred until the first forward data pass at which time the size will be determined.
- **out_size** (*int*) – Dimensionality of output vectors.

Variables

- **upward**(`chainer.links.Linear`) – Linear layer of upward connections.
- **lateral**(`chainer.links.Linear`) – Linear layer of lateral connections.

Example

There are several ways to make a StatelessLSTM link.

Let a two-dimensional input array x , a cell state array h , and the output array of the previous step h be:

```
>>> x = np.zeros((1, 10), dtype=np.float32)
>>> c = np.zeros((1, 20), dtype=np.float32)
>>> h = np.zeros((1, 20), dtype=np.float32)
```

1. Give both `in_size` and `out_size` arguments:

```
>>> l = L.StatelessLSTM(10, 20)
>>> c_new, h_new = l(c, h, x)
>>> c_new.shape
(1, 20)
>>> h_new.shape
(1, 20)
```

2. Omit `in_size` argument or fill it with `None`:

The below two cases are the same.

```
>>> l = L.StatelessLSTM(20)
>>> c_new, h_new = l(c, h, x)
>>> c_new.shape
(1, 20)
>>> h_new.shape
(1, 20)
```

```
>>> l = L.StatelessLSTM(None, 20)
>>> c_new, h_new = l(c, h, x)
>>> c_new.shape
(1, 20)
>>> h_new.shape
(1, 20)
```

Methods

`__call__(c, h, x)`

Returns new cell state and updated output of LSTM.

Parameters

- **c** (`Variable`) – Cell states of LSTM units.
- **h** (`Variable`) – Output at the previous time step.
- **x** (`Variable`) – A new batch from the input sequence.

Returns Returns `(c_new, h_new)`, where `c_new` represents new cell state, and `h_new` is updated output of LSTM units.

Return type tuple of ~chainer.Variable

__getitem__ (*name*)
Equivalent to getattr.

add_link (*name*, *link*)
Registers a child link to this chain.

Deprecated since version v2.0.0: Assign the child link directly to an attribute within `init_scope()` instead. For example, the following code

```
chain.add_link('l1', L.Linear(3, 5))
```

can be replaced by the following line.

```
with chain.init_scope():  
    chain.l1 = L.Linear(3, 5)
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the child link. This name is also used as the attribute name.
- **link** (*Link*) – The link object to be registered.

add_param (*name*, *shape*=None, *dtype*=<class 'numpy.float32'>, *initializer*=None)
Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a `Parameter` object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():  
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not None, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name*, *value*)
Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.

- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** ([Link](#)) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode*='share')

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument *mode* below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameter variables under the returned link object is re-initialized by calling their *initialize()* method, so that all the parameters may have different initial values from the original link. *copy* means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. *share* means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is *share*.

Returns Copied link object.

Return type [Link](#)

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters **link** ([Link](#)) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the [Parameters](#) held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the *enabled* flag of the update rule of each parameter variable to *False*.

enable_update()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a *Parameter* object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (skipself=False)

Returns a generator of all links under the hierarchy.

Parameters `skipself (bool)` – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (skipself=False)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters `skipself (bool)` – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (include_uninit=True)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters `include_uninit (bool)` – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (include_uninit=True)

Returns a generator of all parameters under the link hierarchy.

Parameters `include_uninit (bool)` – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (name)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters **name** (*str*) – Name of the attribute to be registered.

repeat (*n_repeat*, *mode*='init')

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The *mode* argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The *net* object contains 16 blocks, each of which is *ConvBNReLU*. And the *mode* was *init*, so each block is re-initialized with different parameters. If you give *copy* to this argument, each block has same values for its parameters but its object ID is different from others. If it is *share*, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. *copy* means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. *share* means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters `device` – Target device specifier. If omitted, the current device is used.

Returns: self

`to_intel64()`

Copies parameter variables and persistent values to CPU.

`zerograds()`

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

`update_enabled`

True if at least one parameter has an update rule enabled.

`within_init_scope`

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

`xp`

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

5.3.2 Activation/loss/normalization functions with parameters

<code>chainer.links.BatchNormalization</code>	Batch normalization layer on outputs of linear or convolution functions.
<code>chainer.links.BatchRenormalization</code>	Batch renormalization layer on outputs of linear or convolution functions.
<code>chainer.links.LayerNormalization</code>	Layer normalization layer on outputs of linear functions.
<code>chainer.links.BinaryHierarchicalSoftmax</code>	Hierarchical softmax layer over binary tree.
<code>chainer.links.BlackOut</code>	BlackOut loss layer.
<code>chainer.links.CRF1d</code>	Linear-chain conditional random field loss layer.
<code>chainer.links.SimplifiedDropconnect</code>	Fully-connected layer with simplified dropconnect regularization.
<code>chainer.links.PReLU</code>	Parametric ReLU function as a link.
<code>chainer.links.Swish</code>	Swish activation function as a link.
<code>chainer.links.Maxout</code>	Fully-connected maxout layer.
<code>chainer.links.NegativeSampling</code>	Negative sampling loss layer.

chainer.links.BatchNormalization

```
class chainer.links.BatchNormalization(size, decay=0.9, eps=2e-05, dtype=<class
                                     'numpy.float32'>, use_gamma=True,
                                     use_beta=True, initial_gamma=None, ini-
                                     tial_beta=None)
```

Batch normalization layer on outputs of linear or convolution functions.

This link wraps the `batch_normalization()` and `fixed_batch_normalization()` functions.

It runs in three modes: training mode, fine-tuning mode, and testing mode.

In training mode, it normalizes the input by *batch statistics*. It also maintains approximated population statistics by moving averages, which can be used for instant evaluation in testing mode.

In fine-tuning mode, it accumulates the input to compute *population statistics*. In order to correctly compute the population statistics, a user must use this mode to feed mini-batches running through whole training dataset.

In testing mode, it uses pre-computed population statistics to normalize the input variable. The population statistics is approximated if it is computed by training mode, or accurate if it is correctly computed by fine-tuning mode.

Parameters

- **size** (*int* or *tuple of ints*) – Size (or shape) of channel dimensions.
- **decay** (*float*) – Decay rate of moving average. It is used on training.
- **eps** (*float*) – Epsilon value for numerical stability.
- **dtype** (*numpy.dtype*) – Type to use in computing.
- **use_gamma** (*bool*) – If `True`, use scaling parameter. Otherwise, use `unit(1)` which makes no effect.
- **use_beta** (*bool*) – If `True`, use shifting parameter. Otherwise, use `unit(0)` which makes no effect.

See: [Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift](#)

See also:

`batch_normalization()`, `fixed_batch_normalization()`

Variables

- **gamma** (*Variable*) – Scaling parameter.
- **beta** (*Variable*) – Shifting parameter.
- **avg_mean** (*numpy.ndarray* or *cupy.ndarray*) – Population mean.
- **avg_var** (*numpy.ndarray* or *cupy.ndarray*) – Population variance.
- **N** (*int*) – Count of batches given for fine-tuning.
- **decay** (*float*) – Decay rate of moving average. It is used on training.
- **eps** (*float*) – Epsilon value for numerical stability. This value is added to the batch variances.

Example

```
>>> x = np.arange(12).reshape(4, 3).astype(np.float32) ** 2
>>> x
array([[ 0.,  1.,  4.],
       [ 9., 16., 25.],
       [36., 49., 64.],
       [81., 100., 121.]], dtype=float32)
>>> bn = chainer.links.BatchNormalization(3)
>>> bn(x)
variable([[ -1.          , -1.0664359 , -1.1117983 ],
          [-0.71428573, -0.6714596 , -0.6401263 ],
          [ 0.14285715,  0.19748813,  0.23583598],
          [ 1.5714287 ,  1.5404074 ,  1.5160885 ]])
>>> (x - x.mean(axis=0)) / np.sqrt(x.var(axis=0) + 2e-5)
array([[ -1.          , -1.0664359 , -1.1117983 ],
       [-0.71428573, -0.6714596 , -0.6401263 ],
       [ 0.14285715,  0.19748813,  0.235836  ],
       [ 1.5714285 ,  1.5404074 ,  1.5160886 ]], dtype=float32)
```

Consider an input of batched 10 images of 32x32 with 3 channels.

```
>>> x = np.random.randn(10, 3, 32, 32).astype(np.float32)
```

To normalize for each channel, give the number of channels to size.

```
>>> bn = chainer.links.BatchNormalization(3)
>>> bn.avg_mean.shape
(3,)
>>> bn.beta += 2.0
>>> bn.gamma *= 5.0
>>> list(sorted(bn.namedparams()))
[('/beta', variable([2., ...])), ('/gamma', variable([5., ...]))]
>>> y = bn(x)
>>> y.shape
(10, 3, 32, 32)
>>> np.testing.assert_allclose(
...     y.array.mean(axis=(0, 2, 3)), bn.beta.array, atol=1e-6)
>>> np.testing.assert_allclose(
...     y.array.std(axis=(0, 2, 3)),
...     bn.gamma.array, atol=1e-3)
```

To normalize for each channel for each pixel, size should be the tuple of the dimensions.

```
>>> bn = chainer.links.BatchNormalization((3, 32, 32))
>>> bn.avg_mean.shape
(3, 32, 32)
>>> y = bn(x)
>>> y.shape
(10, 3, 32, 32)
>>> np.testing.assert_allclose(
...     y.array.mean(axis=0), bn.beta.array, atol=1e-6)
>>> np.testing.assert_allclose(
...     y.array.std(axis=0),
...     bn.gamma.array, atol=1e-3)
```

By default, channel axis is (or starts from) the 1st axis of the input shape.

Methods

__call__ (*self*, *x*, *finetune=False*)

Invokes the forward propagation of BatchNormalization.

In training mode, the BatchNormalization computes moving averages of mean and variance for evaluation during training, and normalizes the input using batch statistics.

Warning: `test` argument is not supported anymore since v2. Instead, use `chainer.using_config('train', False)`. See `chainer.using_config()`.

Parameters

- **x** (*Variable*) – Input variable.
- **finetune** (*bool*) – If it is in the training mode and `finetune` is `True`, BatchNormalization runs in fine-tuning mode; it accumulates the input array to compute population statistics for normalization, and normalizes the input using batch statistics.

add_param (*name*, *shape=None*, *dtype=<class 'numpy.float32'>*, *initializer=None*)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a `Parameter` object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name*, *value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** ([Link](#)) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode='share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument *mode* below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameter variables under the returned link object is re-initialized by calling their *initialize()* method, so that all the parameters may have different initial values from the original link. *copy* means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. *share* means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is *share*.

Returns Copied link object.

Return type [Link](#)

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters **link** ([Link](#)) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the [Parameters](#) held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the *enabled* flag of the update rule of each parameter variable to *False*.

enable_update()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a *Parameter* object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (skipself=False)

Returns a generator of all links under the hierarchy.

Parameters `skipself (bool)` – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (skipself=False)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters `skipself (bool)` – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (include_uninit=True)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters `include_uninit (bool)` – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (include_uninit=True)

Returns a generator of all parameters under the link hierarchy.

Parameters `include_uninit (bool)` – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (name)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters **name** (*str*) – Name of the attribute to be registered.

repeat (*n_repeat*, *mode*='init')

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The *mode* argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The *net* object contains 16 blocks, each of which is *ConvBNReLU*. And the *mode* was *init*, so each block is re-initialized with different parameters. If you give *copy* to this argument, each block has same values for its parameters but its object ID is different from others. If it is *share*, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. *copy* means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. *share* means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

start_finetuning ()

Resets the population count for collecting population statistics.

This method can be skipped if it is the first time to use the fine-tuning mode. Otherwise, this method should be called before starting the fine-tuning mode again.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

avg_mean = None

avg_var = None

beta = None

gamma = None

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.BatchRenormalization

```
class chainer.links.BatchRenormalization (size, rmax=1, dmax=0, decay=0.9, eps=2e-05,
                                         dtype=<class 'numpy.float32'>,
                                         use_gamma=True, use_beta=True,
                                         initial_gamma=None, initial_beta=None,
                                         freeze_running_statistics=False)
```

Batch renormalization layer on outputs of linear or convolution functions.

This link wraps the `batch_renormalization()` and `fixed_batch_renormalization()` functions.

This is an extension of batch normalization, which ensures that the training and inference models generate the same outputs that depend on individual examples rather than the entire minibatch.

See: [Batch Renormalization: Towards Reducing Minibatch Dependence in Batch-Normalized Models](#)

See also:

```
batch_renormalization(),                               fixed_batch_renormalization()
batch_normalization(),
```

Methods

`__call__(self, x, finetune=False)`

Invokes the forward propagation of BatchNormalization.

In training mode, the BatchNormalization computes moving averages of mean and variance for evaluation during training, and normalizes the input using batch statistics.

Warning: `test` argument is not supported anymore since v2. Instead, use `chainer.using_config('train', False)`. See [chainer.using_config\(\)](#).

Parameters

- **x** (*Variable*) – Input variable.
- **finetune** (*bool*) – If it is in the training mode and `finetune` is `True`, BatchNormalization runs in fine-tuning mode; it accumulates the input array to compute population statistics for normalization, and normalizes the input using batch statistics.

`add_param(name, shape=None, dtype=<class 'numpy.float32'>, initializer=None)`

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a [Parameter](#) object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name*, *value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** ([Link](#)) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode*=*'share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument *mode* below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameter variables under the returned link object is re-initialized by calling their *initialize()* method, so that all the parameters may have different initial values from the original link. *copy* means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. *share* means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is *share*.

Returns Copied link object.

Return type [Link](#)

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters **link** ([Link](#)) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the [Parameters](#) held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for [Chain](#)) by an assignment. A [Parameter](#) object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a [Parameter](#) object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If *name* has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters *name* (*str*) – Name of the attribute to be registered.

repeat (*n_repeat*, *mode*='init')

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The *mode* argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The *net* object contains 16 blocks, each of which is *ConvBNReLU*. And the *mode* was *init*, so each block is re-initialized with different parameters. If you give *copy* to this argument, each block has same values for its parameters but its object ID is different from others. If it is *share*, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. *copy* means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. *share* means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters *serializer* (*AbstractSerializer*) – Serializer object.

start_finetuning ()

Resets the population count for collecting population statistics.

This method can be skipped if it is the first time to use the fine-tuning mode. Otherwise, this method should be called before starting the fine-tuning mode again.

to_cpu()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu(device=None)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64()

Copies parameter variables and persistent values to CPU.

zerograds()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

avg_mean = None

avg_var = None

beta = None

gamma = None

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.LayerNormalization

class `chainer.links.LayerNormalization`(*size=None, eps=1e-06, initial_gamma=None, initial_beta=None*)

Layer normalization layer on outputs of linear functions.

Warning: This feature is experimental. The interface can change in the future.

This link implements a “layer normalization” layer which normalizes the input units by statistics that are computed along the second axis, scales and shifts them. Parameter initialization will be deferred until the first forward data pass at which time the size will be determined.

Parameters

- **size** (*int*) – Size of input units. If `None`, parameter initialization will be deferred until the first forward data pass at which time the size will be determined.
- **eps** (*float*) – Epsilon value for numerical stability of normalization.
- **initial_gamma** (*Initializer*) – Initializer for scaling vector. If `None`, then the vector is filled by 1. If a scalar, the vector is filled by it. If `numpy.ndarray`, the vector is set by it.
- **initial_beta** (*Initializer*) – Initializer for shifting vector. If `None`, then the vector is filled by 0. If a scalar, the vector is filled by it. If `numpy.ndarray`, the vector is set by it.

Variables

- **gamma** (*Parameter*) – Scaling parameter.
- **beta** (*Parameter*) – Shifting parameter.
- **eps** (*float*) – Epsilon value for numerical stability.

See: [Layer Normalization](#)

Methods

__call__ (*x*)

Apply layer normalization to given input.

Parameters **x** (*Variable*) – Batch vectors. Shape of this value must be *(batch_size, unit_size)*, e.g., the output of `linear()`.

Returns Output of the layer normalization.

Return type *Variable*

add_param (*name, shape=None, dtype=<class 'numpy.float32'>, initializer=None*)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a *Parameter* object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute’s type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.

- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name, value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** ([Link](#)) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode='share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type [Link](#)

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters `link` (`Link`) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the `Parameters` held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope ()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for `Chain`) by an assignment. A `Parameter` object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a `Parameter` object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (`skipself=False`)

Returns a generator of all links under the hierarchy.

Parameters `skipself` (`bool`) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (`skipself=False`)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters `skipself` (`bool`) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (`include_uninit=True`)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters `include_uninit` (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters `include_uninit` (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters `name` (*str*) – Name of the attribute to be registered.

repeat (*n_repeat, mode='init'*)

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The `mode` argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the `mode` was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial

parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.BinaryHierarchicalSoftmax

class `chainer.links.BinaryHierarchicalSoftmax` (*in_size, tree*)

Hierarchical softmax layer over binary tree.

In natural language applications, vocabulary size is too large to use softmax loss. Instead, the hierarchical softmax uses product of sigmoid functions. It costs only $O(\log(n))$ time where n is the vocabulary size in average.

At first a user need to prepare a binary tree whose each leaf is corresponding to a word in a vocabulary. When a word x is given, exactly one path from the root of the tree to the leaf of the word exists. Let $\text{path}(x) = ((e_1, b_1), \dots, (e_m, b_m))$ be the path of x , where e_i is an index of i -th internal node, and $b_i \in \{-1, 1\}$ indicates direction to move at i -th internal node (-1 is left, and 1 is right). Then, the probability of x is given as below:

$$\begin{aligned} P(x) &= \prod_{(e_i, b_i) \in \text{path}(x)} P(b_i | e_i) \\ &= \prod_{(e_i, b_i) \in \text{path}(x)} \sigma(b_i x^\top w_{e_i}), \end{aligned}$$

where $\sigma(\cdot)$ is a sigmoid function, and w is a weight matrix.

This function costs $O(\log(n))$ time as an average length of paths is $O(\log(n))$, and $O(n)$ memory as the number of internal nodes equals $n - 1$.

Parameters

- **in_size** (*int*) – Dimension of input vectors.
- **tree** – A binary tree made with tuples like $((1, 2), 3)$.

Variables **w** (*Variable*) – Weight parameter matrix.

See: Hierarchical Probabilistic Neural Network Language Model [Morin+, AISTAT2005].

Methods

__call__ (x, t)

Computes the loss value for given input and ground truth labels.

Parameters

- **x** (*Variable*) – Input to the classifier at each node.
- **t** (*Variable*) – Batch of ground truth labels.

Returns Loss value.

Return type *Variable*

add_param (*name*, *shape=None*, *dtype=<class 'numpy.float32'>*, *initializer=None*)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a *Parameter* object directly to an attribute within *init_scope()* instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.

- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name*, *value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** ([Link](#)) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode*=`'share'`)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type [Link](#)

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters **link** ([Link](#)) – Source link object.

count_params()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

static create_huffman_tree(word_counts)

Makes a Huffman tree from a dictionary containing word counts.

This method creates a binary Huffman tree, that is required for *BinaryHierarchicalSoftmax*. For example, {0: 8, 1: 5, 2: 6, 3: 4} is converted to ((3, 1), (2, 0)).

Parameters *word_counts* (*dict of int key and int or float values*) – Dictionary representing counts of words.

Returns Binary Huffman tree with tuples and keys of *word_counts*.

disable_update()

Disables update rules of all parameters under the link hierarchy.

This method sets the *enabled* flag of the update rule of each parameter variable to *False*.

enable_update()

Enables update rules of all parameters under the link hierarchy.

This method sets the *enabled* flag of the update rule of each parameter variable to *True*.

init_scope()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the *init_scope* method, we can simply assign a *Parameter* object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links(skipself=False)

Returns a generator of all links under the hierarchy.

Parameters *skipself* (*bool*) – If *True*, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks(skipself=False)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters `skipself` (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters `include_uninit` (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters `include_uninit` (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters `name` (*str*) – Name of the attribute to be registered.

repeat (*n_repeat, mode='init'*)

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The `mode` argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the `mode` was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.BlackOut

class `chainer.links.BlackOut` (*in_size*, *counts*, *sample_size*)
 BlackOut loss layer.

See also:

`black_out()` for more detail.

Parameters

- **in_size** (*int*) – Dimension of input vectors.
- **counts** (*int list*) – Number of each identifiers.
- **sample_size** (*int*) – Number of negative samples.

Variables `W` (*Parameter*) – Weight parameter matrix.

Methods

__call__ (*x*, *t*)

Computes the loss value for given input and ground truth labels.

Parameters

- **x** (*Variable*) – Input of the weight matrix multiplication.
- **t** (*Variable*) – Batch of ground truth labels.

Returns Loss value.

Return type *Variable*

add_param (*name*, *shape=None*, *dtype=<class 'numpy.float32'>*, *initializer=None*)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a *Parameter* object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not None, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name*, *value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** ([Link](#)) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode*=*'share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument *mode* below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameter variables under the returned link object is re-initialized by calling their *initialize()* method, so that all the parameters may have different initial values from the original link. *copy* means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. *share* means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is *share*.

Returns Copied link object.

Return type [Link](#)

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters **link** ([Link](#)) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the [Parameters](#) held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for [Chain](#)) by an assignment. A [Parameter](#) object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a [Parameter](#) object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If *name* has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters *name* (*str*) – Name of the attribute to be registered.

repeat (*n_repeat*, *mode*='init')

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The *mode* argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The *net* object contains 16 blocks, each of which is *ConvBNReLU*. And the *mode* was *init*, so each block is re-initialized with different parameters. If you give *copy* to this argument, each block has same values for its parameters but its object ID is different from others. If it is *share*, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. *copy* means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. *share* means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters *serializer* (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

sample_data = None

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.CRF1d

class `chainer.links.CRF1d` (*n_label*)

Linear-chain conditional random field loss layer.

This link wraps the `crfld()` function. It holds a transition cost matrix as a parameter.

Parameters **n_label** (*int*) – Number of labels.

See also:

`crfld()` for more detail.

Variables **cost** (*Variable*) – Transition cost parameter.

Methods

__call__ (*xs*, *ys*, *reduce*='mean')

Call self as a function.

add_param (*name*, *shape*=None, *dtype*=<class 'numpy.float32'>, *initializer*=None)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a *Parameter* object directly to an attribute within *init_scope()* instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not None, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, *dtype* argument is ignored.

add_persistent (*name*, *value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** (*Link*) – Source link object.

argmax (*xs*)

Computes a state that maximizes a joint probability.

Parameters **xs** (*list of Variable*) – Input vector for each label.

Returns A tuple of *Variable* representing each log-likelihood and a list representing the argmax path.

Return type *tuple*

See also:

See `crfld_argmax()` for more detail.

children()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy(mode='share')

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters `mode` (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type *Link*

copyparams(link)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters `link` (*Link*) – Source link object.

count_params()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a *Parameter* object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (skipself=False)

Returns a generator of all links under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (skipself=False)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (include_uninit=True)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (include_uninit=True)

Returns a generator of all parameters under the link hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (name)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If *name* has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters *name* (*str*) – Name of the attribute to be registered.

repeat (n_repeat, mode='init')

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The *mode* argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The *net* object contains 16 blocks, each of which is *ConvBNReLU*. And the *mode* was *init*, so each block is re-initialized with different parameters. If you give *copy* to this argument, each block has same values for its parameters but its object ID is different from others. If it is *share*, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. *copy* means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. *share* means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters *serializer* (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters *device* – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64()

Copies parameter variables and persistent values to CPU.

zerograds()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.SimplifiedDropconnect

class `chainer.links.SimplifiedDropconnect` (*in_size*, *out_size*, *ratio*=0.5, *nobias*=False, *initialW*=None, *initial_bias*=None)

Fully-connected layer with simplified dropconnect regularization.

Notice: This implementation cannot be used for reproduction of the paper. There is a difference between the current implementation and the original one. The original version uses sampling with gaussian distribution before passing activation function, whereas the current implementation averages before activation.

Parameters

- **in_size** (*int*) – Dimension of input vectors. If None, parameter initialization will be deferred until the first forward data pass at which time the size will be determined.
- **out_size** (*int*) – Dimension of output vectors.
- **nobias** (*bool*) – If True, then this link does not use the bias term.
- **initialW** (*initializer*) – Initializer to initialize the weight. When it is `numpy.ndarray`, its `ndim` should be 3.
- **initial_bias** (*initializer*) – Initializer to initialize the bias. If None, the bias will be initialized to zero. When it is `numpy.ndarray`, its `ndim` should be 2.

Variables

- **W** (*Variable*) – Weight parameter.
- **b** (*Variable*) – Bias parameter.

See also:

`simplified_dropconnect()`

See also:

Li, W., Matthew Z., Sixin Z., Yann L., Rob F. (2013). Regularization of Neural Network using DropConnect. International Conference on Machine Learning. [URL](#)

Methods

__call__ (*x*, *train=True*, *mask=None*, *use_batchwise_mask=True*)

Applies the simplified dropconnect layer.

Parameters

- **x** (*chainer.Variable* or *numpy.ndarray* or *cupy.ndarray*) – Batch of input vectors. Its first dimension *n* is assumed to be the *minibatch dimension*.
- **train** (*bool*) – If *True*, executes simplified dropconnect. Otherwise, simplified dropconnect link works as a linear unit.
- **mask** (*None* or *chainer.Variable* or *numpy.ndarray* or *cupy.ndarray*) – If *None*, randomized simplified dropconnect mask is generated. Otherwise, The mask must be (*n*, *M*, *N*) or (*M*, *N*) shaped array, and *use_batchwise_mask* is ignored. Main purpose of this option is debugging. *mask* array will be used as a dropconnect mask.
- **use_batchwise_mask** (*bool*) – If *True*, dropped connections depend on each sample in mini-batch.

Returns Output of the simplified dropconnect layer.

Return type *Variable*

add_param (*name*, *shape=None*, *dtype=<class 'numpy.float32'>*, *initializer=None*)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a *Parameter* object directly to an attribute within *init_scope()* instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int* or *tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not *None*, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, *dtype* argument is ignored.

add_persistent (*name*, *value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** (*Link*) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode='share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument *mode* below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameter variables under the returned link object is re-initialized by calling their *initialize()* method, so that all the parameters may have different initial values from the original link. *copy* means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. *share* means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is *share*.

Returns Copied link object.

Return type *Link*

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters **link** (*Link*) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a *Parameter* object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If *name* has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters *name* (*str*) – Name of the attribute to be registered.

repeat (*n_repeat*, *mode*='init')

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The *mode* argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The *net* object contains 16 blocks, each of which is *ConvBNReLU*. And the *mode* was *init*, so each block is re-initialized with different parameters. If you give *copy* to this argument, each block has same values for its parameters but its object ID is different from others. If it is *share*, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. *copy* means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. *share* means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters *serializer* (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.PReLU

class `chainer.links.PReLU` (*shape=()*, *init=0.25*)

Parametric ReLU function as a link.

Parameters

- **shape** (*tuple of ints*) – Shape of the parameter array.
- **init** (*float*) – Initial parameter value.

See the paper for details: [Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification](#).

See also:

`chainer.functions.prelu()`

Variables **w** (*Parameter*) – Coefficient of parametric ReLU.

Methods

__call__ (*x*)

Applies the parametric ReLU activation function.

Parameters *x* (*Variable*) – Input variable.

Returns Output of the parametric ReLU function.

Return type *Variable*

add_param (*name*, *shape=None*, *dtype=<class 'numpy.float32'>*, *initializer=None*)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a *Parameter* object directly to an attribute within *init_scope()* instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not *None*, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, *dtype* argument is ignored.

add_persistent (*name*, *value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters *link* (*Link*) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (mode='share')

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters `mode (str)` – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type `Link`

copyparams (link)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters `link (Link)` – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the `Parameters` held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope ()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for `Chain`) by an assignment. A `Parameter` object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a `Parameter` object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If *name* has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters *name* (*str*) – Name of the attribute to be registered.

repeat (*n_repeat*, *mode='init'*)

Repeats this link multiple times to make a `Sequential`.

This method returns a `Sequential` object which has the same `Link` multiple times repeatedly. The *mode* argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer `Sequential` block like this:

```

class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')

```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the mode was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes**update_enabled**

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.Swish

class `chainer.links.Swish`(*beta_shape*, *beta_init*=1.0)

Swish activation function as a link.

Parameters

- **beta_shape** (*tuple of ints or None*) – Shape of the parameter variable β . If None, parameter initialization will be deferred until the first forward data pass at which time the shape will be determined.
- **beta_init** (*float*) – Initial value of the parameter variable β .

See the paper for details: [Searching for Activation Functions](#)

To try Swish instead of ReLU, replace `F.ReLU` with individual `Swish` links registered to the model. For example, the model defined in the [MNIST example](#) can be rewritten as follows.

ReLU version (original):

```
class MLP(chainer.Chain):

    def __init__(self, n_units, n_out):
        super(MLP, self).__init__()
        with self.init_scope():
            self.l1 = L.Linear(None, n_units)
            self.l2 = L.Linear(None, n_units)
            self.l3 = L.Linear(None, n_out)

    def __call__(self, x):
        h1 = F.relu(self.l1(x))
        h2 = F.relu(self.l2(h1))
        return self.l3(h2)
```

Swish version:


```

class MLP(chainer.Chain):

    def __init__(self, n_units, n_out):
        super(MLP, self).__init__()
        with self.init_scope():
            self.l1 = L.Linear(None, n_units)
            self.s1 = L.Swish(None)
            self.l2 = L.Linear(None, n_units)
            self.s2 = L.Swish(None)
            self.l3 = L.Linear(None, n_out)

    def __call__(self, x):
        h1 = self.s1(self.l1(x))
        h2 = self.s2(self.l2(h1))
        return self.l3(h2)

```

See also:

See `chainer.functions.swish()` for the definition of Swish activation function.

Variables `beta` (`Parameter`) – Parameter variable β .

Methods

`__call__(x)`

Applies the Swish activation function.

Parameters `x` (`Variable`) – Input variable.

Returns Output of the Swish activation function.

Return type `Variable`

`add_param(name, shape=None, dtype=<class 'numpy.float32'>, initializer=None)`

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a `Parameter` object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```

with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))

```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (`str`) – Name of the parameter. This name is also used as the attribute name.
- **shape** (`int` or `tuple of ints`) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name*, *value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** ([Link](#)) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode*=*'share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument *mode* below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameter variables under the returned link object is re-initialized by calling their *initialize()* method, so that all the parameters may have different initial values from the original link. *copy* means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. *share* means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is *share*.

Returns Copied link object.

Return type [Link](#)

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters **link** ([Link](#)) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the [Parameters](#) held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for [Chain](#)) by an assignment. A [Parameter](#) object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a [Parameter](#) object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If *name* has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters *name* (*str*) – Name of the attribute to be registered.

repeat (*n_repeat*, *mode*='init')

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The *mode* argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The *net* object contains 16 blocks, each of which is *ConvBNReLU*. And the *mode* was *init*, so each block is re-initialized with different parameters. If you give *copy* to this argument, each block has same values for its parameters but its object ID is different from others. If it is *share*, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. *copy* means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. *share* means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters *serializer* (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.Maxout

class `chainer.links.Maxout` (*in_size, out_size, pool_size, initialW=None, initial_bias=0*)

Fully-connected maxout layer.

Let M , P and N be an input dimension, a pool size, and an output dimension, respectively. For an input vector x of size M , it computes

$$Y_i = \max_j (W_{ij} \cdot x + b_{ij}).$$

Here W is a weight tensor of shape (M, P, N) , b an optional bias vector of shape (M, P) and W_{ij} is a sub-vector extracted from W by fixing first and second dimensions to i and j , respectively. Minibatch dimension is omitted in the above equation.

As for the actual implementation, this chain has a Linear link with a $(M * P, N)$ weight matrix and an optional $M * P$ dimensional bias vector.

Parameters

- **in_size** (*int*) – Dimension of input vectors.

- **out_size** (*int*) – Dimension of output vectors.
- **pool_size** (*int*) – Number of channels.
- **initialW** (*initializer*) – Initializer to initialize the weight. When it is `numpy.ndarray`, its `ndim` should be 3.
- **initial_bias** (*initializer*) – Initializer to initialize the bias. If `None`, the bias is omitted. When it is `numpy.ndarray`, its `ndim` should be 2.

Variables **linear** ([Link](#)) – The Linear link that performs affine transformation.

See also:

[maxout\(\)](#)

See also:

Goodfellow, I., Warde-farley, D., Mirza, M., Courville, A., & Bengio, Y. (2013). Maxout Networks. In Proceedings of the 30th International Conference on Machine Learning (ICML-13) (pp. 1319-1327). [URL](#)

Methods

__call__ (*x*)

Applies the maxout layer.

Parameters **x** ([Variable](#)) – Batch of input vectors.

Returns Output of the maxout layer.

Return type [Variable](#)

__getitem__ (*name*)

Equivalent to `getattr`.

add_link (*name, link*)

Registers a child link to this chain.

Deprecated since version v2.0.0: Assign the child link directly to an attribute within `init_scope()` instead. For example, the following code

```
chain.add_link('l1', L.Linear(3, 5))
```

can be replaced by the following line.

```
with chain.init_scope():
    chain.l1 = L.Linear(3, 5)
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the child link. This name is also used as the attribute name.
- **link** ([Link](#)) – The link object to be registered.

add_param (*name, shape=None, dtype=<class 'numpy.float32'>, initializer=None*)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a [Parameter](#) object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name, value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** (*Link*) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode='share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link.

`copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type *Link*

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters **link** (*Link*) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope ()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a *Parameter* object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters **name** (*str*) – Name of the attribute to be registered.

repeat (*n_repeat, mode='init'*)

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The `mode` argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the `mode` was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to

others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.NegativeSampling

class `chainer.links.NegativeSampling` (*in_size*, *counts*, *sample_size*, *power*=0.75)

Negative sampling loss layer.

This link wraps the `negative_sampling()` function. It holds the weight matrix as a parameter. It also builds a sampler internally given a list of word counts.

Parameters

- **in_size** (*int*) – Dimension of input vectors.
- **counts** (*int list*) – Number of each identifiers.
- **sample_size** (*int*) – Number of negative samples.
- **power** (*float*) – Power factor α .

See also:`negative_sampling()` for more detail.**Variables** **W** (*Variable*) – Weight parameter matrix.

Methods

__call__ (*x*, *t*, *reduce*='sum')

Computes the loss value for given input and ground truth labels.

Parameters

- **x** (*Variable*) – Input of the weight matrix multiplication.
- **t** (*Variable*) – Batch of ground truth labels.
- **reduce** (*str*) – Reduction option. Its value must be either 'sum' or 'no'. Otherwise, `ValueError` is raised.

Returns Loss value.**Return type** *Variable***add_param** (*name*, *shape*=None, *dtype*=<class 'numpy.float32'>, *initializer*=None)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a `Parameter` object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name, value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** ([Link](#)) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode='share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type [Link](#)

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters **link** ([Link](#)) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the [Parameters](#) held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope ()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for [Chain](#)) by an assignment. A [Parameter](#) object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a [Parameter](#) object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters **name** (*str*) – Name of the attribute to be registered.

repeat (*n_repeat, mode='init'*)

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The `mode` argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the `mode` was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all

elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

5.3.3 Machine learning models

chainer.links.Classifier

A simple classifier model.

chainer.links.Classifier

```
class chainer.links.Classifier(predictor, lossfun=<function softmax_cross_entropy>, acc-  
                               fun=<function accuracy>, label_key=-1)
```

A simple classifier model.

This is an example of chain that wraps another chain. It computes the loss and accuracy based on a given input/label pair.

Parameters

- **predictor** (*Link*) – Predictor network.
- **lossfun** (*callable*) – Loss function. You can specify one of loss functions from *built-in loss functions*, or your own loss function (see the example below). It should not be an *loss functions with parameters* (i.e., *Link* instance). The function must accept two argument (an output from predictor and its ground truth labels), and return a loss. Returned value must be a Variable derived from the input Variable to perform backpropagation on the variable.
- **accfun** (*callable*) – Function that computes accuracy. You can specify one of evaluation functions from *built-in evaluation functions*, or your own evaluation function. The signature of the function is the same as **lossfun**.
- **label_key** (*int* or *str*) – Key to specify label variable from arguments. When it is *int*, a variable in positional arguments is used. And when it is *str*, a variable in keyword arguments is used.

Variables

- **predictor** (*Link*) – Predictor network.
- **lossfun** (*callable*) – Loss function. See the description in the arguments for details.
- **accfun** (*callable*) – Function that computes accuracy. See the description in the arguments for details.
- **y** (*Variable*) – Prediction for the last minibatch.
- **loss** (*Variable*) – Loss value for the last minibatch.
- **accuracy** (*Variable*) – Accuracy for the last minibatch.
- **compute_accuracy** (*bool*) – If True, compute accuracy on the forward computation. The default value is True.

Note: This link uses `chainer.softmax_cross_entropy()` with default arguments as a loss function (specified by **lossfun**), if users do not explicitly change it. In particular, the loss function does not support double backpropagation. If you need second or higher order differentiation, you need to turn it on with `enable_double_backprop=True`:

```
>>> import chainer.functions as F
>>> import chainer.links as L
>>>
>>> def lossfun(x, t):
...     return F.softmax_cross_entropy(
...         x, t, enable_double_backprop=True)
>>>
>>> predictor = L.Linear(10)
>>> model = L.Classifier(predictor, lossfun=lossfun)
```


Methods

__call__ (*args, **kwargs)

Computes the loss value for an input and label pair.

It also computes accuracy and stores it to the attribute.

Parameters

- **args** (list of `~chainer.Variable`) – Input minibatch.
- **kwargs** (dict of `~chainer.Variable`) – Input minibatch.

When `label_key` is `int`, the corresponding element in `args` is treated as ground truth labels. And when it is `str`, the element in `kwargs` is used. The all elements of `args` and `kwargs` except the ground truth labels are features. It feeds features to the predictor and compare the result with ground truth labels.

Returns Loss value.

Return type `Variable`

__getitem__ (name)

Equivalent to `getattr`.

add_link (name, link)

Registers a child link to this chain.

Deprecated since version v2.0.0: Assign the child link directly to an attribute within `init_scope()` instead. For example, the following code

```
chain.add_link('l1', L.Linear(3, 5))
```

can be replaced by the following line.

```
with chain.init_scope():
    chain.l1 = L.Linear(3, 5)
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (`str`) – Name of the child link. This name is also used as the attribute name.
- **link** (`Link`) – The link object to be registered.

add_param (name, shape=None, dtype=<class 'numpy.float32'>, initializer=None)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a `Parameter` object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (`str`) – Name of the parameter. This name is also used as the attribute name.

- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name, value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** (*Link*) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode='share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type *Link*

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters `link` (`Link`) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the `Parameters` held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope ()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for `Chain`) by an assignment. A `Parameter` object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a `Parameter` object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (`skipself=False`)

Returns a generator of all links under the hierarchy.

Parameters `skipself` (`bool`) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (`skipself=False`)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters `skipself` (`bool`) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (`include_uninit=True`)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters `include_uninit` (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters `include_uninit` (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters `name` (*str*) – Name of the attribute to be registered.

repeat (*n_repeat, mode='init'*)

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The `mode` argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The `net` object contains 16 blocks, each of which is *ConvBNReLU*. And the `mode` was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial

parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

compute_accuracy = True

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

5.3.4 Pre-trained models

Pre-trained models are mainly used to achieve a good performance with a small dataset, or extract a semantic feature vector. Although `CaffeFunction` automatically loads a pre-trained model released as a `caffemodel`, the following link models provide an interface for automatically converting `caffemodels`, and easily extracting semantic feature vectors.

For example, to extract the feature vectors with `VGG16Layers`, which is a common pre-trained model in the field of image recognition, users need to write the following few lines:

```
from chainer.links import VGG16Layers
from PIL import Image

model = VGG16Layers()
img = Image.open("path/to/image.jpg")
feature = model.extract([img], layers=["fc7"])["fc7"]
```

where `fc7` denotes a layer before the last fully-connected layer. Unlike the usual links, these classes automatically load all the parameters from the pre-trained models during initialization.

VGG16Layers

<code>chainer.links.VGG16Layers</code>	A pre-trained CNN model with 16 layers provided by VGG team.
<code>chainer.links.model.vision.vgg.prepare</code>	Converts the given image to the numpy array for VGG models.

chainer.links.VGG16Layers

class `chainer.links.VGG16Layers` (*pretrained_model*='auto')

A pre-trained CNN model with 16 layers provided by VGG team.

During initialization, this chain model automatically downloads the pre-trained caffemodel, convert to another chainer model, stores it on your local directory, and initializes all the parameters with it. This model would be useful when you want to extract a semantic feature vector from a given image, or fine-tune the model on a different dataset. Note that this pre-trained model is released under Creative Commons Attribution License.

If you want to manually convert the pre-trained caffemodel to a chainer model that can be specified in the constructor, please use `convert_caffemodel_to_npz` classmethod instead.

See: K. Simonyan and A. Zisserman, [Very Deep Convolutional Networks for Large-Scale Image Recognition](#)

Parameters `pretrained_model` (*str*) – the destination of the pre-trained chainer model serialized as a `.npz` file. If this argument is specified as `auto`, it automatically downloads the caffemodel from the internet. Note that in this case the converted chainer model is stored on `$CHAINER_DATASET_ROOT/pfnet/chainer/models` directory, where `$CHAINER_DATASET_ROOT` is set as `$HOME/.chainer/dataset` unless you specify another value as an environment variable. The converted chainer model is automatically used from the second time. If the argument is specified as `None`, all the parameters are not initialized by the pre-trained model, but the default initializer used in the original paper, i.e., `chainer.initializers.Normal(scale=0.01)`.

Variables `available_layers` (*list of str*) – The list of available layer names used by `__call__` and `extract` methods.

Methods

`__call__` (*self, x, layers=['prob']*)
Computes all the feature maps specified by `layers`.

Warning: `test` argument is not supported anymore since v2. Instead, use `chainer.using_config('train', train)`. See [chainer.using_config\(\)](#).

Parameters

- **x** (*Variable*) – Input variable. It should be prepared by `prepare` function.
- **layers** (*list of str*) – The list of layer names you want to extract.

Returns A directory in which the key contains the layer name and the value contains the corresponding feature map variable.

Return type Dictionary of `~chainer.Variable`

__getitem__ (*name*)
Equivalent to `getattr`.

add_link (*name, link*)
Registers a child link to this chain.

Deprecated since version v2.0.0: Assign the child link directly to an attribute within `init_scope()` instead. For example, the following code

```
chain.add_link('l1', L.Linear(3, 5))
```

can be replaced by the following line.

```
with chain.init_scope():
    chain.l1 = L.Linear(3, 5)
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the child link. This name is also used as the attribute name.
- **link** (*Link*) – The link object to be registered.

add_param (*name, shape=None, dtype=<class 'numpy.float32'>, initializer=None*)
Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a `Parameter` object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.

- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name*, *value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** ([Link](#)) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

classmethod convert_caffemodel_to_npz (*path_caffemodel*, *path_npz*)

Converts a pre-trained caffemodel to a chainer model.

Parameters

- **path_caffemodel** (*str*) – Path of the pre-trained caffemodel.
- **path_npz** (*str*) – Path of the converted chainer model.

copy (*mode*=`'share'`)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type [Link](#)

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters **link** ([Link](#)) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the [Parameters](#) held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the enabled flag of the update rule of each parameter variable to `False`.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the enabled flag of the update rule of each parameter variable to `True`.

extract (*self, images, layers=['fc7'], size=(224, 224)*)

Extracts all the feature maps of given images.

The difference of directly executing `__call__` is that it directly accepts images as an input and automatically transforms them to a proper variable. That is, it is also interpreted as a shortcut method that implicitly calls `prepare` and `__call__` functions.

Unlike `predict` method, this method does not override `chainer.config.train` and `chainer.config.enable_backprop` configuration. If you want to extract features without updating model parameters, you need to manually set configuration when calling this method as follows:

```
# model is an instance of VGG16Layers
with chainer.using_config('train', False):
    with chainer.using_config('enable_backprop', False):
        feature = model.extract([image])
```

Warning: `test` and `volatile` arguments are not supported anymore since v2. Instead, users should configure training and volatile modes with `train` and `enable_backprop`, respectively.

Note that default behavior of this method is different between v1 and later versions. Specifically, the default values of `test` in v1 were `True` (test mode). But that of `chainer.config.train` is also `True` (train mode). Therefore, users need to explicitly switch `train` to `False` to run the code in test mode and `enable_backprop` to `False` to turn off computational graph construction.

See the [upgrade guide](#).

Parameters

- **images** (*iterable of PIL.Image or numpy.ndarray*) – Input images.
- **layers** (*list of str*) – The list of layer names you want to extract.

- **size** (*pair of ints*) – The resolution of resized images used as an input of CNN. All the given images are not resized if this argument is `None`, but the resolutions of all the images should be the same.

Returns A directory in which the key contains the layer name and the value contains the corresponding feature map variable.

Return type Dictionary of ~chainer.Variable

init_scope()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for [Chain](#)) by an assignment. A [Parameter](#) object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a [Parameter](#) object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

predict (*images, oversample=True*)

Computes all the probabilities of given images.

Parameters

- **images** (*iterable of PIL.Image or numpy.ndarray*) – Input images. When you specify a color image as a `numpy.ndarray`, make sure that color order is RGB.
- **oversample** (*bool*) – If `True`, it averages results across center, corners, and mirrors. Otherwise, it uses only the center.

Returns Output that contains the class probabilities of given images.

Return type *Variable*

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters **name** (*str*) – Name of the attribute to be registered.

repeat (*n_repeat, mode='init'*)

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The `mode` argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the `mode` was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial

parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

available_layers

functions

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.model.vision.vgg.prepare

`chainer.links.model.vision.vgg.prepare` (*image*, *size=(224, 224)*)

Converts the given image to the numpy array for VGG models.

Note that you have to call this method before `__call__` because the pre-trained vgg model requires to resize the given image, convert the RGB to the BGR, subtract the mean, and permute the dimensions before calling.

Parameters

- **image** (*PIL.Image* or *numpy.ndarray*) – Input image. If an input is *numpy.ndarray*, its shape must be (height, width), (height, width, channels), or (channels, height, width), and the order of the channels must be RGB.
- **size** (*pair of ints*) – Size of converted images. If *None*, the given image is not resized.

Returns The converted output array.

Return type *numpy.ndarray*

GoogLeNet

<code>chainer.links.GoogLeNet</code>	A pre-trained GoogLeNet model provided by BVLC.
<code>chainer.links.model.vision.googlenet.prepare</code>	Converts the given image to the numpy array for GoogLeNet.

chainer.links.GoogLeNet

class `chainer.links.GoogLeNet` (*pretrained_model='auto'*)

A pre-trained GoogLeNet model provided by BVLC.

When you specify the path of the pre-trained chainer model serialized as a *.npz* file in the constructor, this chain model automatically initializes all the parameters with it. This model would be useful when you want to extract a semantic feature vector per image, or fine-tune the model on a different dataset.

If you want to manually convert the pre-trained caffemodel to a chainer model that can be specified in the constructor, please use `convert_caffemodel_to_npz` classmethod instead.

GoogLeNet, which is also called Inception-v1, is an architecture of convolutional neural network proposed in 2014. This model is relatively lightweight and requires small memory footprint during training compared with modern architectures such as ResNet. Therefore, if you fine-tune your network based on a model pre-trained by Imagenet and need to train it with large batch size, GoogLeNet may be useful. On the other hand, if you just want an off-the-shelf classifier, we recommend you to use ResNet50 or other models since they are more accurate than GoogLeNet.

The original model is provided here: https://github.com/BVLC/caffe/tree/master/models/bvlc_googlenet

Parameters `pretrained_model` (*str*) – the destination of the pre-trained chainer model serialized as a *.npz* file. If this argument is specified as *auto*, it automatically downloads the caffemodel from the internet. Note that in this case the converted chainer model is stored on `$CHAINER_DATASET_ROOT/pfnet/chainer/models` directory, where `$CHAINER_DATASET_ROOT` is set as `$HOME/.chainer/dataset` unless you specify another value as a environment variable. The converted chainer model is automatically used from the second time. If the argument is specified as *None*, all the parameters are not initialized by the pre-trained model, but the default initializer used in BVLC, i.e., `chainer.initializers.LeCunUniform(scale=1.0)`. Note that, in Caffe, when `weight_filler` is specified as “xavier” type without `variance_norm` parameter, the weights are initialized by `Uniform(-s, s)`, where $s = \sqrt{\frac{3}{fan_{in}}}$ and fan_{in} is the number of input units. This corresponds to `LeCunUniform` in Chainer but not `GlorotUniform`.

Variables **`available_layers`** (*list of str*) – The list of available layer names used by `__call__` and `extract` methods.

Methods

`__call__(self, x, layers=['prob'])`
Computes all the feature maps specified by `layers`.

Warning: `train` argument is not supported anymore since v2. Instead, use `chainer.using_config('train', train)`. See [chainer.using_config\(\)](#).

Parameters

- **`x`** (*Variable*) – Input variable. It should be prepared by `prepare` function.
- **`layers`** (*list of str*) – The list of layer names you want to extract.

Returns A directory in which the key contains the layer name and the value contains the corresponding feature map variable.

Return type Dictionary of `~chainer.Variable`

`__getitem__(name)`
Equivalent to `getattr`.

`add_link` (*name, link*)
Registers a child link to this chain.

Deprecated since version v2.0.0: Assign the child link directly to an attribute within `init_scope()` instead. For example, the following code

```
chain.add_link('l1', L.Linear(3, 5))
```

can be replaced by the following line.

```
with chain.init_scope():
    chain.l1 = L.Linear(3, 5)
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **`name`** (*str*) – Name of the child link. This name is also used as the attribute name.
- **`link`** (*Link*) – The link object to be registered.

`add_param` (*name, shape=None, dtype=<class 'numpy.float32'>, initializer=None*)
Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a `Parameter` object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name, value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** (*Link*) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

classmethod convert_caffemodel_to_npz (*path_caffemodel, path_npz*)

Converts a pre-trained caffemodel to a chainer model.

Parameters

- **path_caffemodel** (*str*) – Path of the pre-trained caffemodel.
- **path_npz** (*str*) – Path of the converted chainer model.

copy (*mode='share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link.

`copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type *Link*

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters **link** (*Link*) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the enabled flag of the update rule of each parameter variable to `False`.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the enabled flag of the update rule of each parameter variable to `True`.

extract (*self, images, layers=['pool5'], size=(224, 224)*)

Extracts all the feature maps of given images.

The difference of directly executing `__call__` is that it directly accepts images as an input and automatically transforms them to a proper variable. That is, it is also interpreted as a shortcut method that implicitly calls `prepare` and `__call__` functions.

Unlike `predict` method, this method does not override `chainer.config.train` and `chainer.config.enable_backprop` configuration. If you want to extract features without updating model parameters, you need to manually set configuration when calling this method as follows:

```
# model is an instance of `GoogLeNet`
with chainer.config('train', False):
    with chainer.config('enable_backprop', False):
        feature = model.extract([image])
```

Warning: `train` and `volatile` arguments are not supported anymore since v2. Instead, users should configure training and volatile modes with `train` and `enable_backprop`, respectively.

Note that default behavior of this method is different between v1 and later versions. Specifically, the default values of `train` arguments in v1 were `False` and `OFF`, while that of `chainer.config.train` are `True`. Therefore, users need to explicitly switch `train` to `False` to run the code in test mode to turn off computational graph construction.

See the [upgrade guide](#).

Parameters

- **images** (*iterable of PIL.Image or numpy.ndarray*) – Input images.
- **layers** (*list of str*) – The list of layer names you want to extract.
- **size** (*pair of ints*) – The resolution of resized images used as an input of CNN. All the given images are not resized if this argument is `None`, but the resolutions of all the images should be the same.

Returns A directory in which the key contains the layer name and the value contains the corresponding feature map variable.

Return type Dictionary of ~chainer.Variable

`init_scope()`

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for [Chain](#)) by an assignment. A [Parameter](#) object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a [Parameter](#) object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

`links(skipself=False)`

Returns a generator of all links under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

`namedlinks(skipself=False)`

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

`namedparams(include_uninit=True)`

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

predict (*images, oversample=True*)

Computes all the probabilities of given images.

Parameters

- **images** (*iterable of PIL.Image or numpy.ndarray*) – Input images. When you specify a color image as a `numpy.ndarray`, make sure that color order is RGB.
- **oversample** (*bool*) – If `True`, it averages results across center, corners, and mirrors. Otherwise, it uses only the center.

Returns Output that contains the class probabilities of given images.

Return type *Variable*

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters **name** (*str*) – Name of the attribute to be registered.

repeat (*n_repeat, mode='init'*)

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The `mode` argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the `mode` was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.**to_cpu** ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes**available_layers****functions****update_enabled**

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

`chainer.links.model.vision.googlenet.prepare`

`chainer.links.model.vision.googlenet.prepare` (*image*, *size*=(224, 224))

Converts the given image to the numpy array for GoogLeNet.

Note that you have to call this method before `__call__` because the pre-trained GoogLeNet model requires to resize the given image, convert the RGB to the BGR, subtract the mean, and permute the dimensions before calling.

Parameters

- **image** (*PIL.Image* or *numpy.ndarray*) – Input image. If an input is *numpy.ndarray*, its shape must be (height, width), (height, width, channels), or (channels, height, width), and the order of the channels must be RGB.
- **size** (*pair of ints*) – Size of converted images. If None, the given image is not resized.

Returns The converted output array.

Return type `numpy.ndarray`

Residual Networks

<code>chainer.links.model.vision.resnet.ResNetLayers</code>	A pre-trained CNN model provided by MSRA.
<code>chainer.links.ResNet50Layers</code>	A pre-trained CNN model with 50 layers provided by MSRA.
<code>chainer.links.ResNet101Layers</code>	A pre-trained CNN model with 101 layers provided by MSRA.
<code>chainer.links.ResNet152Layers</code>	A pre-trained CNN model with 152 layers provided by MSRA.
<code>chainer.links.model.vision.resnet.prepare</code>	Converts the given image to the numpy array for ResNets.

`chainer.links.model.vision.resnet.ResNetLayers`

class `chainer.links.model.vision.resnet.ResNetLayers` (*pretrained_model*, *n_layers*)

A pre-trained CNN model provided by MSRA.

When you specify the path of the pre-trained chainer model serialized as a `.npz` file in the constructor, this chain model automatically initializes all the parameters with it. This model would be useful when you want to extract a semantic feature vector per image, or fine-tune the model on a different dataset. Note that unlike `VGG16Layers`, it does not automatically download a pre-trained caffemodel. This caffemodel can be downloaded at [GitHub](#).

If you want to manually convert the pre-trained caffemodel to a chainer model that can be specified in the constructor, please use `convert_caffemodel_to_npz` classmethod instead.

See: K. He et. al., [Deep Residual Learning for Image Recognition](#)

Parameters

- **pretrained_model** (*str*) – the destination of the pre-trained chainer model serialized as a .npz file. If this argument is specified as `auto`, it automatically loads and converts the caffemodel from `$CHAINER_DATASET_ROOT/pfnet/chainer/models/ResNet-{n-layers}-model.caffemodel`, where `$CHAINER_DATASET_ROOT` is set as `$HOME/.chainer/dataset` unless you specify another value by modifying the environment variable and `{n-layers}` is replaced with the specified number of layers given as the first argument to this constructor. Note that in this case the converted chainer model is stored on the same directory and automatically used from the next time. If this argument is specified as `None`, all the parameters are not initialized by the pre-trained model, but the default initializer used in the original paper, i.e., `chainer.initializers.HeNormal(scale=1.0)`.
- **n_layers** (*int*) – The number of layers of this model. It should be either 50, 101, or 152.

Variables *available_layers* (*list of str*) – The list of available layer names used by `__call__` and `extract` methods.

Methods

`__call__(self, x, layers=['prob'])`
Computes all the feature maps specified by `layers`.

Warning: `test` argument is not supported anymore since v2. Instead, use `chainer.using_config('train', train)`. See `chainer.using_config()`.

Parameters

- **x** (*Variable*) – Input variable. It should be prepared by `prepare` function.
- **layers** (*list of str*) – The list of layer names you want to extract.

Returns A directory in which the key contains the layer name and the value contains the corresponding feature map variable.

Return type Dictionary of ~chainer.Variable

`__getitem__(name)`
Equivalent to `getattr`.

`add_link(name, link)`
Registers a child link to this chain.

Deprecated since version v2.0.0: Assign the child link directly to an attribute within `init_scope()` instead. For example, the following code

```
chain.add_link('l1', L.Linear(3, 5))
```

can be replaced by the following line.

```
with chain.init_scope():
    chain.l1 = L.Linear(3, 5)
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the child link. This name is also used as the attribute name.
- **link** (*Link*) – The link object to be registered.

add_param (*name, shape=None, dtype=<class 'numpy.float32'>, initializer=None*)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a *Parameter* object directly to an attribute within *init_scope()* instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not *None*, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, *dtype* argument is ignored.

add_persistent (*name, value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters *link* (*Link*) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

classmethod convert_caffemodel_to_npz (*path_caffemodel, path_npz, n_layers=50*)

Converts a pre-trained caffemodel to a chainer model.

Parameters

- **path_caffemodel** (*str*) – Path of the pre-trained caffemodel.
- **path_npz** (*str*) – Path of the converted chainer model.

copy (*mode*='share')

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument *mode* below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters *mode* (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default *mode* is `share`.

Returns Copied link object.

Return type *Link*

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters *link* (*Link*) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

extract (*self*, *images*, *layers*=['pool5'], *size*=(224, 224))

Extracts all the feature maps of given images.

The difference of directly executing `__call__` is that it directly accepts images as an input and automatically transforms them to a proper variable. That is, it is also interpreted as a shortcut method that implicitly calls `prepare` and `__call__` functions.

Unlike `predict` method, this method does not override `chainer.config.train` and `chainer.config.enable_backprop` configuration. If you want to extract features without updating model parameters, you need to manually set configuration when calling this method as follows:

```
# model is an instance of ResNetLayers (50 or 101 or 152 layers)
with chainer.using_config('train', False):
    with chainer.using_config('enable_backprop', False):
        feature = model.extract([image])
```

Warning: `test` and `volatile` arguments are not supported anymore since v2. Instead, users should configure training and volatile modes with `train` and `enable_backprop`, respectively.

Note that default behavior of this method is different between v1 and later versions. Specifically, the default values of `test` in v1 were `True` (test mode). But that of `chainer.config.train` is also `True` (train mode). Therefore, users need to explicitly switch `train` to `False` to run the code in test mode and `enable_backprop` to `False` to turn off computational graph construction.

See the [upgrade guide](#).

Parameters

- **images** (*iterable of PIL.Image or numpy.ndarray*) – Input images.
- **layers** (*list of str*) – The list of layer names you want to extract.
- **size** (*pair of ints*) – The resolution of resized images used as an input of CNN. All the given images are not resized if this argument is `None`, but the resolutions of all the images should be the same.

Returns A directory in which the key contains the layer name and the value contains the corresponding feature map variable.

Return type Dictionary of ~chainer.Variable

`init_scope()`

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a *Parameter* object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

`links (skipself=False)`

Returns a generator of all links under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

predict (*images, oversample=True*)

Computes all the probabilities of given images.

Parameters

- **images** (*iterable of PIL.Image or numpy.ndarray*) – Input images. When you specify a color image as a `numpy.ndarray`, make sure that color order is RGB.
- **oversample** (*bool*) – If `True`, it averages results across center, corners, and mirrors. Otherwise, it uses only the center.

Returns Output that contains the class probabilities of given images.

Return type *Variable*

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters **name** (*str*) – Name of the attribute to be registered.

repeat (*n_repeat, mode='init'*)

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The `mode` argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
```

(continues on next page)

(continued from previous page)

```

        self.conv = L.Convolution2D(
            None, 64, 3, 1, 1, nobias=True)
        self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')

```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the `mode` was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: `self`

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: `self`

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

available_layers

functions

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.ResNet50Layers

class `chainer.links.ResNet50Layers` (*pretrained_model*='auto')

A pre-trained CNN model with 50 layers provided by MSRA.

When you specify the path of the pre-trained chainer model serialized as a `.npz` file in the constructor, this chain model automatically initializes all the parameters with it. This model would be useful when you want to extract a semantic feature vector per image, or fine-tune the model on a different dataset. Note that unlike `VGG16Layers`, it does not automatically download a pre-trained caffemodel. This caffemodel can be downloaded at [GitHub](#).

If you want to manually convert the pre-trained caffemodel to a chainer model that can be specified in the constructor, please use `convert_caffemodel_to_npz` classmethod instead.

ResNet50 has 25,557,096 trainable parameters, and it's 58% and 43% fewer than ResNet101 and ResNet152, respectively. On the other hand, the top-5 classification accuracy on ImageNet dataset drops only 0.7% and 1.1% from ResNet101 and ResNet152, respectively. Therefore, ResNet50 may have the best balance between the accuracy and the model size. It would be basically just enough for many cases, but some advanced models for object detection or semantic segmentation use deeper ones as their building blocks, so these deeper ResNets are here for making reproduction work easier.

See: K. He et. al., [Deep Residual Learning for Image Recognition](#)

Parameters `pretrained_model` (*str*) – the destination of the pre-trained chainer model serialized as a `.npz` file. If this argument is specified as `auto`, it automatically loads and converts the caffemodel from `$CHAINER_DATASET_ROOT/pfnet/chainer/models/ResNet-50-model.caffemodel`, where `$CHAINER_DATASET_ROOT` is set as `$HOME/.chainer/dataset` unless you specify another value by modifying the environment variable. Note that in this case the converted chainer model is stored on the same directory and automatically used from the next time. If this argument is specified as `None`, all the parameters are not initialized by the pre-trained model, but the default initializer used in the original paper, i.e., `chainer.initializers.HeNormal(scale=1.0)`.

Variables `available_layers` (*list of str*) – The list of available layer names used by `__call__` and extract methods.

Methods

__call__ (*self*, *x*, *layers*=['prob'])
Computes all the feature maps specified by *layers*.

Warning: `test` argument is not supported anymore since v2. Instead, use `chainer.using_config('train', train)`. See `chainer.using_config()`.

Parameters

- **x** (*Variable*) – Input variable. It should be prepared by `prepare` function.
- **layers** (*list of str*) – The list of layer names you want to extract.

Returns A directory in which the key contains the layer name and the value contains the corresponding feature map variable.

Return type Dictionary of ~chainer.Variable

__getitem__ (*name*)
Equivalent to `getattr`.

add_link (*name*, *link*)
Registers a child link to this chain.

Deprecated since version v2.0.0: Assign the child link directly to an attribute within `init_scope()` instead. For example, the following code

```
chain.add_link('l1', L.Linear(3, 5))
```

can be replaced by the following line.

```
with chain.init_scope():  
    chain.l1 = L.Linear(3, 5)
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the child link. This name is also used as the attribute name.
- **link** (*Link*) – The link object to be registered.

add_param (*name*, *shape*=None, *dtype*=<class 'numpy.float32'>, *initializer*=None)
Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a `Parameter` object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():  
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name, value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** (*Link*) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

classmethod convert_caffemodel_to_npz (*path_caffemodel, path_npz, n_layers=50*)

Converts a pre-trained caffemodel to a chainer model.

Parameters

- **path_caffemodel** (*str*) – Path of the pre-trained caffemodel.
- **path_npz** (*str*) – Path of the converted chainer model.

copy (*mode='share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays

are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type *Link*

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters **link** (*Link*) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the enabled flag of the update rule of each parameter variable to `False`.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the enabled flag of the update rule of each parameter variable to `True`.

extract (*self*, *images*, *layers*=['pool5'], *size*=(224, 224))

Extracts all the feature maps of given images.

The difference of directly executing `__call__` is that it directly accepts images as an input and automatically transforms them to a proper variable. That is, it is also interpreted as a shortcut method that implicitly calls `prepare` and `__call__` functions.

Unlike `predict` method, this method does not override `chainer.config.train` and `chainer.config.enable_backprop` configuration. If you want to extract features without updating model parameters, you need to manually set configuration when calling this method as follows:

```
# model is an instance of ResNetLayers (50 or 101 or 152 layers)
with chainer.using_config('train', False):
    with chainer.using_config('enable_backprop', False):
        feature = model.extract([image])
```

Warning: `test` and `volatile` arguments are not supported anymore since v2. Instead, users should configure training and volatile modes with `train` and `enable_backprop`, respectively.

Note that default behavior of this method is different between v1 and later versions. Specifically, the default values of `test` in v1 were `True` (test mode). But that of `chainer.config.train` is also `True` (train mode). Therefore, users need to explicitly switch `train` to `False` to run the code in test mode and `enable_backprop` to `False` to turn off computational graph construction.

See the [upgrade guide](#).

Parameters

- **images** (*iterable of PIL.Image or numpy.ndarray*) – Input images.
- **layers** (*list of str*) – The list of layer names you want to extract.
- **size** (*pair of ints*) – The resolution of resized images used as an input of CNN. All the given images are not resized if this argument is `None`, but the resolutions of all the images should be the same.

Returns A directory in which the key contains the layer name and the value contains the corresponding feature map variable.

Return type Dictionary of ~chainer.Variable

init_scope()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a *Parameter* object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

predict (*images*, *oversample=True*)

Computes all the probabilities of given images.

Parameters

- **images** (*iterable of PIL.Image or numpy.ndarray*) – Input images. When you specify a color image as a `numpy.ndarray`, make sure that color order is RGB.
- **oversample** (*bool*) – If `True`, it averages results across center, corners, and mirrors. Otherwise, it uses only the center.

Returns Output that contains the class probabilities of given images.

Return type *Variable*

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters **name** (*str*) – Name of the attribute to be registered.

repeat (*n_repeat*, *mode='init'*)

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The `mode` argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the `mode` was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.

- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: `self`

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: `self`

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

available_layers

functions

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.ResNet101Layers

class `chainer.links.ResNet101Layers` (*pretrained_model*='auto')

A pre-trained CNN model with 101 layers provided by MSRA.

When you specify the path of the pre-trained chainer model serialized as a `.npz` file in the constructor, this chain model automatically initializes all the parameters with it. This model would be useful when you want to extract a semantic feature vector per image, or fine-tune the model on a different dataset. Note that unlike `VGG16Layers`, it does not automatically download a pre-trained caffemodel. This caffemodel can be downloaded at [GitHub](#).

If you want to manually convert the pre-trained caffemodel to a chainer model that can be specified in the constructor, please use `convert_caffemodel_to_npz` classmethod instead.

ResNet101 has 44,549,224 trainable parameters, and it's 43% fewer than ResNet152 model, while the top-5 classification accuracy on ImageNet dataset drops 1.1% from ResNet152. For many cases, ResNet50 may have the best balance between the accuracy and the model size.

See: K. He et. al., [Deep Residual Learning for Image Recognition](#)

Parameters `pretrained_model` (*str*) – the destination of the pre-trained chainer model serialized as a `.npz` file. If this argument is specified as `auto`, it automatically loads and converts the caffemodel from `$CHAINER_DATASET_ROOT/pfnet/chainer/models/ResNet-101-model.caffemodel`, where `$CHAINER_DATASET_ROOT` is set as `$HOME/.chainer/dataset` unless you specify another value by modifying the environment variable. Note that in this case the converted chainer model is stored on the same directory and automatically used from the next time. If this argument is specified as `None`, all the parameters are not initialized by the pre-trained model, but the default initializer used in the original paper, i.e., `chainer.initializers.HeNormal(scale=1.0)`.

Variables `available_layers` (*list of str*) – The list of available layer names used by `__call__` and extract methods.

Methods

`__call__` (*self, x, layers=['prob']*)

Computes all the feature maps specified by `layers`.

Warning: `test` argument is not supported anymore since v2. Instead, use `chainer.using_config('train', train)`. See [chainer.using_config\(\)](#).

Parameters

- **x** (*Variable*) – Input variable. It should be prepared by `prepare` function.
- **layers** (*list of str*) – The list of layer names you want to extract.

Returns A directory in which the key contains the layer name and the value contains the corresponding feature map variable.

Return type Dictionary of `~chainer.Variable`

`__getitem__` (*name*)

Equivalent to `getattr`.

add_link (*name*, *link*)

Registers a child link to this chain.

Deprecated since version v2.0.0: Assign the child link directly to an attribute within `init_scope()` instead. For example, the following code

```
chain.add_link('l1', L.Linear(3, 5))
```

can be replaced by the following line.

```
with chain.init_scope():
    chain.l1 = L.Linear(3, 5)
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the child link. This name is also used as the attribute name.
- **link** (*Link*) – The link object to be registered.

add_param (*name*, *shape*=None, *dtype*=<class 'numpy.float32'>, *initializer*=None)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a `Parameter` object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not None, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name*, *value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters `link` ([Link](#)) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

classmethod `convert_caffemodel_to_npz` (`path_caffemodel`, `path_npz`, `n_layers=50`)

Converts a pre-trained caffemodel to a chainer model.

Parameters

- `path_caffemodel` (`str`) – Path of the pre-trained caffemodel.
- `path_npz` (`str`) – Path of the converted chainer model.

copy (`mode='share'`)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters `mode` (`str`) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type [Link](#)

copyparams (`link`)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters `link` ([Link](#)) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the [Parameters](#) held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

extract(self, images, layers=['pool5'], size=(224, 224))

Extracts all the feature maps of given images.

The difference of directly executing `__call__` is that it directly accepts images as an input and automatically transforms them to a proper variable. That is, it is also interpreted as a shortcut method that implicitly calls `prepare` and `__call__` functions.

Unlike `predict` method, this method does not override `chainer.config.train` and `chainer.config.enable_backprop` configuration. If you want to extract features without updating model parameters, you need to manually set configuration when calling this method as follows:

```
# model is an instance of ResNetLayers (50 or 101 or 152 layers)
with chainer.using_config('train', False):
    with chainer.using_config('enable_backprop', False):
        feature = model.extract([image])
```

Warning: `test` and `volatile` arguments are not supported anymore since v2. Instead, users should configure training and volatile modes with `train` and `enable_backprop`, respectively.

Note that default behavior of this method is different between v1 and later versions. Specifically, the default values of `test` in v1 were `True` (test mode). But that of `chainer.config.train` is also `True` (train mode). Therefore, users need to explicitly switch `train` to `False` to run the code in test mode and `enable_backprop` to `False` to turn off computational graph construction.

See the [upgrade guide](#).

Parameters

- **images** (*iterable of PIL.Image or numpy.ndarray*) – Input images.
- **layers** (*list of str*) – The list of layer names you want to extract.
- **size** (*pair of ints*) – The resolution of resized images used as an input of CNN. All the given images are not resized if this argument is `None`, but the resolutions of all the images should be the same.

Returns A dictionary in which the key contains the layer name and the value contains the corresponding feature map variable.

Return type Dictionary of ~chainer.Variable

init_scope()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a `Parameter` object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

predict (*images, oversample=True*)

Computes all the probabilities of given images.

Parameters

- **images** (*iterable of PIL.Image or numpy.ndarray*) – Input images. When you specify a color image as a `numpy.ndarray`, make sure that color order is RGB.
- **oversample** (*bool*) – If `True`, it averages results across center, corners, and mirrors. Otherwise, it uses only the center.

Returns Output that contains the class probabilities of given images.

Return type *Variable*

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters **name** (*str*) – Name of the attribute to be registered.

repeat (*n_repeat*, *mode*='init')

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The *mode* argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The *net* object contains 16 blocks, each of which is *ConvBNReLU*. And the *mode* was *init*, so each block is re-initialized with different parameters. If you give *copy* to this argument, each block has same values for its parameters but its object ID is different from others. If it is *share*, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. *copy* means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. *share* means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters `device` – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64()

Copies parameter variables and persistent values to CPU.

zerograds()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

available_layers

functions

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.ResNet152Layers

class `chainer.links.ResNet152Layers` (*pretrained_model*='auto')

A pre-trained CNN model with 152 layers provided by MSRA.

When you specify the path of the pre-trained chainer model serialized as a `.npz` file in the constructor, this chain model automatically initializes all the parameters with it. This model would be useful when you want to extract a semantic feature vector per image, or fine-tune the model on a different dataset. Note that unlike `VGG16Layers`, it does not automatically download a pre-trained caffemodel. This caffemodel can be downloaded at [GitHub](#).

If you want to manually convert the pre-trained caffemodel to a chainer model that can be specified in the constructor, please use `convert_caffemodel_to_npz` classmethod instead.

ResNet152 has 60,192,872 trainable parameters, and it's the deepest ResNet model and it achieves the best result on ImageNet classification task in [ILSVRC 2015](#).

See: K. He et. al., [Deep Residual Learning for Image Recognition](#)

Parameters `pretrained_model` (*str*) – the destination of the pre-trained chainer model serialized as a `.npz` file. If this argument is specified as `auto`, it automatically loads and converts the caffemodel from `$CHAINER_DATASET_ROOT/pfnet/chainer/models/ResNet-152-model.caffemodel`, where `$CHAINER_DATASET_ROOT` is

set as `$HOME/.chainer/dataset` unless you specify another value by modifying the environment variable. Note that in this case the converted chainer model is stored on the same directory and automatically used from the next time. If this argument is specified as `None`, all the parameters are not initialized by the pre-trained model, but the default initializer used in the original paper, i.e., `chainer.initializers.HeNormal(scale=1.0)`.

Variables `available_layers` (*list of str*) – The list of available layer names used by `__call__` and `extract` methods.

Methods

`__call__(self, x, layers=['prob'])`
Computes all the feature maps specified by `layers`.

Warning: `test` argument is not supported anymore since v2. Instead, use `chainer.using_config('train', train)`. See `chainer.using_config()`.

Parameters

- **x** (*Variable*) – Input variable. It should be prepared by `prepare` function.
- **layers** (*list of str*) – The list of layer names you want to extract.

Returns A directory in which the key contains the layer name and the value contains the corresponding feature map variable.

Return type Dictionary of `~chainer.Variable`

`__getitem__(name)`
Equivalent to `getattr`.

`add_link(name, link)`
Registers a child link to this chain.

Deprecated since version v2.0.0: Assign the child link directly to an attribute within `init_scope()` instead. For example, the following code

```
chain.add_link('l1', L.Linear(3, 5))
```

can be replaced by the following line.

```
with chain.init_scope():
    chain.l1 = L.Linear(3, 5)
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the child link. This name is also used as the attribute name.
- **link** (*Link*) – The link object to be registered.

`add_param(name, shape=None, dtype=<class 'numpy.float32'>, initializer=None)`
Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a `Parameter` object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():  
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name, value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** (*Link*) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

classmethod convert_caffemodel_to_npz (*path_caffemodel, path_npz, n_layers=50*)

Converts a pre-trained caffemodel to a chainer model.

Parameters

- **path_caffemodel** (*str*) – Path of the pre-trained caffemodel.
- **path_npz** (*str*) – Path of the converted chainer model.

copy (*mode='share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters `mode` (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type *Link*

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters `link` (*Link*) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the enabled flag of the update rule of each parameter variable to `False`.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the enabled flag of the update rule of each parameter variable to `True`.

extract (*self*, *images*, *layers*=['pool5'], *size*=(224, 224))

Extracts all the feature maps of given images.

The difference of directly executing `__call__` is that it directly accepts images as an input and automatically transforms them to a proper variable. That is, it is also interpreted as a shortcut method that implicitly calls `prepare` and `__call__` functions.

Unlike `predict` method, this method does not override `chainer.config.train` and `chainer.config.enable_backprop` configuration. If you want to extract features without updating model parameters, you need to manually set configuration when calling this method as follows:

```
# model is an instance of ResNetLayers (50 or 101 or 152 layers)
with chainer.using_config('train', False):
    with chainer.using_config('enable_backprop', False):
        feature = model.extract([image])
```

Warning: `test` and `volatile` arguments are not supported anymore since v2. Instead, users should configure training and volatile modes with `train` and `enable_backprop`, respectively.

Note that default behavior of this method is different between v1 and later versions. Specifically, the default values of `test` in v1 were `True` (test mode). But that of `chainer.config.train` is also `True` (train mode). Therefore, users need to explicitly switch `train` to `False` to run the code in test mode and `enable_backprop` to `False` to turn off computational graph construction.

See the [upgrade guide](#).

Parameters

- **images** (*iterable of `PIL.Image` or `numpy.ndarray`*) – Input images.
- **layers** (*list of `str`*) – The list of layer names you want to extract.
- **size** (*pair of `ints`*) – The resolution of resized images used as an input of CNN. All the given images are not resized if this argument is `None`, but the resolutions of all the images should be the same.

Returns A directory in which the key contains the layer name and the value contains the corresponding feature map variable.

Return type Dictionary of `~chainer.Variable`

`init_scope()`

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for [Chain](#)) by an assignment. A [Parameter](#) object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a [Parameter](#) object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

`links (skipself=False)`

Returns a generator of all links under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

`namedlinks (skipself=False)`

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

predict (*images, oversample=True*)

Computes all the probabilities of given images.

Parameters

- **images** (*iterable of PIL.Image or numpy.ndarray*) – Input images. When you specify a color image as a `numpy.ndarray`, make sure that color order is RGB.
- **oversample** (*bool*) – If `True`, it averages results across center, corners, and mirrors. Otherwise, it uses only the center.

Returns Output that contains the class probabilities of given images.

Return type *Variable*

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters **name** (*str*) – Name of the attribute to be registered.

repeat (*n_repeat, mode='init'*)

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The `mode` argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the mode was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

available_layers

functions

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.model.vision.resnet.prepare

`chainer.links.model.vision.resnet.prepare(image, size=(224, 224))`

Converts the given image to the numpy array for ResNets.

Note that you have to call this method before `__call__` because the pre-trained resnet model requires to resize the given image, convert the RGB to the BGR, subtract the mean, and permute the dimensions before calling.

Parameters

- **image** (*PIL.Image* or `numpy.ndarray`) – Input image. If an input is `numpy.ndarray`, its shape must be (height, width), (height, width, channels), or (channels, height, width), and the order of the channels must be RGB.
- **size** (*pair of ints*) – Size of converted images. If None, the given image is not resized.

Returns The converted output array.

Return type `numpy.ndarray`

Compatibility with other frameworks

<code>chainer.links.TheanoFunction</code>	Theano function wrapper.
<code>chainer.links.caffe.CaffeFunction</code>	Caffe emulator based on the model file of Caffe.

chainer.links.TheanoFunction

class `chainer.links.TheanoFunction(inputs, outputs)`

Theano function wrapper.

Warning: This feature is experimental. The interface can change in the future.

This function wraps Theano function as a `chainer.Link`. A user needs to make input Theano variables and output Theano variables. This function automatically creates Theano function for forward calculation and backward calculation from inputs and outputs. And then, it sends data in `chainer.Variable` to the function and gets results from Theano.

Example

```
>>> import theano
>>> x = theano.tensor.fvector()
>>> y = theano.tensor.fvector()
```

(continues on next page)

(continued from previous page)

```

>>> z = x + y
>>> w = x - y
>>> f = L.TheanoFunction(inputs=[x, y], outputs=[z, w])
>>> a = chainer.Variable(np.array([1, 2], dtype=np.float32))
>>> b = chainer.Variable(np.array([2, 3], dtype=np.float32))
>>> c, d = f(a, b)
>>> c.data
array([3., 5.], dtype=float32)
>>> d.data
array([-1., -1.], dtype=float32)

```

Note: The current implementation always copies `cupy.ndarray` to CPU.

Parameters

- **inputs** (tuple of `theano.tensor.TensorVariable`) – Input variables of Theano. This function accepts the same number of *Variables* in forward computation.
- **outputs** (tuple of `theano.tensor.TensorVariable`) – Output variables of Theano. The function returns the same number of *Variables* as outputs.

Methods

__call__ (*args)
Call self as a function.

add_param (name, shape=None, dtype=<class 'numpy.float32'>, initializer=None)
Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a *Parameter* object directly to an attribute within *init_scope()* instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name*, *value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** ([Link](#)) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode*=*'share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument *mode* below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameter variables under the returned link object is re-initialized by calling their *initialize()* method, so that all the parameters may have different initial values from the original link. *copy* means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. *share* means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is *share*.

Returns Copied link object.

Return type [Link](#)

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters **link** ([Link](#)) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the [Parameters](#) held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for [Chain](#)) by an assignment. A [Parameter](#) object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a [Parameter](#) object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If *name* has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters *name* (*str*) – Name of the attribute to be registered.

repeat (*n_repeat*, *mode*='init')

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The *mode* argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The *net* object contains 16 blocks, each of which is *ConvBNReLU*. And the *mode* was *init*, so each block is re-initialized with different parameters. If you give *copy* to this argument, each block has same values for its parameters but its object ID is different from others. If it is *share*, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. *copy* means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. *share* means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters *serializer* (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.links.caffe.CaffeFunction

class `chainer.links.caffe.CaffeFunction` (*model_path*)

Caffe emulator based on the model file of Caffe.

Given a protocol buffers file of a Caffe model, this class loads and emulates it on `Variable` objects. It supports the official reference models provided by BVLC.

Note: CaffeFunction ignores the following layers:

- Layers that CaffeFunction does not support (including data layers)
 - Layers that have no top blobs
 - Layers whose bottom blobs are incomplete (i.e., some or all of them are not given nor computed)
-

Warning: It does not support full compatibility against Caffe. Some layers and configurations are not implemented in Chainer yet, though the reference models provided by the BVLC team are supported except data layers.

Example

Consider we want to extract the (unnormalized) log class probability of given images using BVLC reference CaffeNet. The model can be downloaded from:

http://dl.caffe.berkeleyvision.org/bvlc_reference_caffenet.caffemodel

We want to compute the `fc8` blob from the `data` blob. It is simply written as follows:

```
# Load the model
func = CaffeFunction('path/to/bvlc_reference_caffenet.caffemodel')

# Minibatch of size 10
x_data = numpy.ndarray((10, 3, 227, 227), dtype=numpy.float32)
... # (Fill the minibatch here)

# Forward the pre-trained net
x = Variable(x_data)
y, = func(inputs={'data': x}, outputs=['fc8'])
```

The result `y` contains the `Variable` corresponding to the `fc8` blob. The computational graph is memorized as a usual forward computation in Chainer, so we can run backprop through this pre-trained net.

Parameters `model_path` (*str*) – Path to the binary-`proto` model file of Caffe.

Variables `forwards` (*dict*) – A mapping from layer names to corresponding functions.

Methods

`__call__(self, inputs, outputs, disable=())`

Executes a sub-network of the network.

This function acts as an interpreter of the network definition for Caffe. On execution, it interprets each layer one by one, and if the bottom blobs are already computed, then emulates the layer and stores output blobs as `Variable` objects.

Warning: `train` argument is not supported anymore since v2. Instead, use `chainer.using_config('train', train)`. See `chainer.using_config()`.

Parameters

- **inputs** (*dict*) – A dictionary whose key-value pairs indicate initial correspondences between blob names and `Variable` objects.
- **outputs** (*Iterable*) – A list of blob names whose corresponding `Variable` objects are returned.
- **disable** (*Iterable*) – A list of layer names that will be ignored during the forward computation.

Returns A tuple of output *Variable* objects corresponding to elements of the *outputs* argument.

Return type `tuple`

`__getitem__` (*name*)
Equivalent to `getattr`.

add_link (*name*, *link*)

Registers a child link to this chain.

Deprecated since version v2.0.0: Assign the child link directly to an attribute within `init_scope()` instead. For example, the following code

```
chain.add_link('l1', L.Linear(3, 5))
```

can be replaced by the following line.

```
with chain.init_scope():  
    chain.l1 = L.Linear(3, 5)
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the child link. This name is also used as the attribute name.
- **link** (*Link*) – The link object to be registered.

add_param (*name*, *shape*=None, *dtype*=<class 'numpy.float32'>, *initializer*=None)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a *Parameter* object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():  
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not None, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, *dtype* argument is ignored.

add_persistent (*name*, *value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** (*Link*) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode*=*'share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument *mode* below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters **mode** (*str*) – It should be either *init*, *copy*, or *share*. *init* means parameter variables under the returned link object is re-initialized by calling their *initialize()* method, so that all the parameters may have different initial values from the original link. *copy* means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. *share* means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is *share*.

Returns Copied link object.

Return type *Link*

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters **link** (*Link*) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for [Chain](#)) by an assignment. A [Parameter](#) object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a [Parameter](#) object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters `skipself` (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters `skipself` (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters `include_uninit` (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters `include_uninit` (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters `name` (*str*) – Name of the attribute to be registered.

repeat (*n_repeat*, *mode*='init')

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The *mode* argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the *mode* was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

5.3.5 Link and Chain base classes

<code>chainer.Link</code>	Building block of model definitions.
<code>chainer.Chain</code>	Composable link with object-like interface.
<code>chainer.ChainList</code>	Composable link with list-like interface.
<code>chainer.Sequential</code>	Sequential model which has a single-stream forward pass.

chainer.Link

class `chainer.Link` (***params*)

Building block of model definitions.

Link is a building block of neural network models that support various features like handling parameters, defining network fragments, serialization, etc.

Link is the primitive structure for the model definitions. It supports management of parameter variables and *persistent values* that should be incorporated to serialization.

Parameter is an instance of `Parameter` registered to a link. A `Parameter` object can be registered as a

parameter of the link by assigning it to an attribute within *an initialization scope*, which is a code surrounded by a `init_scope()` context manager using the `with` statement.

Persistent values are arrays, scalars, or any other serializable values registered via `register_persistent()` or `add_persistent()`.

Note: Whereas arbitrary serializable objects can be registered as persistent values, it is strongly recommended to just register values that should be treated as results of learning. A typical example of persistent values is ones computed during training and required for testing, e.g. running statistics for batch normalization.

Parameters and persistent values are referred by their names. They can be accessed as attributes of the links. Link class itself manages the lists of names of parameters and persistent values to distinguish parameters and persistent values from other attributes.

Link can be composed into more complex models. This composition feature is supported by child classes like `Chain` and `ChainList`. One can create a chain by combining one or more links. See the documents for these classes for details.

As noted above, Link supports the serialization protocol of the `Serializer` class. **Note that only parameters and persistent values are saved and loaded.** Other attributes are considered as a part of user program (i.e. a part of network definition). In order to construct a link from saved file, other attributes must be identically reconstructed by user codes.

Example

This is a simple example of custom link definition. Chainer itself also provides many links defined under the `links` module. They might serve as examples, too.

Consider we want to define a simple primitive link that implements a fully-connected layer based on the `linear()` function. Note that this function takes input units, a weight variable, and a bias variable as arguments. Then, the fully-connected layer can be defined as follows:

```
import chainer
import chainer.functions as F
from chainer import initializers
import numpy as np

class LinearLayer(chainer.Link):

    def __init__(self, n_in, n_out):
        super(LinearLayer, self).__init__()
        with self.init_scope():
            self.W = chainer.Parameter(
                initializers.Normal(), (n_out, n_in))
            self.b = chainer.Parameter(
                initializers.Zero(), (n_out,))

    def __call__(self, x):
        return F.linear(x, self.W, self.b)
```

This example shows that a user can define arbitrary parameters and use them in any methods. Links typically implement the `__call__` operator, although they can also provide other methods to implement the forward propagation.

Parameters `params` – (*deprecated since v2.0.0*) Names, shapes, and optional dtypes of initial parameters. The keywords are used as the parameter names and the corresponding values consist

either of the shape or a tuple of shape and a dtype (`shape, dtype`). If only the shape is supplied, the default dtype will be used.

Variables `name` (*str*) – Name of this link, given by the parent chain (if exists).

Methods

`__call__` (**args, **kwargs*)
Call self as a function.

`add_param` (*name, shape=None, dtype=<class 'numpy.float32'>, initializer=None*)
Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a `Parameter` object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():  
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

`add_persistent` (*name, value*)
Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

`addgrads` (*link*)
Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters `link` (`Link`) – Source link object.

`children` ()
Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (mode='share')

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters `mode` (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type *Link*

copyparams (link)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters `link` (*Link*) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope ()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a `Parameter` object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters *skipself* (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters *include_uninit* (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If *name* has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters *name* (*str*) – Name of the attribute to be registered.

repeat (*n_repeat*, *mode='init'*)

Repeats this link multiple times to make a `Sequential`.

This method returns a `Sequential` object which has the same `Link` multiple times repeatedly. The *mode* argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer `Sequential` block like this:

```

class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')

```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the mode was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes**update_enabled**

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.Chain**class chainer.Chain** (**links)

Composable link with object-like interface.

Composability is one of the most important features of neural nets. Neural net models consist of many reusable fragments, and each model itself might be embedded into a larger learnable system. Chain enables us to write a neural net based on composition, without bothering about routine works like collecting parameters, serialization, copying the structure with parameters shared, etc.

This class actually provides a way to compose one or more links into one structure. A chain can contain one or more *child links*. Child link is a link registered to the chain with its own name. The child link is stored to an attribute of the chain with the name. User can write a whole model or a fragment of neural nets as a child class of Chain.

Each chain itself is also a link. Therefore, one can combine chains into higher-level chains. In this way, links and chains construct a *link hierarchy*. Link hierarchy forms a tree structure, where each node is identified by the path from the root. The path is represented by a string like a file path in UNIX, consisting of names of nodes on the path, joined by slashes `/`.

A child link can be added just by assigning it to an attribute of the chain within `init_scope()`.

The registered child link is saved and loaded on serialization and deserialization, and involved in the optimization. The registered link is called a child. The child link is accessible via `children()` generator, which returns a generator running through the children in registered order.

On registration of a child link, its `name` attribute is also set (or overwritten if the link has already been registered to another chain).

Example

This is a simple example of custom chain definition. Chainer itself also provides some chains defined under the `links` module. They might serve as examples, too.

Consider we want to define a multi-layer perceptron consisting of two hidden layers with rectifiers as activation functions. We can use the `Linear` link as a building block:


```

import chainer
import chainer.functions as F
import chainer.links as L

class MultiLayerPerceptron(chainer.Chain):

    def __init__(self, n_in, n_hidden, n_out):
        super(MultiLayerPerceptron, self).__init__()
        with self.init_scope():
            self.layer1 = L.Linear(n_in, n_hidden)
            self.layer2 = L.Linear(n_hidden, n_hidden)
            self.layer3 = L.Linear(n_hidden, n_out)

    def __call__(self, x):
        # Forward propagation
        h1 = F.relu(self.layer1(x))
        h2 = F.relu(self.layer2(h1))
        return self.layer3(h2)

```

Child links are registered via the assignment within a `with self.init_scope():` block. The forward propagation is often implemented as the `__call__` operator as the above example, though it is not mandatory.

Parameters `links` – Child links. The keywords are used as their names. The names are also set to the links.

Deprecated since version v2.0.0: Assign child links directly to attributes instead.

Methods

`__call__(*args, **kwargs)`
Call self as a function.

`__getitem__(name)`
Equivalent to `getattr`.

`add_link(name, link)`
Registers a child link to this chain.

Deprecated since version v2.0.0: Assign the child link directly to an attribute within `init_scope()` instead. For example, the following code

```
chain.add_link('l1', L.Linear(3, 5))
```

can be replaced by the following line.

```
with chain.init_scope():
    chain.l1 = L.Linear(3, 5)
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the child link. This name is also used as the attribute name.
- **link** (*Link*) – The link object to be registered.

add_param (*name*, *shape=None*, *dtype=<class 'numpy.float32'>*, *initializer=None*)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a `Parameter` object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name*, *value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** (`Link`) – Source link object.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode='share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters `mode` (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type *Link*

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters `link` (*Link*) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope ()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a *Parameter* object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters **name** (*str*) – Name of the attribute to be registered.

repeat (*n_repeat, mode='init'*)

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The `mode` argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))
```

(continues on next page)

(continued from previous page)

```
net = ConvBNReLU().repeat(16, mode='init')
```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the `mode` was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

`update_enabled`

True if at least one parameter has an update rule enabled.

`within_init_scope`

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

`xp`

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

`chainer.ChainList`

class `chainer.ChainList` (*links)

Composable link with list-like interface.

This is another example of compositional link. Unlike `Chain`, this class can be used like a list of child links. Each child link is indexed by a non-negative integer, and it maintains the current number of registered child links. The `add_link()` method inserts a new link at the end of the list. It is useful to write a chain with arbitrary number of child links, e.g. an arbitrarily deep multi-layer perceptron.

Note that this class does not implement all methods of `list`.

Parameters `links` – Initial child links.

Methods

`__call__` (*args, **kwargs)

Call self as a function.

`__getitem__` (index)

Returns the child at given index.

Parameters `index` (`int`) – Index of the child in the list.

Returns The `index`-th child link.

Return type `Link`

`__len__` ()

Returns the number of children.

`__iter__` ()

`add_link` (link)

Registers a child link and adds it to the tail of the list.

Parameters `link` (`Link`) – The link object to be registered.

`add_param` (name, shape=None, dtype=<class 'numpy.float32'>, initializer=None)

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a `Parameter` object directly to an attribute within `init_scope()` instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not `None`, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, `dtype` argument is ignored.

add_persistent (*name, value*)

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** (*Link*) – Source link object.

append (*link*)

Registers a child link and adds it to the tail of the list.

This is equivalent to `add_link()`. This method has been added to emulate the `list` interface.

Parameters **link** (*Link*) – The link object to be registered.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

cleargrads ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode='share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters `mode` (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their `initialize()` method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type *Link*

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters `link` (*Link*) – Source link object.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the *Parameters* held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `False`.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the `enabled` flag of the update rule of each parameter variable to `True`.

init_scope ()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for *Chain*) by an assignment. A *Parameter* object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a *Parameter* object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If `name` has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters **name** (*str*) – Name of the attribute to be registered.

repeat (*n_repeat, mode='init'*)

Repeats this link multiple times to make a *Sequential*.

This method returns a *Sequential* object which has the same *Link* multiple times repeatedly. The `mode` argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer *Sequential* block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the `mode` was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (*int*) – Number of times to repeat.
- **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned *Sequential* will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting *Sequential* object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (*serializer*)

Serializes the link object.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

Attributes

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

chainer.Sequential

class `chainer.Sequential(*layers)`

Sequential model which has a single-stream forward pass.

Warning: This feature is experimental. The interface can change in the future.

This class enables to construct a network which has sequential structure easily. While `Chain` and `ChainList` can only take `Link` object as input to their constructor, this `Sequential` can take arbitrary number of any callable objects for the forward pass computation. A `Sequential` calls the given callable objects sequentially inside of the `__call__()` method in the same order as the given arguments. Therefore, you do not need to write the forward pass computation explicitly.

Example

The below example code shows how to use this class to construct a simple sequential network:

```
import chainer
import chainer.functions as F
import link.Links as L
from chainer import Sequential

# Model definition without writing __call__ function
model = Sequential(
    L.Linear(n_in, n_hidden),
    F.relu,
    L.Linear(n_hidden, n_hidden),
    F.relu,
    L.Linear(n_hidden, n_out)
)

# Compute the forward pass
y = model(x)
```

where `x` denotes a mini-batch of `n_in`-dimensional input vectors.

Furthermore, `Sequential` supports built-in list APIs, so you can concatenate `Sequential` objects to create a longer `Sequential` model easily with the same ways as Python lists:

```
model_A = Sequential(L.Linear(10, 10), F.relu)
model_B = Sequential(L.Linear(10, 10), F.sigmoid)
model_C = model_A + model_B
```

To repeat a `Sequential` object multiple times, you can use `repeat()` method.

```
model_D = model_A.repeat(3)
```

You can also add your own functions or any callable objects to a `Sequential` object:

```
from link.Links.model.vision.vgg import VGG16Layers()

model = Sequential()
model.append(L.Linear(n_out, n_hidden))
model.append(F.relu)
model.append(F.Reshape((1, 3, 224, 224)))
model.append(VGG16Layers())
model.append(lambda x: x['prob'])

y = model(x)
```

The above code example shows how to add some layers to the `model` using `append()` method and then add a large network (`VGG16Layers`) and finally add a lambda function to extract the `prob` output.

You can check the structure of your model briefly using `print` as following:

```
>>> print(model_C)
0      Linear      W(10, 10)      b(10,)
1      relu
2      Linear      W(10, 10)      b(10,)
3      sigmoid
```

Note: Note that a `Sequential` link which has at least one lambda function as its member cannot be pickled. So, please use partial method from `functools` package instead:

```
from functools import partial

# This is not pickable
model = Sequential(
    L.Convolution2D(None, 64, 3, 1, 1),
    lambda x: F.max_pooling_2d(x, 2)
)

# This is pickable
model = Sequential(
    L.Convolution2D(None, 64, 3, 1, 1),
    partial(F.max_pooling_2d, ksize=2)
)
```

Parameters `layers` – The layers which are called in its order. Each component should be a callable object including `Link` object and functions defined under the `chainer.functions`, e.g., `relu`, etc.

Methods

`__call__(*x)`

Forward pass computation.

This method performs the forward pass computation by giving the input variable `x` to the layers registered in the constructor in the same order as the order in which the arguments are given to the constructor.

It should be noted that the input variable is given directly to the first layer and all intermediate outputs generated during the forward pass are also directly fed to the next layer. Therefore, the number of outputs

at a layer should be the same as the number of inputs at the next layer.

Parameters **x** – Input variables.

Returns The output of the final layer in the given layers.

`__getitem__(i)`

Returns the child at given index.

Parameters **index** (*int*) – Index of the child in the list.

Returns The *index*-th child link.

Return type *Link*

`__setitem__(i, layer)`

`__len__()`

Returns the number of children.

`__iter__()`

`add_link(link)`

Registers a child link and adds it to the tail of the list.

Parameters **link** (*Link*) – The link object to be registered.

`add_param(name, shape=None, dtype=<class 'numpy.float32'>, initializer=None)`

Registers a parameter to the link.

Deprecated since version v2.0.0: Assign a *Parameter* object directly to an attribute within *init_scope()* instead. For example, the following code

```
link.add_param('W', shape=(5, 3))
```

can be replaced by the following assignment.

```
with link.init_scope():
    link.W = chainer.Parameter(None, (5, 3))
```

The latter is easier for IDEs to keep track of the attribute's type.

Parameters

- **name** (*str*) – Name of the parameter. This name is also used as the attribute name.
- **shape** (*int or tuple of ints*) – Shape of the parameter array. If it is omitted, the parameter variable is left uninitialized.
- **dtype** – Data type of the parameter array.
- **initializer** – If it is not *None*, the data is initialized with the given initializer. If it is an array, the data is directly initialized by it. If it is callable, it is used as a weight initializer. Note that in these cases, *dtype* argument is ignored.

`add_persistent(name, value)`

Registers a persistent value to the link.

The registered value is saved and loaded on serialization and deserialization. The value is set to an attribute of the link.

Parameters

- **name** (*str*) – Name of the persistent value. This name is also used for the attribute name.
- **value** – Value to be registered.

addgrads (*link*)

Accumulates gradient values from given link.

This method adds each gradient array of the given link to corresponding gradient array of this link. The accumulation is even done across host and different devices.

Parameters **link** ([Link](#)) – Source link object.

append (*layer*)

Registers a child link and adds it to the tail of the list.

This is equivalent to [add_link\(\)](#). This method has been added to emulate the `list` interface.

Parameters **link** ([Link](#)) – The link object to be registered.

children ()

Returns a generator of all child links.

Returns A generator object that generates all child links.

clear ()**cleargrads** ()

Clears all gradient arrays.

This method should be called before the backward computation at every iteration of the optimization.

copy (*mode='share'*)

Copies the link hierarchy to new one.

The whole hierarchy rooted by this link is copied. There are three modes to perform copy. Please see the document for the argument `mode` below.

The name of the link is reset on the copy, since the copied instance does not belong to the original parent chain (even if exists).

Parameters **mode** (*str*) – It should be either `init`, `copy`, or `share`. `init` means parameter variables under the returned link object is re-initialized by calling their [initialize\(\)](#) method, so that all the parameters may have different initial values from the original link. `copy` means that the link object is deeply copied, so that its parameters are not re-initialized but are also deeply copied. Thus, all parameters have same initial values but can be changed independently. `share` means that the link is shallowly copied, so that its parameters' arrays are shared with the original one. Thus, their values are changed synchronously. The default mode is `share`.

Returns Copied link object.

Return type [Link](#)

copyparams (*link*)

Copies all parameters from given link.

This method copies data arrays of all parameters in the hierarchy. The copy is even done across the host and devices. Note that this method does not copy the gradient arrays.

Parameters **link** ([Link](#)) – Source link object.

count (*layer*)**count_by_layer_type** (*type_name*)

Count the number of layers by layer type.

This method counts the number of layers which have the name given by the argument `type_name`. For example, if you want to know the number of [Linear](#) layers included in this model, `type_name` should

be `Linear`. If you want to know the number of `Function` classes or user-defined functions which have a specific name, `type_name` should be the function name, e.g., `relu` or `reshape`, etc.

Parameters `type_name` (`str`) – The class or function name of a layer you want to enumerate.

count_params ()

Counts the total number of parameters.

This method counts the total number of scalar values included in all the `Parameters` held by this link and its descendants.

If the link contains uninitialized parameters, this method raises a warning.

Returns The total size of parameters (int)

disable_update ()

Disables update rules of all parameters under the link hierarchy.

This method sets the enabled flag of the update rule of each parameter variable to `False`.

enable_update ()

Enables update rules of all parameters under the link hierarchy.

This method sets the enabled flag of the update rule of each parameter variable to `True`.

extend (`sequential`)

flatten ()

Flatten nested `Sequential` links.

This method flattens all the nested `Sequential` links inside this `Sequential` link.

Returns A flattened `Sequential` object.

Example

```
>>> import chainer
>>> import chainer.functions as F
>>> import chainer.links as L
>>> a = chainer.Sequential(L.Linear(None, 10), F.relu)
>>> b = chainer.Sequential(L.Linear(None, 10), F.relu)
>>> a.append(b)
>>> print(a) # Without flatten
0      Linear  W(None) b(10,)
1      relu
2      Sequential      which has 2 layers
>>> print(a.flatten()) # With flatten
0      Linear  W(None) b(10,)
1      relu
2      Linear  W(None) b(10,)
3      relu
```

index (`layer`, `start=None`, `end=None`)

init_scope ()

Creates an initialization scope.

This method returns a context manager object that enables registration of parameters (and links for `Chain`) by an assignment. A `Parameter` object can be automatically registered by assigning it to an attribute under this context manager.

Example

In most cases, the parameter registration is done in the initializer method. Using the `init_scope` method, we can simply assign a `Parameter` object to register it to the link.

```
class MyLink(chainer.Link):
    def __init__(self):
        super().__init__()
        with self.init_scope():
            self.W = chainer.Parameter(0, (10, 5))
            self.b = chainer.Parameter(0, (5,))
```

insert (*i*, *layer*)

links (*skipself=False*)

Returns a generator of all links under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all links.

namedlinks (*skipself=False*)

Returns a generator of all (path, link) pairs under the hierarchy.

Parameters **skipself** (*bool*) – If `True`, then the generator skips this link and starts with the first child link.

Returns A generator object that generates all (path, link) pairs.

namedparams (*include_uninit=True*)

Returns a generator of all (path, param) pairs under the hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all (path, parameter) pairs. The paths are relative from this link.

params (*include_uninit=True*)

Returns a generator of all parameters under the link hierarchy.

Parameters **include_uninit** (*bool*) – If `True`, it also generates uninitialized parameters.

Returns A generator object that generates all parameters.

pop (*i=-1*)

register_persistent (*name*)

Registers an attribute of a given name as a persistent value.

This is a convenient method to register an existing attribute as a persistent value. If *name* has been already registered as a parameter, this method removes it from the list of parameter names and re-registers it as a persistent value.

Parameters **name** (*str*) – Name of the attribute to be registered.

remove (*layer*)

remove_by_layer_type (*type_name*)

Remove layers by layer type.

This method removes layers from the Sequential object by the layer's class name or function name. If you want to remove a `Link`, the argument *type_name* should be its class name, e.g., `Linear`

or `Convolution2D`, etc. If you want to remove a `Function` class or any other callable objects, `type_name` should be the function name, e.g., `relu` or `reshape`, etc.

Parameters `type_name` (`str`) – The name of a layer you want to remove.

repeat (`n_repeat`, `mode='init'`)

Repeats this link multiple times to make a `Sequential`.

This method returns a `Sequential` object which has the same `Link` multiple times repeatedly. The `mode` argument means how to copy this link to repeat.

Example

You can repeat the same link multiple times to create a longer `Sequential` block like this:

```
class ConvBNReLU(chainer.Chain):

    def __init__(self):
        super(ConvBNReLU, self).__init__()
        with self.init_scope():
            self.conv = L.Convolution2D(
                None, 64, 3, 1, 1, nobias=True)
            self.bn = L.BatchNormalization(64)

    def __call__(self, x):
        return F.relu(self.bn(self.conv(x)))

net = ConvBNReLU().repeat(16, mode='init')
```

The `net` object contains 16 blocks, each of which is `ConvBNReLU`. And the `mode` was `init`, so each block is re-initialized with different parameters. If you give `copy` to this argument, each block has same values for its parameters but its object ID is different from others. If it is `share`, each block is same to others in terms of not only parameters but also the object IDs because they are shallow-copied, so that when the parameter of one block is changed, all the parameters in the others also change.

Parameters

- **n_repeat** (`int`) – Number of times to repeat.
- **mode** (`str`) – It should be either `init`, `copy`, or `share`. `init` means parameters of each repeated element in the returned `Sequential` will be re-initialized, so that all elements have different initial parameters. `copy` means that the parameters will not be re-initialized but object itself will be deep-copied, so that all elements have same initial parameters but can be changed independently. `share` means all the elements which consist the resulting `Sequential` object are same object because they are shallow-copied, so that all parameters of elements are shared with each other.

serialize (`serializer`)

Serializes the link object.

Parameters **serializer** (`AbstractSerializer`) – Serializer object.

to_cpu ()

Copies parameter variables and persistent values to CPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to CPU, the link implementation must override this method to do so.

Returns: self

to_gpu (*device=None*)

Copies parameter variables and persistent values to GPU.

This method does not handle non-registered attributes. If some of such attributes must be copied to GPU, the link implementation must override this method to do so.

Parameters **device** – Target device specifier. If omitted, the current device is used.

Returns: self

to_intel64 ()

Copies parameter variables and persistent values to CPU.

zerograds ()

Initializes all gradient arrays by zero.

This method can be used for the same purpose of `cleargrads`, but less efficient. This method is left for backward compatibility.

Deprecated since version v1.15: Use `cleargrads()` instead.

__add__ (*other*)

Attributes

update_enabled

True if at least one parameter has an update rule enabled.

within_init_scope

True if the current code is inside of an initialization scope.

See `init_scope()` for the details of the initialization scope.

xp

Array module for this link.

Depending on which of CPU/GPU this link is on, this property returns `numpy` or `cupy`.

5.4 Optimizers

<code>chainer.optimizers.AdaDelta</code>	Zeiler's ADADELTA.
<code>chainer.optimizers.AdaGrad</code>	AdaGrad optimizer.
<code>chainer.optimizers.Adam</code>	Adam optimizer.
<code>chainer.optimizers.MomentumSGD</code>	Momentum SGD optimizer.
<code>chainer.optimizers.NesterovAG</code>	Nesterov's Accelerated Gradient.
<code>chainer.optimizers.RMSprop</code>	RMSprop optimizer.
<code>chainer.optimizers.RMSpropGraves</code>	Alex Graves's RMSprop.
<code>chainer.optimizers.SGD</code>	Vanilla Stochastic Gradient Descent.
<code>chainer.optimizers.SMORMS3</code>	Simon Funk's SMORMS3.

5.4.1 `chainer.optimizers.AdaDelta`

class `chainer.optimizers.AdaDelta` (*rho=0.95, eps=1e-06*)

Zeiler's ADADELTA.

See: <http://www.matthewzeiler.com/pubs/googleTR2012/googleTR2012.pdf>

Parameters

- **rho** (*float*) – Exponential decay rate of the first and second order moments.
- **eps** (*float*) – Small value for the numerical stability.

Methods

add_hook (*hook*, *name=None*, *timing='auto'*)

Registers a hook function.

Hook function is typically called right after the gradient computation, though the timing depends on the optimization method, and the timing attribute.

Parameters

- **hook** (*callable*) – Hook function. If `hook.call_for_each_param` is true, this hook function is called for each parameter by passing the update rule and the parameter. Otherwise, this hook function is called only once each iteration by passing the optimizer.
- **name** (*str*) – Name of the registration. If omitted, `hook.name` is used by default.
- **timing** (*str*) – Specifies when the hook is called. If 'auto', the timing property of the hook will decide the timing. If 'pre', the hook will be called before any updates. If 'post', the hook will be called after any updates.

call_hooks (*timing='pre'*)

Invokes hook functions in registration order.

create_update_rule ()

Creates a new update rule object.

This method creates an update rule object. It is called by `setup()` to set up an update rule of each parameter. Each implementation of the gradient method should override this method to provide the default update rule implementation.

Returns Update rule object.

Return type *UpdateRule*

new_epoch ()

Starts a new epoch.

This method increments the *epoch* count. Note that if the optimizer depends on the epoch count, then user should call this method appropriately at the beginning of each epoch.

reallocate_cleared_grads ()

Reallocate gradients cleared by `cleargrad()`.

This method allocates arrays for all gradients which have `None`. This method is called before and after every optimizer hook. If an inheriting optimizer does not require this allocation, the optimizer can override this method with a blank function.

remove_hook (*name*)

Removes a hook function.

Parameters **name** (*str*) – Registered name of the hook function to remove.

serialize (*serializer*)

Serializes or deserializes the optimizer.

It only saves or loads the following things:

- Optimizer states
- Global states (`t` and `epoch`)

It does not save nor load the parameters of the target link. They should be separately saved or loaded.

Parameters `serializer` (`AbstractSerializer`) – Serializer or deserializer object.

set_loss_scale (`loss_scale`)
Sets loss scaling factor.

setup (`link`)
Sets a target link and initializes the optimizer states.

Given link is set to the `target` attribute. It also prepares the optimizer state dictionaries corresponding to all parameters in the link hierarchy. The existing states are discarded.

Parameters `link` (`Link`) – Target link object.

Returns The optimizer instance.

Note: As of v4.0.0, this function returns the optimizer instance itself so that you can instantiate and setup the optimizer in one line, e.g., `optimizer = SomeOptimizer().setup(link)`.

update (`lossfun=None`, `*args`, `**kwargs`)
Updates parameters based on a loss function or computed gradients.

This method runs in two ways.

- If `lossfun` is given, then it is used as a loss function to compute gradients.
- Otherwise, this method assumes that the gradients are already computed.

In both cases, the computed gradients are used to update parameters. The actual update routines are defined by the update rule of each parameter.

use_cleargrads (`use=True`)
Enables or disables use of `cleargrads()` in `update`.

Parameters `use` (`bool`) – If `True`, this function enables use of `cleargrads`. If `False`, disables use of `cleargrads` (`zerograds` is used).

Deprecated since version v2.0: Note that `update()` calls `cleargrads()` by default. `cleargrads()` is more efficient than `zerograds()`, so one does not have to call `use_cleargrads()`. This method remains for backward compatibility.

use_fp32_update (`flag=True`)
Enables use of parameter update in fp32.

Attributes

`epoch = 0`

`eps`
Alias to `self.hyperparam.eps`

`rho`
Alias to `self.hyperparam.rho`

`t = 0`

`target = None`

5.4.2 chainer.optimizers.AdaGrad

class `chainer.optimizers.AdaGrad` (*lr=0.001, eps=1e-08*)
AdaGrad optimizer.

See: <http://jmlr.org/papers/v12/duchi11a.html>

Parameters

- **lr** (*float*) – Learning rate.
- **eps** (*float*) – Small value for the numerical stability.

Methods

add_hook (*hook, name=None, timing='auto'*)
Registers a hook function.

Hook function is typically called right after the gradient computation, though the timing depends on the optimization method, and the timing attribute.

Parameters

- **hook** (*callable*) – Hook function. If `hook.call_for_each_param` is true, this hook function is called for each parameter by passing the update rule and the parameter. Otherwise, this hook function is called only once each iteration by passing the optimizer.
- **name** (*str*) – Name of the registration. If omitted, `hook.name` is used by default.
- **timing** (*str*) – Specifies when the hook is called. If 'auto', the timing property of the hook will decide the timing. If 'pre', the hook will be called before any updates. If 'post', the hook will be called after any updates.

call_hooks (*timing='pre'*)
Invokes hook functions in registration order.

create_update_rule ()
Creates a new update rule object.

This method creates an update rule object. It is called by `setup()` to set up an update rule of each parameter. Each implementation of the gradient method should override this method to provide the default update rule implementation.

Returns Update rule object.

Return type *UpdateRule*

new_epoch ()
Starts a new epoch.

This method increments the *epoch* count. Note that if the optimizer depends on the epoch count, then user should call this method appropriately at the beginning of each epoch.

reallocate_cleared_grads ()
Reallocate gradients cleared by `cleargrad()`.

This method allocates arrays for all gradients which have `None`. This method is called before and after every optimizer hook. If an inheriting optimizer does not require this allocation, the optimizer can override this method with a blank function.

remove_hook (*name*)
Removes a hook function.

Parameters `name` (*str*) – Registered name of the hook function to remove.

serialize (*serializer*)

Serializes or deserializes the optimizer.

It only saves or loads the following things:

- Optimizer states
- Global states (*t* and *epoch*)

It does not save nor load the parameters of the target link. They should be separately saved or loaded.

Parameters `serializer` (*AbstractSerializer*) – Serializer or deserializer object.

set_loss_scale (*loss_scale*)

Sets loss scaling factor.

setup (*link*)

Sets a target link and initializes the optimizer states.

Given link is set to the *target* attribute. It also prepares the optimizer state dictionaries corresponding to all parameters in the link hierarchy. The existing states are discarded.

Parameters `link` (*Link*) – Target link object.

Returns The optimizer instance.

Note: As of v4.0.0, this function returns the optimizer instance itself so that you can instantiate and setup the optimizer in one line, e.g., `optimizer = SomeOptimizer().setup(link)`.

update (*lossfun=None, *args, **kws*)

Updates parameters based on a loss function or computed gradients.

This method runs in two ways.

- If *lossfun* is given, then it is used as a loss function to compute gradients.
- Otherwise, this method assumes that the gradients are already computed.

In both cases, the computed gradients are used to update parameters. The actual update routines are defined by the update rule of each parameter.

use_cleargrads (*use=True*)

Enables or disables use of *cleargrads()* in *update*.

Parameters `use` (*bool*) – If *True*, this function enables use of *cleargrads*. If *False*, disables use of *cleargrads* (*zerograds* is used).

Deprecated since version v2.0: Note that *update()* calls *cleargrads()* by default. *cleargrads()* is more efficient than *zerograds()*, so one does not have to call *use_cleargrads()*. This method remains for backward compatibility.

use_fp32_update (*flag=True*)

Enables use of parameter update in fp32.

Attributes

`epoch = 0`

`eps`

Alias to `self.hyperparam.eps`

```

lr
    Alias to self.hyperparam.lr

t = 0

target = None

```

5.4.3 chainer.optimizers.Adam

```

class chainer.optimizers.Adam(alpha=0.001, beta1=0.9, beta2=0.999, eps=1e-08, eta=1.0,
                             weight_decay_rate=0, amsgrad=False)

```

Adam optimizer.

See: [Adam: A Method for Stochastic Optimization](#)

Modified for proper weight decay (also called AdamW). AdamW introduces the additional parameters `eta` and `weight_decay_rate`, which can be used to properly scale the learning rate, and decouple the weight decay rate from `alpha`, as shown in the below paper.

Note that with the default values `eta = 1` and `weight_decay_rate = 0`, this implementation is identical to the standard Adam method.

See: [Fixing Weight Decay Regularization in Adam](#)

A flag `amsgrad` to use the AMSGrad variant of Adam from the paper: [On the Convergence of Adam and Beyond](#)

Parameters

- **alpha** (*float*) – Coefficient of learning rate.
- **beta1** (*float*) – Exponential decay rate of the first order moment.
- **beta2** (*float*) – Exponential decay rate of the second order moment.
- **eps** (*float*) – Small value for the numerical stability.
- **eta** (*float*) – Schedule multiplier, can be used for warm restarts.
- **weight_decay_rate** (*float*) – Weight decay rate.
- **amsgrad** (*bool*) – Whether to use AMSGrad variant of Adam.

Methods

```

add_hook(hook, name=None, timing='auto')

```

Registers a hook function.

Hook function is typically called right after the gradient computation, though the timing depends on the optimization method, and the timing attribute.

Parameters

- **hook** (*callable*) – Hook function. If `hook.call_for_each_param` is true, this hook function is called for each parameter by passing the update rule and the parameter. Otherwise, this hook function is called only once each iteration by passing the optimizer.
- **name** (*str*) – Name of the registration. If omitted, `hook.name` is used by default.
- **timing** (*str*) – Specifies when the hook is called. If 'auto', the timing property of the hook will decide the timing. If 'pre', the hook will be called before any updates. If 'post', the hook will be called after any updates.

call_hooks (*timing='pre'*)

Invokes hook functions in registration order.

create_update_rule ()

Creates a new update rule object.

This method creates an update rule object. It is called by `setup()` to set up an update rule of each parameter. Each implementation of the gradient method should override this method to provide the default update rule implementation.

Returns Update rule object.

Return type *UpdateRule*

new_epoch ()

Starts a new epoch.

This method increments the *epoch* count. Note that if the optimizer depends on the epoch count, then user should call this method appropriately at the beginning of each epoch.

reallocate_cleared_grads ()

Reallocate gradients cleared by `cleargrad()`.

This method allocates arrays for all gradients which have `None`. This method is called before and after every optimizer hook. If an inheriting optimizer does not require this allocation, the optimizer can override this method with a blank function.

remove_hook (*name*)

Removes a hook function.

Parameters *name* (*str*) – Registered name of the hook function to remove.

serialize (*serializer*)

Serializes or deserializes the optimizer.

It only saves or loads the following things:

- Optimizer states
- Global states (*t* and *epoch*)

It does not saves nor loads the parameters of the target link. They should be separately saved or loaded.

Parameters *serializer* (*AbstractSerializer*) – Serializer or deserializer object.

set_loss_scale (*loss_scale*)

Sets loss scaling factor.

setup (*link*)

Sets a target link and initializes the optimizer states.

Given link is set to the *target* attribute. It also prepares the optimizer state dictionaries corresponding to all parameters in the link hierarchy. The existing states are discarded.

Parameters *link* (*Link*) – Target link object.

Returns The optimizer instance.

Note: As of v4.0.0, this function returns the optimizer instance itself so that you can instantiate and setup the optimizer in one line, e.g., `optimizer = SomeOptimizer().setup(link)`.

update (*lossfun=None, *args, **kwargs*)

Updates parameters based on a loss function or computed gradients.

This method runs in two ways.

- If `lossfun` is given, then it is used as a loss function to compute gradients.
- Otherwise, this method assumes that the gradients are already computed.

In both cases, the computed gradients are used to update parameters. The actual update routines are defined by the update rule of each parameter.

use_cleargrads (*use=True*)

Enables or disables use of `cleargrads()` in `update`.

Parameters *use* (*bool*) – If `True`, this function enables use of `cleargrads`. If `False`, disables use of `cleargrads` (`zerograds` is used).

Deprecated since version v2.0: Note that `update()` calls `cleargrads()` by default. `cleargrads()` is more efficient than `zerograds()`, so one does not have to call `use_cleargrads()`. This method remains for backward compatibility.

use_fp32_update (*flag=True*)

Enables use of parameter update in fp32.

Attributes

alpha

Alias to `self.hyperparam.alpha`

amsgrad

Alias to `self.hyperparam.amsgrad`

beta1

Alias to `self.hyperparam.beta1`

beta2

Alias to `self.hyperparam.beta2`

epoch = 0

eps

Alias to `self.hyperparam.eps`

eta

Alias to `self.hyperparam.eta`

lr

t = 0

target = None

weight_decay_rate

Alias to `self.hyperparam.weight_decay_rate`

5.4.4 chainer.optimizers.MomentumSGD

class `chainer.optimizers.MomentumSGD` (*lr=0.01, momentum=0.9*)

Momentum SGD optimizer.

Parameters

- **lr** (*float*) – Learning rate.

- **momentum** (*float*) – Exponential decay rate of the first order moment.

Methods

add_hook (*hook*, *name=None*, *timing='auto'*)

Registers a hook function.

Hook function is typically called right after the gradient computation, though the timing depends on the optimization method, and the timing attribute.

Parameters

- **hook** (*callable*) – Hook function. If `hook.call_for_each_param` is true, this hook function is called for each parameter by passing the update rule and the parameter. Otherwise, this hook function is called only once each iteration by passing the optimizer.
- **name** (*str*) – Name of the registration. If omitted, `hook.name` is used by default.
- **timing** (*str*) – Specifies when the hook is called. If 'auto', the timing property of the hook will decide the timing. If 'pre', the hook will be called before any updates. If 'post', the hook will be called after any updates.

call_hooks (*timing='pre'*)

Invokes hook functions in registration order.

create_update_rule ()

Creates a new update rule object.

This method creates an update rule object. It is called by `setup()` to set up an update rule of each parameter. Each implementation of the gradient method should override this method to provide the default update rule implementation.

Returns Update rule object.

Return type *UpdateRule*

new_epoch ()

Starts a new epoch.

This method increments the *epoch* count. Note that if the optimizer depends on the epoch count, then user should call this method appropriately at the beginning of each epoch.

reallocate_cleared_grads ()

Reallocate gradients cleared by `cleargrad()`.

This method allocates arrays for all gradients which have `None`. This method is called before and after every optimizer hook. If an inheriting optimizer does not require this allocation, the optimizer can override this method with a blank function.

remove_hook (*name*)

Removes a hook function.

Parameters **name** (*str*) – Registered name of the hook function to remove.

serialize (*serializer*)

Serializes or deserializes the optimizer.

It only saves or loads the following things:

- Optimizer states
- Global states (*t* and *epoch*)

It does not save nor load the parameters of the target link. They should be separately saved or loaded.

Parameters `serializer` (`AbstractSerializer`) – Serializer or deserializer object.

set_loss_scale (`loss_scale`)

Sets loss scaling factor.

setup (`link`)

Sets a target link and initializes the optimizer states.

Given link is set to the `target` attribute. It also prepares the optimizer state dictionaries corresponding to all parameters in the link hierarchy. The existing states are discarded.

Parameters `link` (`Link`) – Target link object.

Returns The optimizer instance.

Note: As of v4.0.0, this function returns the optimizer instance itself so that you can instantiate and setup the optimizer in one line, e.g., `optimizer = SomeOptimizer().setup(link)`.

update (`lossfun=None, *args, **kws`)

Updates parameters based on a loss function or computed gradients.

This method runs in two ways.

- If `lossfun` is given, then it is used as a loss function to compute gradients.
- Otherwise, this method assumes that the gradients are already computed.

In both cases, the computed gradients are used to update parameters. The actual update routines are defined by the update rule of each parameter.

use_cleargrads (`use=True`)

Enables or disables use of `cleargrads()` in `update`.

Parameters `use` (`bool`) – If `True`, this function enables use of `cleargrads`. If `False`, disables use of `cleargrads` (`zerograds` is used).

Deprecated since version v2.0: Note that `update()` calls `cleargrads()` by default. `cleargrads()` is more efficient than `zerograds()`, so one does not have to call `use_cleargrads()`. This method remains for backward compatibility.

use_fp32_update (`flag=True`)

Enables use of parameter update in fp32.

Attributes

`epoch = 0`

`lr`

Alias to `self.hyperparam.lr`

`momentum`

Alias to `self.hyperparam.momentum`

`t = 0`

`target = None`

5.4.5 chainer.optimizers.NesterovAG

class `chainer.optimizers.NesterovAG` (*lr=0.01, momentum=0.9*)
Nesterov's Accelerated Gradient.

See: <https://arxiv.org/abs/1212.0901>

Parameters

- **lr** (*float*) – Learning rate.
- **momentum** (*float*) – Exponential decay rate of the first order moment.

Methods

add_hook (*hook, name=None, timing='auto'*)
Registers a hook function.

Hook function is typically called right after the gradient computation, though the timing depends on the optimization method, and the timing attribute.

Parameters

- **hook** (*callable*) – Hook function. If `hook.call_for_each_param` is `true`, this hook function is called for each parameter by passing the update rule and the parameter. Otherwise, this hook function is called only once each iteration by passing the optimizer.
- **name** (*str*) – Name of the registration. If omitted, `hook.name` is used by default.
- **timing** (*str*) – Specifies when the hook is called. If 'auto', the timing property of the hook will decide the timing. If 'pre', the hook will be called before any updates. If 'post', the hook will be called after any updates.

call_hooks (*timing='pre'*)
Invokes hook functions in registration order.

create_update_rule ()
Creates a new update rule object.

This method creates an update rule object. It is called by `setup()` to set up an update rule of each parameter. Each implementation of the gradient method should override this method to provide the default update rule implementation.

Returns Update rule object.

Return type *UpdateRule*

new_epoch ()
Starts a new epoch.

This method increments the *epoch* count. Note that if the optimizer depends on the epoch count, then user should call this method appropriately at the beginning of each epoch.

reallocate_cleared_grads ()
Reallocate gradients cleared by `cleargrad()`.

This method allocates arrays for all gradients which have `None`. This method is called before and after every optimizer hook. If an inheriting optimizer does not require this allocation, the optimizer can override this method with a blank function.

remove_hook (*name*)
Removes a hook function.

Parameters `name` (*str*) – Registered name of the hook function to remove.

serialize (*serializer*)

Serializes or deserializes the optimizer.

It only saves or loads the following things:

- Optimizer states
- Global states (*t* and *epoch*)

It does not save nor load the parameters of the target link. They should be separately saved or loaded.

Parameters `serializer` (*AbstractSerializer*) – Serializer or deserializer object.

set_loss_scale (*loss_scale*)

Sets loss scaling factor.

setup (*link*)

Sets a target link and initializes the optimizer states.

Given link is set to the *target* attribute. It also prepares the optimizer state dictionaries corresponding to all parameters in the link hierarchy. The existing states are discarded.

Parameters `link` (*Link*) – Target link object.

Returns The optimizer instance.

Note: As of v4.0.0, this function returns the optimizer instance itself so that you can instantiate and setup the optimizer in one line, e.g., `optimizer = SomeOptimizer().setup(link)`.

update (*lossfun=None, *args, **kws*)

Updates parameters based on a loss function or computed gradients.

This method runs in two ways.

- If *lossfun* is given, then it is used as a loss function to compute gradients.
- Otherwise, this method assumes that the gradients are already computed.

In both cases, the computed gradients are used to update parameters. The actual update routines are defined by the update rule of each parameter.

use_cleargrads (*use=True*)

Enables or disables use of *cleargrads()* in *update*.

Parameters `use` (*bool*) – If *True*, this function enables use of *cleargrads*. If *False*, disables use of *cleargrads* (*zerograds* is used).

Deprecated since version v2.0: Note that *update()* calls *cleargrads()* by default. *cleargrads()* is more efficient than *zerograds()*, so one does not have to call *use_cleargrads()*. This method remains for backward compatibility.

use_fp32_update (*flag=True*)

Enables use of parameter update in fp32.

Attributes

`epoch = 0`

`lr`

Alias to `self.hyperparam.lr`

```
momentum
    Alias to self.hyperparam.momentum

t = 0

target = None
```

5.4.6 chainer.optimizers.RMSprop

class `chainer.optimizers.RMSprop` (*lr=0.01, alpha=0.99, eps=1e-08*)
RMSprop optimizer.

See: T. Tieleman and G. Hinton (2012). Lecture 6.5 - rmsprop, COURSERA: Neural Networks for Machine Learning.

Parameters

- **lr** (*float*) – Learning rate.
- **alpha** (*float*) – Exponential decay rate of the second order moment.
- **eps** (*float*) – Small value for the numerical stability.

Methods

add_hook (*hook, name=None, timing='auto'*)
Registers a hook function.

Hook function is typically called right after the gradient computation, though the timing depends on the optimization method, and the timing attribute.

Parameters

- **hook** (*callable*) – Hook function. If `hook.call_for_each_param` is true, this hook function is called for each parameter by passing the update rule and the parameter. Otherwise, this hook function is called only once each iteration by passing the optimizer.
- **name** (*str*) – Name of the registration. If omitted, `hook.name` is used by default.
- **timing** (*str*) – Specifies when the hook is called. If 'auto', the timing property of the hook will decide the timing. If 'pre', the hook will be called before any updates. If 'post', the hook will be called after any updates.

call_hooks (*timing='pre'*)
Invokes hook functions in registration order.

create_update_rule ()
Creates a new update rule object.

This method creates an update rule object. It is called by `setup()` to set up an update rule of each parameter. Each implementation of the gradient method should override this method to provide the default update rule implementation.

Returns Update rule object.

Return type *UpdateRule*

new_epoch ()
Starts a new epoch.

This method increments the *epoch* count. Note that if the optimizer depends on the epoch count, then user should call this method appropriately at the beginning of each epoch.

reallocate_cleared_grads()

Reallocate gradients cleared by `cleargrad()`.

This method allocates arrays for all gradients which have `None`. This method is called before and after every optimizer hook. If an inheriting optimizer does not require this allocation, the optimizer can override this method with a blank function.

remove_hook(name)

Removes a hook function.

Parameters `name` (`str`) – Registered name of the hook function to remove.

serialize(serializer)

Serializes or deserializes the optimizer.

It only saves or loads the following things:

- Optimizer states
- Global states (`t` and `epoch`)

It does not save nor load the parameters of the target link. They should be separately saved or loaded.

Parameters `serializer` (`AbstractSerializer`) – Serializer or deserializer object.

set_loss_scale(loss_scale)

Sets loss scaling factor.

setup(link)

Sets a target link and initializes the optimizer states.

Given link is set to the `target` attribute. It also prepares the optimizer state dictionaries corresponding to all parameters in the link hierarchy. The existing states are discarded.

Parameters `link` (`Link`) – Target link object.

Returns The optimizer instance.

Note: As of v4.0.0, this function returns the optimizer instance itself so that you can instantiate and setup the optimizer in one line, e.g., `optimizer = SomeOptimizer().setup(link)`.

update(lossfun=None, *args, **kws)

Updates parameters based on a loss function or computed gradients.

This method runs in two ways.

- If `lossfun` is given, then it is used as a loss function to compute gradients.
- Otherwise, this method assumes that the gradients are already computed.

In both cases, the computed gradients are used to update parameters. The actual update routines are defined by the update rule of each parameter.

use_cleargrads(use=True)

Enables or disables use of `cleargrads()` in `update`.

Parameters `use` (`bool`) – If `True`, this function enables use of `cleargrads`. If `False`, disables use of `cleargrads` (`zerograd` is used).

Deprecated since version v2.0: Note that `update()` calls `cleargrads()` by default. `cleargrads()` is more efficient than `zerograd()`, so one does not have to call `use_cleargrads()`. This method remains for backward compatibility.

use_fp32_update (*flag=True*)
Enables use of parameter update in fp32.

Attributes

alpha
Alias to `self.hyperparam.alpha`

epoch = 0

eps
Alias to `self.hyperparam.eps`

lr
Alias to `self.hyperparam.lr`

t = 0

target = None

5.4.7 chainer.optimizers.RMSpropGraves

class `chainer.optimizers.RMSpropGraves` (*lr=0.0001*, *alpha=0.95*, *momentum=0.9*,
eps=0.0001)

Alex Graves's RMSprop.

See: <https://arxiv.org/abs/1308.0850>

Parameters

- **lr** (*float*) – Learning rate.
- **alpha** (*float*) – Exponential decay rate of the first and second order moments of the raw gradient.
- **momentum** (*float*) – Exponential decay rate of the first order moment of the adjusted gradient.
- **eps** (*float*) – Small value for the numerical stability.

Methods

add_hook (*hook*, *name=None*, *timing='auto'*)

Registers a hook function.

Hook function is typically called right after the gradient computation, though the timing depends on the optimization method, and the timing attribute.

Parameters

- **hook** (*callable*) – Hook function. If `hook.call_for_each_param` is true, this hook function is called for each parameter by passing the update rule and the parameter. Otherwise, this hook function is called only once each iteration by passing the optimizer.
- **name** (*str*) – Name of the registration. If omitted, `hook.name` is used by default.
- **timing** (*str*) – Specifies when the hook is called. If 'auto', the timing property of the hook will decide the timing. If 'pre', the hook will be called before any updates. If 'post', the hook will be called after any updates.

call_hooks (*timing='pre'*)

Invokes hook functions in registration order.

create_update_rule ()

Creates a new update rule object.

This method creates an update rule object. It is called by `setup()` to set up an update rule of each parameter. Each implementation of the gradient method should override this method to provide the default update rule implementation.

Returns Update rule object.

Return type `UpdateRule`

new_epoch ()

Starts a new epoch.

This method increments the `epoch` count. Note that if the optimizer depends on the epoch count, then user should call this method appropriately at the beginning of each epoch.

reallocate_cleared_grads ()

Reallocate gradients cleared by `cleargrad()`.

This method allocates arrays for all gradients which have `None`. This method is called before and after every optimizer hook. If an inheriting optimizer does not require this allocation, the optimizer can override this method with a blank function.

remove_hook (*name*)

Removes a hook function.

Parameters **name** (*str*) – Registered name of the hook function to remove.

serialize (*serializer*)

Serializes or deserializes the optimizer.

It only saves or loads the following things:

- Optimizer states
- Global states (`t` and `epoch`)

It does not saves nor loads the parameters of the target link. They should be separately saved or loaded.

Parameters **serializer** (`AbstractSerializer`) – Serializer or deserializer object.

set_loss_scale (*loss_scale*)

Sets loss scaling factor.

setup (*link*)

Sets a target link and initializes the optimizer states.

Given link is set to the `target` attribute. It also prepares the optimizer state dictionaries corresponding to all parameters in the link hierarchy. The existing states are discarded.

Parameters **link** (`Link`) – Target link object.

Returns The optimizer instance.

Note: As of v4.0.0, this function returns the optimizer instance itself so that you can instantiate and setup the optimizer in one line, e.g., `optimizer = SomeOptimizer().setup(link)`.

update (*lossfun=None, *args, **kwargs*)

Updates parameters based on a loss function or computed gradients.

This method runs in two ways.

- If `lossfun` is given, then it is used as a loss function to compute gradients.
- Otherwise, this method assumes that the gradients are already computed.

In both cases, the computed gradients are used to update parameters. The actual update routines are defined by the update rule of each parameter.

use_cleargrads (*use=True*)

Enables or disables use of `cleargrads()` in `update`.

Parameters *use* (*bool*) – If `True`, this function enables use of `cleargrads`. If `False`, disables use of `cleargrads` (`zerograds` is used).

Deprecated since version v2.0: Note that `update()` calls `cleargrads()` by default. `cleargrads()` is more efficient than `zerograds()`, so one does not have to call `use_cleargrads()`. This method remains for backward compatibility.

use_fp32_update (*flag=True*)

Enables use of parameter update in fp32.

Attributes

alpha

Alias to `self.hyperparam.alpha`

epoch = 0

eps

Alias to `self.hyperparam.eps`

lr

Alias to `self.hyperparam.lr`

momentum

Alias to `self.hyperparam.momentum`

t = 0

target = None

5.4.8 chainer.optimizers.SGD

class `chainer.optimizers.SGD` (*lr=0.01*)

Vanilla Stochastic Gradient Descent.

Parameters *lr* (*float*) – Learning rate.

Methods

add_hook (*hook*, *name=None*, *timing='auto'*)

Registers a hook function.

Hook function is typically called right after the gradient computation, though the timing depends on the optimization method, and the timing attribute.

Parameters

- **hook** (*callable*) – Hook function. If `hook.call_for_each_param` is true, this hook function is called for each parameter by passing the update rule and the parameter. Otherwise, this hook function is called only once each iteration by passing the optimizer.
- **name** (*str*) – Name of the registration. If omitted, `hook.name` is used by default.
- **timing** (*str*) – Specifies when the hook is called. If ‘auto’, the timing property of the hook will decide the timing. If ‘pre’, the hook will be called before any updates. If ‘post’, the hook will be called after any updates.

call_hooks (*timing='pre'*)

Invokes hook functions in registration order.

create_update_rule ()

Creates a new update rule object.

This method creates an update rule object. It is called by `setup()` to set up an update rule of each parameter. Each implementation of the gradient method should override this method to provide the default update rule implementation.

Returns Update rule object.

Return type *UpdateRule*

new_epoch ()

Starts a new epoch.

This method increments the *epoch* count. Note that if the optimizer depends on the epoch count, then user should call this method appropriately at the beginning of each epoch.

reallocate_cleared_grads ()

Reallocate gradients cleared by `cleargrad()`.

This method allocates arrays for all gradients which have `None`. This method is called before and after every optimizer hook. If an inheriting optimizer does not require this allocation, the optimizer can override this method with a blank function.

remove_hook (*name*)

Removes a hook function.

Parameters **name** (*str*) – Registered name of the hook function to remove.

serialize (*serializer*)

Serializes or deserializes the optimizer.

It only saves or loads the following things:

- Optimizer states
- Global states (*t* and *epoch*)

It does not save nor load the parameters of the target link. They should be separately saved or loaded.

Parameters **serializer** (*AbstractSerializer*) – Serializer or deserializer object.

set_loss_scale (*loss_scale*)

Sets loss scaling factor.

setup (*link*)

Sets a target link and initializes the optimizer states.

Given link is set to the *target* attribute. It also prepares the optimizer state dictionaries corresponding to all parameters in the link hierarchy. The existing states are discarded.

Parameters **link** (*Link*) – Target link object.

Returns The optimizer instance.

Note: As of v4.0.0, this function returns the optimizer instance itself so that you can instantiate and setup the optimizer in one line, e.g., `optimizer = SomeOptimizer().setup(link)`.

update (*lossfun=None, *args, **kwargs*)

Updates parameters based on a loss function or computed gradients.

This method runs in two ways.

- If `lossfun` is given, then it is used as a loss function to compute gradients.
- Otherwise, this method assumes that the gradients are already computed.

In both cases, the computed gradients are used to update parameters. The actual update routines are defined by the update rule of each parameter.

use_cleargrads (*use=True*)

Enables or disables use of `cleargrads()` in `update`.

Parameters `use` (*bool*) – If `True`, this function enables use of `cleargrads`. If `False`, disables use of `cleargrads` (`zerograds` is used).

Deprecated since version v2.0: Note that `update()` calls `cleargrads()` by default. `cleargrads()` is more efficient than `zerograds()`, so one does not have to call `use_cleargrads()`. This method remains for backward compatibility.

use_fp32_update (*flag=True*)

Enables use of parameter update in fp32.

Attributes

`epoch = 0`

`lr`

Alias to `self.hyperparam.lr`

`t = 0`

`target = None`

5.4.9 chainer.optimizers.SMORMS3

class `chainer.optimizers.SMORMS3` (*lr=0.001, eps=1e-16*)

Simon Funk's SMORMS3.

See <http://sifter.org/~simon/journal/20150420.html>.

Parameters

- `lr` (*float*) – Learning rate.
- `eps` (*float*) – Small value for the numerical stability.

Methods

add_hook (*hook*, *name=None*, *timing='auto'*)

Registers a hook function.

Hook function is typically called right after the gradient computation, though the timing depends on the optimization method, and the timing attribute.

Parameters

- **hook** (*callable*) – Hook function. If `hook.call_for_each_param` is true, this hook function is called for each parameter by passing the update rule and the parameter. Otherwise, this hook function is called only once each iteration by passing the optimizer.
- **name** (*str*) – Name of the registration. If omitted, `hook.name` is used by default.
- **timing** (*str*) – Specifies when the hook is called. If 'auto', the timing property of the hook will decide the timing. If 'pre', the hook will be called before any updates. If 'post', the hook will be called after any updates.

call_hooks (*timing='pre'*)

Invokes hook functions in registration order.

create_update_rule ()

Creates a new update rule object.

This method creates an update rule object. It is called by `setup()` to set up an update rule of each parameter. Each implementation of the gradient method should override this method to provide the default update rule implementation.

Returns Update rule object.

Return type *UpdateRule*

new_epoch ()

Starts a new epoch.

This method increments the `epoch` count. Note that if the optimizer depends on the epoch count, then user should call this method appropriately at the beginning of each epoch.

reallocate_cleared_grads ()

Reallocate gradients cleared by `cleargrad()`.

This method allocates arrays for all gradients which have `None`. This method is called before and after every optimizer hook. If an inheriting optimizer does not require this allocation, the optimizer can override this method with a blank function.

remove_hook (*name*)

Removes a hook function.

Parameters **name** (*str*) – Registered name of the hook function to remove.

serialize (*serializer*)

Serializes or deserializes the optimizer.

It only saves or loads the following things:

- Optimizer states
- Global states (*t* and *epoch*)

It does not save nor load the parameters of the target link. They should be separately saved or loaded.

Parameters **serializer** (*AbstractSerializer*) – Serializer or deserializer object.

set_loss_scale (*loss_scale*)

Sets loss scaling factor.

setup (*link*)

Sets a target link and initializes the optimizer states.

Given link is set to the *target* attribute. It also prepares the optimizer state dictionaries corresponding to all parameters in the link hierarchy. The existing states are discarded.

Parameters *link* (*Link*) – Target link object.

Returns The optimizer instance.

Note: As of v4.0.0, this function returns the optimizer instance itself so that you can instantiate and setup the optimizer in one line, e.g., `optimizer = SomeOptimizer().setup(link)`.

update (*lossfun=None, *args, **kwargs*)

Updates parameters based on a loss function or computed gradients.

This method runs in two ways.

- If *lossfun* is given, then it is used as a loss function to compute gradients.
- Otherwise, this method assumes that the gradients are already computed.

In both cases, the computed gradients are used to update parameters. The actual update routines are defined by the update rule of each parameter.

use_cleargrads (*use=True*)

Enables or disables use of *cleargrads()* in *update*.

Parameters *use* (*bool*) – If *True*, this function enables use of *cleargrads*. If *False*, disables use of *cleargrads* (*zerograds* is used).

Deprecated since version v2.0: Note that *update()* calls *cleargrads()* by default. *cleargrads()* is more efficient than *zerograds()*, so one does not have to call *use_cleargrads()*. This method remains for backward compatibility.

use_fp32_update (*flag=True*)

Enables use of parameter update in fp32.

Attributes

epoch = 0

eps

Alias to `self.hyperparam.eps`

lr

Alias to `self.hyperparam.lr`

t = 0

target = None

5.4.10 Optimizer base classes

<code>chainer.Optimizer</code>	Base class of all numerical optimizers.
<code>chainer.UpdateRule</code>	Base class of all update rules.
<code>chainer.optimizer.Hyperparameter</code>	Set of hyperparameter entries of an optimizer.
<code>chainer.GradientMethod</code>	Base class of all single gradient-based optimizers.

chainer.Optimizer

class chainer.Optimizer

Base class of all numerical optimizers.

This class provides basic features for all optimization methods. It optimizes parameters of a *target link*. The target link is registered via the `setup()` method, and then the `update()` method updates its parameters based on a given loss function.

Each optimizer implementation must be defined as a child class of Optimizer. It must override `update()` method.

If the optimizer is based on single gradient computation (like most first-order methods), then it should inherit `GradientMethod`, which adds some features dedicated for the first order methods, including the support of `UpdateRule`.

Optimizer instance also supports *hook functions*. Hook function is registered by the `add_hook()` method. Each hook function is called in registration order before of after the actual parameter update (configurable). If the hook function has an attribute `call_for_each_param` and its value is `True`, the hook function is used as a hook function of all update rules (i.e., it is invoked for every parameter by passing the corresponding update rule and the parameter).

Variables

- **`target`** – Target link object. It is set by the `setup()` method.
- **`t`** – Number of update steps. It must be incremented by the `update()` method.
- **`epoch`** – Current epoch. It is incremented by the `new_epoch()` method.

Methods

`add_hook(hook, name=None, timing='auto')`

Registers a hook function.

Hook function is typically called right after the gradient computation, though the timing depends on the optimization method, and the timing attribute.

Parameters

- **`hook`** (*callable*) – Hook function. If `hook.call_for_each_param` is `true`, this hook function is called for each parameter by passing the update rule and the parameter. Otherwise, this hook function is called only once each iteration by passing the optimizer.
- **`name`** (*str*) – Name of the registration. If omitted, `hook.name` is used by default.
- **`timing`** (*str*) – Specifies when the hook is called. If 'auto', the timing property of the hook will decide the timing. If 'pre', the hook will be called before any updates. If 'post', the hook will be called after any updates.

`call_hooks(timing='pre')`

Invokes hook functions in registration order.

new_epoch()

Starts a new epoch.

This method increments the `epoch` count. Note that if the optimizer depends on the epoch count, then user should call this method appropriately at the beginning of each epoch.

remove_hook(*name*)

Removes a hook function.

Parameters **name** (*str*) – Registered name of the hook function to remove.

serialize(*serializer*)

Serializes or deserializes the optimizer.

It only saves or loads the following things:

- Optimizer states
- Global states (`t` and `epoch`)

It does not save nor load the parameters of the target link. They should be separately saved or loaded.

Parameters **serializer** (*AbstractSerializer*) – Serializer or deserializer object.

set_loss_scale(*loss_scale*)

Sets loss scaling factor.

setup(*link*)

Sets a target link and initializes the optimizer states.

Given link is set to the `target` attribute. It also prepares the optimizer state dictionaries corresponding to all parameters in the link hierarchy. The existing states are discarded.

Parameters **link** (*Link*) – Target link object.

Returns The optimizer instance.

Note: As of v4.0.0, this function returns the optimizer instance itself so that you can instantiate and setup the optimizer in one line, e.g., `optimizer = SomeOptimizer().setup(link)`.

update(*lossfun=None, *args, **kws*)

Updates the parameters.

This method updates the parameters of the target link. The behavior of this method is different for the cases either `lossfun` is given or not.

If `lossfun` is given, this method typically clears the gradients, calls the loss function with given extra arguments, and calls the `backward()` method of its output to compute the gradients. The actual implementation might call `lossfun` more than once.

If `lossfun` is not given, then this method assumes that the gradients of all parameters are already computed. An implementation that requires multiple gradient computations might raise an error on this case.

In both cases, this method invokes the update procedure for all parameters.

Parameters

- **lossfun** (*callable*) – Loss function. You can specify one of loss functions from *built-in loss functions*, or your own loss function. It should not be an *loss functions with parameters* (i.e., *Link* instance). The function must accept arbitrary arguments and return one *Variable* object that represents the loss (or objective) value. Returned value must be a *Variable* derived from the input *Variable* object. `lossfun` can be omitted for single gradient-based methods. In this case, this method assumes gradient arrays computed.

- **kwds** (*args*,) – Arguments for the loss function.

Attributes

```
epoch = 0
t = 0
target = None
```

chainer.UpdateRule

class `chainer.UpdateRule` (*parent_hyperparam=None*)

Base class of all update rules.

Update rule is an object that implements how to update one parameter variable using the gradient of a loss function. This class provides the interface and the common features of any update rules.

An update rule can be set to a *Variable* object that represents a parameter array of a model. An *Optimizer* instance defines which parameters to update, and the update rule instance of each parameter defines how to update it.

Hook functions can be set to any update rule instance. The hook function is called just before or after any updates (configurable) in the order of registrations.

An implementation of update rule should override `update_core()` or its device-dependent variants (i.e., `update_core_cpu()` and `update_core_gpu()`).

The state (e.g. a moving average of the gradient) of the update rule is stored into the state dictionary. An implementation of update rule using state should also override `init_state()` to initialize the state at the first update. The values of the state dictionary are automatically copied to the appropriate device before the update based on the data and grad arrays.

Parameters `parent_hyperparam` (*Hyperparameter*) – Hyperparameter that provides the default values.

Variables

- **enabled** (*bool*) – Flag to configure if this update rule is active. If the update rule is not active (i.e., `enabled = False`), the `update()` method does not update the parameter.
- **hyperparam** (*Hyperparameter*) – Hyperparameter of the update rule.
- **t** (*int*) – Number of updates made by this update rule.

Methods

add_hook (*hook*, *name=None*, *timing='auto'*)

Adds a hook function.

The hook function is called before or after any updates (see the timing attribute).

Parameters

- **hook** (*callable*) – Hook function to be added. It takes two arguments: the update rule object and the parameter variable.
- **name** (*str*) – Name of the hook function. The name attribute of the hook function is used by default.

- **timing** (*str*) – Specifies when the hook is called. If ‘auto’, the timing property of the hook will decide the timing. If ‘pre’, the hook will be called before any updates. If ‘post’, the hook will be called after any updates. If ‘auto’ and the timing property of the hook is not available, timing will default to ‘pre’.

init_state (*param*)

Initializes the state.

Any implementations that use the state should override this method. This method is called at the first update.

Parameters **param** (*Variable*) – Parameter variable. It can be used to extract the shape and the data type of the parameter.

remove_hook (*name*)

Removes the specified hook function.

Parameters **name** (*str*) – Name of the hook function to be removed. The hook function registered with this name will be removed.

serialize (*serializer*)

Serializes the update rule state.

Be careful that this method only saves/loads the state of the update rule. The parameters of the target link is not saved/loaded by this method, and so you need to serialize the target link separately if you want to fully recover the training state including parameters.

Parameters **serializer** (*AbstractSerializer*) – Serializer object.

update (*param*)

Invokes hook functions and updates the parameter.

Parameters **param** (*Variable*) – Variable to be updated.

update_core (*param*)

Updates the parameter.

Implementation of UpdateRule should override this method or both of `update_core_cpu()` and `update_core_gpu()`.

Parameters **param** (*Variable*) – Variable to be updated.

update_core_cpu (*param*)

Updates the parameter on CPU.

See `update_core()` for details.

Parameters **param** (*Variable*) – Variable to be updated.

update_core_gpu (*param*)

Updates the parameter on GPU.

See `update_core()` for details.

Parameters **param** (*Variable*) – Variable to be updated.

use_fp32_update (*flag=True*)

Enables use of parameter update in fp32.

This method enables use of parameter update in fp32. When it is enabled and data type of original parameter variable is fp16, fp32 copy of parameter variable is automatically created and retained at `self.fp32_param`. And the parameter is update in fp32 in the following way.

1. copies the grad of original parameter variable to the grad of fp32 parameter variable, converting its data type from fp16 to fp32.

2. updates the parameter in fp32.
3. copies the data of fp32 parameter variable to the data of original parameter variable, converting its data type from fp32 to fp16.

See meth:*update* for details.

Attributes

state

State dictionary.

chainer.optimizer.Hyperparameter

class chainer.optimizer.Hyperparameter (parent=None)

Set of hyperparameter entries of an optimizer.

This is a utility class to provide a set of hyperparameter entries for update rules and an optimizer. Each entry can be set as an attribute of a hyperparameter object.

A hyperparameter object can hold a reference to its parent hyperparameter object. When an attribute does not exist in the child hyperparameter, it automatically refers to the parent. We typically set the hyperparameter of the gradient method as the parent of the hyperparameter of each update rule. It enables us to centralize the management of hyperparameters (e.g. we can change the learning rate of all update rules just by modifying the hyperparameter of the central optimizer object), while users can freely customize the hyperparameter of each update rule if needed.

Parameters **parent** ([Hyperparameter](#)) – Parent hyperparameter.

Methods

get_dict ()

Converts the hyperparameter into a dictionary.

Returns Dictionary containing all entries that can be referred by this hyperparameter object.

Attributes

parent

Parent hyperparameter object.

chainer.GradientMethod

class chainer.GradientMethod

Base class of all single gradient-based optimizers.

This is an extension of the [Optimizer](#) class. Typical gradient methods that just require the gradient at the current parameter vector on an update can be implemented as its child class.

This class uses [UpdateRule](#) to manage the update rule of each parameter. A child class of GradientMethod should override [create_update_rule\(\)](#) to create the default update rule of each parameter.

This class also provides `hyperparam`, which is the hyperparameter used as the default configuration of each update rule. All built-in gradient method implementations also provide proxy properties that act as aliases to

the attributes of `hyperparam`. It is recommended to provide such an alias to each attribute. It can be done by only adding one line for each attribute using `HyperparameterProxy`.

Variables `hyperparam` (`Hyperparameter`) – The hyperparameter of the gradient method. It is used as the default configuration of each update rule (i.e., the hyperparameter of each update rule refers this hyperparameter as its parent).

Methods

add_hook (*hook*, *name=None*, *timing='auto'*)

Registers a hook function.

Hook function is typically called right after the gradient computation, though the timing depends on the optimization method, and the timing attribute.

Parameters

- **hook** (*callable*) – Hook function. If `hook.call_for_each_param` is `true`, this hook function is called for each parameter by passing the update rule and the parameter. Otherwise, this hook function is called only once each iteration by passing the optimizer.
- **name** (*str*) – Name of the registration. If omitted, `hook.name` is used by default.
- **timing** (*str*) – Specifies when the hook is called. If `'auto'`, the timing property of the hook will decide the timing. If `'pre'`, the hook will be called before any updates. If `'post'`, the hook will be called after any updates.

call_hooks (*timing='pre'*)

Invokes hook functions in registration order.

create_update_rule ()

Creates a new update rule object.

This method creates an update rule object. It is called by `setup()` to set up an update rule of each parameter. Each implementation of the gradient method should override this method to provide the default update rule implementation.

Returns Update rule object.

Return type `UpdateRule`

new_epoch ()

Starts a new epoch.

This method increments the `epoch` count. Note that if the optimizer depends on the epoch count, then user should call this method appropriately at the beginning of each epoch.

reallocate_cleared_grads ()

Reallocate gradients cleared by `cleargrad()`.

This method allocates arrays for all gradients which have `None`. This method is called before and after every optimizer hook. If an inheriting optimizer does not require this allocation, the optimizer can override this method with a blank function.

remove_hook (*name*)

Removes a hook function.

Parameters **name** (*str*) – Registered name of the hook function to remove.

serialize (*serializer*)

Serializes or deserializes the optimizer.

It only saves or loads the following things:

- Optimizer states
- Global states (`t` and `epoch`)

It does not save nor load the parameters of the target link. They should be separately saved or loaded.

Parameters `serializer` (`AbstractSerializer`) – Serializer or deserializer object.

set_loss_scale (`loss_scale`)

Sets loss scaling factor.

setup (`link`)

Sets a target link and initializes the optimizer states.

Given link is set to the `target` attribute. It also prepares the optimizer state dictionaries corresponding to all parameters in the link hierarchy. The existing states are discarded.

Parameters `link` (`Link`) – Target link object.

Returns The optimizer instance.

Note: As of v4.0.0, this function returns the optimizer instance itself so that you can instantiate and setup the optimizer in one line, e.g., `optimizer = SomeOptimizer().setup(link)`.

update (`lossfun=None, *args, **kwargs`)

Updates parameters based on a loss function or computed gradients.

This method runs in two ways.

- If `lossfun` is given, then it is used as a loss function to compute gradients.
- Otherwise, this method assumes that the gradients are already computed.

In both cases, the computed gradients are used to update parameters. The actual update routines are defined by the update rule of each parameter.

use_cleargrads (`use=True`)

Enables or disables use of `cleargrads()` in `update`.

Parameters `use` (`bool`) – If `True`, this function enables use of `cleargrads`. If `False`, disables use of `cleargrads` (`zerograds` is used).

Deprecated since version v2.0: Note that `update()` calls `cleargrads()` by default. `cleargrads()` is more efficient than `zerograds()`, so one does not have to call `use_cleargrads()`. This method remains for backward compatibility.

use_fp32_update (`flag=True`)

Enables use of parameter update in fp32.

Attributes

`epoch = 0`

`t = 0`

`target = None`

5.4.11 Hook functions

<code>chainer.optimizer_hooks.WeightDecay</code>	Optimizer/UpdateRule hook function for weight decay regularization.
<code>chainer.optimizer_hooks.Lasso</code>	Optimizer/UpdateRule hook function for Lasso regularization.
<code>chainer.optimizer_hooks.GradientClipping</code>	Optimizer hook function for gradient clipping.
<code>chainer.optimizer_hooks.GradientHardClipping</code>	Optimizer/UpdateRule hook function for gradient clipping.
<code>chainer.optimizer_hooks.GradientNoise</code>	Optimizer/UpdateRule hook function for adding gradient noise.
<code>chainer.optimizer_hooks.GradientLARS</code>	Optimizer/UpdateRule hook function for layer wise adaptive rate scaling.

chainer.optimizer_hooks.WeightDecay

class `chainer.optimizer_hooks.WeightDecay` (*rate*)

Optimizer/UpdateRule hook function for weight decay regularization.

This hook function adds a scaled parameter to the corresponding gradient. It can be used as a regularization.

Parameters *rate* (*float*) – Coefficient for the weight decay.

Variables

- *rate* (*float*) – Coefficient for the weight decay.
- *timing* (*string*) – Specifies when this hook should be called by the Optimizer/UpdateRule. Valid values are ‘pre’ (before any updates) and ‘post’ (after any updates).
- *call_for_each_param* (*bool*) – Specifies if this hook is called for each parameter (True) or only once (False) by an optimizer to which this hook is registered. This function does not expect users to switch the value from default one, which is *True*.

New in version 4.0.0: The *timing* parameter.

Methods

`__call__` (*rule*, *param*)
Call self as a function.

Attributes

```
call_for_each_param = True
name = 'WeightDecay'
timing = 'pre'
```

chainer.optimizer_hooks.Lasso

class `chainer.optimizer_hooks.Lasso` (*rate*)

Optimizer/UpdateRule hook function for Lasso regularization.

This hook function adds a scaled parameter to the sign of each weight. It can be used as a regularization.

Parameters `rate` (*float*) – Coefficient for the weight decay.

Variables

- `rate` (*float*) – Coefficient for the weight decay.
- `timing` (*string*) – Specifies when this hook should be called by the Optimizer/UpdateRule. Valid values are ‘pre’ (before any updates) and ‘post’ (after any updates).
- `call_for_each_param` (*bool*) – Specifies if this hook is called for each parameter (`True`) or only once (`False`) by an optimizer to which this hook is registered. This function does not expect users to switch the value from default one, which is `True`.

New in version 4.0.0: The `timing` parameter.

Methods

`__call__` (*rule, param*)
Call self as a function.

Attributes

```
call_for_each_param = True
name = 'Lasso'
timing = 'pre'
```

chainer.optimizer_hooks.GradientClipping

class `chainer.optimizer_hooks.GradientClipping` (*threshold*)

Optimizer hook function for gradient clipping.

This hook function scales all gradient arrays to fit to the defined L2 norm threshold.

Parameters `threshold` (*float*) – L2 norm threshold.

Variables

- `threshold` (*float*) – L2 norm threshold of gradient norm.
- `timing` (*string*) – Specifies when this hook should be called by the Optimizer/UpdateRule. Valid values are ‘pre’ (before any updates) and ‘post’ (after any updates).

New in version 4.0.0: The `timing` parameter.

Methods

`__call__` (*opt*)
Call self as a function.

Attributes

```
name = 'GradientClipping'
timing = 'pre'
```

chainer.optimizer_hooks.GradientHardClipping

class chainer.optimizer_hooks.**GradientHardClipping**(*lower_bound*, *upper_bound*)
Optimizer/UpdateRule hook function for gradient clipping.

This hook function clips all gradient arrays to be within a lower and upper bound.

Parameters

- **lower_bound** (*float*) – The lower bound of the gradient value.
- **upper_bound** (*float*) – The upper bound of the gradient value.

Variables

- **lower_bound** (*float*) – The lower bound of the gradient value.
- **upper_bound** (*float*) – The upper bound of the gradient value.
- **timing** (*string*) – Specifies when this hook should be called by the Optimizer/UpdateRule. Valid values are ‘pre’ (before any updates) and ‘post’ (after any updates).
- **call_for_each_param** (*bool*) – Specifies if this hook is called for each parameter (True) or only once (False) by an optimizer to which this hook is registered. This function does not expect users to switch the value from default one, which is *True*.

New in version 4.0.0: The *timing* parameter.

Methods

```
__call__(rule, param)
    Call self as a function.
```

Attributes

```
call_for_each_param = True
name = 'GradientHardClipping'
timing = 'pre'
```

chainer.optimizer_hooks.GradientNoise

class chainer.optimizer_hooks.**GradientNoise**(*eta*, *noise_func*=<function *exponential_decay_noise*>)
Optimizer/UpdateRule hook function for adding gradient noise.

This hook function simply adds noise generated by the *noise_func* to the gradient. By default it adds time-dependent annealed Gaussian noise to the gradient at every training step:

$$g_t \leftarrow g_t + N(0, \sigma_t^2)$$

where

$$\sigma_t^2 = \frac{\eta}{(1+t)^\gamma}$$

with η selected from $\{0.01, 0.3, 1.0\}$ and $\gamma = 0.55$.

Parameters

- **eta** (*float*) – Parameter that defines the scale of the noise, which for the default noise function is recommended to be either 0.01, 0.3 or 1.0.
- **noise_func** (*function*) – Noise generating function which by default is given by [Adding Gradient Noise Improves Learning for Very Deep Networks](#).

Variables

- **timing** (*string*) – Specifies when this hook should be called by the Optimizer/UpdateRule. Valid values are ‘pre’ (before any updates) and ‘post’ (after any updates).
- **call_for_each_param** (*bool*) – Specifies if this hook is called for each parameter (True) or only once (False) by an optimizer to which this hook is registered. This function does not expect users to switch the value from default one, which is *True*.

New in version 4.0.0: The *timing* parameter.

Methods

__call__ (*rule, param*)
Call self as a function.

Attributes

```
call_for_each_param = True
name = 'GradientNoise'
timing = 'pre'
```

chainer.optimizer_hooks.GradientLARS

class chainer.optimizer_hooks.**GradientLARS** (*threshold=0.01, weight_decay=0.0, eps=1e-09*)

Optimizer/UpdateRule hook function for layer wise adaptive rate scaling.

See: [Large Batch Training of Convolutional Networks](#).

See: [Convergence Analysis of Gradient Descent Algorithms with Proportional Updates](#).

This hook function scales all gradient arrays to fit to the weight norm.

In <https://arxiv.org/abs/1708.03888>,

$$\begin{aligned} v_{t+1} &= m * v_t + \gamma * \lambda * (\nabla L(w_t) + \beta w_t), \\ w_{t+1} &= w_t - v_{t+1}, \end{aligned}$$

where

- γ : learning_rate

- m : momentum
- β : weight_decay
- η : lars_coeficient
- λ : local_lr = $\eta * \frac{\|w_t\|}{\|\nabla L(w_t)\| + \beta * \|w_t\|}$.

As lr in `chainer.optimizers.SGD` or `chainer.optimizers.MomentumSGD` corresponds to $\gamma * \eta$, we define `clip_rate` as $\frac{\|w_t\|}{\|\nabla L(w_t)\| + \beta * \|w_t\|}$ and reformulate the aforementioned formula as: $v_{t+1} = m * v_t + lr * clip_rate * (\nabla L(w_t) + \beta w_t)$ and implement in this way. So you do not set `lars_coeficient`.

Parameters

- **threshold** (*float*) – If weight norm is more than threshold, this function scales all gradient arrays to fit weight norm. (See <<https://arxiv.org/abs/1801.03137>>)
- **weight_decay** (*float*) – Coefficient for the weight decay.
- **eps** (*float*) – Small value for the numerical stability. (See <<https://arxiv.org/abs/1801.03137>>)

Variables

- **threshold** (*float*) – If weight norm is more than threshold, this function scales all gradient arrays to fit weight norm. (See <<https://arxiv.org/abs/1801.03137>>)
- **weight_decay** (*float*) – Coefficient for the weight decay.
- **eps** (*float*) – Small value for the numerical stability. (See <<https://arxiv.org/abs/1801.03137>>)
- **timing** (*string*) – Specifies when this hook should be called by the Optimizer/UpdateRule. Valid values are ‘pre’ (before any updates) and ‘post’ (after any updates).
- **call_for_each_param** (*bool*) – Specifies if this hook is called for each parameter (`True`) or only once (`False`) by an optimizer to which this hook is registered. This function does not expect users to switch the value from default one, which is `True`.

Methods

`__call__` (*rule, param*)
Call self as a function.

Attributes

```
call_for_each_param = True
name = 'GradientLARS'
timing = 'pre'
```

5.5 Weight Initializers

Weight initializers are used to initialize arrays. They destructively modify the content of `numpy.ndarray` or `cupy.ndarray`. Typically, weight initializers are passed to `Links` to initialize their weights and biases.

A weight initializer can be any of the following objects.

- `chainer.Initializer` class instance.
- Python or NumPy scalar or `numpy.ndarray`.
- A callable that takes an array (`numpy.ndarray` or `cupy.ndarray`) and feeds the initial data into it.
- `None`, in which case *the default initializer* is used. Unless explicitly specified, it is `LeCunNormal` with scale value 1.

5.5.1 Base class

<code>chainer.Initializer</code>	Initializes array.
----------------------------------	--------------------

`chainer.Initializer`

class `chainer.Initializer` (*dtype=None*)

Initializes array.

It initializes the given array.

Variables `dtype` – Data type specifier. It is for type check in `__call__` function.

Methods

`__call__` (*array*)

Initializes given array.

This method destructively changes the value of array. The derived class is required to implement this method. The algorithms used to make the new values depend on the concrete derived classes.

Parameters `array` (*`numpy.ndarray` or `cupy.ndarray`*) – An array to be initialized by this initializer.

5.5.2 Concrete initializers

<code>chainer.initializers.Identity</code>	Initializes array with the identity matrix.
<code>chainer.initializers.Constant</code>	Initializes array with constant value.
<code>chainer.initializers.Zero</code>	Initializes array to all-zero.
<code>chainer.initializers.One</code>	Initializes array to all-one.
<code>chainer.initializers.NaN</code>	Initializes array to all-NaN.
<code>chainer.initializers.Normal</code>	Initializes array with a normal distribution.
<code>chainer.initializers.LeCunNormal</code>	Initializes array with scaled Gaussian distribution.
<code>chainer.initializers.GlorotNormal</code>	Initializes array with scaled Gaussian distribution.
<code>chainer.initializers.HeNormal</code>	Initializes array with scaled Gaussian distribution.
<code>chainer.initializers.Orthogonal</code>	Initializes array with an orthogonal system.
<code>chainer.initializers.Uniform</code>	Initializes array with a scaled uniform distribution.
<code>chainer.initializers.LeCunUniform</code>	Initializes array with a scaled uniform distribution.
<code>chainer.initializers.GlorotUniform</code>	Initializes array with a scaled uniform distribution.
<code>chainer.initializers.HeUniform</code>	Initializes array with scaled uniform distribution.

chainer.initializers.Identity

class chainer.initializers.**Identity** (*scale=1.0, dtype=None*)

Initializes array with the identity matrix.

It initializes the given array with the constant multiple of the identity matrix. Note that arrays to be passed must be 2D squared matrices.

Variables **scale** (*scalar*) – A constant to be multiplied to identity matrices.

Methods

__call__ (*array*)

Initializes given array.

This method destructively changes the value of array. The derived class is required to implement this method. The algorithms used to make the new values depend on the concrete derived classes.

Parameters **array** (*numpy.ndarray or cupy.ndarray*) – An array to be initialized by this initializer.

chainer.initializers.Constant

class chainer.initializers.**Constant** (*fill_value, dtype=None*)

Initializes array with constant value.

Variables

- **fill_value** (*scalar or numpy.ndarray or cupy.ndarray*) – A constant to be assigned to the initialized array. Broadcast is allowed on this assignment.
- **dtype** – Data type specifier.

Methods

__call__ (*array*)

Initializes given array.

This method destructively changes the value of array. The derived class is required to implement this method. The algorithms used to make the new values depend on the concrete derived classes.

Parameters **array** (*numpy.ndarray or cupy.ndarray*) – An array to be initialized by this initializer.

Attributes

fill_value = None

chainer.initializers.Zero

class chainer.initializers.**Zero** (*dtype=None*)

Initializes array to all-zero.

Variables **dtype** – Data type specifier.

Methods

`__call__(array)`

Initializes given array.

This method destructively changes the value of array. The derived class is required to implement this method. The algorithms used to make the new values depend on the concrete derived classes.

Parameters **array** (*numpy.ndarray* or *cupy.ndarray*) – An array to be initialized by this initializer.

Attributes

`fill_value = 0.0`

chainer.initializers.One

class `chainer.initializers.One(dtype=None)`

Initializes array to all-one.

Variables **dtype** – Data type specifier.

Methods

`__call__(array)`

Initializes given array.

This method destructively changes the value of array. The derived class is required to implement this method. The algorithms used to make the new values depend on the concrete derived classes.

Parameters **array** (*numpy.ndarray* or *cupy.ndarray*) – An array to be initialized by this initializer.

Attributes

`fill_value = 1.0`

chainer.initializers.NaN

class `chainer.initializers.NaN(dtype=None)`

Initializes array to all-NaN.

Variables **dtype** – Data type specifier.

Methods

`__call__(array)`

Initializes given array.

This method destructively changes the value of array. The derived class is required to implement this method. The algorithms used to make the new values depend on the concrete derived classes.

Parameters **array** (*numpy.ndarray* or *cupy.ndarray*) – An array to be initialized by this initializer.

Attributes

fill_value = nan

chainer.initializers.Normal

class chainer.initializers.**Normal** (*scale=0.05, dtype=None*)

Initializes array with a normal distribution.

Each element of the array is initialized by the value drawn independently from Gaussian distribution whose mean is 0, and standard deviation is *scale*.

Parameters

- **scale** (*float*) – Standard deviation of Gaussian distribution.
- **dtype** – Data type specifier.

Methods

__call__ (*array*)

Initializes given array.

This method destructively changes the value of array. The derived class is required to implement this method. The algorithms used to make the new values depend on the concrete derived classes.

Parameters **array** (*numpy.ndarray* or *cupy.ndarray*) – An array to be initialized by this initializer.

chainer.initializers.LeCunNormal

class chainer.initializers.**LeCunNormal** (*scale=1.0, dtype=None*)

Initializes array with scaled Gaussian distribution.

Each element of the array is initialized by the value drawn independently from Gaussian distribution whose mean is 0, and standard deviation is $scale \times \sqrt{\frac{1}{fan_{in}}}$, where fan_{in} is the number of input units.

Reference: LeCun 98, Efficient Backprop <http://yann.lecun.com/exdb/publis/pdf/lecun-98b.pdf>

Parameters

- **scale** (*float*) – A constant that determines the scale of the standard deviation.
- **dtype** – Data type specifier.

Methods

__call__ (*array*)

Initializes given array.

This method destructively changes the value of array. The derived class is required to implement this method. The algorithms used to make the new values depend on the concrete derived classes.

Parameters **array** (*numpy.ndarray* or *cupy.ndarray*) – An array to be initialized by this initializer.

chainer.initializers.GlorotNormal

class chainer.initializers.**GlorotNormal** (*scale=1.0, dtype=None*)

Initializes array with scaled Gaussian distribution.

Each element of the array is initialized by the value drawn independently from Gaussian distribution whose mean is 0, and standard deviation is $scale \times \sqrt{\frac{2}{fan_{in} + fan_{out}}}$, where fan_{in} and fan_{out} are the number of input and output units, respectively.

Reference: Glorot & Bengio, AISTATS 2010

Parameters

- **scale** (*float*) – A constant that determines the scale of the standard deviation.
- **dtype** – Data type specifier.

Methods

__call__ (*array*)

Initializes given array.

This method destructively changes the value of array. The derived class is required to implement this method. The algorithms used to make the new values depend on the concrete derived classes.

Parameters **array** (*numpy.ndarray* or *cupy.ndarray*) – An array to be initialized by this initializer.

chainer.initializers.HeNormal

class chainer.initializers.**HeNormal** (*scale=1.0, dtype=None, fan_option='fan_in'*)

Initializes array with scaled Gaussian distribution.

Each element of the array is initialized by the value drawn independently from Gaussian distribution whose mean is 0, and standard deviation is $scale \times \sqrt{\frac{2}{fan}}$. If `fan_option == 'fan_in'`, *fan* is the number of input units. If `fan_option == 'fan_out'`, *fan* is the number of output units.

Reference: He et al., <https://arxiv.org/abs/1502.01852>

Parameters

- **scale** (*float*) – A constant that determines the scale of the standard deviation.
- **dtype** – Data type specifier.
- **fan_option** (*{'fan_in', 'fan_out'}*) – Decides how to compute the standard deviation. The default value is 'fan_in'.

Methods

__call__ (*array*)

Initializes given array.

This method destructively changes the value of array. The derived class is required to implement this method. The algorithms used to make the new values depend on the concrete derived classes.

Parameters **array** (*numpy.ndarray* or *cupy.ndarray*) – An array to be initialized by this initializer.

chainer.initializers.Orthogonal

class chainer.initializers.**Orthogonal** (*scale=1.1, dtype=None*)

Initializes array with an orthogonal system.

This initializer first makes a matrix of the same shape as the array to be initialized whose elements are drawn independently from standard Gaussian distribution. Next, it applies Singular Value Decomposition (SVD) to the matrix. Then, it initializes the array with either side of resultant orthogonal matrices, depending on the shape of the input array. Finally, the array is multiplied by the constant *scale*.

If the *ndim* of the input array is more than 2, we consider the array to be a matrix by concatenating all axes except the first one.

The number of vectors consisting of the orthogonal system (i.e. first element of the shape of the array) must be equal to or smaller than the dimension of each vector (i.e. second element of the shape of the array).

Variables

- **scale** (*float*) – A constant to be multiplied by.
- **dtype** – Data type specifier.

Reference: Saxe et al., <https://arxiv.org/abs/1312.6120>

Methods

__call__ (*array*)

Initializes given array.

This method destructively changes the value of array. The derived class is required to implement this method. The algorithms used to make the new values depend on the concrete derived classes.

Parameters **array** (*numpy.ndarray* or *cupy.ndarray*) – An array to be initialized by this initializer.

chainer.initializers.Uniform

class chainer.initializers.**Uniform** (*scale=0.05, dtype=None*)

Initializes array with a scaled uniform distribution.

Each element of the array is initialized by the value drawn independently from uniform distribution $[-scale, scale]$.

Variables

- **scale** (*float*) – A constant that determines the scale of the uniform distribution.
- **dtype** – Data type specifier.

Methods

`__call__(array)`

Initializes given array.

This method destructively changes the value of array. The derived class is required to implement this method. The algorithms used to make the new values depend on the concrete derived classes.

Parameters **array** (*numpy.ndarray* or *cupy.ndarray*) – An array to be initialized by this initializer.

chainer.initializers.LeCunUniform

class `chainer.initializers.LeCunUniform` (*scale=1.0, dtype=None*)

Initializes array with a scaled uniform distribution.

Each element of the array is initialized by the value drawn independently from uniform distribution $[-s, s]$ where $s = scale \times \sqrt{\frac{3}{fan_{in}}}$. Here fan_{in} is the number of input units.

Reference: LeCun 98, Efficient Backprop <http://yann.lecun.com/exdb/publis/pdf/lecun-98b.pdf>

Variables

- **scale** (*float*) – A constant that determines the scale of the uniform distribution.
- **dtype** – Data type specifier.

Methods

`__call__(array)`

Initializes given array.

This method destructively changes the value of array. The derived class is required to implement this method. The algorithms used to make the new values depend on the concrete derived classes.

Parameters **array** (*numpy.ndarray* or *cupy.ndarray*) – An array to be initialized by this initializer.

chainer.initializers.GlorotUniform

class `chainer.initializers.GlorotUniform` (*scale=1.0, dtype=None*)

Initializes array with a scaled uniform distribution.

Each element of the array is initialized by the value drawn independently from uniform distribution $[-s, s]$ where $s = scale \times \sqrt{\frac{6}{fan_{in} + fan_{out}}}$. Here, fan_{in} and fan_{out} are the number of input and output units, respectively.

Variables

- **scale** (*float*) – A constant that determines the scale of the uniform distribution.
- **dtype** – Data type specifier.

Methods

`__call__(array)`

Initializes given array.

This method destructively changes the value of array. The derived class is required to implement this method. The algorithms used to make the new values depend on the concrete derived classes.

Parameters **array** (*numpy.ndarray* or *cupy.ndarray*) – An array to be initialized by this initializer.

chainer.initializers.HeUniform

class `chainer.initializers.HeUniform(scale=1.0, dtype=None)`

Initializes array with scaled uniform distribution.

Each element of the array is initialized by the value drawn independently from uniform distribution $[-s, s]$ where $s = scale \times \sqrt{\frac{6}{fan_{in}}}$. Here, fan_{in} is the number of input units.

Variables

- **scale** (*float*) – A constant that determines the scale of the uniform distribution.
- **dtype** – Data type specifier.

Methods

`__call__(array)`

Initializes given array.

This method destructively changes the value of array. The derived class is required to implement this method. The algorithms used to make the new values depend on the concrete derived classes.

Parameters **array** (*numpy.ndarray* or *cupy.ndarray*) – An array to be initialized by this initializer.

5.5.3 Helper function

<code>chainer.initializers.generate_array</code>	Return initialized array.
--	---------------------------

chainer.initializers.generate_array

`chainer.initializers.generate_array(initializer, shape, xp)`

Return initialized array.

The algorithms used to make the new values depend on the concrete derived classes. The dtype of a generated array depends on `initializer.dtype`.

Parameters

- **initializer** – A callable object that takes *numpy.ndarray* or *cupy.ndarray* and edits its value.
- **shape** (*tuple*) – Shape of a return array.
- **xp** (*module*) – *cupy* or *numpy*.

Returns An initialized array.

Return type `numpy.ndarray` or `cupy.ndarray`

5.6 Training Tools

Chainer provides a standard implementation of the training loops under the `chainer.training` module. It is built on top of many other core features of Chainer, including Variable and Function, Link/Chain/ChainList, Optimizer, Dataset, and Reporter/Summary. Compared to the training loop abstraction of other machine learning tool kits, Chainer's training framework aims at maximal flexibility, while keeps the simplicity for the typical usages. Most components are pluggable, and users can overwrite the definition.

The core of the training loop abstraction is `Trainer`, which implements the training loop itself. The training loop consists of two parts: one is `Updater`, which actually updates the parameters to train, and the other is `Extension` for arbitrary functionalities other than the parameter update.

Updater and some extensions use `chainer.dataset` and `Iterator` to scan the datasets and load mini-batches. The trainer also uses `Reporter` to collect the observed values, and some extensions use `DictSummary` to accumulate them and computes the statistics.

You can find many examples for the usage of this training utilities from the official examples. You can also search the extension implementations from `Extensions`.

5.6.1 Trainer

`chainer.training.Trainer`

The standard training loop in Chainer.

`chainer.training.Trainer`

class `chainer.training.Trainer` (*updater*, *stop_trigger=None*, *out='result'*, *extensions=None*)

The standard training loop in Chainer.

Trainer is an implementation of a training loop. Users can invoke the training by calling the `run()` method.

Each iteration of the training loop proceeds as follows.

- Update of the parameters. It includes the mini-batch loading, forward and backward computations, and an execution of the update formula. These are all done by the update object held by the trainer.
- Invocation of trainer extensions in the descending order of their priorities. A trigger object is attached to each extension, and it decides at each iteration whether the extension should be executed. Trigger objects are callable objects that take the trainer object as the argument and return a boolean value indicating whether the extension should be called or not.

Extensions are callable objects that take the trainer object as the argument. There are three ways to define custom extensions: inheriting the `Extension` class, decorating functions by `make_extension()`, and defining any callable including lambda functions. See `Extension` for more details on custom extensions and how to configure them.

Users can register extensions to the trainer by calling the `extend()` method, where some configurations can be added.

- Trigger object, which is also explained above. In most cases, `IntervalTrigger` is used, in which case users can simply specify a tuple of the interval length and its unit, like `(1000, 'iteration')` or `(1, 'epoch')`.

- The order of execution of extensions is determined by their priorities. Extensions of higher priorities are invoked earlier. There are three standard values for the priorities:
 - `PRIORITY_WRITER`. This is the priority for extensions that write some records to the `observation` dictionary. It includes cases that the extension directly adds values to the `observation` dictionary, or the extension uses the `chainer.report()` function to report values to the `observation` dictionary.
 - `PRIORITY_EDITOR`. This is the priority for extensions that edit the `observation` dictionary based on already reported values.
 - `PRIORITY_READER`. This is the priority for extensions that only read records from the `observation` dictionary. This is also suitable for extensions that do not use the `observation` dictionary at all.

The current state of the trainer object and objects handled by the trainer can be serialized through the standard serialization protocol of Chainer. It enables us to easily suspend and resume the training loop.

Note: The serialization does not recover everything of the training loop. It only recovers the states which change over the training (e.g. parameters, optimizer states, the batch iterator state, extension states, etc.). You must initialize the objects correctly before deserializing the states.

On the other hand, it means that users can change the settings on deserialization. For example, the exit condition can be changed on the deserialization, so users can train the model for some iterations, suspend it, and then resume it with larger number of total iterations.

During the training, it also creates a `Reporter` object to store observed values on each update. For each iteration, it creates a fresh `observation` dictionary and stores it in the `observation` attribute.

Links of the target model of each optimizer are registered to the reporter object as observers, where the name of each observer is constructed as the format `<optimizer name><link name>`. The link name is given by the `chainer.Link.namedlink()` method, which represents the path to each link in the hierarchy. Other observers can be registered by accessing the reporter object via the `reporter` attribute.

The default trainer is *plain*, i.e., it does not contain any extensions.

Parameters

- **updater** (`Updater`) – Updater object. It defines how to update the models.
- **stop_trigger** – Trigger that determines when to stop the training loop. If it is not callable, it is passed to `IntervalTrigger`.
- **out** – Output directory.
- **extensions** – Extensions registered to the trainer.

Variables

- **updater** – The updater object for this trainer.
- **stop_trigger** – Trigger that determines when to stop the training loop. The training loop stops at the iteration on which this trigger returns `True`.
- **observation** – Observation of values made at the last update. See the `Reporter` class for details.
- **out** – Output directory.
- **reporter** – Reporter object to report observed values.

Methods

extend (*extension*, *name=None*, *trigger=None*, *priority=None*, ***kwargs*)

Registers an extension to the trainer.

Extension is a callable object which is called after each update unless the corresponding trigger object decides to skip the iteration. The order of execution is determined by priorities: extensions with higher priorities are called earlier in each iteration. Extensions with the same priority are invoked in the order of registrations.

If two or more extensions with the same name are registered, suffixes are added to the names of the second to last extensions. The suffix is `_N` where `N` is the ordinal of the extensions.

See *Extension* for the interface of extensions.

Parameters

- **extension** – Extension to register.
- **name** (*str*) – Name of the extension. If it is omitted, the `default_name` attribute of the extension is used instead. Note that the name would be suffixed by an ordinal in case of duplicated names as explained above.
- **trigger** (*tuple or Trigger*) – Trigger object that determines when to invoke the extension. If it is `None`, `extension.trigger` is used instead. If it is `None` and the extension does not have the trigger attribute, the extension is triggered at every iteration by default. If the trigger is not callable, it is passed to `IntervalTrigger` to build an interval trigger.
- **priority** (*int*) – Invocation priority of the extension. Extensions are invoked in the descending order of priorities in each iteration. If this is `None`, `extension.priority` is used instead.

get_extension (*name*)

Returns the extension of a given name.

Parameters **name** (*str*) – Name of the extension.

Returns Extension.

run (*show_loop_exception_msg=True*)

Executes the training loop.

This method is the core of `Trainer`. It executes the whole loop of training the models.

Note that this method cannot run multiple times for one trainer object.

serialize (*serializer*)

Attributes

elapsed_time

Total time used for the training.

The time is in seconds. If the training is resumed from snapshot, it includes the time of all the previous training to get the current state of the trainer.

5.6.2 Updaters

<code>chainer.training.Updater</code>	Interface of updater objects for trainers.
<code>chainer.training.updaters. StandardUpdater</code>	Standard implementation of Updater.
<code>chainer.training.updaters. ParallelUpdater</code>	Implementation of a parallel GPU Updater.
<code>chainer.training.updaters. MultiprocessParallelUpdater</code>	Implementation of a multiprocess parallel GPU Updater.

chainer.training.Updater

class `chainer.training.Updater`

Interface of updater objects for trainers.

Updater implements a training iteration as `update()`. Typically, the updating iteration proceeds as follows.

- Fetch a minibatch from *dataset* via *Iterator*.
- Run forward and backward process of *Chain*.
- Update parameters according to their *UpdateRule*.

The first line is processed by `Iterator.__next__`. The second and third are processed by `Optimizer.update`. Users can also implement their original updating iteration by overriding `Updater.update`.

Methods

connect_trainer (*trainer*)

Connects the updater to the trainer that will call it.

The typical usage of this method is to register additional links to the reporter of the trainer. This method is called at the end of the initialization of *Trainer*. The default implementation does nothing.

Parameters *trainer* (*Trainer*) – Trainer object to which the updater is registered.

finalize ()

Finalizes the updater object.

This method is called at the end of training loops. It should finalize each dataset iterator used in this updater.

get_all_optimizers ()

Gets a dictionary of all optimizers for this updater.

Returns Dictionary that maps names to optimizers.

Return type *dict*

get_optimizer (*name*)

Gets the optimizer of given name.

Updater holds one or more optimizers with names. They can be retrieved by this method.

Parameters *name* (*str*) – Name of the optimizer.

Returns Optimizer of the name.

Return type *Optimizer*

serialize (*serializer*)

Serializes the current state of the updater object.

update()

Updates the parameters of the target model.

This method implements an update formula for the training task, including data loading, forward/backward computations, and actual updates of parameters.

This method is called once at each iteration of the training loop.

chainer.training.updaters.StandardUpdater

```
class chainer.training.updaters.StandardUpdater(iterator, optimizer, con-
                                              verter=<function concat_examples>,
                                              device=None, loss_func=None,
                                              loss_scale=None)
```

Standard implementation of Updater.

This is the standard implementation of *Updater*. It accepts one or more training datasets and one or more optimizers. The default update routine assumes that there is only one training dataset and one optimizer. Users can override this update routine by inheriting this class and overriding the *update_core()* method. Each batch is converted to input arrays by *chainer.dataset.concat_examples()* by default, which can also be manually set by *converter* argument.

Parameters

- **iterator** – Dataset iterator for the training dataset. It can also be a dictionary that maps strings to iterators. If this is just an iterator, then the iterator is registered by the name 'main'.
- **optimizer** – Optimizer to update parameters. It can also be a dictionary that maps strings to optimizers. If this is just an optimizer, then the optimizer is registered by the name 'main'.
- **converter** – Converter function to build input arrays. Each batch extracted by the main iterator and the *device* option are passed to this function. *chainer.dataset.concat_examples()* is used by default.
- **device** – Device to which the training data is sent. Negative value indicates the host memory (CPU).
- **loss_func** – Loss function. The target link of the main optimizer is used by default.
- **loss_scale** (*float*) – Loss scaling factor. Loss scaling is a usefull technique to mitigate vanishing gradient issue that tends to happen when low precision data type like float16 is used during training. If you set loss scaling factor, gradients of loss values are to be multiplied by the factor before backprop starts. The factor is propagated to whole gradients in a computational graph along the backprop. The gradients of parameters are divided by the factor just before the parameters are to be updated.

Variables

- **converter** – Converter function.
- **loss_func** – Loss function. If it is *None*, the target link of the main optimizer is used instead.
- **device** – Device to which the training data is sent.
- **iteration** – Current number of completed updates.

Methods

connect_trainer (*trainer*)

Connects the updater to the trainer that will call it.

The typical usage of this method is to register additional links to the reporter of the trainer. This method is called at the end of the initialization of *Trainer*. The default implementation does nothing.

Parameters *trainer* (*Trainer*) – Trainer object to which the updater is registered.

finalize ()

Finalizes the updater object.

This method calls the *finalize* method of each iterator that this updater has. It is called at the end of training loops.

get_all_optimizers ()

Gets a dictionary of all optimizers for this updater.

Returns Dictionary that maps names to optimizers.

Return type *dict*

get_iterator (*name*)

Gets the dataset iterator of given name.

Parameters *name* (*str*) – Name of the dataset iterator.

Returns Corresponding dataset iterator.

Return type *Iterator*

get_optimizer (*name*)

Gets the optimizer of given name.

Parameters *name* (*str*) – Name of the optimizer.

Returns Corresponding optimizer.

Return type *Optimizer*

serialize (*serializer*)

Serializes the current state of the updater object.

update ()

Updates the parameters of the target model.

This method implements an update formula for the training task, including data loading, forward/backward computations, and actual updates of parameters.

This method is called once at each iteration of the training loop.

update_core ()

Attributes

epoch

epoch_detail

is_new_epoch

previous_epoch_detail

chainer.training.updaters.ParallelUpdater

```
class chainer.training.updaters.ParallelUpdater(iterator, optimizer, converter=<function concat_examples>,
models=None, devices=None, loss_func=None, loss_scale=None)
```

Implementation of a parallel GPU Updater.

This is an implementation of `Updater` that uses multiple GPUs. It behaves similarly to `StandardUpdater`. The update routine is modified to support data-parallel computation on multiple GPUs in one machine. It is based on synchronous parallel SGD: it parallelizes the gradient computation over a mini-batch, and updates the parameters only in the main device.

Parameters

- **iterator** – Dataset iterator for the training dataset. It can also be a dictionary that maps strings to iterators. If this is just an iterator, then the iterator is registered by the name 'main'.
- **optimizer** – Optimizer to update parameters. It can also be a dictionary that maps strings to optimizers. If this is just an optimizer, then the optimizer is registered by the name 'main'.
- **converter** – Converter function to build input arrays. Each batch extracted by the main iterator is split equally between the devices and then passed with corresponding `device` option to this function. `concat_examples()` is used by default.
- **models** – Dictionary of models. The main model should be the same model attached to the 'main' optimizer.
- **devices** – Dictionary of devices to which the training data is sent. The devices should be arranged in a dictionary with the same structure as `models`.
- **loss_func** – Loss function. The model is used as a loss function by default.
- **loss_scale** (*float*) – Loss scaling factor. Loss scaling is a useful technique to mitigate vanishing gradient issue that tends to happen when low precision data type like float16 is used during training. If you set loss scaling factor, gradients of loss values are to be multiplied by the factor before backprop starts. The factor is propagated to whole gradients in a computational graph along the backprop. The gradients of parameters are divided by the factor just before the parameters are to be updated.

Methods

connect_trainer (*trainer*)

Connects the updater to the trainer that will call it.

The typical usage of this method is to register additional links to the reporter of the trainer. This method is called at the end of the initialization of `Trainer`. The default implementation does nothing.

Parameters **trainer** (`Trainer`) – Trainer object to which the updater is registered.

finalize ()

Finalizes the updater object.

This method calls the `finalize` method of each iterator that this updater has. It is called at the end of training loops.

get_all_optimizers ()

Gets a dictionary of all optimizers for this updater.

Returns Dictionary that maps names to optimizers.

Return type `dict`

get_iterator (*name*)

Gets the dataset iterator of given name.

Parameters **name** (*str*) – Name of the dataset iterator.

Returns Corresponding dataset iterator.

Return type *Iterator*

get_optimizer (*name*)

Gets the optimizer of given name.

Parameters **name** (*str*) – Name of the optimizer.

Returns Corresponding optimizer.

Return type *Optimizer*

serialize (*serializer*)

Serializes the current state of the updater object.

update ()

Updates the parameters of the target model.

This method implements an update formula for the training task, including data loading, forward/backward computations, and actual updates of parameters.

This method is called once at each iteration of the training loop.

update_core ()

Attributes

epoch

epoch_detail

is_new_epoch

previous_epoch_detail

chainer.training.updaters.MultiprocessParallelUpdater

```
class chainer.training.updaters.MultiprocessParallelUpdater (iterators, optimizer,  
                                                             converter=<function  
                                                             concat_examples>,  
                                                             devices=None)
```

Implementation of a multiprocess parallel GPU Updater.

This is an implementation of `Updater` that uses multiple GPUs with multi-process data parallelism. It uses Nvidia NCCL for communication between multiple GPUs.

It behaves similarly to `StandardUpdater`. The update routine is modified to support data-parallel computation on multiple GPUs in one machine. It is based on synchronous parallel SGD: it parallelizes the gradient computation over a mini-batch, and updates the parameters only in the main device.

It does not transfer the values collected by `Reporter` in the sub devices to the main device. So you can only see the reported values in the main device.

Parameters

- **iterators** – List of dataset iterator for the training dataset. The number of the iterators must be same to the number of GPUs you use.
- **optimizer** – Optimizer to update parameters. The model should be attached to the optimizer.
- **converter** – Converter function to build input arrays. Each batch extracted by the iterator is split equally between the devices and then passed with corresponding `device` option to this function. `concat_examples()` is used by default.
- **devices** – Dictionary or list of devices to which the training data is sent. The master device will be the first one in the list or the value attached to the key 'main'.

Methods

static available()

connect_trainer(*trainer*)

Connects the updater to the trainer that will call it.

The typical usage of this method is to register additional links to the reporter of the trainer. This method is called at the end of the initialization of `Trainer`. The default implementation does nothing.

Parameters **trainer** (`Trainer`) – Trainer object to which the updater is registered.

finalize()

Finalizes the updater object.

This method calls the `finalize` method of each iterator that this updater has. It is called at the end of training loops.

get_all_optimizers()

Gets a dictionary of all optimizers for this updater.

Returns Dictionary that maps names to optimizers.

Return type `dict`

get_iterator(*name*)

Gets the dataset iterator of given name.

Parameters **name** (`str`) – Name of the dataset iterator.

Returns Corresponding dataset iterator.

Return type `Iterator`

get_optimizer(*name*)

Gets the optimizer of given name.

Parameters **name** (`str`) – Name of the optimizer.

Returns Corresponding optimizer.

Return type `Optimizer`

serialize(*serializer*)

Serializes the current state of the updater object.

setup_workers()

update()

Updates the parameters of the target model.

This method implements an update formula for the training task, including data loading, forward/backward computations, and actual updates of parameters.

This method is called once at each iteration of the training loop.

update_core()

Attributes

epoch

epoch_detail

is_new_epoch

previous_epoch_detail

5.6.3 Extensions

An extension is a callable object that can perform arbitrary actions during the training loop. Extensions can be registered to *Trainer* by using *Trainer.extend()* method, and they are invoked when the *Trigger* condition is satisfied.

In addition to the built-in extensions listed below, you can define your own extension by implementing *Extension* or using the *make_extension()* decorator. See *Trainer Extensions* for details.

Common

<i>chainer.training.Extension</i>	Base class of trainer extensions.
<i>chainer.training.make_extension</i>	Decorator to make given functions into trainer extensions.

chainer.training.Extension

class chainer.training.**Extension**

Base class of trainer extensions.

Extension of *Trainer* is a callable object that takes the trainer object as the argument. It also provides some default configurations as its attributes, e.g. the default trigger and the default priority. This class provides a set of typical default values for these attributes.

There are three ways to define users' own extensions: inheriting this class, decorating closures by *make_extension()*, or using any callable including lambda functions as extensions. Decorator can slightly reduce the overhead and is much easier to use, while this class provides more flexibility (for example, it can have methods to configure the behavior). Using a lambda function allows one-line coding for simple purposes, but users have to specify the configurations as arguments to *Trainer.extend()*. For a callable not inheriting this class, the default configurations of this class are used unless the user explicitly specifies them in *Trainer.extend()* method.

Variables

- **trigger** – Default value of trigger for this extension. It is set to (1, 'iteration')

by default.

- **priority** – Default priority of the extension. It is set to `PRIORITY_READER` by default.

Methods

__call__ (*trainer*)

Invokes the extension.

Implementations should override this operator. This method is called at iterations which the corresponding trigger accepts.

Parameters **trainer** (*Trainer*) – Trainer object that calls this operator.

finalize ()

Finalizes the extension.

This method is called at the end of the training loop.

initialize (*trainer*)

Initializes up the trainer state.

This method is called before entering the training loop. An extension that modifies the state of *Trainer* can override this method to initialize it.

When the trainer has been restored from a snapshot, this method has to recover an appropriate part of the state of the trainer.

For example, *ExponentialShift* extension changes the optimizer’s hyperparameter at each invocation. Note that the hyperparameter is not saved to the snapshot; it is the responsibility of the extension to recover the hyperparameter. The *ExponentialShift* extension recovers it in its `initialize` method if it has been loaded from a snapshot, or just setting the initial value otherwise.

Parameters **trainer** (*Trainer*) – Trainer object that runs the training loop.

serialize (*serializer*)

Serializes the extension state.

It is called when a trainer that owns this extension is serialized. It serializes nothing by default.

Attributes

default_name

Default name of the extension.

It is the name of the class by default. Implementation can override this property, or provide a class attribute to hide it.

priority = 100

trigger = (1, 'iteration')

chainer.training.make_extension

`chainer.training.make_extension` (*trigger=None, default_name=None, priority=None, finalizer=None, initializer=None, **kwargs*)

Decorator to make given functions into trainer extensions.

This decorator just adds some attributes to a given function. The value of the attributes are given by the arguments of this decorator.

See [Extension](#) for details of trainer extensions. Most of the default values of arguments also follow those for this class.

Parameters

- **trigger** – Default trigger of the extension.
- **default_name** – Default name of the extension. The name of a given function is used by default.
- **priority** (*int*) – Default priority of the extension.
- **finalizer** – Finalizer function of this extension. It is called at the end of the training loop.
- **initializer** – Initializer function of this extension. It is called at the beginning of the training loop.

Evaluation and Metrics Collection

These extensions provide features to collect additional metrics. The typical use case is to use [Evaluator](#) to perform evaluation with a validation dataset to compute validation loss/accuracy.

<code>chainer.training.extensions.Evaluator</code>	Trainer extension to evaluate models on a validation set.
<code>chainer.training.extensions.MicroAverage</code>	Calculates micro-average ratio.
<code>chainer.training.extensions.FailOnNonNumber</code>	Trainer extension to raise <code>RuntimeError</code> if parameters contain NaN or Inf.
<code>chainer.training.extensions.ParameterStatistics</code>	Trainer extension to report parameter statistics.
<code>chainer.training.extensions.observe_lr</code>	Returns a trainer extension to record the learning rate.
<code>chainer.training.extensions.observe_value</code>	Returns a trainer extension to continuously record a value.

chainer.training.extensions.Evaluator

```
class chainer.training.extensions.Evaluator(iterator, target, converter=<function  
concat_examples>, device=None,  
eval_hook=None, eval_func=None)
```

Trainer extension to evaluate models on a validation set.

This extension evaluates the current models by a given evaluation function. It creates a [Reporter](#) object to store values observed in the evaluation function on each iteration. The report for all iterations are aggregated to [DictSummary](#). The collected mean values are further reported to the reporter object of the trainer, where the name of each observation is prefixed by the evaluator name. See [Reporter](#) for details in naming rules of the reports.

Evaluator has a structure to customize similar to that of [StandardUpdater](#). The main differences are:

- There are no optimizers in an evaluator. Instead, it holds links to evaluate.
- An evaluation loop function is used instead of an update function.

- Preparation routine can be customized, which is called before each evaluation. It can be used, e.g., to initialize the state of stateful recurrent networks.

There are two ways to modify the evaluation behavior besides setting a custom evaluation function. One is by setting a custom evaluation loop via the `eval_func` argument. The other is by inheriting this class and overriding the `evaluate()` method. In latter case, users have to create and handle a reporter object manually. Users also have to copy the iterators before using them, in order to reuse them at the next time of evaluation. In both cases, the functions are called in testing mode (i.e., `chainer.config.train` is set to `False`).

This extension is called at the end of each epoch by default.

Parameters

- **iterator** – Dataset iterator for the validation dataset. It can also be a dictionary of iterators. If this is just an iterator, the iterator is registered by the name 'main'.
- **target** – Link object or a dictionary of links to evaluate. If this is just a link object, the link is registered by the name 'main'.
- **converter** – Converter function to build input arrays. `concat_examples()` is used by default.
- **device** – Device to which the validation data is sent. Negative value indicates the host memory (CPU).
- **eval_hook** – Function to prepare for each evaluation process. It is called at the beginning of the evaluation. The evaluator extension object is passed at each call.
- **eval_func** – Evaluation function called at each iteration. The target link to evaluate as a callable is used by default.

Variables

- **converter** – Converter function.
- **device** – Device to which the validation data is sent.
- **eval_hook** – Function to prepare for each evaluation process.
- **eval_func** – Evaluation function called at each iteration.

Methods

`__call__(trainer=None)`

Executes the evaluator extension.

Unlike usual extensions, this extension can be executed without passing a trainer object. This extension reports the performance on validation dataset using the `report()` function. Thus, users can use this extension independently from any trainer by manually configuring a `Reporter` object.

Parameters **trainer** (`Trainer`) – Trainer object that invokes this extension. It can be omitted in case of calling this extension manually.

Returns Result dictionary that contains mean statistics of values reported by the evaluation function.

Return type `dict`

`evaluate()`

Evaluates the model and returns a result dictionary.

This method runs the evaluation loop over the validation dataset. It accumulates the reported values to `DictSummary` and returns a dictionary whose values are means computed by the summary.

Note that this function assumes that the main iterator raises `StopIteration` or code in the evaluation loop raises an exception. So, if this assumption is not held, the function could be caught in an infinite loop.

Users can override this method to customize the evaluation routine.

Note: This method encloses `eval_func` calls with `function.no_backprop_mode()` context, so all calculations using `FunctionNodes` inside `eval_func` do not make computational graphs. It is for reducing the memory consumption.

Returns Result dictionary. This dictionary is further reported via `report()` without specifying any observer.

Return type `dict`

finalize()

Finalizes the evaluator object.

This method calls the `finalize` method of each iterator that this evaluator has. It is called at the end of training loops.

get_all_iterators()

Returns a dictionary of all iterators.

get_all_targets()

Returns a dictionary of all target links.

get_iterator(name)

Returns the iterator of the given name.

get_target(name)

Returns the target link of the given name.

initialize(trainer)

Initializes up the trainer state.

This method is called before entering the training loop. An extension that modifies the state of `Trainer` can override this method to initialize it.

When the trainer has been restored from a snapshot, this method has to recover an appropriate part of the state of the trainer.

For example, `ExponentialShift` extension changes the optimizer's hyperparameter at each invocation. Note that the hyperparameter is not saved to the snapshot; it is the responsibility of the extension to recover the hyperparameter. The `ExponentialShift` extension recovers it in its `initialize` method if it has been loaded from a snapshot, or just setting the initial value otherwise.

Parameters `trainer` (`Trainer`) – Trainer object that runs the training loop.

serialize(serializer)

Serializes the extension state.

It is called when a trainer that owns this extension is serialized. It serializes nothing by default.

Attributes

`default_name = 'validation'`

`name = None`


```
priority = 300
trigger = (1, 'epoch')
```

chainer.training.extensions.MicroAverage

```
class chainer.training.extensions.MicroAverage(numerator_key, denominator_key, result_key, trigger=(1, 'epoch'))
```

Calculates micro-average ratio.

Give N batches and values $\{n_1, \dots, n_N\}$ and $\{d_1, \dots, d_N\}$, this extension calculates micro-average of these ratio defined as:

$$\frac{\sum_i^N n_i}{\sum_i^N d_i}.$$

A user usually uses the number of examples which a system correctly predict as n_i and the number of total examples in i -th batch as d_i . This value is called macro-average of precision.

Note that macro-average is defined as:

$$\frac{1}{N} \sum_i^N (n_i/d_i),$$

It is same to the micro-average when each mini-batch has the same d_i .

You need to report numerator value (the number of correct examples) and denominator value (the number of examples) in your model.

```
>>> class MyModel(chainer.Link):
...     def __call__(self, x, y):
...         loss = F.softmax_cross_entropy(x, y)
...         correct = (x.data.argmax(axis=1) == y.data).sum()
...         total = len(y.data)
...         reporter.report({'correct': correct, 'total': total}, self)
...         return loss
```

And then, make an extension with corresponding reporting keys and register it.

```
>>> ext = extensions.MicroAverage(
...     'main/correct', 'main/total', 'main/accuracy')
```

Parameters

- **numerator_key** (*str*) – Key string of obserbation storing a numerator value.
- **denominator_key** (*str*) – Key string of obserbation storing a denominator value.
- **result_key** (*str*) – Key string of obserbation to store a result.
- **trigger** – Trigger that decides when to calculate average. This is distinct from the trigger of this extension itself. If it is a tuple in the form `<int>, 'epoch'` or `<int>, 'iteration'`, it is passed to `IntervalTrigger`.

Methods

__call__ (*trainer*)
Invokes the extension.

Implementations should override this operator. This method is called at iterations which the corresponding trigger accepts.

Parameters `trainer` (`Trainer`) – Trainer object that calls this operator.

finalize ()

Finalizes the extension.

This method is called at the end of the training loop.

initialize (`trainer`)

Initializes up the trainer state.

This method is called before entering the training loop. An extension that modifies the state of `Trainer` can override this method to initialize it.

When the trainer has been restored from a snapshot, this method has to recover an appropriate part of the state of the trainer.

For example, `ExponentialShift` extension changes the optimizer's hyperparameter at each invocation. Note that the hyperparameter is not saved to the snapshot; it is the responsibility of the extension to recover the hyperparameter. The `ExponentialShift` extension recovers it in its `initialize` method if it has been loaded from a snapshot, or just setting the initial value otherwise.

Parameters `trainer` (`Trainer`) – Trainer object that runs the training loop.

serialize (`serializer`)

Serializes the extension state.

It is called when a trainer that owns this extension is serialized. It serializes nothing by default.

Attributes

default_name

Default name of the extension.

It is the name of the class by default. Implementation can override this property, or provide a class attribute to hide it.

priority = 200

trigger = (1, 'iteration')

chainer.training.extensions.FailOnNonNumber

class `chainer.training.extensions.FailOnNonNumber`

Trainer extension to raise `RuntimeError` if parameters contain NaN or Inf.

Although parameters including non-number such as NaN and Inf are unnecessary in most cases, `Trainer` will continue to compute even if the parameters in a given optimizer diverge. This extension is aimed to reduce unnecessary computations by throwing `RuntimeError` if the parameters contain NaN or Inf.

Methods

__call__ (`trainer`)

Invokes the extension.

Implementations should override this operator. This method is called at iterations which the corresponding trigger accepts.

Parameters `trainer` (`Trainer`) – Trainer object that calls this operator.

finalize ()

Finalizes the extension.

This method is called at the end of the training loop.

initialize (`trainer`)

Initializes up the trainer state.

This method is called before entering the training loop. An extension that modifies the state of `Trainer` can override this method to initialize it.

When the trainer has been restored from a snapshot, this method has to recover an appropriate part of the state of the trainer.

For example, `ExponentialShift` extension changes the optimizer’s hyperparameter at each invocation. Note that the hyperparameter is not saved to the snapshot; it is the responsibility of the extension to recover the hyperparameter. The `ExponentialShift` extension recovers it in its `initialize` method if it has been loaded from a snapshot, or just setting the initial value otherwise.

Parameters `trainer` (`Trainer`) – Trainer object that runs the training loop.

serialize (`serializer`)

Serializes the extension state.

It is called when a trainer that owns this extension is serialized. It serializes nothing by default.

Attributes

default_name

Default name of the extension.

It is the name of the class by default. Implementation can override this property, or provide a class attribute to hide it.

priority = 100

trigger = (1, 'iteration')

chainer.training.extensions.ParameterStatistics

```
class chainer.training.extensions.ParameterStatistics(links, statistics={'max':
    <function ParameterStatistics.<lambda> at
    0x7f8e0f8fe620>, 'mean':
    <function ParameterStatistics.<lambda> at
    0x7f8e0f8fe488>, 'min':
    <function ParameterStatistics.<lambda> at
    0x7f8e0f8fe598>, 'percentile': <function ParameterStatistics.<lambda> at
    0x7f8e0f8fe730>, 'std':
    <function ParameterStatistics.<lambda> at
    0x7f8e0f8fe510>, 'zeros': <function ParameterStatistics.<lambda>
    at 0x7f8e0f8fe6a8>},
    report_params=True,
    report_grads=True,
    prefix=None, trigger=(1, 'epoch'),
    skip_nan_params=False)
```

Trainer extension to report parameter statistics.

Statistics are collected and reported for a given *Link* or an iterable of *Links*. If a link contains child links, the statistics are reported separately for each child.

Any function that takes a one-dimensional `numpy.ndarray` or a `cupy.ndarray` and outputs a single or multiple real numbers can be registered to handle the collection of statistics, e.g. `numpy.ndarray.mean()`.

The keys of reported statistics follow the convention of link name followed by parameter name, attribute name and function name, e.g. `VGG16Layers/conv1_1/W/data/mean`. They are prepended with an optional prefix and appended with integer indices if the statistics generating function return multiple values.

Parameters

- **links** (*Link* or *iterable of ~chainer.Link*) – Link(s) containing the parameters to observe. The link is expected to have a `name` attribute which is used as a part of the report key.
- **statistics** (*dict*) – Dictionary with function name to function mappings. The name is a string and is used as a part of the report key. The function is responsible for generating the statistics.
- **report_params** (*bool*) – If `True`, report statistics for parameter values such as weights and biases.
- **report_grads** (*bool*) – If `True`, report statistics for parameter gradients.
- **prefix** (*str*) – Optional prefix to prepend to the report keys.
- **trigger** – Trigger that decides when to aggregate the results and report the values.
- **skip_nan_params** (*bool*) – If `True`, statistics are not computed for parameters including NaNs and a single NaN value is immediately reported instead. Otherwise, this extension

will simply try to compute the statistics without performing any checks for NaNs.

Methods

__call__ (*trainer*)

Execute the statistics extension.

Collect statistics for the current state of parameters.

Note that this method will merely update its statistic summary, unless the internal trigger is fired. If the trigger is fired, the summary will also be reported and then reset for the next accumulation.

Parameters *trainer* (*Trainer*) – Associated trainer that invoked this extension.

finalize ()

Finalizes the extension.

This method is called at the end of the training loop.

initialize (*trainer*)

Initializes up the trainer state.

This method is called before entering the training loop. An extension that modifies the state of *Trainer* can override this method to initialize it.

When the trainer has been restored from a snapshot, this method has to recover an appropriate part of the state of the trainer.

For example, *ExponentialShift* extension changes the optimizer’s hyperparameter at each invocation. Note that the hyperparameter is not saved to the snapshot; it is the responsibility of the extension to recover the hyperparameter. The *ExponentialShift* extension recovers it in its *initialize* method if it has been loaded from a snapshot, or just setting the initial value otherwise.

Parameters *trainer* (*Trainer*) – Trainer object that runs the training loop.

register_statistics (*name*, *function*)

Register a function to compute a certain statistic.

The registered function will be called each time the extension runs and the results will be included in the report.

Parameters

- **name** (*str*) – Name of the statistic.
- **function** – Function to generate the statistic. Any function that takes a one-dimensional `numpy.ndarray` or a `cupy.ndarray` and outputs a single or multiple real numbers is allowed.

serialize (*serializer*)

Serializes the extension state.

It is called when a trainer that owns this extension is serialized. It serializes nothing by default.

Attributes

default_name = 'parameter_statistics'

default_statistics = {'max': <function ParameterStatistics.<lambda> at 0x7f8e0f8fe620>

priority = 300

```
report_key_template = '{prefix}{link_name}{param_name}/{attr_name}/{function_name}'
trigger = (1, 'iteration')
```

chainer.training.extensions.observe_lr

`chainer.training.extensions.observe_lr` (*optimizer_name*='main', *observation_key*='lr')

Returns a trainer extension to record the learning rate.

Parameters

- **optimizer_name** (*str*) – Name of optimizer whose learning rate is recorded.
- **observation_key** (*str*) – Key of observation to record.

Returns The extension function.

chainer.training.extensions.observe_value

`chainer.training.extensions.observe_value` (*observation_key*, *target_func*)

Returns a trainer extension to continuously record a value.

Parameters

- **observation_key** (*str*) – Key of observation to record.
- **target_func** (*function*) – Function that returns the value to record. It must take one argument: :class:`~chainer.training.Trainer` object.

Returns The extension function.

Optimizer Behavior Control

These extensions provide features to adjust optimizer behavior. The typical use case is to change the learning rate of the optimizer over time.

<code>chainer.training.extensions.ExponentialShift</code>	Trainer extension to exponentially shift an optimizer attribute.
<code>chainer.training.extensions.LinearShift</code>	Trainer extension to change an optimizer attribute linearly.

chainer.training.extensions.ExponentialShift

class `chainer.training.extensions.ExponentialShift` (*attr*, *rate*, *init*=None, *target*=None, *optimizer*=None)

Trainer extension to exponentially shift an optimizer attribute.

This extension exponentially increases or decreases the specified attribute of the optimizer. The typical use case is an exponential decay of the learning rate.

This extension is also called before the training loop starts by default.

Parameters

- **attr** (*str*) – Name of the attribute to shift.

- **rate** (*float*) – Rate of the exponential shift. This value is multiplied to the attribute at each call.
- **init** (*float*) – Initial value of the attribute. If it is `None`, the extension extracts the attribute at the first call and uses it as the initial value.
- **target** (*float*) – Target value of the attribute. If the attribute reaches this value, the shift stops.
- **optimizer** (*Optimizer*) – Target optimizer to adjust the attribute. If it is `None`, the main optimizer of the updater is used.

Methods

__call__ (*trainer*)

Invokes the extension.

Implementations should override this operator. This method is called at iterations which the corresponding trigger accepts.

Parameters **trainer** (*Trainer*) – Trainer object that calls this operator.

finalize ()

Finalizes the extension.

This method is called at the end of the training loop.

initialize (*trainer*)

Initializes up the trainer state.

This method is called before entering the training loop. An extension that modifies the state of *Trainer* can override this method to initialize it.

When the trainer has been restored from a snapshot, this method has to recover an appropriate part of the state of the trainer.

For example, *ExponentialShift* extension changes the optimizer’s hyperparameter at each invocation. Note that the hyperparameter is not saved to the snapshot; it is the responsibility of the extension to recover the hyperparameter. The *ExponentialShift* extension recovers it in its `initialize` method if it has been loaded from a snapshot, or just setting the initial value otherwise.

Parameters **trainer** (*Trainer*) – Trainer object that runs the training loop.

serialize (*serializer*)

Serializes the extension state.

It is called when a trainer that owns this extension is serialized. It serializes nothing by default.

Attributes

default_name

Default name of the extension.

It is the name of the class by default. Implementation can override this property, or provide a class attribute to hide it.

priority = 100

trigger = (1, 'iteration')

chainer.training.extensions.LinearShift

```
class chainer.training.extensions.LinearShift(attr, value_range, time_range, optimizer=None)
```

Trainer extension to change an optimizer attribute linearly.

This extension changes an optimizer attribute from the first value to the last value linearly within a specified duration. The typical use case is warming up of the momentum coefficient.

For example, suppose that this extension is called at every iteration, and `value_range == (x, y)` and `time_range == (i, j)`. Then, this extension keeps the attribute to be `x` up to the `i`-th iteration, linearly shifts the value to `y` by the `j`-th iteration, and then keeps the value to be `y` after the `j`-th iteration.

This extension is also called before the training loop starts by default.

Parameters

- **attr** (*str*) – Name of the optimizer attribute to adjust.
- **value_range** (*tuple of float*) – The first and the last values of the attribute.
- **time_range** (*tuple of ints*) – The first and last counts of calls in which the attribute is adjusted.
- **optimizer** (*Optimizer*) – Target optimizer object. If it is `None`, the main optimizer of the trainer is used.

Methods

```
__call__(trainer)
```

Invokes the extension.

Implementations should override this operator. This method is called at iterations which the corresponding trigger accepts.

Parameters **trainer** (*Trainer*) – Trainer object that calls this operator.

```
finalize()
```

Finalizes the extension.

This method is called at the end of the training loop.

```
initialize(trainer)
```

Initializes up the trainer state.

This method is called before entering the training loop. An extension that modifies the state of *Trainer* can override this method to initialize it.

When the trainer has been restored from a snapshot, this method has to recover an appropriate part of the state of the trainer.

For example, *ExponentialShift* extension changes the optimizer's hyperparameter at each invocation. Note that the hyperparameter is not saved to the snapshot; it is the responsibility of the extension to recover the hyperparameter. The *ExponentialShift* extension recovers it in its `initialize` method if it has been loaded from a snapshot, or just setting the initial value otherwise.

Parameters **trainer** (*Trainer*) – Trainer object that runs the training loop.

```
serialize(serializer)
```

Serializes the extension state.

It is called when a trainer that owns this extension is serialized. It serializes nothing by default.

Attributes

`default_name`

Default name of the extension.

It is the name of the class by default. Implementation can override this property, or provide a class attribute to hide it.

`priority = 100`

`trigger = (1, 'iteration')`

Reporting

These extensions provide features to perform reporting of metrics and various statistics to the console or files.

<code>chainer.training.extensions.PrintReport</code>	Trainer extension to print the accumulated results.
<code>chainer.training.extensions.ProgressBar</code>	Trainer extension to print a progress bar and recent training status.
<code>chainer.training.extensions.LogReport</code>	Trainer extension to output the accumulated results to a log file.
<code>chainer.training.extensions.PlotReport</code>	Trainer extension to output plots.
<code>chainer.training.extensions.VariableStatisticsPlot</code>	Trainer extension to plot statistics for Variables.
<code>chainer.training.extensions.dump_graph</code>	Returns a trainer extension to dump a computational graph.

`chainer.training.extensions.PrintReport`

```
class chainer.training.extensions.PrintReport (entries, log_report='LogReport',
                                              out=<_io.TextIOWrapper
                                              name='<stdout>' mode='w'
                                              encoding='UTF-8'>)
```

Trainer extension to print the accumulated results.

This extension uses the log accumulated by a `LogReport` extension to print specified entries of the log in a human-readable format.

Parameters

- **entries** (*list of str*) – List of keys of observations to print.
- **log_report** (*str or LogReport*) – Log report to accumulate the observations. This is either the name of a `LogReport` extensions registered to the trainer, or a `LogReport` instance to use internally.
- **out** – Stream to print the bar. Standard output is used by default.

Methods

`__call__ (trainer)`

Invokes the extension.

Implementations should override this operator. This method is called at iterations which the corresponding trigger accepts.

Parameters `trainer` (`Trainer`) – Trainer object that calls this operator.

finalize ()

Finalizes the extension.

This method is called at the end of the training loop.

initialize (`trainer`)

Initializes up the trainer state.

This method is called before entering the training loop. An extension that modifies the state of `Trainer` can override this method to initialize it.

When the trainer has been restored from a snapshot, this method has to recover an appropriate part of the state of the trainer.

For example, `ExponentialShift` extension changes the optimizer's hyperparameter at each invocation. Note that the hyperparameter is not saved to the snapshot; it is the responsibility of the extension to recover the hyperparameter. The `ExponentialShift` extension recovers it in its `initialize` method if it has been loaded from a snapshot, or just setting the initial value otherwise.

Parameters `trainer` (`Trainer`) – Trainer object that runs the training loop.

serialize (`serializer`)

Serializes the extension state.

It is called when a trainer that owns this extension is serialized. It serializes nothing by default.

Attributes

default_name

Default name of the extension.

It is the name of the class by default. Implementation can override this property, or provide a class attribute to hide it.

priority = 100

trigger = (1, 'iteration')

chainer.training.extensions.ProgressBar

```
class chainer.training.extensions.ProgressBar(training_length=None,          up-
                                             date_interval=100,      bar_length=50,
                                             out=<_io.TextIOWrapper
                                             name='<stdout>'          mode='w'
                                             encoding='UTF-8'>)
```

Trainer extension to print a progress bar and recent training status.

This extension prints a progress bar at every call. It watches the current iteration and epoch to print the bar.

Parameters

- **training_length** (`tuple`) – Length of whole training. It consists of an integer and either 'epoch' or 'iteration'. If this value is omitted and the stop trigger of the trainer is `IntervalTrigger`, this extension uses its attributes to determine the length of the training.

- **update_interval** (*int*) – Number of iterations to skip printing the progress bar.
- **bar_length** (*int*) – Length of the progress bar in characters.
- **out** – Stream to print the bar. Standard output is used by default.

Methods

__call__ (*trainer*)

Invokes the extension.

Implementations should override this operator. This method is called at iterations which the corresponding trigger accepts.

Parameters **trainer** (*Trainer*) – Trainer object that calls this operator.

finalize ()

Finalizes the extension.

This method is called at the end of the training loop.

initialize (*trainer*)

Initializes up the trainer state.

This method is called before entering the training loop. An extension that modifies the state of *Trainer* can override this method to initialize it.

When the trainer has been restored from a snapshot, this method has to recover an appropriate part of the state of the trainer.

For example, *ExponentialShift* extension changes the optimizer’s hyperparameter at each invocation. Note that the hyperparameter is not saved to the snapshot; it is the responsibility of the extension to recover the hyperparameter. The *ExponentialShift* extension recovers it in its `initialize` method if it has been loaded from a snapshot, or just setting the initial value otherwise.

Parameters **trainer** (*Trainer*) – Trainer object that runs the training loop.

serialize (*serializer*)

Serializes the extension state.

It is called when a trainer that owns this extension is serialized. It serializes nothing by default.

Attributes

default_name

Default name of the extension.

It is the name of the class by default. Implementation can override this property, or provide a class attribute to hide it.

priority = 100

trigger = (1, 'iteration')

chainer.training.extensions.LogReport

class `chainer.training.extensions.LogReport` (*keys=None, trigger=(1, 'epoch'), postprocess=None, log_name='log'*)

Trainer extension to output the accumulated results to a log file.

This extension accumulates the observations of the trainer to `DictSummary` at a regular interval specified by a supplied trigger, and writes them into a log file in JSON format.

There are two triggers to handle this extension. One is the trigger to invoke this extension, which is used to handle the timing of accumulating the results. It is set to 1, 'iteration' by default. The other is the trigger to determine when to emit the result. When this trigger returns True, this extension appends the summary of accumulated values to the list of past summaries, and writes the list to the log file. Then, this extension makes a new fresh summary object which is used until the next time that the trigger fires.

It also adds some entries to each result dictionary.

- 'epoch' and 'iteration' are the epoch and iteration counts at the output, respectively.
- 'elapsed_time' is the elapsed time in seconds since the training begins. The value is taken from `Trainer.elapsed_time`.

Parameters

- **keys** (*iterable of strs*) – Keys of values to accumulate. If this is None, all the values are accumulated and output to the log file.
- **trigger** – Trigger that decides when to aggregate the result and output the values. This is distinct from the trigger of this extension itself. If it is a tuple in the form `<int>, 'epoch'` or `<int>, 'iteration'`, it is passed to `IntervalTrigger`.
- **postprocess** – Callback to postprocess the result dictionaries. Each result dictionary is passed to this callback on the output. This callback can modify the result dictionaries, which are used to output to the log file.
- **log_name** (*str*) – Name of the log file under the output directory. It can be a format string: the last result dictionary is passed for the formatting. For example, users can use `'{iteration}'` to separate the log files for different iterations. If the log name is None, it does not output the log to any file.

Methods

`__call__(trainer)`

Invokes the extension.

Implementations should override this operator. This method is called at iterations which the corresponding trigger accepts.

Parameters `trainer` (`Trainer`) – Trainer object that calls this operator.

`finalize()`

Finalizes the extension.

This method is called at the end of the training loop.

`initialize(trainer)`

Initializes up the trainer state.

This method is called before entering the training loop. An extension that modifies the state of `Trainer` can override this method to initialize it.

When the trainer has been restored from a snapshot, this method has to recover an appropriate part of the state of the trainer.

For example, `ExponentialShift` extension changes the optimizer's hyperparameter at each invocation. Note that the hyperparameter is not saved to the snapshot; it is the responsibility of the extension

to recover the hyperparameter. The *ExponentialShift* extension recovers it in its `initialize` method if it has been loaded from a snapshot, or just setting the initial value otherwise.

Parameters `trainer` (*Trainer*) – Trainer object that runs the training loop.

serialize (*serializer*)

Serializes the extension state.

It is called when a trainer that owns this extension is serialized. It serializes nothing by default.

Attributes

default_name

Default name of the extension.

It is the name of the class by default. Implementation can override this property, or provide a class attribute to hide it.

log

The current list of observation dictionaries.

priority = 100

trigger = (1, 'iteration')

chainer.training.extensions.PlotReport

```
class chainer.training.extensions.PlotReport (y_keys,      x_key='iteration',      trig-
                                             ger=(1, 'epoch'), postprocess=None,
                                             file_name='plot.png',      marker='x',
                                             grid=True)
```

Trainer extension to output plots.

This extension accumulates the observations of the trainer to *DictSummary* at a regular interval specified by a supplied trigger, and plot a graph with using them.

There are two triggers to handle this extension. One is the trigger to invoke this extension, which is used to handle the timing of accumulating the results. It is set to 1, 'iteration' by default. The other is the trigger to determine when to emit the result. When this trigger returns True, this extension appends the summary of accumulated values to the list of past summaries, and writes the list to the log file. Then, this extension makes a new fresh summary object which is used until the next time that the trigger fires.

It also adds 'epoch' and 'iteration' entries to each result dictionary, which are the epoch and iteration counts at the output.

Warning: If your environment needs to specify a backend of matplotlib explicitly, please call `matplotlib.use` before calling `trainer.run`. For example:

```
import matplotlib
matplotlib.use('Agg')

trainer.extend(
    extensions.PlotReport(['main/loss', 'validation/main/loss'],
                          'epoch', file_name='loss.png'))
trainer.run()
```

Then, once one of instances of this extension is called, `matplotlib.use` will have no effect.

For the details, please see here: https://matplotlib.org/faq/usage_faq.html#what-is-a-backend

Parameters

- **y_keys** (*iterable of strs*) – Keys of values regarded as y. If this is None, nothing is output to the graph.
- **x_key** (*str*) – Keys of values regarded as x. The default value is 'iteration'.
- **trigger** – Trigger that decides when to aggregate the result and output the values. This is distinct from the trigger of this extension itself. If it is a tuple in the form <int>, 'epoch' or <int>, 'iteration', it is passed to IntervalTrigger.
- **postprocess** – Callback to postprocess the result dictionaries. Figure object, Axes object, and all plot data are passed to this callback in this order. This callback can modify the figure.
- **file_name** (*str*) – Name of the figure file under the output directory. It can be a format string.
- **marker** (*str*) – The marker used to plot the graph. Default is 'x'. If None is given, it draws with no markers.
- **grid** (*bool*) – Set the axis grid on if True. Default is True.

Methods

__call__ (*trainer*)

Invokes the extension.

Implementations should override this operator. This method is called at iterations which the corresponding trigger accepts.

Parameters **trainer** (*Trainer*) – Trainer object that calls this operator.

static available ()

finalize ()

Finalizes the extension.

This method is called at the end of the training loop.

initialize (*trainer*)

Initializes up the trainer state.

This method is called before entering the training loop. An extension that modifies the state of *Trainer* can override this method to initialize it.

When the trainer has been restored from a snapshot, this method has to recover an appropriate part of the state of the trainer.

For example, *ExponentialShift* extension changes the optimizer's hyperparameter at each invocation. Note that the hyperparameter is not saved to the snapshot; it is the responsibility of the extension to recover the hyperparameter. The *ExponentialShift* extension recovers it in its *initialize* method if it has been loaded from a snapshot, or just setting the initial value otherwise.

Parameters **trainer** (*Trainer*) – Trainer object that runs the training loop.

serialize (*serializer*)

Serializes the extension state.

It is called when a trainer that owns this extension is serialized. It serializes nothing by default.

Attributes

default_name

Default name of the extension.

It is the name of the class by default. Implementation can override this property, or provide a class attribute to hide it.

priority = 100

trigger = (1, 'iteration')

chainer.training.extensions.VariableStatisticsPlot

```
class chainer.training.extensions.VariableStatisticsPlot(targets,
                                                    max_sample_size=1000,
                                                    report_data=True,
                                                    report_grad=True,
                                                    plot_mean=True,
                                                    plot_std=True, per-
                                                    centile_sigmas=(0, 0.13,
                                                    2.28, 15.87, 50, 84.13,
                                                    97.72, 99.87, 100),
                                                    trigger=(1, 'epoch'),
                                                    file_name='statistics.png',
                                                    figsize=None,
                                                    marker=None,
                                                    grid=True)
```

Trainer extension to plot statistics for Variables.

This extension collects statistics for a single Variable, a list of Variables or similarly a single or a list of Links containing one or more Variables. In case multiple Variables are found, the means are computed. The collected statistics are plotted and saved as an image in the directory specified by the Trainer.

Statistics include mean, standard deviation and percentiles.

This extension uses reservoir sampling to preserve memory, using a fixed size running sample. This means that collected items in the sample are discarded uniformly at random when the number of items becomes larger than the maximum sample size, but each item is expected to occur in the sample with equal probability.

Parameters

- **targets** (Variable, Link or list of either) – Parameters for which statistics are collected.
- **max_sample_size** (*int*) – Maximum number of running samples.
- **report_data** (*bool*) – If True, data (e.g. weights) statistics are plotted. If False, they are neither computed nor plotted.
- **report_grad** (*bool*) – If True, gradient statistics are plotted. If False, they are neither computed nor plotted.
- **plot_mean** (*bool*) – If True, means are plotted. If False, they are neither computed nor plotted.
- **plot_std** (*bool*) – If True, standard deviations are plotted. If False, they are neither computed nor plotted.

- **percentile_sigmas** (*float or tuple of floats*) – Percentiles to plot in the range [0, 100].
- **trigger** – Trigger that decides when to save the plots as an image. This is distinct from the trigger of this extension itself. If it is a tuple in the form <int>, 'epoch' or <int>, 'iteration', it is passed to IntervalTrigger.
- **file_name** (*str*) – Name of the output image file under the output directory.
- **figsize** (*tuple of int*) – Matplotlib `figsize` argument that specifies the size of the output image.
- **marker** (*str*) – Matplotlib `marker` argument that specified the marker style of the plots.
- **grid** (*bool*) – Matplotlib `grid` argument that specifies whether grids are rendered in in the plots or not.

Methods

__call__ (*trainer*)

Invokes the extension.

Implementations should override this operator. This method is called at iterations which the corresponding trigger accepts.

Parameters **trainer** (*Trainer*) – Trainer object that calls this operator.

static available ()

finalize ()

Finalizes the extension.

This method is called at the end of the training loop.

initialize (*trainer*)

Initializes up the trainer state.

This method is called before entering the training loop. An extension that modifies the state of *Trainer* can override this method to initialize it.

When the trainer has been restored from a snapshot, this method has to recover an appropriate part of the state of the trainer.

For example, *ExponentialShift* extension changes the optimizer's hyperparameter at each invocation. Note that the hyperparameter is not saved to the snapshot; it is the responsibility of the extension to recover the hyperparameter. The *ExponentialShift* extension recovers it in its `initialize` method if it has been loaded from a snapshot, or just setting the initial value otherwise.

Parameters **trainer** (*Trainer*) – Trainer object that runs the training loop.

save_plot_using_module (*file_path, plt*)

serialize (*serializer*)

Serializes the extension state.

It is called when a trainer that owns this extension is serialized. It serializes nothing by default.

Attributes

default_name

Default name of the extension.

It is the name of the class by default. Implementation can override this property, or provide a class attribute to hide it.

```
priority = 100
trigger = (1, 'iteration')
```

chainer.training.extensions.dump_graph

```
chainer.training.extensions.dump_graph(root_name, out_name='cg.dot', variable_style=None, function_style=None)
```

Returns a trainer extension to dump a computational graph.

This extension dumps a computational graph. The graph is output in DOT language.

It only dumps a graph at the first invocation.

Note: As of v2.0.0, the computational graph is not kept by default. This extension changes this behavior until the first invocation. **It is strongly recommended to use it with the default trigger setting.**

The detailed behavior of this extension since v2.0.0 is as follows.

1. In its initializer, it turns on the `chainer.config.keep_graph_on_report` flag.
2. At the first iteration, it dumps the graph using the graph held by the reported variable.
3. After dumping the graph, it turns off the flag (if it was originally turned off) so that any variable reported afterward does not hold a computational graph.

When the `keep_graph_on_report` flag is turned on, the computational graph created by the updater is kept during the invocation of extensions. It will cause an unnecessarily large memory consumption when an extension also uses a large amount of memory, e.g. *Evaluator*.

With the default setting, the `dump_graph` extension is called at the first iteration. Since *Evaluator* is not called at the first iteration in most cases, it does not cause any memory problem.

Parameters

- **root_name** (*str*) – Name of the root of the computational graph. The root variable is retrieved by this name from the observation dictionary of the trainer.
- **out_name** (*str*) – Output file name.
- **variable_style** (*dict*) – Dot node style for variables. Each variable is rendered by an octagon by default.
- **function_style** (*dict*) – Dot node style for functions. Each function is rendered by a rectangular by default.

See also:

See `build_computational_graph()` for the `variable_style` and `function_style` arguments.

Snapshot

These extensions provide features to take snapshots of models.

<code>chainer.training.extensions.snapshot</code>	Returns a trainer extension to take snapshots of the trainer.
<code>chainer.training.extensions.snapshot_object</code>	Returns a trainer extension to take snapshots of a given object.

chainer.training.extensions.snapshot

`chainer.training.extensions.snapshot` (*savefun*=<function *save_npz*>, *filename*=*'snapshot_iter_{.updater.iteration}'*)

Returns a trainer extension to take snapshots of the trainer.

This extension serializes the trainer object and saves it to the output directory. It is used to support resuming the training loop from the saved state.

This extension is called once per epoch by default. To take a snapshot at a different interval, a trigger object specifying the required interval can be passed along with this extension to the *extend()* method of the trainer.

The default priority is -100, which is lower than that of most built-in extensions.

Note: This extension first writes the serialized object to a temporary file and then rename it to the target file name. Thus, if the program stops right before the renaming, the temporary file might be left in the output directory.

Parameters

- **savefun** – Function to save the trainer. It takes two arguments: the output file path and the trainer object.
- **filename** (*str*) – Name of the file into which the trainer is serialized. It can be a format string, where the trainer object is passed to the `str.format()` method.

chainer.training.extensions.snapshot_object

`chainer.training.extensions.snapshot_object` (*target*, *filename*, *savefun*=<function *save_npz*>)

Returns a trainer extension to take snapshots of a given object.

This extension serializes the given object and saves it to the output directory.

This extension is called once per epoch by default. To take a snapshot at a different interval, a trigger object specifying the required interval can be passed along with this extension to the *extend()* method of the trainer.

The default priority is -100, which is lower than that of most built-in extensions.

Parameters

- **target** – Object to serialize.
- **filename** (*str*) – Name of the file into which the object is serialized. It can be a format string, where the trainer object is passed to the `str.format()` method. For example, `'snapshot_{.updater.iteration}'` is converted to `'snapshot_10000'` at the 10,000th iteration.
- **savefun** – Function to save the object. It takes two arguments: the output file path and the object to serialize.

Returns An extension function.

5.6.4 Triggers

A trigger is a callable object to decide when to process some specific event within the training loop. It takes a `Trainer` object as the argument, and returns `True` if some event should be fired.

It is mainly used to determine when to call an extension. It is also used to determine when to quit the training loop.

<code>chainer.training.get_trigger</code>	Gets a trigger object.
<code>chainer.training.triggers. BestValueTrigger</code>	Trigger invoked when specific value becomes best.
<code>chainer.training.triggers. EarlyStoppingTrigger</code>	Trigger for Early Stopping
<code>chainer.training.triggers. IntervalTrigger</code>	Trigger based on a fixed interval.
<code>chainer.training.triggers. ManualScheduleTrigger</code>	Trigger invoked at specified point(s) of iterations or epochs.
<code>chainer.training.triggers. MaxValueTrigger</code>	Trigger invoked when specific value becomes maximum.
<code>chainer.training.triggers. MinValueTrigger</code>	Trigger invoked when specific value becomes minimum.
<code>chainer.training.triggers. TimeTrigger</code>	Trigger based on a fixed time interval.

chainer.training.get_trigger

`chainer.training.get_trigger(trigger)`

Gets a trigger object.

Trigger object is a callable that accepts a `Trainer` object as an argument and returns a boolean value. When it returns `True`, various kinds of events can occur depending on the context in which the trigger is used. For example, if the trigger is passed to the `Trainer` as the *stop trigger*, the training loop breaks when the trigger returns `True`. If the trigger is passed to the `extend()` method of a trainer, then the registered extension is invoked only when the trigger returns `True`.

This function returns a trigger object based on the argument. If `trigger` is already a callable, it just returns the trigger. If `trigger` is `None`, it returns a trigger that never fires. Otherwise, it passes the value to `IntervalTrigger`.

Parameters `trigger` – Trigger object. It can be either an already built trigger object (i.e., a callable object that accepts a trainer object and returns a bool value), or a tuple. In latter case, the tuple is passed to `IntervalTrigger`.

Returns `trigger` if it is a callable, otherwise a `IntervalTrigger` object made from `trigger`.

chainer.training.triggers.BestValueTrigger

class `chainer.training.triggers.BestValueTrigger(key, compare, trigger=(1, 'epoch'))`

Trigger invoked when specific value becomes best.

Parameters

- **key** (*str*) – Key of value.
- **compare** (*callable*) – Compare function which takes current best value and new value and returns whether new value is better than current best.
- **trigger** – Trigger that decides the comparison interval between current best value and new value. This must be a tuple in the form of `<int>, 'epoch'` or `<int>, 'iteration'` which is passed to *IntervalTrigger*.

Methods

`__call__(trainer)`

Decides whether the extension should be called on this iteration.

Parameters **trainer** (*Trainer*) – Trainer object that this trigger is associated with. The observation of this trainer is used to determine if the trigger should fire.

Returns `True` if the corresponding extension should be invoked in this iteration.

Return type `bool`

`chainer.training.triggers.EarlyStoppingTrigger`

```
class chainer.training.triggers.EarlyStoppingTrigger(check_trigger=(1, 'epoch'),
                                                    monitor='main/loss',
                                                    patients=3,
                                                    mode='auto',
                                                    verbose=False,
                                                    max_trigger=(100, 'epoch'))
```

Trigger for Early Stopping

It can be used as a stop trigger of *Trainer* to realize *early stopping* technique.

This trigger works as follows. Within each *check interval* defined by the `check_trigger` argument, it monitors and accumulates the reported value at each iteration. At the end of each interval, it computes the mean of the accumulated values and compares it to the previous ones to maintain the *best* value. When it finds that the best value is not updated for some periods (defined by *patients*), this trigger fires.

Parameters

- **monitor** (*str*) – The metric you want to monitor
- **check_trigger** – Trigger that decides the comparison interval between current best value and new value. This must be a tuple in the form of `<int>, 'epoch'` or `<int>, 'iteration'` which is passed to *IntervalTrigger*.
- **patients** (*int*) – Counts to let the trigger be patient. The trigger will not fire until the condition is met for successive patient checks.
- **mode** (*str*) – 'max', 'min', or 'auto'. It is used to determine how to compare the monitored values.
- **verbose** (*bool*) – Enable verbose output. If verbose is true, you can get more information
- **max_trigger** – Upper bound of the number of training loops

Methods

`__call__(trainer)`

Decides whether the training loop should be stopped.

Parameters **trainer** (`Trainer`) – Trainer object that this trigger is associated with. The observation of this trainer is used to determine if the trigger should fire.

Returns `True` if the training loop should be stopped.

Return type `bool`

`get_training_length()`

`chainer.training.triggers.IntervalTrigger`

class `chainer.training.triggers.IntervalTrigger` (*period, unit*)

Trigger based on a fixed interval.

This trigger accepts iterations divided by a given interval. There are two ways to specify the interval: per iterations and epochs. *Iteration* means the number of updates, while *epoch* means the number of sweeps over the training dataset. Fractional values are allowed if the interval is a number of epochs; the trigger uses the *iteration* and *epoch_detail* attributes defined by the updater.

For the description of triggers, see `get_trigger()`.

Parameters

- **period** (*int or float*) – Length of the interval. Must be an integer if unit is 'iteration'.
- **unit** (*str*) – Unit of the length specified by period. It must be either 'iteration' or 'epoch'.

Methods

`__call__` (*trainer*)

Decides whether the extension should be called on this iteration.

Parameters **trainer** (`Trainer`) – Trainer object that this trigger is associated with. The updater associated with this trainer is used to determine if the trigger should fire.

Returns `True` if the corresponding extension should be invoked in this iteration.

Return type `bool`

`get_training_length()`

`serialize` (*serializer*)

`chainer.training.triggers.ManualScheduleTrigger`

class `chainer.training.triggers.ManualScheduleTrigger` (*points, unit*)

Trigger invoked at specified point(s) of iterations or epochs.

This trigger accepts iterations or epochs indicated by given point(s). There are two ways to specify the point(s): iteration and epoch. *iteration* means the number of updates, while *epoch* means the number of sweeps over the training dataset. Fractional values are allowed if the point is a number of epochs; the trigger uses the *iteration* and *epoch_detail* attributes defined by the updater.

Parameters

- **points** (*int, float, or list of int or float*) – time of the trigger. Must be an integer or list of integer if unit is 'iteration'.

- **unit** (*str*) – Unit of the time specified by `points`. It must be either `'iteration'` or `'epoch'`.

Methods

`__call__` (*trainer*)

Decides whether the extension should be called on this iteration.

Parameters **trainer** (*Trainer*) – Trainer object that this trigger is associated with. The updater associated with this trainer is used to determine if the trigger should fire.

Returns `True` if the corresponding extension should be invoked in this iteration.

Return type `bool`

serialize (*serializer*)

`chainer.training.triggers.MaxValueTrigger`

class `chainer.training.triggers.MaxValueTrigger` (*key, trigger=(1, 'epoch')*)

Trigger invoked when specific value becomes maximum.

For example you can use this trigger to take snapshot on the epoch the validation accuracy is maximum.

Parameters

- **key** (*str*) – Key of value. The trigger fires when the value associated with this key becomes maximum.
- **trigger** – Trigger that decides the comparison interval between current best value and new value. This must be a tuple in the form of `<int>, 'epoch'` or `<int>, 'iteration'` which is passed to `IntervalTrigger`.

Methods

`__call__` (*trainer*)

Decides whether the extension should be called on this iteration.

Parameters **trainer** (*Trainer*) – Trainer object that this trigger is associated with. The observation of this trainer is used to determine if the trigger should fire.

Returns `True` if the corresponding extension should be invoked in this iteration.

Return type `bool`

`chainer.training.triggers.MinValueTrigger`

class `chainer.training.triggers.MinValueTrigger` (*key, trigger=(1, 'epoch')*)

Trigger invoked when specific value becomes minimum.

For example you can use this trigger to take snapshot on the epoch the validation loss is minimum.

Parameters

- **key** (*str*) – Key of value. The trigger fires when the value associated with this key becomes minimum.

- **trigger** – Trigger that decides the comparison interval between current best value and new value. This must be a tuple in the form of `<int>`, `'epoch'` or `<int>`, `'iteration'` which is passed to `IntervalTrigger`.

Methods

`__call__(trainer)`

Decides whether the extension should be called on this iteration.

Parameters **trainer** (`Trainer`) – Trainer object that this trigger is associated with. The observation of this trainer is used to determine if the trigger should fire.

Returns `True` if the corresponding extension should be invoked in this iteration.

Return type `bool`

chainer.training.triggers.TimeTrigger

`class chainer.training.triggers.TimeTrigger(period)`

Trigger based on a fixed time interval.

This trigger accepts iterations with a given interval time.

Parameters **period** (`float`) – Interval time. It is given in seconds.

Methods

`__call__(trainer)`

Call self as a function.

`serialize(serializer)`

5.7 Datasets

5.7.1 Dataset Abstraction (chainer.dataset)

Chainer supports a common interface for training and validation of datasets. The dataset support consists of three components: datasets, iterators, and batch conversion functions.

Dataset represents a set of examples. The interface is only determined by combination with iterators you want to use on it. The built-in iterators of Chainer require the dataset to support `__getitem__` and `__len__` methods. In particular, the `__getitem__` method should support indexing by both an integer and a slice. We can easily support slice indexing by inheriting `DatasetMixin`, in which case users only have to implement `get_example()` method for indexing. Basically, datasets are considered as *stateless* objects, so that we do not need to save the dataset as a checkpoint of the training procedure.

Iterator iterates over the dataset, and at each iteration, it yields a mini-batch of examples as a list. Iterators should support the `Iterator` interface, which includes the standard iterator protocol of Python. Iterators manage where to read next, which means they are *stateful*.

Batch conversion function converts the mini-batch into arrays to feed to the neural nets. They are also responsible to send each array to an appropriate device. Chainer currently provides two implementations:

- `concat_examples()` is a plain implementation which is used as the default choice.

- `ConcatWithAsyncTransfer` is a variant which is basically same as `concat_examples()` except that it overlaps other GPU computations and data transfer for the next iteration.

These components are all customizable, and designed to have a minimum interface to restrict the types of datasets and ways to handle them. In most cases, though, implementations provided by Chainer itself are enough to cover the usages.

Chainer also has a light system to download, manage, and cache concrete examples of datasets. All datasets managed through the system are saved under *the dataset root directory*, which is determined by the `CHAINER_DATASET_ROOT` environment variable, and can also be set by the `set_dataset_root()` function.

Dataset Representation

See *Dataset Examples* (`chainer.datasets`) for dataset implementations.

`chainer.dataset.DatasetMixin`Default implementation of dataset indexing.

`chainer.dataset.DatasetMixin`

class `chainer.dataset.DatasetMixin`

Default implementation of dataset indexing.

`DatasetMixin` provides the `__getitem__()` operator. The default implementation uses `get_example()` to extract each example, and combines the results into a list. This mixin makes it easy to implement a new dataset that does not support efficient slicing.

Dataset implementation using `DatasetMixin` still has to provide the `__len__()` operator explicitly.

Methods

`__getitem__(index)`

Returns an example or a sequence of examples.

It implements the standard Python indexing and one-dimensional integer array indexing. It uses the `get_example()` method by default, but it may be overridden by the implementation to, for example, improve the slicing performance.

Parameters `index` (`int`, `slice`, `list` or `numpy.ndarray`) – An index of an example or indexes of examples.

Returns If `index` is `int`, returns an example created by `get_example`. If `index` is either `slice` or one-dimensional `list` or `numpy.ndarray`, returns a list of examples created by `get_example`.

Example

```
>>> import numpy
>>> from chainer import dataset
>>> class SimpleDataset(dataset.DatasetMixin):
...     def __init__(self, values):
...         self.values = values
...     def __len__(self):
...         return len(self.values)
...     def get_example(self, i):
...         return self.values[i]
```

(continues on next page)

(continued from previous page)

```

...
>>> ds = SimpleDataset([0, 1, 2, 3, 4, 5])
>>> ds[1]    # Access by int
1
>>> ds[1:3]  # Access by slice
[1, 2]
>>> ds[[4, 0]] # Access by one-dimensional integer list
[4, 0]
>>> index = numpy.arange(3)
>>> ds[index] # Access by one-dimensional integer numpy.ndarray
[0, 1, 2]

```

__len__()

Returns the number of data points.

get_example(i)

Returns the i-th example.

Implementations should override it. It should raise `IndexError` if the index is invalid.

Parameters *i* (*int*) – The index of the example.

Returns The i-th example.

Iterator Interface

See [Iterator](#) for dataset iterator implementations.

chainer.dataset.Iterator

Base class of all dataset iterators.

chainer.dataset.Iterator

class `chainer.dataset.Iterator`

Base class of all dataset iterators.

Iterator iterates over the dataset, yielding a minibatch at each iteration. Minibatch is a list of examples. Each implementation should implement an iterator protocol (e.g., the `__next__()` method).

Note that, even if the iterator supports setting the batch size, it does not guarantee that each batch always contains the same number of examples. For example, if you let the iterator to stop at the end of the sweep, the last batch may contain a fewer number of examples.

The interface between the iterator and the underlying dataset is not fixed, and up to the implementation.

Each implementation should provide the following attributes (not needed to be writable).

- `batch_size`: Number of examples within each minibatch.
- `epoch`: Number of completed sweeps over the dataset.
- `epoch_detail`: Floating point number version of the epoch. For example, if the iterator is at the middle of the dataset at the third epoch, then this value is 2.5.
- `previous_epoch_detail`: The value of `epoch_detail` at the previous iteration. This value is `None` before the first iteration.
- `is_new_epoch`: True if the epoch count was incremented at the last update.

Each implementation should also support serialization to resume/suspend the iteration.

Methods

`__enter__()`

With statement context manager method

This method does nothing by default. Implementation may override it to better handle the internal resources by with statement.

`__exit__(exc_type, exc_value, traceback)`

With statement context manager method

This method does nothing by default. Implementation may override it to better handle the internal resources by with statement.

`__next__()`

Returns the next batch.

This is a part of the iterator protocol of Python. It may raise the `StopIteration` exception when it stops the iteration.

`__iter__()`

Returns self.

`finalize()`

Finalizes the iterator and possibly releases the resources.

This method does nothing by default. Implementation may override it to better handle the internal resources.

`next()`

Python2 alternative of `__next__`.

It calls `__next__()` by default.

`serialize(serializer)`

Serializes the internal state of the iterator.

This is a method to support serializer protocol of Chainer.

Note: It should only serialize the internal state that changes over the iteration. It should not serializes what is set manually by users such as the batch size.

Batch Conversion Function

<code>chainer.dataset.concat_examples</code>	Concatenates a list of examples into array(s).
<code>chainer.dataset.ConcatWithAsyncTransfer</code>	Interface to concatenate data and transfer them to GPU asynchronously.
<code>chainer.dataset.to_device</code>	Send an array to a given device.

`chainer.dataset.concat_examples`

`chainer.dataset.concat_examples(batch, device=None, padding=None)`

Concatenates a list of examples into array(s).

This function converts an “array of tuples” into a “tuple of arrays”. Specifically, given a list of examples each of which consists of a list of elements, this function first makes an array by taking the element in the same position from each example and concatenates them along the newly-inserted first axis (called *batch dimension*) into one array. It repeats this for all positions and returns the resulting arrays.

The output type depends on the type of examples in *batch*. For instance, consider each example consists of two arrays (*x*, *y*). Then, this function concatenates *x* ‘s into one array, and *y* ‘s into another array, and returns a tuple of these two arrays. Another example: consider each example is a dictionary of two entries whose keys are ‘*x*’ and ‘*y*’, respectively, and values are arrays. Then, this function concatenates *x* ‘s into one array, and *y* ‘s into another array, and returns a dictionary with two entries *x* and *y* whose values are the concatenated arrays.

When the arrays to concatenate have different shapes, the behavior depends on the *padding* value. If *padding* is *None* (default), it raises an error. Otherwise, it builds an array of the minimum shape that the contents of all arrays can be substituted to. The padding value is then used to the extra elements of the resulting arrays.

Example

```
>>> import numpy as np
>>> from chainer import dataset
>>> x = ([1, 2], 1),
...      ([3, 4], 2),
...      ([5, 6], 3)]
>>> dataset.concat_examples(x)
(array([[1, 2],
        [3, 4],
        [5, 6]]), array([1, 2, 3]))
>>>
>>> y = (np.array([1, 2]), 0),
...      (np.array([3]), 1),
...      (np.array([ ]), 2)]
>>> dataset.concat_examples(y, padding=100)
(array([[ 1,  2],
        [ 3, 100],
        [100, 100]]), array([0, 1, 2]))
>>>
>>> z = (np.array([1, 2]), np.array([0])),
...      (np.array([3]), np.array([ ])),
...      (np.array([ ]), np.array([2])))
>>> dataset.concat_examples(z, padding=(100, 200))
(array([[ 1,  2],
        [ 3, 100],
        [100, 100]]), array([[ 0],
        [200],
        [ 2]]))
>>> w = [{'feature': np.array([1, 2]), 'label': 0},
...      {'feature': np.array([3, 4]), 'label': 1},
...      {'feature': np.array([5, 6]), 'label': 2}]
>>> dataset.concat_examples(w)
{'feature': array([[1, 2],
        [3, 4],
        [5, 6]]), 'label': array([0, 1, 2])}
```

Parameters

- **batch** (*list*) – A list of examples. This is typically given by a dataset iterator.

- **device** (*int*) – Device ID to which each array is sent. Negative value indicates the host memory (CPU). If it is omitted, all arrays are left in the original device.
- **padding** – Scalar value for extra elements. If this is None (default), an error is raised on shape mismatch. Otherwise, an array of minimum dimensionalities that can accommodate all arrays is created, and elements outside of the examples are padded by this value.

Returns Array, a tuple of arrays, or a dictionary of arrays. The type depends on the type of each example in the batch.

chainer.dataset.ConcatWithAsyncTransfer

class chainer.dataset.ConcatWithAsyncTransfer (*stream=None*)

Interface to concatenate data and transfer them to GPU asynchronously.

It enables to transfer next batch of input data to GPU while GPU is running kernels for training using current batch of input data.

An instance of this class is mainly intended to be used as a converter function of an updater like below.

```
from chainer.dataset import convert
...
updater = chainer.training.updaters.StandardUpdater(
    ...,
    converter=convert.ConcatWithAsyncTransfer(),
    ...)
```

Parameters **stream** (*cupy.cuda.Stream*) – CUDA stream. If None, a stream is automatically created on the first call. Data transfer operation is launched asynchronously using the stream.

Methods

__call__ (*batch, device=None, padding=None*)

Concatenate data and transfer them to GPU asynchronously.

See also [chainer.dataset.concat_examples\(\)](#).

Parameters

- **batch** (*list*) – A list of examples.
- **device** (*int*) – Device ID to which each array is sent.
- **padding** – Scalar value for extra elements.

Returns Array, a tuple of arrays, or a dictionary of arrays. The type depends on the type of each example in the batch.

chainer.dataset.to_device

chainer.dataset.to_device (*device, x*)

Send an array to a given device.

This method sends a given array to a given device. This method is used in [concat_examples\(\)](#). You can also use this method in a custom converter method used in [Updater](#) and [Extension](#) such as [StandardUpdater](#) and [Evaluator](#).

See also `chainer.dataset.concat_examples()`.

Parameters

- **device** (*int* or *None*) – Device ID to which an array is sent. If it is negative value, an array is sent to CPU. If it is positive, an array is sent to GPU with the given ID. If it is *None*, an array is left in the original device.
- **x** (*numpy.ndarray* or *cupy.ndarray*) – An array to send.

Returns Converted array.

Dataset Management

<code>chainer.dataset.get_dataset_root</code>	Gets the path to the root directory to download and cache datasets.
<code>chainer.dataset.set_dataset_root</code>	Sets the root directory to download and cache datasets.
<code>chainer.dataset.cached_download</code>	Downloads a file and caches it.
<code>chainer.dataset.cache_or_load_file</code>	Caches a file if it does not exist, or loads it otherwise.

chainer.dataset.get_dataset_root

`chainer.dataset.get_dataset_root()`

Gets the path to the root directory to download and cache datasets.

Returns The path to the dataset root directory.

Return type `str`

chainer.dataset.set_dataset_root

`chainer.dataset.set_dataset_root(path)`

Sets the root directory to download and cache datasets.

There are two ways to set the dataset root directory. One is by setting the environment variable `CHAINER_DATASET_ROOT`. The other is by using this function. If both are specified, one specified via this function is used. The default dataset root is `$HOME/.chainer/dataset`.

Parameters `path` (*str*) – Path to the new dataset root directory.

chainer.dataset.cached_download

`chainer.dataset.cached_download(url)`

Downloads a file and caches it.

It downloads a file from the URL if there is no corresponding cache. After the download, this function stores a cache to the directory under the dataset root (see `set_dataset_root()`). If there is already a cache for the given URL, it just returns the path to the cache without downloading the same file.

Note: This function raises `OSError` when it fails to create the cache directory. In older version, it raised `RuntimeError`.

Parameters `url` (*str*) – URL to download from.

Returns Path to the downloaded file.

Return type `str`

`chainer.dataset.cache_or_load_file`

`chainer.dataset.cache_or_load_file(path, creator, loader)`

Caches a file if it does not exist, or loads it otherwise.

This is a utility function used in dataset loading routines. The `creator` creates the file to given path, and returns the content. If the file already exists, the `loader` is called instead, and it loads the file and returns the content.

Note that the path passed to the creator is temporary one, and not same as the path given to this function. This function safely renames the file created by the creator to a given path, even if this function is called simultaneously by multiple threads or processes.

Parameters

- **path** (`str`) – Path to save the cached file.
- **creator** – Function to create the file and returns the content. It takes a path to temporary place as the argument. Before calling the creator, there is no file at the temporary path.
- **loader** – Function to load the cached file and returns the content.

Returns It returns the returned values by the creator or the loader.

5.7.2 Dataset Examples (`chainer.datasets`)

The most basic `dataset` implementation is an array. Both NumPy and CuPy arrays can be used directly as datasets.

In many cases, though, the simple arrays are not enough to write the training procedure. In order to cover most of such cases, Chainer provides many built-in implementations of datasets.

These built-in datasets are divided into two groups. One is a group of general datasets. Most of them are wrapper of other datasets to introduce some structures (e.g., tuple or dict) to each data point. The other one is a group of concrete, popular datasets. These concrete examples use the downloading utilities in the `chainer.dataset` module to cache downloaded and converted datasets.

5.7.3 General Datasets

General datasets are further divided into four types.

The first one is `DictDataset` and `TupleDataset`, both of which combine other datasets and introduce some structures on them.

The second one is `ConcatenatedDataset` and `SubDataset`. `ConcatenatedDataset` represents a concatenation of existing datasets. It can be used to merge datasets and make a larger dataset. `SubDataset` represents a subset of an existing dataset. It can be used to separate a dataset for hold-out validation or cross validation. Convenient functions to make random splits are also provided.

The third one is `TransformDataset`, which wraps around a dataset by applying a function to data indexed from the underlying dataset. It can be used to modify behavior of a dataset that is already prepared.

The last one is a group of domain-specific datasets. Currently, `ImageDataset` and `LabeledImageDataset` are provided for datasets of images.

DictDataset

*chainer.datasets.DictDataset*Dataset of a dictionary of datasets.

chainer.datasets.DictDataset

class `chainer.datasets.DictDataset` (***datasets*)

Dataset of a dictionary of datasets.

It combines multiple datasets into one dataset. Each example is represented by a dictionary mapping a key to an example of the corresponding dataset.

Parameters **datasets** – Underlying datasets. The keys are used as the keys of each example. All datasets must have the same length.

Methods

`__getitem__` (*index*)`__len__` ()

TupleDataset

*chainer.datasets.TupleDataset*Dataset of tuples from multiple equal-length datasets.

chainer.datasets.TupleDataset

class `chainer.datasets.TupleDataset` (**datasets*)

Dataset of tuples from multiple equal-length datasets.

A `TupleDataset` combines multiple equal-length datasets into a single dataset of tuples. The *i*-th tuple contains the *i*-th example from each of the argument datasets, in the same order that they were supplied.

Recall that in Chainer, a dataset is defined as an iterable that supports both `__getitem__` and `__len__`. The `__getitem__` method should support indexing by both an integer and a slice.

As an example, consider creating a `TupleDataset` from two argument datasets `d1 = [8, 0, 5, 1]` and `d2 = [3, 1, 7, 4]` as `tuple_dataset = TupleDataset(d1, d2)`. The `tuple_dataset` will then contain the examples `(8, 3)`, `(0, 1)`, `(5, 7)`, `(1, 4)`. Note that this behavior is similar to that of the built-in `zip()` function.

Parameters **datasets** – Underlying datasets that will be aggregated. Each dataset must be an iterable that implements `__getitem__` and `__len__`. The *j*-th dataset will be used for the *j*-th item of each example tuple. All datasets must have the same length.

Methods

`__getitem__` (*index*)`__len__` ()

ConcatenatedDataset

<code>chainer.datasets.ConcatenatedDataset</code>	Dataset which concatenates some base datasets.
---	--

chainer.datasets.ConcatenatedDataset

class `chainer.datasets.ConcatenatedDataset` (*datasets)

Dataset which concatenates some base datasets.

This dataset wraps some base datasets and works as a concatenated dataset. For example, if a base dataset with 10 samples and another base dataset with 20 samples are given, this dataset works as a dataset which has 30 samples.

Parameters **datasets** – The underlying datasets. Each dataset has to support `__len__()` and `__getitem__()`.

Methods

`__getitem__` (index)

Returns an example or a sequence of examples.

It implements the standard Python indexing and one-dimensional integer array indexing. It uses the `get_example()` method by default, but it may be overridden by the implementation to, for example, improve the slicing performance.

Parameters **index** (*int*, *slice*, *list* or *numpy.ndarray*) – An index of an example or indexes of examples.

Returns If index is int, returns an example created by `get_example`. If index is either slice or one-dimensional list or `numpy.ndarray`, returns a list of examples created by `get_example`.

Example

```
>>> import numpy
>>> from chainer import dataset
>>> class SimpleDataset(dataset.DatasetMixin):
...     def __init__(self, values):
...         self.values = values
...     def __len__(self):
...         return len(self.values)
...     def get_example(self, i):
...         return self.values[i]
...
>>> ds = SimpleDataset([0, 1, 2, 3, 4, 5])
>>> ds[1]      # Access by int
1
>>> ds[1:3]    # Access by slice
[1, 2]
>>> ds[[4, 0]] # Access by one-dimensional integer list
[4, 0]
>>> index = numpy.arange(3)
>>> ds[index]  # Access by one-dimensional integer numpy.ndarray
[0, 1, 2]
```

`__len__()`

Returns the number of data points.

`get_example(i)`

Returns the i -th example.

Implementations should override it. It should raise `IndexError` if the index is invalid.

Parameters `i` (*int*) – The index of the example.

Returns The i -th example.

SubDataset

<code>chainer.datasets.SubDataset</code>	Subset of a base dataset.
<code>chainer.datasets.split_dataset</code>	Splits a dataset into two subsets.
<code>chainer.datasets.split_dataset_random</code>	Splits a dataset into two subsets randomly.
<code>chainer.datasets.get_cross_validation_datasets</code>	Creates a set of training/test splits for cross validation.
<code>chainer.datasets.get_cross_validation_datasets_random</code>	Creates a set of training/test splits for cross validation randomly.

chainer.datasets.SubDataset

class `chainer.datasets.SubDataset` (*dataset, start, finish, order=None*)

Subset of a base dataset.

`SubDataset` defines a subset of a given base dataset. The subset is defined as an interval of indexes, optionally with a given permutation.

If `order` is given, then the i -th example of this dataset is the `order[start + i]`-th example of the base dataset, where i is a non-negative integer. If `order` is not given, then the i -th example of this dataset is the `start + i`-th example of the base dataset. Negative indexing is also allowed: in this case, the term `start + i` is replaced by `finish + i`.

`SubDataset` is often used to split a dataset into training and validation subsets. The training set is used for training, while the validation set is used to track the generalization performance, i.e. how the learned model works well on unseen data. We can tune hyperparameters (e.g. number of hidden units, weight initializers, learning rate, etc.) by comparing the validation performance. Note that we often use another set called test set to measure the quality of the tuned hyperparameter, which can be made by nesting multiple `SubDatasets`.

There are two ways to make training-validation splits. One is a single split, where the dataset is split just into two subsets. It can be done by `split_dataset()` or `split_dataset_random()`. The other one is a k -fold cross validation, in which the dataset is divided into k subsets, and k different splits are generated using each of the k subsets as a validation set and the rest as a training set. It can be done by `get_cross_validation_datasets()`.

Parameters

- **dataset** – Base dataset.
- **start** (*int*) – The first index in the interval.
- **finish** (*int*) – The next-to-the-last index in the interval.
- **order** (*sequence of ints*) – Permutation of indexes in the base dataset. If this is `None`, then the ascending order of indexes is used.

Methods

__getitem__ (*index*)

Returns an example or a sequence of examples.

It implements the standard Python indexing and one-dimensional integer array indexing. It uses the `get_example()` method by default, but it may be overridden by the implementation to, for example, improve the slicing performance.

Parameters *index* (*int*, *slice*, *list* or *numpy.ndarray*) – An index of an example or indexes of examples.

Returns If *index* is *int*, returns an example created by `get_example`. If *index* is either *slice* or one-dimensional *list* or *numpy.ndarray*, returns a list of examples created by `get_example`.

Example

```
>>> import numpy
>>> from chainer import dataset
>>> class SimpleDataset(dataset.DatasetMixin):
...     def __init__(self, values):
...         self.values = values
...     def __len__(self):
...         return len(self.values)
...     def get_example(self, i):
...         return self.values[i]
...
>>> ds = SimpleDataset([0, 1, 2, 3, 4, 5])
>>> ds[1]      # Access by int
1
>>> ds[1:3]    # Access by slice
[1, 2]
>>> ds[[4, 0]] # Access by one-dimensional integer list
[4, 0]
>>> index = numpy.arange(3)
>>> ds[index]  # Access by one-dimensional integer numpy.ndarray
[0, 1, 2]
```

__len__ ()

Returns the number of data points.

get_example (*i*)

Returns the *i*-th example.

Implementations should override it. It should raise `IndexError` if the index is invalid.

Parameters *i* (*int*) – The index of the example.

Returns The *i*-th example.

chainer.datasets.split_dataset

`chainer.datasets.split_dataset` (*dataset*, *split_at*, *order=None*)

Splits a dataset into two subsets.

This function creates two instances of `SubDataset`. These instances do not share any examples, and they together cover all examples of the original dataset.

Parameters

- **dataset** – Dataset to split.
- **split_at** (*int*) – Position at which the base dataset is split.
- **order** (*sequence of ints*) – Permutation of indexes in the base dataset. See the document of *SubDataset* for details.

Returns Two *SubDataset* objects. The first subset represents the examples of indexes `order[:split_at]` while the second subset represents the examples of indexes `order[split_at:]`.

Return type `tuple`

chainer.datasets.split_dataset_random

`chainer.datasets.split_dataset_random(dataset, first_size, seed=None)`

Splits a dataset into two subsets randomly.

This function creates two instances of *SubDataset*. These instances do not share any examples, and they together cover all examples of the original dataset. The split is automatically done randomly.

Parameters

- **dataset** – Dataset to split.
- **first_size** (*int*) – Size of the first subset.
- **seed** (*int*) – Seed the generator used for the permutation of indexes. If an integer being convertible to 32 bit unsigned integers is specified, it is guaranteed that each sample in the given dataset always belongs to a specific subset. If `None`, the permutation is changed randomly.

Returns Two *SubDataset* objects. The first subset contains `first_size` examples randomly chosen from the dataset without replacement, and the second subset contains the rest of the dataset.

Return type `tuple`

chainer.datasets.get_cross_validation_datasets

`chainer.datasets.get_cross_validation_datasets(dataset, n_fold, order=None)`

Creates a set of training/test splits for cross validation.

This function generates `n_fold` splits of the given dataset. The first part of each split corresponds to the training dataset, while the second part to the test dataset. No pairs of test datasets share any examples, and all test datasets together cover the whole base dataset. Each test dataset contains almost same number of examples (the numbers may differ up to 1).

Parameters

- **dataset** – Dataset to split.
- **n_fold** (*int*) – Number of splits for cross validation.
- **order** (*sequence of ints*) – Order of indexes with which each split is determined. If it is `None`, then no permutation is used.

Returns List of dataset splits.

Return type list of tuples

`chainer.datasets.get_cross_validation_datasets_random`

`chainer.datasets.get_cross_validation_datasets_random(dataset, n_fold, seed=None)`

Creates a set of training/test splits for cross validation randomly.

This function acts almost same as `get_cross_validation_dataset()`, except automatically generating random permutation.

Parameters

- **dataset** – Dataset to split.
- **n_fold** (*int*) – Number of splits for cross validation.
- **seed** (*int*) – Seed the generator used for the permutation of indexes. If an integer beging convertible to 32 bit unsigned integers is specified, it is guaranteed that each sample in the given dataset always belongs to a specific subset. If `None`, the permutation is changed randomly.

Returns List of dataset splits.

Return type list of tuples

TransformDataset

<code>chainer.datasets.TransformDataset</code>	Dataset that indexes the base dataset and transforms the data.
--	--

`chainer.datasets.TransformDataset`

class `chainer.datasets.TransformDataset(dataset, transform)`

Dataset that indexes the base dataset and transforms the data.

This dataset wraps the base dataset by modifying the behavior of the base dataset's `__getitem__()`. Arrays returned by `__getitem__()` of the base dataset with integer as an argument are transformed by the given function `transform`. Also, `__len__()` returns the integer returned by the base dataset's `__len__()`.

The function `transform` takes, as an argument, `in_data`, which is the output of the base dataset's `__getitem__()`, and returns the transformed arrays as output. Please see the following example. Since `in_data` directly refers to the item in the dataset, take care that `transform` not modify it. For example, note that the line `img = img - 0.5` bellow is correct since it makes a copy of `img`. However, it would be incorrect to use `img -= 0.5` since that would update the contents of the item in the dataset in place, corrupting it.

```
>>> from chainer.datasets import get_mnist
>>> from chainer.datasets import TransformDataset
>>> dataset, _ = get_mnist()
>>> def transform(in_data):
...     img, label = in_data
...     img = img - 0.5 # scale to [-0.5, -0.5]
...     return img, label
>>> dataset = TransformDataset(dataset, transform)
```

Parameters

- **dataset** – The underlying dataset. The index of this dataset corresponds to the index of the base dataset. This object needs to support functions `__getitem__()` and `__len__()` as described above.
- **transform** (*callable*) – A function that is called to transform values returned by the underlying dataset's `__getitem__()`.

Methods

`__getitem__` (*index*)

Returns an example or a sequence of examples.

It implements the standard Python indexing and one-dimensional integer array indexing. It uses the `get_example()` method by default, but it may be overridden by the implementation to, for example, improve the slicing performance.

Parameters *index* (*int*, *slice*, *list* or *numpy.ndarray*) – An index of an example or indexes of examples.

Returns If *index* is *int*, returns an example created by `get_example`. If *index* is either *slice* or one-dimensional *list* or *numpy.ndarray*, returns a list of examples created by `get_example`.

Example

```
>>> import numpy
>>> from chainer import dataset
>>> class SimpleDataset(dataset.DatasetMixin):
...     def __init__(self, values):
...         self.values = values
...     def __len__(self):
...         return len(self.values)
...     def get_example(self, i):
...         return self.values[i]
...
>>> ds = SimpleDataset([0, 1, 2, 3, 4, 5])
>>> ds[1]      # Access by int
1
>>> ds[1:3]    # Access by slice
[1, 2]
>>> ds[[4, 0]] # Access by one-dimensional integer list
[4, 0]
>>> index = numpy.arange(3)
>>> ds[index]  # Access by one-dimensional integer numpy.ndarray
[0, 1, 2]
```

`__len__` ()

Returns the number of data points.

`get_example` (*i*)

Returns the *i*-th example.

Implementations should override it. It should raise `IndexError` if the index is invalid.

Parameters *i* (*int*) – The index of the example.

Returns The *i*-th example.

ImageDataset

<code>chainer.datasets.ImageDataset</code>	Dataset of images built from a list of paths to image files.
<code>chainer.datasets.ZippedImageDataset</code>	Dataset of images built from a zip file.
<code>chainer.datasets.MultiZippedImageDataset</code>	Dataset of images built from a list of paths to zip files.

chainer.datasets.ImageDataset

class `chainer.datasets.ImageDataset` (*paths*, *root*='.', *dtype*=<class 'numpy.float32'>)
 Dataset of images built from a list of paths to image files.

This dataset reads an external image file on every call of the `__getitem__()` operator. The paths to the image to retrieve is given as either a list of strings or a text file that contains paths in distinct lines.

Each image is automatically converted to arrays of shape `channels, height, width`, where `channels` represents the number of channels in each pixel (e.g., 1 for grey-scale images, and 3 for RGB-color images).

Note: This dataset requires the Pillow package being installed. In order to use this dataset, install Pillow (e.g. by using the command `pip install Pillow`). Be careful to prepare appropriate libraries for image formats you want to use (e.g. `libpng` for PNG images, and `libjpeg` for JPG images).

Warning: You are responsible for preprocessing the images before feeding them to a model. For example, if your dataset contains both RGB and grayscale images, make sure that you convert them to the same format. Otherwise you will get errors because the input dimensions are different for RGB and grayscale images.

Parameters

- **paths** (*str* or *list of strs*) – If it is a string, it is a path to a text file that contains paths to images in distinct lines. If it is a list of paths, the *i*-th element represents the path to the *i*-th image. In both cases, each path is a relative one from the root path given by another argument.
- **root** (*str*) – Root directory to retrieve images from.
- **dtype** – Data type of resulting image arrays.

Methods

`__getitem__` (*index*)

Returns an example or a sequence of examples.

It implements the standard Python indexing and one-dimensional integer array indexing. It uses the `get_example()` method by default, but it may be overridden by the implementation to, for example, improve the slicing performance.

Parameters *index* (*int*, *slice*, *list* or *numpy.ndarray*) – An index of an example or indexes of examples.

Returns If index is int, returns an example created by `get_example`. If index is either slice or one-dimensional list or `numpy.ndarray`, returns a list of examples created by `get_example`.

Example

```
>>> import numpy
>>> from chainer import dataset
>>> class SimpleDataset(dataset.DatasetMixin):
...     def __init__(self, values):
...         self.values = values
...     def __len__(self):
...         return len(self.values)
...     def get_example(self, i):
...         return self.values[i]
...
>>> ds = SimpleDataset([0, 1, 2, 3, 4, 5])
>>> ds[1]      # Access by int
1
>>> ds[1:3]    # Access by slice
[1, 2]
>>> ds[[4, 0]] # Access by one-dimensional integer list
[4, 0]
>>> index = numpy.arange(3)
>>> ds[index]  # Access by one-dimensional integer numpy.ndarray
[0, 1, 2]
```

`__len__()`

Returns the number of data points.

`get_example(i)`

Returns the i-th example.

Implementations should override it. It should raise `IndexError` if the index is invalid.

Parameters `i` (*int*) – The index of the example.

Returns The i-th example.

chainer.datasets.ZippedImageDataset

class `chainer.datasets.ZippedImageDataset` (*zipfilename*, *dtype*=<class 'numpy.float32'>)

Dataset of images built from a zip file.

This dataset reads an external image file in the given zipfile. The zipfile shall contain only image files. This shall be able to replace `ImageDataset` and works better on NFS and other networked file systems. If zipfile becomes too large you may consider `MultiZippedImageDataset` as a handy alternative.

Known issue: pickle and unpickle on same process may cause race condition on `ZipFile`. Pickle of this class is expected to be sent to different processes via `ChainerMN`.

Parameters

- **zipfilename** (*str*) – a string to point zipfile path
- **dtype** – Data type of resulting image arrays

Methods

`__getitem__(index)`

Returns an example or a sequence of examples.

It implements the standard Python indexing and one-dimensional integer array indexing. It uses the `get_example()` method by default, but it may be overridden by the implementation to, for example, improve the slicing performance.

Parameters `index` (*int*, *slice*, *list* or *numpy.ndarray*) – An index of an example or indexes of examples.

Returns If index is int, returns an example created by `get_example`. If index is either slice or one-dimensional list or `numpy.ndarray`, returns a list of examples created by `get_example`.

Example

```
>>> import numpy
>>> from chainer import dataset
>>> class SimpleDataset(dataset.DatasetMixin):
...     def __init__(self, values):
...         self.values = values
...     def __len__(self):
...         return len(self.values)
...     def get_example(self, i):
...         return self.values[i]
...
>>> ds = SimpleDataset([0, 1, 2, 3, 4, 5])
>>> ds[1]      # Access by int
1
>>> ds[1:3]    # Access by slice
[1, 2]
>>> ds[[4, 0]] # Access by one-dimensional integer list
[4, 0]
>>> index = numpy.arange(3)
>>> ds[index]  # Access by one-dimensional integer numpy.ndarray
[0, 1, 2]
```

`__len__()`

Returns the number of data points.

`get_example(i)`

Returns the i-th example.

Implementations should override it. It should raise `IndexError` if the index is invalid.

Parameters `i` (*int*) – The index of the example.

Returns The i-th example.

`chainer.datasets.MultiZippedImageDataset`

```
class chainer.datasets.MultiZippedImageDataset(zipfilenames, dtype=<class
                                              'numpy.float32'>)
```

Dataset of images built from a list of paths to zip files.

This dataset reads an external image file in given zipfiles. The zipfiles shall contain only image files. This shall be able to replace ImageDataset and works better on NFS and other networked file systems. The user shall find good balance between zipfile size and number of zipfiles (e.g. granularity)

Parameters

- **zipfilenames** (*list of strings*) – List of zipped archive filename.
- **dtype** – Data type of resulting image arrays.

Methods

`__getitem__` (*index*)

Returns an example or a sequence of examples.

It implements the standard Python indexing and one-dimensional integer array indexing. It uses the `get_example()` method by default, but it may be overridden by the implementation to, for example, improve the slicing performance.

Parameters *index* (*int*, *slice*, *list* or *numpy.ndarray*) – An index of an example or indexes of examples.

Returns If index is int, returns an example created by `get_example`. If index is either slice or one-dimensional list or `numpy.ndarray`, returns a list of examples created by `get_example`.

Example

```
>>> import numpy
>>> from chainer import dataset
>>> class SimpleDataset(dataset.DatasetMixin):
...     def __init__(self, values):
...         self.values = values
...     def __len__(self):
...         return len(self.values)
...     def get_example(self, i):
...         return self.values[i]
...
>>> ds = SimpleDataset([0, 1, 2, 3, 4, 5])
>>> ds[1]      # Access by int
1
>>> ds[1:3]    # Access by slice
[1, 2]
>>> ds[[4, 0]] # Access by one-dimensional integer list
[4, 0]
>>> index = numpy.arange(3)
>>> ds[index]  # Access by one-dimensional integer numpy.ndarray
[0, 1, 2]
```

`__len__` ()

Returns the number of data points.

`get_example` (*i*)

Returns the *i*-th example.

Implementations should override it. It should raise `IndexError` if the index is invalid.

Parameters *i* (*int*) – The index of the example.

Returns The *i*-th example.

LabeledImageDataset

<code>chainer.datasets.LabeledImageDataset</code>	Dataset of image and label pairs built from a list of paths and labels.
---	---

chainer.datasets.LabeledImageDataset

```
class chainer.datasets.LabeledImageDataset (pairs,          root='',          dtype=<class
                                             'numpy.float32'>,      label_dtype=<class
                                             'numpy.int32'>)
```

Dataset of image and label pairs built from a list of paths and labels.

This dataset reads an external image file like *ImageDataset*. The difference from *ImageDataset* is that this dataset also returns a label integer. The paths and labels are given as either a list of pairs or a text file contains paths/labels pairs in distinct lines. In the latter case, each path and corresponding label are separated by white spaces. This format is same as one used in Caffe.

Note: This dataset requires the Pillow package being installed. In order to use this dataset, install Pillow (e.g. by using the command `pip install Pillow`). Be careful to prepare appropriate libraries for image formats you want to use (e.g. `libpng` for PNG images, and `libjpeg` for JPG images).

Warning: You are responsible for preprocessing the images before feeding them to a model. For example, if your dataset contains both RGB and grayscale images, make sure that you convert them to the same format. Otherwise you will get errors because the input dimensions are different for RGB and grayscale images.

Parameters

- **pairs** (*str* or *list of tuples*) – If it is a string, it is a path to a text file that contains paths to images in distinct lines. If it is a list of pairs, the *i*-th element represents a pair of the path to the *i*-th image and the corresponding label. In both cases, each path is a relative one from the root path given by another argument.
- **root** (*str*) – Root directory to retrieve images from.
- **dtype** – Data type of resulting image arrays.
- **label_dtype** – Data type of the labels.

Methods

`__getitem__` (*index*)

Returns an example or a sequence of examples.

It implements the standard Python indexing and one-dimensional integer array indexing. It uses the `get_example()` method by default, but it may be overridden by the implementation to, for example, improve the slicing performance.

Parameters **index** (*int*, *slice*, *list* or *numpy.ndarray*) – An index of an example or indexes of examples.

Returns If index is int, returns an example created by `get_example`. If index is either slice or one-dimensional list or `numpy.ndarray`, returns a list of examples created by `get_example`.

Example

```
>>> import numpy
>>> from chainer import dataset
>>> class SimpleDataset(dataset.DatasetMixin):
...     def __init__(self, values):
...         self.values = values
...     def __len__(self):
...         return len(self.values)
...     def get_example(self, i):
...         return self.values[i]
...
>>> ds = SimpleDataset([0, 1, 2, 3, 4, 5])
>>> ds[1]      # Access by int
1
>>> ds[1:3]    # Access by slice
[1, 2]
>>> ds[[4, 0]] # Access by one-dimensional integer list
[4, 0]
>>> index = numpy.arange(3)
>>> ds[index]  # Access by one-dimensional integer numpy.ndarray
[0, 1, 2]
```

`__len__()`

Returns the number of data points.

`get_example(i)`

Returns the i-th example.

Implementations should override it. It should raise `IndexError` if the index is invalid.

Parameters *i* (*int*) – The index of the example.

Returns The i-th example.

5.7.4 Concrete Datasets

<code>chainer.datasets.get_mnist</code>	Gets the MNIST dataset.
<code>chainer.datasets.get_fashion_mnist</code>	Gets the Fashion-MNIST dataset.
<code>chainer.datasets.get_cifar10</code>	Gets the CIFAR-10 dataset.
<code>chainer.datasets.get_cifar100</code>	Gets the CIFAR-100 dataset.
<code>chainer.datasets.get_ptb_words</code>	Gets the Penn Tree Bank dataset as long word sequences.
<code>chainer.datasets.get_ptb_words_vocabulary</code>	Gets the Penn Tree Bank word vocabulary.
<code>chainer.datasets.get_svhn</code>	Gets the SVHN dataset.

chainer.datasets.get_mnist

```
chainer.datasets.get_mnist (withlabel=True, ndim=1, scale=1.0, dtype=<class 'numpy.float32'>,
                             label_dtype=<class 'numpy.int32'>, rgb_format=False)
```

Gets the MNIST dataset.

MNIST is a set of hand-written digits represented by grey-scale 28x28 images. In the original images, each pixel is represented by one-byte unsigned integer. This function scales the pixels to floating point values in the interval `[0, scale]`.

This function returns the training set and the test set of the official MNIST dataset. If `withlabel` is `True`, each dataset consists of tuples of images and labels, otherwise it only consists of images.

Parameters

- **withlabel** (*bool*) – If `True`, it returns datasets with labels. In this case, each example is a tuple of an image and a label. Otherwise, the datasets only contain images.
- **ndim** (*int*) – Number of dimensions of each image. The shape of each image is determined depending on `ndim` as follows:
 - `ndim == 1`: the shape is `(784,)`
 - `ndim == 2`: the shape is `(28, 28)`
 - `ndim == 3`: the shape is `(1, 28, 28)`
- **scale** (*float*) – Pixel value scale. If it is 1 (default), pixels are scaled to the interval `[0, 1]`.
- **dtype** – Data type of resulting image arrays.
- **label_dtype** – Data type of the labels.
- **rgb_format** (*bool*) – if `ndim == 3` and `rgb_format` is `True`, the image will be converted to rgb format by duplicating the channels so the image shape is `(3, 28, 28)`. Default is `False`.

Returns A tuple of two datasets. If `withlabel` is `True`, both datasets are *TupleDataset* instances. Otherwise, both datasets are arrays of images.

chainer.datasets.get_fashion_mnist

```
chainer.datasets.get_fashion_mnist (withlabel=True, ndim=1, scale=1.0, dtype=<class
                                     'numpy.float32'>, label_dtype=<class 'numpy.int32'>,
                                     rgb_format=False)
```

Gets the Fashion-MNIST dataset.

Fashion-MNIST is a set of fashion articles represented by grey-scale 28x28 images. In the original images, each pixel is represented by one-byte unsigned integer. This function scales the pixels to floating point values in the interval `[0, scale]`.

This function returns the training set and the test set of the official Fashion-MNIST dataset. If `withlabel` is `True`, each dataset consists of tuples of images and labels, otherwise it only consists of images.

Parameters

- **withlabel** (*bool*) – If `True`, it returns datasets with labels. In this case, each example is a tuple of an image and a label. Otherwise, the datasets only contain images.
- **ndim** (*int*) – Number of dimensions of each image. The shape of each image is determined depending on `ndim` as follows:

- `ndim == 1`: the shape is (784,)
- `ndim == 2`: the shape is (28, 28)
- `ndim == 3`: the shape is (1, 28, 28)
- **scale** (*float*) – Pixel value scale. If it is 1 (default), pixels are scaled to the interval [0, 1].
- **dtype** – Data type of resulting image arrays.
- **label_dtype** – Data type of the labels.
- **rgb_format** (*bool*) – if `ndim == 3` and `rgb_format` is `True`, the image will be converted to rgb format by duplicating the channels so the image shape is (3, 28, 28). Default is `False`.

Returns A tuple of two datasets. If `withlabel` is `True`, both datasets are `TupleDataset` instances. Otherwise, both datasets are arrays of images.

chainer.datasets.get_cifar10

`chainer.datasets.get_cifar10(withlabel=True, ndim=3, scale=1.0)`

Gets the CIFAR-10 dataset.

CIFAR-10 is a set of small natural images. Each example is an RGB color image of size 32x32, classified into 10 groups. In the original images, each component of pixels is represented by one-byte unsigned integer. This function scales the components to floating point values in the interval [0, `scale`].

This function returns the training set and the test set of the official CIFAR-10 dataset. If `withlabel` is `True`, each dataset consists of tuples of images and labels, otherwise it only consists of images.

Parameters

- **withlabel** (*bool*) – If `True`, it returns datasets with labels. In this case, each example is a tuple of an image and a label. Otherwise, the datasets only contain images.
- **ndim** (*int*) – Number of dimensions of each image. The shape of each image is determined depending on `ndim` as follows:
 - `ndim == 1`: the shape is (3072,)
 - `ndim == 3`: the shape is (3, 32, 32)
- **scale** (*float*) – Pixel value scale. If it is 1 (default), pixels are scaled to the interval [0, 1].

Returns A tuple of two datasets. If `withlabel` is `True`, both datasets are `TupleDataset` instances. Otherwise, both datasets are arrays of images.

chainer.datasets.get_cifar100

`chainer.datasets.get_cifar100(withlabel=True, ndim=3, scale=1.0)`

Gets the CIFAR-100 dataset.

CIFAR-100 is a set of small natural images. Each example is an RGB color image of size 32x32, classified into 100 groups. In the original images, each component pixels is represented by one-byte unsigned integer. This function scales the components to floating point values in the interval [0, `scale`].

This function returns the training set and the test set of the official CIFAR-100 dataset. If `withlabel` is `True`, each dataset consists of tuples of images and labels, otherwise it only consists of images.

Parameters

- **withlabel** (*bool*) – If `True`, it returns datasets with labels. In this case, each example is a tuple of an image and a label. Otherwise, the datasets only contain images.
- **ndim** (*int*) – Number of dimensions of each image. The shape of each image is determined depending on `ndim` as follows:
 - `ndim == 1`: the shape is `(3072,)`
 - `ndim == 3`: the shape is `(3, 32, 32)`
- **scale** (*float*) – Pixel value scale. If it is 1 (default), pixels are scaled to the interval `[0, 1]`.

Returns A tuple of two datasets. If `withlabel` is `True`, both are `TupleDataset` instances. Otherwise, both datasets are arrays of images.

chainer.datasets.get_ptb_words

`chainer.datasets.get_ptb_words()`

Gets the Penn Tree Bank dataset as long word sequences.

Penn Tree Bank is originally a corpus of English sentences with linguistic structure annotations. This function uses a variant distributed at <https://github.com/wojzaremba/lstm>, which omits the annotation and splits the dataset into three parts: training, validation, and test.

This function returns the training, validation, and test sets, each of which is represented as a long array of word IDs. All sentences in the dataset are concatenated by End-of-Sentence mark ‘<eos>’, which is treated as one of the vocabulary.

Returns `Int32` vectors of word IDs.

Return type tuple of `numpy.ndarray`

See also:

Use `get_ptb_words_vocabulary()` to get the mapping between the words and word IDs.

chainer.datasets.get_ptb_words_vocabulary

`chainer.datasets.get_ptb_words_vocabulary()`

Gets the Penn Tree Bank word vocabulary.

Returns Dictionary that maps words to corresponding word IDs. The IDs are used in the Penn Tree Bank long sequence datasets.

Return type `dict`

See also:

See `get_ptb_words()` for the actual datasets.

chainer.datasets.get_svhn

`chainer.datasets.get_svhn(withlabel=True, scale=1.0, dtype=<class 'numpy.float32'>, label_dtype=<class 'numpy.int32'>)`

Gets the SVHN dataset.

The [Street View House Numbers \(SVHN\)](#) dataset is a dataset similar to MNIST but composed of cropped images of house numbers. The functionality of this function is identical to the counterpart for the MNIST dataset (`get_mnist()`), with the exception that there is no `ndim` argument.

Note: SciPy is required to use this feature.

Parameters

- **withlabel** (*bool*) – If `True`, it returns datasets with labels. In this case, each example is a tuple of an image and a label. Otherwise, the datasets only contain images.
- **scale** (*float*) – Pixel value scale. If it is 1 (default), pixels are scaled to the interval `[0, 1]`.
- **dtype** – Data type of resulting image arrays.
- **label_dtype** – Data type of the labels.

Returns A tuple of two datasets. If `withlabel` is `True`, both datasets are [TupleDataset](#) instances. Otherwise, both datasets are arrays of images.

5.8 Iterator

Chainer provides some iterators that implement typical strategies to create mini-batches by iterating over datasets. [SerialIterator](#) is the simplest one, which extract mini-batches in the main thread. [MultiprocessIterator](#) and [MultithreadIterator](#) are a parallelized version of [SerialIterator](#). It maintains worker subprocesses and subthreads to load the next mini-batch in parallel.

chainer.iterators.SerialIterator	Dataset iterator that serially reads the examples.
chainer.iterators.MultiprocessIterator	Dataset iterator that loads examples in parallel.
chainer.iterators.MultithreadIterator	Dataset iterator that loads examples in parallel.

5.8.1 chainer.iterators.SerialIterator

class `chainer.iterators.SerialIterator` (*dataset, batch_size, repeat=True, shuffle=True*)

Dataset iterator that serially reads the examples.

This is a simple implementation of [Iterator](#) that just visits each example in either the order of indexes or a shuffled order.

To avoid unintentional performance degradation, the `shuffle` option is set to `True` by default. For validation, it is better to set it to `False` when the underlying dataset supports fast slicing. If the order of examples has an important meaning and the updater depends on the original order, this option should be set to `False`.

This iterator saves `-1` instead of `None` in snapshots since some serializers do not support `None`.

Parameters

- **dataset** – Dataset to iterate.
- **batch_size** (*int*) – Number of examples within each batch.

- **repeat** (*bool*) – If `True`, it infinitely loops over the dataset. Otherwise, it stops iteration at the end of the first epoch.
- **shuffle** (*bool*) – If `True`, the order of examples is shuffled at the beginning of each epoch. Otherwise, examples are extracted in the order of indexes.

Methods

`__enter__()`

With statement context manager method

This method does nothing by default. Implementation may override it to better handle the internal resources by with statement.

`__exit__(exc_type, exc_value, traceback)`

With statement context manager method

This method does nothing by default. Implementation may override it to better handle the internal resources by with statement.

`__next__()`

Returns the next batch.

This is a part of the iterator protocol of Python. It may raise the `StopIteration` exception when it stops the iteration.

`__iter__()`

Returns self.

`finalize()`

Finalizes the iterator and possibly releases the resources.

This method does nothing by default. Implementation may override it to better handle the internal resources.

`next()`

Returns the next batch.

This is a part of the iterator protocol of Python. It may raise the `StopIteration` exception when it stops the iteration.

`reset()`

`serialize(serializer)`

Serializes the internal state of the iterator.

This is a method to support serializer protocol of Chainer.

Note: It should only serialize the internal state that changes over the iteration. It should not serializes what is set manually by users such as the batch size.

Attributes

`epoch_detail`

`previous_epoch_detail`

`repeat`

5.8.2 chainer.iterators.MultiprocessIterator

```
class chainer.iterators.MultiprocessIterator (dataset,      batch_size,      repeat=True,
                                              shuffle=True,      n_processes=None,
                                              n_prefetch=1, shared_mem=None)
```

Dataset iterator that loads examples in parallel.

This is an implementation of `Iterator` that loads examples with worker processes. It uses the standard `multiprocessing` module to parallelize the loading. The dataset is sent to the worker processes in the standard way using pickle.

Note that this iterator effectively prefetches the examples for the next batch asynchronously after the current batch is returned.

This iterator saves `-1` instead of `None` in snapshots since some serializers do not support `None`.

Parameters

- **dataset** (`Dataset`) – Dataset to iterate.
- **batch_size** (`int`) – Number of examples within each batch.
- **repeat** (`bool`) – If `True`, it infinitely loops over the dataset. Otherwise, it stops iteration at the end of the first epoch.
- **shuffle** (`bool`) – If `True`, the order of examples is shuffled at the beginning of each epoch. Otherwise, examples are extracted in the order of indexes.
- **n_processes** (`int`) – Number of worker processes. The number of CPUs is used by default.
- **n_prefetch** (`int`) – Number of prefetch batches.
- **shared_mem** (`int`) – The size of using shared memory per data. If `None`, size is adjusted automatically.

Methods

__enter__ ()

With statement context manager method

This method does nothing by default. Implementation may override it to better handle the internal resources by with statement.

__exit__ (`exc_type`, `exc_value`, `traceback`)

With statement context manager method

This method does nothing by default. Implementation may override it to better handle the internal resources by with statement.

__next__ ()

Returns the next batch.

This is a part of the iterator protocol of Python. It may raise the `StopIteration` exception when it stops the iteration.

__iter__ ()

Returns self.

__copy__ ()

finalize ()

next ()

Returns the next batch.

This is a part of the iterator protocol of Python. It may raise the `StopIteration` exception when it stops the iteration.

reset ()

serialize (*serializer*)

Serializes the internal state of the iterator.

This is a method to support serializer protocol of Chainer.

Note: It should only serialize the internal state that changes over the iteration. It should not serializes what is set manually by users such as the batch size.

Attributes

`epoch_detail`

`previous_epoch_detail`

5.8.3 chainer.iterators.MultithreadIterator

class `chainer.iterators.MultithreadIterator` (*dataset*, *batch_size*, *repeat=True*, *shuffle=True*, *n_threads=1*)

Dataset iterator that loads examples in parallel.

This is an implementation of `Iterator` that loads examples with worker threads. It uses the standard `threading` module to parallelize the loading.

Note that this iterator effectively prefetches the examples for the next batch asynchronously after the current batch is returned.

This iterator saves `-1` instead of `None` in snapshots since some serializers do not support `None`.

Parameters

- **dataset** (*Dataset*) – Dataset to iterate.
- **batch_size** (*int*) – Number of examples within each batch.
- **repeat** (*bool*) – If `True`, it infinitely loops over the dataset. Otherwise, it stops iteration at the end of the first epoch.
- **shuffle** (*bool*) – If `True`, the order of examples is shuffled at the beginning of each epoch. Otherwise, examples are extracted in the order of indexes.
- **n_threads** (*int*) – Number of worker threads.

Methods

`__enter__` ()

With statement context manager method

This method does nothing by default. Implementation may override it to better handle the internal resources by with statement.

__exit__ (*exc_type, exc_value, traceback*)

With statement context manager method

This method does nothing by default. Implementation may override it to better handle the internal resources by with statement.

__next__ ()

Returns the next batch.

This is a part of the iterator protocol of Python. It may raise the `StopIteration` exception when it stops the iteration.

__iter__ ()

Returns self.

finalize ()

Finalizes the iterator and possibly releases the resources.

This method does nothing by default. Implementation may override it to better handle the internal resources.

next ()

Returns the next batch.

This is a part of the iterator protocol of Python. It may raise the `StopIteration` exception when it stops the iteration.

reset ()

serialize (*serializer*)

Serializes the internal state of the iterator.

This is a method to support serializer protocol of Chainer.

Note: It should only serialize the internal state that changes over the iteration. It should not serializes what is set manually by users such as the batch size.

Attributes

epoch_detail

previous_epoch_detail

repeat

5.9 Serializers

5.9.1 Serialization in NumPy NPZ format

NumPy serializers can be used in arbitrary environments that Chainer runs with. It consists of asymmetric serializer/deserializer due to the fact that `numpy.savez()` does not support online serialization. Therefore, serialization requires two-step manipulation: first packing the objects into a flat dictionary, and then serializing it into npz format.

<code>chainer.serializers.</code>	Serializer for dictionary.
<code>DictionarySerializer</code>	
<code>chainer.serializers.NpzDeserializer</code>	Deserializer for NPZ format.
<code>chainer.serializers.save_npz</code>	Saves an object to the file in NPZ format.
<code>chainer.serializers.load_npz</code>	Loads an object from the file in NPZ format.

chainer.serializers.DictionarySerializer

class `chainer.serializers.DictionarySerializer` (*target=None, path=""*)

Serializer for dictionary.

This is the standard serializer in Chainer. The hierarchy of objects are simply mapped to a flat dictionary with keys representing the paths to objects in the hierarchy.

Note: Despite of its name, this serializer DOES NOT serialize the object into external files. It just build a flat dictionary of arrays that can be fed into `numpy.savez()` and `numpy.savez_compressed()`. If you want to use this serializer directly, you have to manually send a resulting dictionary to one of these functions.

Parameters

- **target** (*dict*) – The dictionary that this serializer saves the objects to. If target is None, then a new dictionary is created.
- **path** (*str*) – The base path in the hierarchy that this serializer indicates.

Variables **target** (*dict*) – The target dictionary. Once the serialization completes, this dictionary can be fed into `numpy.savez()` or `numpy.savez_compressed()` to serialize it in the NPZ format.

Methods

__call__ (*key, value*)

Serializes or deserializes a value by given name.

This operator saves or loads a value by given name.

If this is a serializer, then the value is simply saved at the key. Note that some type information might be missed depending on the implementation (and the target file format).

If this is a deserializer, then the value is loaded by the key. The deserialization differently works on scalars and arrays. For scalars, the *value* argument is used just for determining the type of restored value to be converted, and the converted value is returned. For arrays, the restored elements are directly copied into the *value* argument. String values are treated like scalars.

Note: As of v2.0.0, serializers and deserializers are required to correctly handle the None value. When *value* is None, serializers save it in format-dependent ways, and deserializers just return the loaded value. When the saved None value is loaded by a deserializer, it should quietly return the None value without modifying the *value* object.

Parameters

- **key** (*str*) – Name of the serialization entry.

- **value** (*scalar, numpy.ndarray, cupy.ndarray, None, or str*) – Object to be (de)serialized. None is only supported by deserializers.

Returns Serialized or deserialized value.

__getitem__ (*key*)

Gets a child serializer.

This operator creates a `_child_` serializer represented by the given key.

Parameters **key** (*str*) – Name of the child serializer.

save (*obj*)

Saves an object by this serializer.

This is equivalent to `obj.serialize(self)`.

Parameters **obj** – Target object to be serialized.

chainer.serializers.NpzDeserializer

```
class chainer.serializers.NpzDeserializer(npz, path="", strict=True, ignore_names=None)
```

Deserializer for NPZ format.

This is the standard deserializer in Chainer. This deserializer can be used to read an object serialized by `save_npz()`.

Parameters

- **npz** – *npz* file object.
- **path** – The base path that the deserialization starts from.
- **strict** (*bool*) – If `True`, the deserializer raises an error when an expected value is not found in the given NPZ file. Otherwise, it ignores the value and skip deserialization.
- **ignore_names** (*string, callable or list of them*) – If callable, it is a function that takes a name of a parameter and a persistent and returns `True` when it needs to be skipped. If string, this is a name of a parameter or persistent that are going to be skipped. This can also be a list of callables and strings that behave as described above.

Methods

__call__ (*key, value*)

Serializes or deserializes a value by given name.

This operator saves or loads a value by given name.

If this is a serializer, then the value is simply saved at the key. Note that some type information might be missed depending on the implementation (and the target file format).

If this is a deserializer, then the value is loaded by the key. The deserialization differently works on scalars and arrays. For scalars, the `value` argument is used just for determining the type of restored value to be converted, and the converted value is returned. For arrays, the restored elements are directly copied into the `value` argument. String values are treated like scalars.

Note: As of v2.0.0, serializers and deserializers are required to correctly handle the `None` value. When `value` is `None`, serializers save it in format-dependent ways, and deserializers just return the loaded

value. When the saved `None` value is loaded by a deserializer, it should quietly return the `None` value without modifying the `value` object.

Parameters

- **key** (*str*) – Name of the serialization entry.
- **value** (*scalar*, *numpy.ndarray*, *cupy.ndarray*, *None*, or *str*) – Object to be (de)serialized. `None` is only supported by deserializers.

Returns Serialized or deserialized value.

`__getitem__` (*key*)

Gets a child serializer.

This operator creates a `_child_` serializer represented by the given key.

Parameters **key** (*str*) – Name of the child serializer.

load (*obj*)

Loads an object from this deserializer.

This is equivalent to `obj.serialize(self)`.

Parameters **obj** – Target object to be serialized.

chainer.serializers.save_npz

`chainer.serializers.save_npz` (*file*, *obj*, *compression=True*)

Saves an object to the file in NPZ format.

This is a short-cut function to save only one object into an NPZ file.

Parameters

- **file** (*str* or *file-like*) – Target file to write to.
- **obj** – Object to be serialized. It must support serialization protocol.
- **compression** (*bool*) – If `True`, compression in the resulting zip file is enabled.

See also:

`chainer.serializers.load_npz()`

chainer.serializers.load_npz

`chainer.serializers.load_npz` (*file*, *obj*, *path=""*, *strict=True*)

Loads an object from the file in NPZ format.

This is a short-cut function to load from an *.npz* file that contains only one object.

Parameters

- **file** (*str* or *file-like*) – File to be loaded.
- **obj** – Object to be deserialized. It must support serialization protocol.
- **path** (*str*) – The path in the hierarchy of the serialized data under which the data is to be loaded. The default behavior (blank) will load all data under the root path.

- **strict** (*bool*) – If True, the deserializer raises an error when an expected value is not found in the given NPZ file. Otherwise, it ignores the value and skip deserialization.

See also:

`chainer.serializers.save_npz()`

5.9.2 Serialization in HDF5 format

<code>chainer.serializers.HDF5Serializer</code>	Serializer for HDF5 format.
<code>chainer.serializers.HDF5Deserializer</code>	Deserializer for HDF5 format.
<code>chainer.serializers.save_hdf5</code>	Saves an object to the file in HDF5 format.
<code>chainer.serializers.load_hdf5</code>	Loads an object from the file in HDF5 format.

chainer.serializers.HDF5Serializer

class `chainer.serializers.HDF5Serializer` (*group*, *compression=4*)

Serializer for HDF5 format.

This is the standard serializer in Chainer. The chain hierarchy is simply mapped to HDF5 hierarchical groups.

Parameters

- **group** (*h5py.Group*) – The group that this serializer represents.
- **compression** (*int*) – Gzip compression level.

Methods

`__call__` (*key*, *value*)

Serializes or deserializes a value by given name.

This operator saves or loads a value by given name.

If this is a serializer, then the value is simply saved at the key. Note that some type information might be missed depending on the implementation (and the target file format).

If this is a deserializer, then the value is loaded by the key. The deserialization differently works on scalars and arrays. For scalars, the `value` argument is used just for determining the type of restored value to be converted, and the converted value is returned. For arrays, the restored elements are directly copied into the `value` argument. String values are treated like scalars.

Note: As of v2.0.0, serializers and deserializers are required to correctly handle the `None` value. When `value` is `None`, serializers save it in format-dependent ways, and deserializers just return the loaded value. When the saved `None` value is loaded by a deserializer, it should quietly return the `None` value without modifying the `value` object.

Parameters

- **key** (*str*) – Name of the serialization entry.
- **value** (*scalar*, *numpy.ndarray*, *cupy.ndarray*, *None*, or *str*) – Object to be (de)serialized. `None` is only supported by deserializers.

Returns Serialized or deserialized value.

`__getitem__` (*key*)

Gets a child serializer.

This operator creates a `_child_` serializer represented by the given key.

Parameters `key` (*str*) – Name of the child serializer.

`save` (*obj*)

Saves an object by this serializer.

This is equivalent to `obj.serialize(self)`.

Parameters `obj` – Target object to be serialized.

chainer.serializers.HDF5Deserializer

class `chainer.serializers.HDF5Deserializer` (*group*, *strict=True*)

Deserializer for HDF5 format.

This is the standard deserializer in Chainer. This deserializer can be used to read an object serialized by `HDF5Serializer`.

Parameters

- **group** (*h5py.Group*) – The group that the deserialization starts from.
- **strict** (*bool*) – If `True`, the deserializer raises an error when an expected value is not found in the given HDF5 file. Otherwise, it ignores the value and skip deserialization.

Methods

`__call__` (*key*, *value*)

Serializes or deserializes a value by given name.

This operator saves or loads a value by given name.

If this is a serializer, then the value is simply saved at the key. Note that some type information might be missed depending on the implementation (and the target file format).

If this is a deserializer, then the value is loaded by the key. The deserialization differently works on scalars and arrays. For scalars, the `value` argument is used just for determining the type of restored value to be converted, and the converted value is returned. For arrays, the restored elements are directly copied into the `value` argument. String values are treated like scalars.

Note: As of v2.0.0, serializers and deserializers are required to correctly handle the `None` value. When `value` is `None`, serializers save it in format-dependent ways, and deserializers just return the loaded value. When the saved `None` value is loaded by a deserializer, it should quietly return the `None` value without modifying the `value` object.

Parameters

- **key** (*str*) – Name of the serialization entry.
- **value** (*scalar*, *numpy.ndarray*, *cupy.ndarray*, *None*, or *str*) – Object to be (de)serialized. `None` is only supported by deserializers.

Returns Serialized or deserialized value.

`__getitem__(key)`

Gets a child serializer.

This operator creates a `_child_` serializer represented by the given key.

Parameters `key` (*str*) – Name of the child serializer.

`load(obj)`

Loads an object from this deserializer.

This is equivalent to `obj.serialize(self)`.

Parameters `obj` – Target object to be serialized.

`chainer.serializers.save_hdf5`

`chainer.serializers.save_hdf5(filename, obj, compression=4)`

Saves an object to the file in HDF5 format.

This is a short-cut function to save only one object into an HDF5 file. If you want to save multiple objects to one HDF5 file, use `HDF5Serializer` directly by passing appropriate `h5py.Group` objects.

Parameters

- **filename** (*str*) – Target file name.
- **obj** – Object to be serialized. It must support serialization protocol.
- **compression** (*int*) – Gzip compression level.

Note: Currently `save_hdf5()` only supports writing to an actual file on file system due to a limitation of HD5F library. See [h5py/h5py#687](#) for details.

See also:

`chainer.serializers.load_hdf5()`

`chainer.serializers.load_hdf5`

`chainer.serializers.load_hdf5(filename, obj)`

Loads an object from the file in HDF5 format.

This is a short-cut function to load from an HDF5 file that contains only one object. If you want to load multiple objects from one HDF5 file, use `HDF5Deserializer` directly by passing appropriate `h5py.Group` objects.

Parameters

- **filename** (*str*) – Name of the file to be loaded.
- **obj** – Object to be deserialized. It must support serialization protocol.

Note: Currently `load_hdf5()` only supports loading an actual file on file system due to a limitation of HD5F library. See [h5py/h5py#687](#) for details.

See also:

`chainer.serializers.save_hdf5()`

5.9.3 Serializers base classes

<code>chainer.Serializer</code>	Base class of all serializers.
<code>chainer.AbstractSerializer</code>	Abstract base class of all serializers and deserializers.
<code>chainer.Deserializer</code>	Base class of all deserializers.

`chainer.Serializer`

class `chainer.Serializer`

Base class of all serializers.

Methods

`__call__` (*key*, *value*)

Serializes or deserializes a value by given name.

This operator saves or loads a value by given name.

If this is a serializer, then the value is simply saved at the key. Note that some type information might be missed depending on the implementation (and the target file format).

If this is a deserializer, then the value is loaded by the key. The deserialization differently works on scalars and arrays. For scalars, the `value` argument is used just for determining the type of restored value to be converted, and the converted value is returned. For arrays, the restored elements are directly copied into the `value` argument. String values are treated like scalars.

Note: As of v2.0.0, serializers and deserializers are required to correctly handle the `None` value. When `value` is `None`, serializers save it in format-dependent ways, and deserializers just return the loaded value. When the saved `None` value is loaded by a deserializer, it should quietly return the `None` value without modifying the `value` object.

Parameters

- **key** (*str*) – Name of the serialization entry.
- **value** (*scalar*, *numpy.ndarray*, *cupy.ndarray*, *None*, or *str*) – Object to be (de)serialized. `None` is only supported by deserializers.

Returns Serialized or deserialized value.

`__getitem__` (*key*)

Gets a child serializer.

This operator creates a `_child_` serializer represented by the given key.

Parameters **key** (*str*) – Name of the child serializer.

save (*obj*)

Saves an object by this serializer.

This is equivalent to `obj.serialize(self)`.

Parameters **obj** – Target object to be serialized.

chainer.AbstractSerializer

class chainer.**AbstractSerializer**

Abstract base class of all serializers and deserializers.

Methods

__call__(*key*, *value*)

Serializes or deserializes a value by given name.

This operator saves or loads a value by given name.

If this is a serializer, then the value is simply saved at the key. Note that some type information might be missed depending on the implementation (and the target file format).

If this is a deserializer, then the value is loaded by the key. The deserialization differently works on scalars and arrays. For scalars, the *value* argument is used just for determining the type of restored value to be converted, and the converted value is returned. For arrays, the restored elements are directly copied into the *value* argument. String values are treated like scalars.

Note: As of v2.0.0, serializers and deserializers are required to correctly handle the `None` value. When *value* is `None`, serializers save it in format-dependent ways, and deserializers just return the loaded value. When the saved `None` value is loaded by a deserializer, it should quietly return the `None` value without modifying the *value* object.

Parameters

- **key** (*str*) – Name of the serialization entry.
- **value** (*scalar*, *numpy.ndarray*, *cupy.ndarray*, *None*, or *str*) – Object to be (de)serialized. `None` is only supported by deserializers.

Returns Serialized or deserialized value.

__getitem__(*key*)

Gets a child serializer.

This operator creates a `_child_` serializer represented by the given key.

Parameters **key** (*str*) – Name of the child serializer.

chainer.Deserializer

class chainer.**Deserializer**

Base class of all deserializers.

Methods

__call__(*key*, *value*)

Serializes or deserializes a value by given name.

This operator saves or loads a value by given name.

If this is a serializer, then the value is simply saved at the key. Note that some type information might be missed depending on the implementation (and the target file format).

If this is a deserializer, then the value is loaded by the key. The deserialization differently works on scalars and arrays. For scalars, the `value` argument is used just for determining the type of restored value to be converted, and the converted value is returned. For arrays, the restored elements are directly copied into the `value` argument. String values are treated like scalars.

Note: As of v2.0.0, serializers and deserializers are required to correctly handle the `None` value. When `value` is `None`, serializers save it in format-dependent ways, and deserializers just return the loaded value. When the saved `None` value is loaded by a deserializer, it should quietly return the `None` value without modifying the `value` object.

Parameters

- **key** (*str*) – Name of the serialization entry.
- **value** (*scalar, numpy.ndarray, cupy.ndarray, None, or str*) – Object to be (de)serialized. `None` is only supported by deserializers.

Returns Serialized or deserialized value.

`__getitem__` (*key*)

Gets a child serializer.

This operator creates a `_child_` serializer represented by the given key.

Parameters **key** (*str*) – Name of the child serializer.

`load` (*obj*)

Loads an object from this deserializer.

This is equivalent to `obj.serialize(self)`.

Parameters **obj** – Target object to be serialized.

5.10 Utilities

5.10.1 Convolution/Deconvolution utilities

<code>chainer.utils.get_conv_outsize</code>	Calculates output size of convolution.
<code>chainer.utils.get_deconv_outsize</code>	Calculates output size of deconvolution.

`chainer.utils.get_conv_outsize`

`chainer.utils.get_conv_outsize` (*size, k, s, p, cover_all=False, d=1*)

Calculates output size of convolution.

This function takes the size of input feature map, kernel, stride, and pooling of one particular dimension, then calculates the output feature map size of that dimension.

See also:

`get_deconv_outsize()`

Parameters

- **size** (*int*) – The size of input feature map. It usually is the length of a side of feature

map.

- **k** (*int*) – The size of convolution kernel.
- **s** (*int*) – The size of stride.
- **p** (*int*) – The size of padding.
- **cover_all** (*bool*) – Use `cover_all` option or not.
- **d** (*int*) – The size of dilation.

Returns The expected output size of the convolution operation.

Return type `int`

`chainer.utils.get_deconv_outsize`

`chainer.utils.get_deconv_outsize(size, k, s, p, cover_all=False, d=1)`

Calculates output size of deconvolution.

This function takes the size of input feature map, kernel, stride, and pooling of one particular dimension, then calculates the output feature map size of that dimension.

See also:

`get_conv_outsize()`

Parameters

- **size** (*int*) – The size of input feature map. It usually is the length of a side of feature map.
- **k** (*int*) – The size of deconvolution kernel.
- **s** (*int*) – The size of stride.
- **p** (*int*) – The size of padding.
- **cover_all** (*bool*) – Use `cover_all` option or not.
- **d** (*int*) – The size of dilation.

Returns The expected output size of the deconvolution operation.

Return type `int`

5.10.2 CUDA utilities

Device, context and memory management on CuPy.

Note: The package `chainer.cuda` has been renamed to `chainer.backends.cuda` as of v4.0.0, but the previous module path `chainer.cuda` is also available.

Chainer uses [CuPy](#) (with very thin wrapper) to exploit the speed of GPU computation. Following modules and classes defined in CuPy are imported to `chainer.backends.cuda` module for convenience (refer to this table when reading chainer's source codes).

imported name	original name
<code>chainer.backends.cuda.cupy</code>	<code>cupy</code>
<code>chainer.backends.cuda.cupyx</code>	<code>cupyx</code>
<code>chainer.backends.cuda.ndarray</code>	<code>cupy.ndarray</code>
<code>chainer.backends.cuda.cupy.cuda</code>	<code>cupy.cuda</code>
<code>chainer.backends.cuda.Device</code>	<code>cupy.cuda.Device</code>
<code>chainer.backends.cuda.Event</code>	<code>cupy.cuda.Event</code>
<code>chainer.backends.cuda.Stream</code>	<code>cupy.cuda.Stream</code>

Chainer replaces the default allocator of CuPy by its memory pool implementation. It enables us to reuse the device memory over multiple forward/backward computations, and temporary arrays for consecutive elementwise operations.

Devices

<code>chainer.backends.cuda.get_device</code>	Gets the device from a device object, an ID integer or an array object.
<code>chainer.backends.cuda.get_device_from_id</code>	Gets the device from an ID integer.
<code>chainer.backends.cuda.get_device_from_array</code>	Gets the device from a list of CuPy array or a single CuPy array.

`chainer.backends.cuda.get_device`

`chainer.backends.cuda.get_device(*args)`

Gets the device from a device object, an ID integer or an array object.

Note: This API is deprecated. Please use `get_device_from_id()` or `get_device_from_array()` instead.

This is a convenient utility to select a correct device if the type of `arg` is unknown (i.e., one can use this function on arrays that may be on CPU or GPU). The returned device object supports the context management protocol of Python for the *with* statement.

Parameters `args` – Values to specify a GPU device. The first device object, integer or `cupy.ndarray` object is used to select a device. If it is a device object, it is returned. If it is an integer, the corresponding device is returned. If it is a CuPy array, the device on which this array reside is returned. If any arguments are neither integers nor CuPy arrays, a dummy device object representing CPU is returned.

Returns Device object specified by given `args`.

See also:

See `cupy.cuda.Device` for the device selection not by arrays.

`chainer.backends.cuda.get_device_from_id`

`chainer.backends.cuda.get_device_from_id(device_id)`

Gets the device from an ID integer.

Parameters `device_id` (*int* or *None*) – The ID of the device which this function returns.

chainer.backends.cuda.get_device_from_array

chainer.backends.cuda.get_device_from_array(*arrays)

Gets the device from a list of CuPy array or a single CuPy array.

The device on which the given CuPy array reside is returned.

Note: This method only recognizes `cupy.ndarrays` in arguments. Especially note that, unlike `get_array_module()`, this method does not recognize `Variable` objects. If you need to get device from the `Variable` instance `v`, you need to use `get_device_from_array(v.array)`.

Parameters **arrays** (`cupy.ndarray` or list of `cupy.ndarray`) – A CuPy array which this function returns the device corresponding to. If a list of `cupy.ndarrays` are given, it returns the first device object of an array in the list.

CuPy array allocation and copy

<code>chainer.backends.cuda.copy</code>	Copies a <code>cupy.ndarray</code> object using the default stream.
<code>chainer.backends.cuda.to_cpu</code>	Copies the given GPU array to host CPU.
<code>chainer.backends.cuda.to_gpu</code>	Copies the given CPU array to the specified device.

chainer.backends.cuda.copy

chainer.backends.cuda.copy(array, out=None, out_device=None, stream=None)

Copies a `cupy.ndarray` object using the default stream.

This function can copy the device array to the destination array on another device.

Parameters

- **array** (`cupy.ndarray`) – Array to be copied.
- **out** (`cupy.ndarray`) – Destination array. If it is not `None`, then `out_device` argument is ignored.
- **out_device** – Destination device specifier. Actual device object is obtained by passing this value to `get_device()`.
- **stream** (`cupy.cuda.Stream`) – CUDA stream.

Returns

Copied array.

If `out` is not specified, then the array is allocated on the device specified by `out_device` argument.

Return type `cupy.ndarray`

chainer.backends.cuda.to_cpu

chainer.backends.cuda.to_cpu(array, stream=None)

Copies the given GPU array to host CPU.

Parameters

- **array** (*array*, *None*, list or tuple) – Array or arrays to be sent to CPU.
- **stream** (*cupy.cuda.Stream*) – CUDA stream.

Returns

Array on CPU.

If some of the arrays are already on CPU, then this function just returns those arrays without performing any copy.

If input arrays include *None*, it is returned as *None* as is.

Return type *numpy.ndarray*, list or tuple

chainer.backends.cuda.to_gpu

`chainer.backends.cuda.to_gpu` (*array*, *device=None*, *stream=None*)

Copies the given CPU array to the specified device.

Parameters

- **array** (*array*, *None*, list or tuple) – Array or arrays to be sent to GPU.
- **device** – Device specifier.
- **stream** (*Stream*) – (*deprecated since v3.0.0*) CUDA stream. If not *None*, the copy runs asynchronously.

Returns

Array or arrays on GPU.

If some of the arrays are already on GPU, then this function just returns those arrays without performing any copy.

If input arrays include *None*, it is returned as *None* as is.

Return type *cupy.ndarray*, list or tuple

Kernel definition utilities

<code>chainer.backends.cuda.memoize</code>	Makes a function memoizing the result for each argument and device.
<code>chainer.backends.cuda.clear_memo</code>	Clears the memoized results for all functions decorated by <code>memoize</code> .
<code>chainer.backends.cuda.elementwise</code>	Creates an elementwise kernel function.
<code>chainer.backends.cuda.reduce</code>	Creates a global reduction kernel function.

chainer.backends.cuda.memoize

`chainer.backends.cuda.memoize` (*for_each_device=False*)

Makes a function memoizing the result for each argument and device.

This is a similar version of `cupy.memoize()`. The difference is that this function can be used in the global scope even if CUDA is not available. In such case, this function does nothing.

Note: This decorator acts as a dummy if CUDA is not available. It cannot be used for general purpose memoization even if `for_each_device` is set to `False`.

chainer.backends.cuda.clear_memo

`chainer.backends.cuda.clear_memo()`

Clears the memoized results for all functions decorated by `memoize`.

This function works like `cupy.clear_memo()` as a counterpart for `chainer.backends.cuda.memoize()`. It can be used even if CUDA is not available. In such a case, this function does nothing.

chainer.backends.cuda.elementwise

`chainer.backends.cuda.elementwise(in_params, out_params, operation, name, **kwargs)`

Creates an elementwise kernel function.

This function uses `memoize()` to cache the kernel object, i.e. the resulting kernel object is cached for each argument combination and CUDA device.

The arguments are the same as those for `cupy.ElementwiseKernel`, except that the `name` argument is mandatory.

chainer.backends.cuda.reduce

`chainer.backends.cuda.reduce(in_params, out_params, map_expr, reduce_expr, post_map_expr, identity, name, **kwargs)`

Creates a global reduction kernel function.

This function uses `memoize()` to cache the resulting kernel object, i.e. the resulting kernel object is cached for each argument combination and CUDA device.

The arguments are the same as those for `cupy.ReductionKernel`, except that the `name` argument is mandatory.

CPU/GPU generic code support

`chainer.backends.cuda.
get_array_module`

Gets an appropriate one from `numpy` or `cupy`.

chainer.backends.cuda.get_array_module

`chainer.backends.cuda.get_array_module(*args)`

Gets an appropriate one from `numpy` or `cupy`.

This is almost equivalent to `cupy.get_array_module()`. The differences are that this function can be used even if CUDA is not available and that it will return their data arrays' array module for *Variable* arguments.

Parameters `args` – Values to determine whether NumPy or CuPy should be used.

Returns `cupy` or `numpy` is returned based on the types of the arguments.

Return type module

cuDNN support

<code>chainer.backends.cuda.set_max_workspace_size</code>	Sets the workspace size for cuDNN.
<code>chainer.backends.cuda.get_max_workspace_size</code>	Gets the workspace size for cuDNN.

chainer.backends.cuda.set_max_workspace_size

`chainer.backends.cuda.set_max_workspace_size(size)`
Sets the workspace size for cuDNN.

Check “cuDNN Library User Guide” for detail.

Parameters `size` – The workspace size for cuDNN.

chainer.backends.cuda.get_max_workspace_size

`chainer.backends.cuda.get_max_workspace_size()`
Gets the workspace size for cuDNN.

Check “cuDNN Library User Guide” for detail.

Returns The workspace size for cuDNN.

Return type int

5.10.3 Common algorithms

<code>chainer.utils.WalkerAlias</code>	Implementation of Walker’s alias method.
--	--

chainer.utils.WalkerAlias

class `chainer.utils.WalkerAlias(probs)`
Implementation of Walker’s alias method.

This method generates a random sample from given probabilities p_1, \dots, p_n in $O(1)$ time. It is more efficient than `choice()`. This class works on both CPU and GPU.

Parameters `probs` (*float list*) – Probabilities of entries. They are normalized with $\text{sum}(probs)$.

See: [Wikipedia article](#)

Methods

sample (*shape*)
Generates a random sample based on given probabilities.

Parameters `shape` (*tuple of int*) – Shape of a return value.

Returns Returns a generated array with the given shape. If a sampler is in CPU mode the return value is a `numpy.ndarray` object, and if it is in GPU mode the return value is a `cupy.ndarray` object.

`sample_cpu(shape)`

`sample_gpu(shape)`

`to_cpu()`

Make a sampler CPU mode.

`to_gpu()`

Make a sampler GPU mode.

5.10.4 Reporter

Reporter

<code>chainer.Reporter</code>	Object to which observed values are reported.
<code>chainer.get_current_reporter</code>	Returns the current reporter object.
<code>chainer.report</code>	Reports observed values with the current reporter object.
<code>chainer.report_scope</code>	Returns a report scope with the current reporter.

chainer.Reporter

class `chainer.Reporter`

Object to which observed values are reported.

Reporter is used to collect values that users want to watch. The reporter object holds a mapping from value names to the actually observed values. We call this mapping *observations*.

When a value is passed to the reporter, an object called *observer* can be optionally attached. In this case, the name of the observer is added as the prefix of the value name. The observer name should be registered beforehand.

See the following example:

```
>>> from chainer import Reporter, report, report_scope
>>>
>>> reporter = Reporter()
>>> observer = object() # it can be an arbitrary (reference) object
>>> reporter.add_observer('my_observer', observer)
>>> observation = {}
>>> with reporter.scope(observation):
...     reporter.report({'x': 1}, observer)
...
>>> observation
{'my_observer/x': 1}
```

There are also a global API to add values:

```
>>> observation = {}
>>> with report_scope(observation):
...     report({'x': 1}, observer)
...

```

(continues on next page)

(continued from previous page)

```
>>> observation
{'my_observer/x': 1}
```

The most important application of Reporter is to report observed values from each link or chain in the training and validation procedures. *Trainer* and some extensions prepare their own Reporter object with the hierarchy of the target link registered as observers. We can use *report()* function inside any links and chains to report the observed values (e.g., training loss, accuracy, activation statistics, etc.).

Variables *observation* – Dictionary of observed values.

Methods

__enter__()

Makes this reporter object current.

__exit__(*exc_type, exc_value, traceback*)

Recovers the previous reporter object to the current.

add_observer(*name, observer*)

Registers an observer of values.

Observer defines a scope of names for observed values. Values observed with the observer are registered with names prefixed by the observer name.

Parameters

- **name** (*str*) – Name of the observer.
- **observer** – The observer object. Note that the reporter distinguishes the observers by their object ids (i.e., `id(owner)`), rather than the object equality.

add_observers(*prefix, observers*)

Registers multiple observers at once.

This is a convenient method to register multiple objects at once.

Parameters

- **prefix** (*str*) – Prefix of each name of observers.
- **observers** – Iterator of name and observer pairs.

report(*values, observer=None*)

Reports observed values.

The values are written with the key, prefixed by the name of the observer object if given.

Note: As of v2.0.0, if a value is of type *Variable*, the variable is copied without preserving the computational graph and the new variable object purged from the graph is stored to the observer. This behavior can be changed by setting `chainer.config.keep_graph_on_report` to `True`.

Parameters

- **values** (*dict*) – Dictionary of observed values.
- **observer** – Observer object. Its object ID is used to retrieve the observer name, which is used as the prefix of the registration name of the observed value.

scope (*observation*)

Creates a scope to report observed values to *observation*.

This is a context manager to be passed to `with` statements. In this scope, the observation dictionary is changed to the given one.

It also makes this reporter object current.

Parameters *observation* (*dict*) – Observation dictionary. All observations reported inside of the `with` statement are written to this dictionary.

chainer.get_current_reporter

`chainer.get_current_reporter()`

Returns the current reporter object.

chainer.report

`chainer.report(values, observer=None)`

Reports observed values with the current reporter object.

Any reporter object can be set current by the `with` statement. This function calls the `Report.report()` method of the current reporter. If no reporter object is current, this function does nothing.

Example

The most typical example is a use within links and chains. Suppose that a link is registered to the current reporter as an observer (for example, the target link of the optimizer is automatically registered to the reporter of the *Trainer*). We can report some values from the link as follows:

```
class MyRegressor(chainer.Chain):
    def __init__(self, predictor):
        super(MyRegressor, self).__init__(predictor=predictor)

    def __call__(self, x, y):
        # This chain just computes the mean absolute and squared
        # errors between the prediction and y.
        pred = self.predictor(x)
        abs_error = F.sum(F.abs(pred - y)) / len(x)
        loss = F.mean_squared_error(pred, y)

        # Report the mean absolute and squared errors.
        report({'abs_error': abs_error, 'squared_error': loss}, self)

    return loss
```

If the link is named 'main' in the hierarchy (which is the default name of the target link in the *StandardUpdater*), these reported values are named 'main/abs_error' and 'main/squared_error'. If these values are reported inside the Evaluator extension, 'validation/' is added at the head of the link name, thus the item names are changed to 'validation/main/abs_error' and 'validation/main/squared_error' ('validation' is the default name of the Evaluator extension).

Parameters

- **values** (*dict*) – Dictionary of observed values.
- **observer** – Observer object. Its object ID is used to retrieve the observer name, which is used as the prefix of the registration name of the observed value.

chainer.report_scope

`chainer.report_scope(observation)`

Returns a report scope with the current reporter.

This is equivalent to `get_current_reporter().scope(observation)`, except that it does not make the reporter current redundantly.

Summary and DictSummary

<code>chainer.Summary</code>	Online summarization of a sequence of scalars.
<code>chainer.DictSummary</code>	Online summarization of a sequence of dictionaries.

chainer.Summary

class `chainer.Summary`

Online summarization of a sequence of scalars.

Summary computes the statistics of given scalars online.

Methods

add (*value*)

Adds a scalar value.

Parameters *value* – Scalar value to accumulate. It is either a NumPy scalar or a zero-dimensional array (on CPU or GPU).

compute_mean ()

Computes the mean.

make_statistics ()

Computes and returns the mean and standard deviation values.

Returns Mean and standard deviation values.

Return type *tuple*

serialize (*serializer*)

chainer.DictSummary

class `chainer.DictSummary`

Online summarization of a sequence of dictionaries.

`DictSummary` computes the statistics of a given set of scalars online. It only computes the statistics for scalar values and variables of scalar values in the dictionaries.

Methods

`add(d)`

Adds a dictionary of scalars.

Parameters `d` (*dict*) – Dictionary of scalars to accumulate. Only elements of scalars, zero-dimensional arrays, and variables of zero-dimensional arrays are accumulated.

`compute_mean()`

Creates a dictionary of mean values.

It returns a single dictionary that holds a mean value for each entry added to the summary.

Returns Dictionary of mean values.

Return type *dict*

`make_statistics()`

Creates a dictionary of statistics.

It returns a single dictionary that holds mean and standard deviation values for every entry added to the summary. For an entry of name 'key', these values are added to the dictionary by names 'key' and 'key.std', respectively.

Returns Dictionary of statistics of all entries.

Return type *dict*

`serialize(serializer)`

5.10.5 Experimental feature annotation

chainer.utils.experimental

Declares that user is using an experimental feature.

`chainer.utils.experimental`

`chainer.utils.experimental(api_name)`

Declares that user is using an experimental feature.

The developer of an API can mark it as *experimental* by calling this function. When users call experimental APIs, `FutureWarning` is issued. The presentation of `FutureWarning` is disabled by setting `chainer.disable_experimental_feature_warning` to `True`, which is `False` by default.

The basic usage is to call it in the function or method we want to mark as experimental along with the API name.

```
from chainer import utils

def f(x):
    utils.experimental('chainer.foo.bar.f')
    # concrete implementation of f follows

f(1)
```

```
... FutureWarning: chainer.foo.bar.f is experimental. The interface can change in
↳ the future. ...
```

We can also make a whole class experimental. In that case, we should call this function in its `__init__` method.

```
class C():
    def __init__(self):
        utils.experimental('chainer.foo.C')

C()
```

```
... FutureWarning: chainer.foo.C is experimental. The interface can change in the
↪future. ...
```

If we want to mark `__init__` method only, rather than class itself, it is recommended that we explicitly feed its API name.

```
class D():
    def __init__(self):
        utils.experimental('D.__init__')

D()
```

```
... FutureWarning: D.__init__ is experimental. The interface can change in the
↪future. ...
```

Currently, we do not have any sophisticated way to mark some usage of non-experimental function as experimental. But we can support such usage by explicitly branching it.

```
def g(x, experimental_arg=None):
    if experimental_arg is not None:
        utils.experimental('experimental_arg of chainer.foo.g')
```

Parameters `api_name` (*str*) – The name of an API marked as experimental.

5.11 Configuring Chainer

Chainer provides some global settings that affect the behavior of some functionalities. Such settings can be configured using the *unified configuration system*. The system provides a transparent way to manage the configuration for each process and for each thread.

The configuration is managed by two global objects: `chainer.global_config` and `chainer.config`.

- The `global_config` object maintains the configuration shared in the Python process. This is an instance of the `GlobalConfig` class. It can be used just as a plain object, and users can freely set any attributes on it.
- The `config` object, on the other hand, maintains the configuration for the current thread. This is an instance of the `LocalConfig` class. It behaves like a thread-local object, and any attribute modifications are only visible to the current thread.

If no value is set to `config` for a given key, `global_config` is transparently referred. Thanks to this transparent lookup, users can always use `config` to read any configuration so that the thread-local configuration is used if available and otherwise the default global setting is used.

The following entries of the configuration are currently provided by Chainer. Some entries support environment variables to set the default values. Note that the default values are set in the global config.

5.11.1 Configuration Keys

- **cudnn_deterministic (default: False)** Flag to configure deterministic computations in cuDNN APIs.
If it is `True`, convolution functions that use cuDNN use the deterministic mode (i.e, the computation is reproducible). Otherwise, the results of convolution functions using cuDNN may be non-deterministic in exchange for better performance.
- **debug (default: False)** Debug mode flag.
If it is `True`, Chainer runs in debug mode. Enabling debug mode may introduce some performance overhead. See [Debug Mode](#) for more information of the debug mode.
You can change the default value to `True` by setting `CHAINER_DEBUG` environment variable to 1.
- **enable_backprop (default: True)** Flag to enable backpropagation support.
If it is `True`, computational graphs are created during forward passes by `FunctionNode`s, allowing backpropagation to start from any `Variable` in the graph. Otherwise, computational graphs are not created but memory consumptions are reduced. So calling `backward()` on the results of a function will not compute any gradients of any input.
- **keep_graph_on_report (default: False)** Flag to configure whether or not to let `report()` keep the computational graph.
If it is `False`, `report()` does not keep the computational graph when a `Variable` object is reported. It means that `report()` stores a copy of the `Variable` object which is purged from the computational graph. If it is `True`, `report()` just stores the `Variable` object as is with the computational graph left attached.
You can change the default value to `True` by setting `CHAINER_KEEP_GRAPH_ON_REPORT` environment variable to 1.
- **train (default: True)** Training mode flag.
If it is `True`, Chainer runs in training mode. Otherwise, it runs in the testing (evaluation) mode.
This configuration is used by Functions and Links that need to behave differently between training phase and evaluation (inference) phase. One example is `chainer.links.BatchNormization` updates statistics using input data only when `train` is set to `True`. The other example is `chainer.functions.dropout()`, which does nothing when `train` is set to `False`.
Generally, you are responsible to change the configuration to `False` during evaluation. If you are using `Trainer` with `Evaluator` extension, `train` configuration will automatically be switched to `False` during evaluation in the training loop.
Note that this parameter does not reduce memory consumption or affect the creation of computational graphs required in order to compute gradients.
- **type_check (default: True)** Type checking mode flag.
If it is `True`, Chainer checks the types (data types and shapes) of inputs on `Function` applications. Otherwise, it skips type checking.
You can change the default value to `False` by setting `CHAINER_TYPE_CHECK` environment variable to 0.
- **use_cudnn (default: 'auto')** Flag to configure whether or not to use cuDNN.
This is a ternary flag with 'always', 'auto', and 'never' as its allowed values. The meaning of each flag is as follows.
 - If it is 'always', Chainer will try to use cuDNN everywhere if possible.

- If it is 'auto', Chainer will use cuDNN only if it is known that the usage does not degrade the performance.
- If it is 'never', Chainer will never use cuDNN anywhere.

You can change the default value by setting `CHAINER_USE_CUDNN` environment variable to any of 'always', 'auto' or 'never'.

- **use_ideep (default: 'never')** Flag to configure whether or not to use iDeep.

This is a ternary flag with 'always', 'auto', and 'never' as its allowed values. The meaning of each flag is as follows.

- If it is 'always', Chainer will try to use iDeep everywhere if possible.
- If it is 'auto', Chainer will use iDeep only if it is known that the usage does not degrade the performance.
- If it is 'never', Chainer will never use iDeep anywhere.

You can change the default value by setting `CHAINER_USE_IDEEP` environment variable to any of 'always', 'auto' or 'never'.

Note that in spite of the configuration, optimizers will use iDeep if and only if the link is converted manually to iDeep (e.g., `model.to_intel64()`).

- **lazy_grad_sum (default: False)** Flag to control the behavior of gradient accumulation.

If it is `True`, gradients are accumulated in batch for performance. Otherwise gradients are accumulated one by one.

You can change the default value to `True` by setting `CHAINER_LAZY_GRAD_SUM` environment variable to 1.

- **use_cudnn_tensor_core (default: 'auto')** Flag to configure whether or not to enable Tensor Core operations in cuDNN.

This is a ternary flag with 'always', 'auto', and 'never' as its allowed values. The meaning of each flag is as follows.

- If it is `always`, Chainer uses cuDNN's Tensor Core operations.
- If it is `never`, Chainer does not use cuDNN's Tensor Core operations.
- If it is `auto`, Chainer checks cuDNN version, the data type of input, the compute capability of the GPU used, and configures whether or not to use cuDNN's Tensor Core operations.

- **autotune (default: False)** Autotune for convolutional networks flag.

If it is `True`, Chainer uses the cuDNN autotune feature to find the fastest calculation process for `chainer.links.Convolution2D`, `ConvolutionND`, `Deconvolution2D`, or `DeconvolutionND` links.

5.11.2 User-defined Keys

Users can also define their own configurations. There are two ways:

1. Use Chainer's configuration objects. In this case, **it is strongly recommended to prefix the name by "user_"** to avoid name conflicts with configurations introduced to Chainer in the future.
2. Use your own configuration objects. Users can define their own configuration objects using `chainer.configuration.GlobalConfig` and `chainer.configuration.LocalConfig`. In this case, there is no need to take care of the name conflicts.

5.11.3 Changing Configuration

If you want to share a setting within the process, set an attribute to the global configuration. This value is automatically extracted by referring to the local config.

```
>>> chainer.global_config.train
True
>>> chainer.config.train
True

>>> chainer.global_config.train = False

>>> chainer.global_config.train
False
>>> chainer.config.train
False
```

If you set an attribute to the local configuration, the value is only visible to the current thread.

```
>>> chainer.global_config.train
True
>>> chainer.config.train
True

>>> chainer.config.train = False

>>> chainer.global_config.train
True
>>> chainer.config.train
False
```

If you want to temporarily modify the configuration for the specific scope, you can use `using_config()`. For example, if you only want to enable debug mode in a fragment of code, write as follows.

```
>>> with chainer.using_config('debug', True):
...     pass # code running in debug mode
```

If you want to switch to the test mode for an evaluation, you can do that in the same way.

```
>>> # Do training here
>>> with chainer.using_config('train', False):
...     pass # Perform evaluation here
```

Note that `Evaluator` automatically switches to the test mode, and thus you do not need to manually switch in the loss function for the evaluation.

You can also make your own code behave differently in training and test modes as follows.

```
if chainer.config.train:
    pass # code only running in the training mode
else:
    pass # code only running in the test mode
```

`chainer.global_config`

Global configuration of Chainer.

`chainer.config`

Thread-local configuration of Chainer.

Continued on next page

Table 64 – continued from previous page

<code>chainer.using_config</code>	Context manager to temporarily change the thread-local configuration.
<code>chainer.configuration.GlobalConfig</code>	The plain object that represents the global configuration of Chainer.
<code>chainer.configuration.LocalConfig</code>	Thread-local configuration of Chainer.

`chainer.global_config`

`chainer.global_config = <chainer.configuration.GlobalConfig object>`

Global configuration of Chainer.

It is an instance of `chainer.configuration.GlobalConfig`. See *Configuring Chainer* for details.

`chainer.config`

`chainer.config = <chainer.configuration.LocalConfig object>`

Thread-local configuration of Chainer.

It is an instance of `chainer.configuration.LocalConfig`, and is referring to `global_config` as its default configuration. See *Configuring Chainer* for details.

`chainer.using_config`

`chainer.using_config(name, value, config=chainer.config)`

Context manager to temporarily change the thread-local configuration.

Parameters

- **name** (*str*) – Name of the configuration to change.
- **value** – Temporary value of the configuration entry.
- **config** (*LocalConfig*) – Configuration object. Chainer’s thread-local configuration is used by default.

See also:

Configuring Chainer

`chainer.configuration.GlobalConfig`

class `chainer.configuration.GlobalConfig`

The plain object that represents the global configuration of Chainer.

Methods

show (*file=sys.stdout*)

Prints the global config entries.

The entries are sorted in the lexicographical order of the entry name.

Parameters **file** – Output file-like object.

chainer.configuration.LocalConfig

class chainer.configuration.LocalConfig(*global_config*)

Thread-local configuration of Chainer.

This class implements the local configuration. When a value is set to this object, the configuration is only updated in the current thread. When a user tries to access an attribute and there is no local value, it automatically retrieves a value from the global configuration.

Methods

show (*file=sys.stdout*)

Prints the config entries.

The entries are sorted in the lexicographical order of the entry names.

Parameters *file* – Output file-like object.

Example

You can easily print the list of configurations used in the current thread.

```
>>> chainer.config.show()
debug                False
enable_backprop      True
train                True
type_check            True
```

5.11.4 Environment Variables

Here are the environment variables Chainer uses.

CHAINER_SEED	Default seed value of random number generators for CUDA. If it is not set, the seed value is generated from Python random module. Set an integer value in decimal format.
CHAINER_DATA_DIRECTORY	Default directory path to store the downloaded datasets. See Datasets for details.
CHAINER_CUDNN	Set 0 to completely disable cuDNN in Chainer. In this case, cuDNN will not be used regardless of CHAINER_USE_CUDNN and <code>chainer.config.use_cudnn</code> configuration. Otherwise cuDNN is enabled automatically.
CHAINER_USE_CUDNN	Used as the default value for <code>chainer.config.use_cudnn</code> configuration. The value must be any of 'always', 'auto' or 'never'. If CHAINER_CUDNN is set to 0, this environment variable has no effect. See Configuring Chainer for details.
CHAINER_USE_IDEEP	Used as the default value for <code>chainer.config.use_idEEP</code> configuration. The value must be any of 'always', 'auto' or 'never'. See Configuring Chainer for details.
CHAINER_LAZY_GGRADSUM	Used as the default value for <code>chainer.config.lazy_grad_sum</code> configuration. Set 1 to enable batch accumulation of gradients. See Configuring Chainer for details.
CHAINER_TYPE_CHECK	Used as the default value for <code>chainer.config.type_check</code> configuration. Set 0 to disable type checking. Otherwise type checking is enabled automatically. See Configuring Chainer and Type checking utilities for details.
CHAINER_DEBUG	Used as the default value for <code>chainer.config.debug</code> configuration. Set 1 to enable debug mode. It is disabled by default. In debug mode, Chainer performs various runtime checks that can help debug user's code at the cost of some overhead. See Configuring Chainer and Debug Mode for details.
CHAINER_KEEP_GRAPH_ON_REPORT	Used as the default value for <code>chainer.config.keep_graph_on_report</code> configuration. Set 1 to let <code>report()</code> keep the computational graph. See Configuring Chainer for details.
CHAINER_PYTHON_FORCE_VERSION	Set 1 to force using Chainer with Python 3.5.0. Note that Chainer does not work with Python 3.5.0. Use Python 3.5.1+ or other supported versions (see Installation).

The following environment variables are only effective when running unit tests.

CHAINER_TEST_GPU	Number of GPUs available for unit tests. When running unit test, test cases that require more GPUs than the specified value will be skipped. Set 0 to skip all test cases that require GPU. See Unit Testing for details.
CHAINER_TEST_RANDOM_SEED	Set non-fixed seed for random number generators, even for test cases annotated with <code>fix_random</code> .

5.12 Debug Mode

In debug mode, Chainer checks values of variables on runtime and shows more detailed error messages. It helps you to debug your programs. However, it requires some additional overhead time.

If you want to enable debug mode for the entire code, you can set CHAINER_DEBUG environment variable to 1.

You can also enable or disable debug mode for the specific scope of code with `chainer.using_config()` or by changing `chainer.config.debug` configuration.

```
with chainer.using_config('debug', True):
    ...
```

See [Configuring Chainer](#) for the details of Chainer's configuration mechanism.

In debug mode, Chainer checks all results of forward and backward computation, and if it finds a NaN value, it raises `RuntimeError`. Some functions and links also check validity of input values more strictly.

You can check if debug mode is enabled with `chainer.is_debug()` function.

<code>chainer.is_debug</code>	Returns if the debug mode is enabled or not in the current thread.
<code>chainer.set_debug</code>	Enables or disables the debug mode in the current thread.

5.12.1 `chainer.is_debug`

`chainer.is_debug()`

Returns if the debug mode is enabled or not in the current thread.

Returns `True` if the debug mode is enabled.

Return type `bool`

5.12.2 `chainer.set_debug`

`chainer.set_debug(debug)`

Enables or disables the debug mode in the current thread.

Note: `chainer.set_debug(value)` is equivalent to `chainer.config.debug = value`.

Parameters `debug` (`bool`) – New debug mode.

5.12.3 Deprecated interface

As of v2.0.0, it is recommended to turn on the debug mode using `chainer.config.debug`. See [Configuring Chainer](#) for the way to use the config object. We leave the reference of the conventional way (which has been available since Chainer v1) as follows.

<code>chainer.DebugMode</code>	Debug mode context.
--------------------------------	---------------------

`chainer.DebugMode`

class `chainer.DebugMode(debug)`

Debug mode context.

This class provides a context manager for debug mode. When entering the context, it sets the debug mode to the value of `debug` parameter with memorizing its original value. When exiting the context, it sets the debug mode back to the original value.

Deprecated since version v2.0.0: Use `chainer.using_config()` instead. See [Debug Mode](#) for details.

Parameters `debug` (`bool`) – Debug mode used in the context.

Methods

`__enter__()`

`__exit__(*args)`

5.13 Visualization of Computational Graph

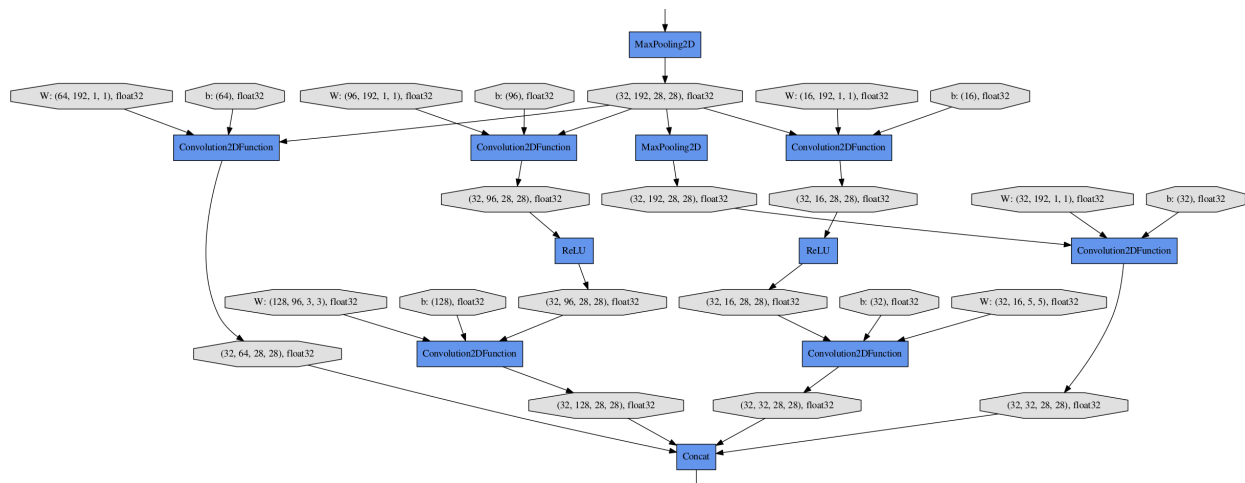
As neural networks get larger and complicated, it gets much harder to confirm if their architectures are constructed properly. Chainer supports visualization of computational graphs. Users can generate computational graphs by invoking `build_computational_graph()`. Generated computational graphs are dumped to specified format (Currently Dot Language is supported).

Basic usage is as follows:

```
import chainer.computational_graph as c
...
g = c.build_computational_graph(vs)
with open('path/to/output/file', 'w') as o:
    o.write(g.dump())
```

where `vs` is list of `Variable` instances and `g` is an instance of `ComputationalGraph`. This code generates the computational graph that are backward-reachable (i.e. reachable by repetition of steps backward) from at least one of `vs`.

Here is an example of (a part of) the generated graph (inception(3a) in GoogLeNet). This example is from `example/imagenet`.



`chainer.computational_graph.
build_computational_graph`
`chainer.computational_graph.
ComputationalGraph`

Builds a graph of functions and variables backward-reachable from outputs.

Class that represents computational graph.

5.13.1 `chainer.computational_graph.build_computational_graph`

```
chainer.computational_graph.build_computational_graph(outputs, remove_split=True,
                                                       variable_style={'fillcolor':
                                                           '#E0E0E0', 'shape': 'octagon',
                                                           'style': 'filled'},
                                                       function_style={'fillcolor':
                                                           '#6495ED', 'shape':
                                                           'record', 'style': 'filled'},
                                                       rankdir='TB',
                                                       remove_variable=False,
                                                       show_name=True)
```

Builds a graph of functions and variables backward-reachable from outputs.

Parameters

- **outputs** (*list*) – nodes from which the graph is constructed. Each element of outputs must be either *Variable* object, *VariableNode* object, or *FunctionNode* object.
- **remove_split** (*bool*) – It must be `True`. This argument is left for backward compatibility.
- **variable_style** (*dict*) – Dot node style for variable. Possible keys are 'shape', 'color', 'fillcolor', 'style', and etc.
- **function_style** (*dict*) – Dot node style for function.
- **rankdir** (*str*) – Direction of the graph that must be TB (top to bottom), BT (bottom to top), LR (left to right) or RL (right to left).
- **remove_variable** (*bool*) – If `True`, *VariableNodes* are removed from the resulting computational graph. Only *FunctionNodes* are shown in the output.
- **show_name** (*bool*) – If `True`, the name attribute of each node is added to the label of the node. Default is `True`.

Returns

A graph consisting of nodes and edges that are backward-reachable from at least one of outputs.

If `unchain_backward` was called in some variable in the computational graph before this function, backward step is stopped at this variable.

For example, suppose that computational graph is as follows:

```

      |--> f ----> y
x  --+
      |--> g ----> z
```

Let `outputs = [y, z]`. Then the full graph is emitted.

Next, let `outputs = [y]`. Note that `z` and `g` are not backward-reachable from `y`. The resulting graph would be following:

```
x ----> f ----> y
```

See `TestGraphBuilder` for details.

Return type *ComputationalGraph*

Note: The default behavior of `ComputationalGraph` has been changed from v1.23.0, so that it outputs the richest representation of a graph as default, namely, styles are set and names of functions and variables are shown. To reproduce the same result as previous versions (\leq v1.22.0), please specify `variable_style=None`, `function_style=None`, and `show_name=False` explicitly.

5.13.2 chainer.computational_graph.ComputationalGraph

```
class chainer.computational_graph.ComputationalGraph(nodes, edges, variable_style={'fillcolor': '#E0E0E0', 'shape': 'octagon', 'style': 'filled'}, function_style={'fillcolor': '#6495ED', 'shape': 'record', 'style': 'filled'}, rankdir='TB', remove_variable=False, show_name=True)
```

Class that represents computational graph.

Note: We assume that the computational graph is directed and acyclic.

Parameters

- **nodes** (*list*) – List of nodes. Each node is either `VariableNode` object or `FunctionNode` object.
- **edges** (*list*) – List of edges. Each edge consists of pair of nodes.
- **variable_style** (*dict*) – Dot node style for variable.
- **function_style** (*dict*) – Dot node style for function.
- **rankdir** (*str*) – Direction of the graph that must be TB (top to bottom), BT (bottom to top), LR (left to right) or RL (right to left).
- **remove_variable** (*bool*) – If `True`, `VariableNodes` are removed from the resulting computational graph. Only `FunctionNodes` are shown in the output.
- **show_name** (*bool*) – If `True`, the name attribute of each node is added to the label of the node. Default is `True`.

Note: The default behavior of `ComputationalGraph` has been changed from v1.23.0, so that it outputs the richest representation of a graph as default, namely, styles are set and names of functions and variables are shown. To reproduce the same result as previous versions (\leq v1.22.0), please specify `variable_style=None`, `function_style=None`, and `show_name=False` explicitly.

Methods

dump (*format='dot'*)
Dumps graph as a text.

Parameters

- **format** (*str*) – The graph language name of the output.
- **it must be 'dot'.** (*Currently,*) –

Returns The graph in specified format.

Return type *str*

5.14 Caffe Model Support

Caffe is a popular framework maintained by BVLC at UC Berkeley. It is widely used by computer vision communities, and aims at fast computation and easy usage without any programming. The BVLC team provides trained reference models in their Model Zoo, which can reduce training time required for a new task.

5.14.1 Import

Chainer can import the reference models and emulate the network by *Link* implementations. This functionality is provided by the `chainer.links.caffe.CaffeFunction` class.

<code>chainer.links.caffe.CaffeFunction</code>	Caffe emulator based on the model file of Caffe.
--	--

5.14.2 Export

Chainer can export a model from *Link*.

<code>chainer.exporters.caffe.export</code>	(Experimental) Export a computational graph as Caffe format.
---	--

chainer.exporters.caffe.export

`chainer.exporters.caffe.export` (*model*, *args*, *directory=None*, *export_params=True*,
graph_name='Graph')
 (Experimental) Export a computational graph as Caffe format.

Parameters

- **model** (*Chain*) – The model object you want to export in Caffe format. It should have `__call__()` method because the second argument *args* is directly given to the model by the `()` accessor.
- **args** (*list of ~chainer.Variable*) – The arguments which are given to the model directly.
- **directory** (*str*) – The directory used for saving the resulting Caffe model. If *None*, nothing is saved to the disk.
- **export_params** (*bool*) – If *True*, this function exports all the parameters included in the given model at the same time. If *False*, the exported Caffe model doesn't include any parameter values.
- **graph_name** (*str*) – A string to be used for the `name` field of the graph in the exported Caffe model.

Note: Currently, this function supports networks that created by following layer functions.

- `linear()`
- `convolution_2d()`
- `deconvolution_2d()`
- `max_pooling_2d()`
- `average_pooling_2d()`
- `batch_normalization()`
- `local_response_normalization()`
- `relu()`
- `concat()`
- `softmax()`
- `reshape()`
- `add()`

This function can export at least following networks.

- GoogLeNet
- ResNet
- VGG

And, this function use testing (evaluation) mode.

Example

```
>>> from chainer.exporters import caffe
>>>
>>> class Model(chainer.Chain):
...     def __init__(self):
...         super(Model, self).__init__()
...         with self.init_scope():
...             self.l1 = L.Convolution2D(None, 1, 1, 1, 0)
...             self.b2 = L.BatchNormalization(1)
...             self.l3 = L.Linear(None, 1)
...
...     def __call__(self, x):
...         h = F.relu(self.l1(x))
...         h = self.b2(h)
...         return self.l3(h)
...
>>> x = chainer.Variable(np.zeros((1, 10, 10, 10), np.float32))
>>> caffe.export(Model(), [x], None, True, 'test')
```

5.15 Assertion and Testing

Chainer provides some facilities to make debugging easy.

5.15.1 Type checking utilities

FunctionNode uses a systematic type checking of the `chainer.utils.type_check` module. It enables users to easily find bugs of forward and backward implementations. You can find examples of type checking in some function implementations.

<code>chainer.utils.type_check.Expr</code>	Abstract syntax tree of an expression.
<code>chainer.utils.type_check.expect</code>	Evaluates and tests all given expressions.
<code>chainer.utils.type_check.TypeInfo</code>	Type information of an input/gradient array.
<code>chainer.utils.type_check.TypeInfoTuple</code>	Type information of input/gradient tuples.

`chainer.utils.type_check.Expr`

class `chainer.utils.type_check.Expr` (*priority*)

Abstract syntax tree of an expression.

It represents an abstract syntax tree, and isn't a value. You can get its actual value with `eval()` function, and get syntax representation with the `__str__()` method. Each comparison operator (e.g. `==`) generates a new *Expr* object which represents the result of comparison between two expressions.

Example

Let `x` and `y` be instances of *Expr*, then

```
>>> x = Variable(1, 'x')
>>> y = Variable(1, 'y')
>>> c = (x == y)
```

is also an instance of *Expr*. To evaluate and get its value, call `eval()` method:

```
>>> c.eval()
True
```

Call `str` function to get a representation of the original equation:

```
>>> str(c)
'x == y'
```

You can actually compare an expression with a value:

```
>>> (x == 1).eval()
True
```

Note that you can't use boolean operators such as `and`, as they try to cast expressions to boolean values:

```
>>> z = Variable(1, 'z')
>>> x == y and y == z # raises an error
```

(continues on next page)

(continued from previous page)

```
Traceback (most recent call last):
RuntimeError: Don't convert Expr to bool. Please call Expr.eval method to
↳ evaluate expression.
```

Methods

`__call__ (*args)`

Call self as a function.

`__getitem__ (key)`

`eval ()`

Evaluates the tree to get actual value.

Behavior of this function depends on an implementation class. For example, a binary operator `+` calls the `__add__` function with the two results of `eval ()` function.

`__eq__ (y)`

`__ne__ (y)`

`__lt__ (y)`

`__le__ (y)`

`__gt__ (y)`

`__ge__ (y)`

`__nonzero__ ()`

`__bool__ ()`

`__neg__ ()`

`__add__ (y)`

`__radd__ (y)`

`__sub__ (y)`

`__rsub__ (y)`

`__mul__ (y)`

`__rmul__ (y)`

`__truediv__ (y)`

`__rtruediv__ (y)`

`__floordiv__ (y)`

`__rfloordiv__ (y)`

`__pow__ (y)`

chainer.utils.type_check.expect

`chainer.utils.type_check.expect` (**bool_exprs*)

Evaluates and tests all given expressions.

This function evaluates given boolean expressions in order. When at least one expression is evaluated as `False`, that means the given condition is not satisfied. You can check conditions with this function.

Parameters *bool_exprs* (*tuple of Bool expressions*) – Bool expressions you want to evaluate.

chainer.utils.type_check.TypeInfo

class `chainer.utils.type_check.TypeInfo` (*shape, dtype*)

Type information of an input/gradient array.

It contains type information of an array, such as the shape of array and the number of dimensions. This information is independent of CPU or GPU array.

Methods

Attributes

size

chainer.utils.type_check.TypeInfoTuple

class `chainer.utils.type_check.TypeInfoTuple`

Type information of input/gradient tuples.

It is a sub-class of tuple containing *TypeInfo*. The *i*-th element of this object contains type information of the *i*-th input/gradient data. As each element is *Expr*, you can easily check its validity.

Methods

count (*value*) → integer – return number of occurrences of value

index (*value*[, *start*[, *stop*]]) → integer – return first index of value.

Raises `ValueError` if the value is not present.

size ()

Returns an expression representing its length.

Returns An expression object representing length of the tuple.

Return type *Expr*

5.15.2 Gradient checking utilities

Most function implementations are numerically tested by *gradient checking*. This method computes numerical gradients of forward routines and compares their results with the corresponding backward routines. It enables us to make the source of issues clear when we hit an error of gradient computations. The `chainer.gradient_check` module makes it easy to implement the gradient checking.

<code>chainer.gradient_check. check_backward</code>	Test backward procedure of a given function.
<code>chainer.gradient_check. numerical_grad</code>	Computes numerical gradient by finite differences.

chainer.gradient_check.check_backward

`chainer.gradient_check.check_backward`(*func*, *x_data*, *y_grad*, *params*=(), *eps*=0.001, *atol*=1e-05, *rtol*=0.0001, *no_grads*=None, *dtype*=None, *detect_nondifferentiable*=False)

Test backward procedure of a given function.

This function automatically checks the backward-process of a given function to ensure that the computed gradients are approximately correct. For example, assuming you've defined a `FunctionNode` class `MyFunc`, that takes two arguments and returns one value, you can wrap it in an ordinary function and check its gradient computations as follows:

```
def func(xs):
    y, = MyFunc().apply(xs)
    return y

x1_data = xp.array(...)
x2_data = xp.array(...)
gy_data = xp.array(...)
check_backward(func, (x1_data, x2_data), gy_data)
```

This method creates `Variable` objects with `x_data` and calls `func` with the `Variables` to get its result as `Variable`. Then, it sets `y_grad` array to `grad` attribute of the result and calls `backward` method to get gradients of the inputs. To check correctness of the gradients, the function calls `numerical_grad()` to calculate numerically the gradients and compares the types of gradients with `chainer.testing.assert_allclose()`.

To reduce computational time, it uses directional derivative along a random vector. A function $g : \mathbb{R} \rightarrow \mathbb{R}^n$ is defined as $g(\delta) = f(x + \delta r)$, where $\delta \in \mathbb{R}$, $r \in \mathbb{R}^n$ is a random vector and f is a function which you want to test. Its gradient is

$$g'(\delta) = f'(x + \delta r) \cdot r.$$

Therefore, $g'(0) = f'(x) \cdot r$. So we can check the correctness of back propagation of f indirectly by comparing this equation with the gradient of g numerically calculated and that of f computed by backprop. If r is chosen from uniform distribution, we can conclude with high probability that the gradient of f itself is correct.

If input objects (`x1_data` or/and `x2_data` in this example) represent integer variables, their gradients are ignored.

You can simplify a test when `MyFunc` gets only one argument:

```
check_backward(func, x1_data, gy_data)
```

If `MyFunc` is a loss function which returns a zero-dimensional array, pass `None` to `gy_data`. In this case, it sets 1 to `grad` attribute of the result:

```
check_backward(my_loss_func,
               (x1_data, x2_data), None)
```

If `MyFunc` returns multiple outputs, pass all gradients for outputs as a tuple:


```
gy1_data = xp.array(...)
gy2_data = xp.array(...)
check_backward(func, x1_data, (gy1_data, gy2_data))
```

You can also test a [Link](#). To check gradients of parameters of the link, set a tuple of the parameters to `params` arguments:

```
check_backward(my_link, (x1_data, x2_data), gy_data,
                (my_link.W, my_link.b))
```

Note that `params` are not `ndarrays`, but `Variables`.

Function objects are acceptable as `func` argument:

```
check_backward(lambda x1, x2: f(x1, x2),
                (x1_data, x2_data), gy_data)
```

Note: `func` is called many times to get numerical gradients for all inputs. This function doesn't work correctly when `func` behaves randomly as it gets different gradients.

Parameters

- **func** (*callable*) – A function which gets *Variables* and returns *Variables*. `func` must return a tuple of *Variables* or one *Variable*. You can use a *Function*, *FunctionNode* or a *Link* object or any other function satisfying the condition.
- **x_data** (*ndarray or tuple of ndarrays*) – A set of `ndarrays` to be passed to `func`. If `x_data` is one `ndarray` object, it is treated as `(x_data,)`.
- **y_grad** (*ndarray or tuple of ndarrays or None*) – A set of `ndarrays` representing gradients of return-values of `func`. If `y_grad` is one `ndarray` object, it is treated as `(y_grad,)`. If `func` is a loss-function, `y_grad` should be set to `None`.
- **params** (*Variable or tuple of ~chainer.Variable*) – A set of *Variables* whose gradients are checked. When `func` is a *Link* object, set its parameters as `params`. If `params` is one *Variable* object, it is treated as `(params,)`.
- **eps** (*float*) – Epsilon value to be passed to `numerical_grad()`.
- **atol** (*float*) – Absolute tolerance to be passed to `chainer.testing.assert_allclose()`.
- **rtol** (*float*) – Relative tolerance to be passed to `chainer.testing.assert_allclose()`.
- **no_grads** (*list of bool*) – Flag to skip variable for gradient assertion. It should be same length as `x_data`.
- **dtype** (*dtype*) – `x_data`, `y_grad` and `params` are casted to this `dtype` when calculating numerical gradients. Only float types and `None` are allowed.
- **detect_nondifferentiable** (*bool*) – If `True`, check for non-differentiable inputs is enabled. If `func` is non-differentiable at `x_data`, `check_backward` raises `NondifferentiableError`.

See also:

`numerical_grad()`

chainer.gradient_check.numerical_grad

```
chainer.gradient_check.numerical_grad(f, inputs, grad_outputs, eps=0.001,
                                     detect_nondifferentiable=False, diff_atol=0,
                                     diff_rtol=0.01, center_outputs=None)
```

Computes numerical gradient by finite differences.

This function is used to implement gradient check. For usage example, see unit tests of `chainer.functions`.

By default, `numerical_grad` computes the gradient to the first order of `eps`.

Parameters

- **f** (*callable*) – Python function with no arguments that runs forward computation and returns the result.
- **inputs** (*tuple of arrays*) – Tuple of arrays that should be treated as inputs. Each element of them is slightly modified to realize numerical gradient by finite differences.
- **grad_outputs** (*tuple of arrays*) – Tuple of arrays that are treated as output gradients.
- **eps** (*float*) – Epsilon value of finite differences.
- **detect_nondifferentiable** (*bool*) – False by default. If True, `numerical_grad` checks whether `f` is differentiable at `inputs`. It requires evaluation of `f` at 5 points instead of 2. As a side effect, the accuracy of numerical gradient will be increased to the third order of `eps`. If it turns out that `f` is non-differentiable at input, `numerical_grad` raises `NondifferentiableError`.
- **diff_atol** (*float*) – Absolute tolerance of fitting error of non-differentiable point detection.
- **diff_rtol** (*float*) – Tolerance of fitting error of non-differentiable point detection relative to the output values of `f`.
- **center_outputs** (*tuple of arrays or None*) – Only used if `detect_nondifferentiable` is True. If specified, these arrays are used as the outputs of `f` at `inputs`. Otherwise, it is calculated. It can be used to reduce the computation if these arrays are already calculated before calling `numerical_grad`.

Returns Numerical gradient arrays corresponding to `inputs`.

Return type `tuple`

5.15.3 Standard Assertions

The assertions have same names as NumPy's ones. The difference from NumPy is that they can accept both `numpy.ndarray` and `cupy.ndarray`.

<code>chainer.testing.assert_allclose</code>
--

Asserts if some corresponding element of <code>x</code> and <code>y</code> differs too much.
--

chainer.testing.assert_allclose

```
chainer.testing.assert_allclose(x, y, atol=1e-05, rtol=0.0001, verbose=True)
```

Asserts if some corresponding element of `x` and `y` differs too much.

This function can handle both CPU and GPU arrays simultaneously.

Parameters

- **x** – Left-hand-side array.
- **y** – Right-hand-side array.
- **atol** (*float*) – Absolute tolerance.
- **rtol** (*float*) – Relative tolerance.
- **verbose** (*bool*) – If `True`, it outputs verbose messages on error.

5.15.4 Function testing utilities

Chainer provides some utilities for testing its functions.

<code>chainer.testing.unary_math_function_unittest</code>	Decorator for testing unary mathematical Chainer functions.
---	---

chainer.testing.unary_math_function_unittest

```
chainer.testing.unary_math_function_unittest(func, func_expected=None, label_expected=None, make_data=None, is_linear=False, forward_options=None, backward_options=None, double_backward_options=None)
```

Decorator for testing unary mathematical Chainer functions.

This decorator makes test classes test unary mathematical Chainer functions. Tested are forward and backward, including double backward, computations on CPU and GPU across parameterized shape and dtype.

Parameters

- **func** (*function or Function*) – Chainer function to be tested by the decorated test class. Taking *Function* is for backward compatibility.
- **func_expected** – Function used to provide expected values for testing forward computation. If not given, a corresponding numpy function for `func` is implicitly picked up by its name.
- **label_expected** (*string*) – String used to test labels of Chainer functions. If not given, the name of `func` is implicitly used.
- **make_data** – Function to customize input and gradient data used in the tests. It takes shape and dtype as its arguments, and returns a tuple of input, gradient and double gradient data. By default, uniform distribution ranged `[-1, 1]` is used for all of them.
- **is_linear** (*bool*) – Tells the decorator that `func` is a linear function so that it wraps `func` as a non-linear function to perform double backward test. The default value is `False`.
- **forward_options** (*dict*) – Options to be specified as an argument of `chainer.testing.assert_allclose()` function. If not given, preset tolerance values are automatically selected.
- **backward_options** (*dict*) – Options to be specified as an argument of `chainer.gradient_check.check_backward()` function. If not given, preset tolerance values are automatically selected depending on dtype.

- **double_backward_options** (*dict*) – Options to be specified as an argument of `chainer.gradient_check.check_double_backward()` function. If not given, preset tolerance values are automatically selected depending on `dtype`.

The decorated test class tests forward, backward and double backward computations on CPU and GPU across the following `parameterize()` ed parameters:

- `shape`: rank of zero, and rank of more than zero
- `dtype`: `numpy.float16`, `numpy.float32` and `numpy.float64`

Additionally, it tests the label of the Chainer function.

Chainer functions tested by the test class decorated with the decorator should have the following properties:

- Unary, taking one parameter and returning one value
- `dtype` of input and output are the same
- Elementwise operation for the supplied ndarray

Example

The following code defines a test class that tests `sin()` Chainer function, which takes a parameter with `dtype` of float and returns a value with the same `dtype`.

```
>>> import unittest
>>> from chainer import testing
>>> from chainer import functions as F
>>>
>>> @testing.unary_math_function_unittest(F.sin)
... class TestSin(unittest.TestCase):
...     pass
```

Because the test methods are implicitly injected to `TestSin` class by the decorator, it is enough to place `pass` in the class definition.

To customize test data, `make_data` optional parameter can be used. The following is an example of testing `sqrt` Chainer function, which is tested in positive value domain here instead of the default input.

```
>>> import numpy
>>>
>>> def make_data(shape, dtype):
...     x = numpy.random.uniform(0.1, 1, shape).astype(dtype)
...     gy = numpy.random.uniform(-1, 1, shape).astype(dtype)
...     ggx = numpy.random.uniform(-1, 1, shape).astype(dtype)
...     return x, gy, ggx
...
>>> @testing.unary_math_function_unittest(F.sqrt,
...                                     make_data=make_data)
... class TestSqrt(unittest.TestCase):
...     pass
```

`make_data` function which returns input, gradient and double gradient data generated in proper value domains with given `shape` and `dtype` parameters is defined, then passed to the decorator's `make_data` parameter.

API Compatibility Policy

This document explains the design policy on compatibilities of Chainer APIs. Development team should follow this policy on deciding to add, extend, and change APIs and their behaviors.

This document is written for both users and developers. Users can decide the level of dependencies on Chainer's implementations in their codes based on this document. Developers should read through this document before creating pull requests that contain changes on the interface. Note that this document may contain ambiguities on the level of supported compatibilities.

6.1 Targeted Versions

This policy is applied to Chainer v2.0.0 and higher. Note that this policy is not applied to Chainer of lower versions. For older versions of Chainer, see [the old version of API Compatibility Policy](#).

6.2 Versioning and Backward Compatibility

The versioning of Chainer follows the [PEP 440](#) and a part of [Semantic versioning](#). See [Contribution Guide](#) for details of versioning.

The backward compatibility is kept for **revision updates** and **minor updates**, which are applied to the stable version. A **major update** from the latest release candidate basically keeps the backward compatibility, although it is not guaranteed. Any **pre-releases** may break the backward compatibility.

6.3 Breaking the Compatibility

We sometimes need to break the backward compatibility to improve the framework design and to support new kinds of machine learning methods. Such a change is only made into pre-releases (alpha, beta, and release candidate) and sometimes into the major update.

A change that breaks the compatibility affects user codes. We try to lower the cost of adapting your code to the newer version. The following list shows an example of what we can do to reduce the cost (*Note: this is not a promise; what kind of actions we can take depends on the situation*).

- When an argument is removed from an existing API, passing the argument to the updated API will emit an error with a special error message. The error message tells you how to fix your code.
- When a function or a class is removed, we make the current stable version emit a deprecation warning. **Note that the deprecation warning is not printed by default in Python.** You have to manually turn on the deprecation warning by `warnings.simplefilter('always', DeprecationWarning)`.
- When a definition of a link is changed, we try to enable it to deserialize a model dumped with an older version of Chainer. In most cases, we cannot guarantee that a model serialized with a newer version of Chainer is loadable by an older version of Chainer.

Note: Since Chainer v2, we have stopped adopting any solid processes to break backward compatibilities (e.g. a solid schedule for deprecating and removing a feature) in order to keep the development fast enough to support the cutting-edge research. **It does not mean we stop taking care of maintainability of user codes.** We are still paying much attention to not breaking user codes.

6.4 Experimental APIs

Thanks to many contributors, we have introduced many new features to Chainer.

However, we have sometimes released new features only to later notice that their APIs are not appropriate. In particular, we sometimes know that the API is likely to be modified in the near future because we do not have enough knowledge about how well the current design fits to the real usages. **The objective of experimental APIs is to declare that the APIs are likely to be updated in the near future so that users can decide if they can(not) use them.**

Any newly added API can be marked as *experimental*. Any API that is not experimental is called *stable* in this document.

Note: Undocumented behaviors are not considered as APIs, so they can be changed at any time (even in a revision update). The treatment of undocumented behaviors are described in [Undocumented behaviors](#) section.

When users use experimental APIs for the first time, warnings are raised once for each experimental API, unless users explicitly disable the emission of the warnings in advance.

See the document of `chainer.utils.experimental()` to know how developers mark APIs as experimental and how users enable or disable the warnings practically.

Note: It is up to developers if APIs should be annotated as experimental or not. We recommend to make the APIs experimental if they implement large modules or make a decision from several design choices.

6.5 Supported Backward Compatibility

This section defines backward compatibilities that revision updates must maintain.

6.5.1 Documented Interface

Chainer has the official API documentation. Many applications can be written based on the documented features. We support backward compatibilities of documented features. In other words, codes only based on the documented features run correctly with revision-updated versions.

Developers are encouraged to use apparent names for objects of implementation details. For example, attributes outside of the documented APIs should have one or more underscores at the prefix of their names.

Note: Although it is not stated as a rule, we also try to keep the compatibility for any interface that looks like a stable feature. For example, if the name of a symbol (function, class, method, attribute, etc.) is not prefixed by an underscore and the API is not experimental, the API should be kept over revision updates even if it is not documented.

6.5.2 Undocumented behaviors

Behaviors of Chainer implementation not stated in the documentation are undefined. Undocumented behaviors are not guaranteed to be stable between different revision versions.

Even revision updates may contain changes to undefined behaviors. One of the typical examples is a bug fix. Another example is an improvement on implementation, which may change the internal object structures not shown in the documentation. As a consequence, **even revision updates do not support compatibility of pickling, unless the full layout of pickled objects is clearly documented.**

6.5.3 Documentation Error

Compatibility is basically determined based on the documentation, although it sometimes contains errors. It may make the APIs confusing to assume the documentation always stronger than the implementations. We therefore may fix the documentation errors in any updates that may break the compatibility in regard to the documentation.

Note: Developers should not fix the documentation and implementation of the same functionality at the same time in revision updates as a “bug fix” unless the bug is so critical that no users are expected to be using the old version correctly.

6.5.4 Object Attributes and Properties

Object attributes and properties are sometimes replaced by each other. It does not break the user codes, except the codes depend on how the attributes and properties are implemented.

6.5.5 Functions and Methods

Methods may be replaced by callable attributes keeping the compatibility of parameters and return values. It does not break the user codes, except the codes depend on how the methods and callable attributes are implemented.

6.5.6 Exceptions and Warnings

The specifications of raising exceptions are considered as a part of standard backward compatibilities. No exception is raised in the future revision versions with correct usages that the documentation allows.

On the other hand, warnings may be added at any revision updates for any APIs. It means revision updates do not keep backward compatibility of warnings.

6.6 Model Format Compatibility

Links and chains serialized by official serializers that Chainer provides are correctly loaded with the future versions. They might not be correctly loaded with Chainer of the lower versions.

Note: Current serialization APIs do not support versioning. It prevents us from introducing changes in the layout of objects that support serialization. We are discussing versioning in serialization APIs.

6.7 Installation Compatibility

The installation process is another concern of compatibilities.

Any changes on the set of dependent libraries that force modifications on the existing environments should be done in pre-releases and major updates. Such changes include following cases:

- dropping supported versions of dependent libraries (e.g. dropping cuDNN v2)
- adding new mandatory dependencies (e.g. adding h5py to setup_requires)

Note: We sometimes have to narrow the supported versions due to bugs in the specific versions of libraries. In such a case, we may drop the support of those versions even in revision updates unless a workaround is found for the issue.

This is a guide for all contributions to Chainer. The development of Chainer is running on [the official repository at GitHub](#). Anyone that wants to register an issue or to send a pull request should read through this document.

Note: Many points of this document are updated at v2. We strongly recommend all contributors of v1 to read through the document again.

7.1 Classification of Contributions

There are several ways to contribute to Chainer community:

1. Registering an issue
2. Sending a pull request (PR)
3. Sending a question/reply to [StackOverflow](#) (with `chainer` tag) or [Chainer User Group](#)
4. Open-sourcing an external example
5. Writing a post about Chainer

This document mainly focuses on 1 and 2, though other contributions are also appreciated.

7.2 Development Cycle

This section explains the development process of Chainer. Before contributing to Chainer, it is strongly recommended to understand the development cycle.

7.2.1 Versioning

The versioning of Chainer follows [PEP 440](#) and a part of [Semantic versioning](#). The version number consists of three or four parts: $X.Y.Z_w$ where X denotes the **major version**, Y denotes the **minor version**, Z denotes the **revision number**, and the optional w denotes the pre-release suffix. While the major, minor, and revision numbers follow the rule of semantic versioning, the pre-release suffix follows PEP 440 so that the version string is much friendly with Python eco-system.

Note that a major update basically does not contain compatibility-breaking changes from the last release candidate (RC). This is not a strict rule, though; if there is a critical API bug that we have to fix for the major version, we may add breaking changes to the major version up.

As for the backward compatibility, see [API Compatibility Policy](#).

7.2.2 Release Cycle

Starting from v2.0.0, we are developing two tracks of versions at the same time. The first one is the track of **stable versions**, which is a series of revision updates for the latest major version. The second one is the track of **development versions**, which is a series of pre-releases for the upcoming major version.

Consider that $X.0.0$ is the latest major version and $Y.0.0$, $Z.0.0$ are the succeeding major versions. Then, the timeline of the updates is depicted by the following table.

Date	ver X	ver Y	ver Z
0 weeks	X.0.0rc1	–	–
4 weeks	X.0.0	Y.0.0a1	–
8 weeks	X.1.0*	Y.0.0b1	–
12 weeks	X.2.0*	Y.0.0rc1	–
16 weeks	–	Y.0.0	Z.0.0a1

(* These might be revision releases)

The dates shown in the left-most column are relative to the release of $X.0.0rc1$. In particular, each revision/minor release is made four weeks after the previous one of the same major version, and the pre-release of the upcoming major version is made at the same time. Whether these releases are revision or minor is determined based on the contents of each update.

Note that there are only three stable releases for the versions $X.x.x$. During the parallel development of $Y.0.0$ and $Z.0.0a1$, the version Y is treated as an **almost-stable version** and Z is treated as a development version.

If there is a critical bug found in $X.x.x$ after stopping the development of version X , we may release a hot-fix for this version at any time.

We create a milestone for each upcoming release at GitHub. The GitHub milestone is basically used for collecting the issues and PRs resolved in the release.

7.2.3 Git Branches

The `master` branch is used to develop pre-release versions. It means that **alpha, beta, and RC updates are developed at the `master` branch**. This branch contains the most up-to-date source tree that includes features newly added after the latest major version.

The stable version is developed at the individual branch named as vN where “ N ” reflects the version number (we call it a *versioned branch*). For example, v3.0.0, v3.0.1, and v3.0.2 will be developed at the `v3` branch.

Notes for contributors: When you send a pull request, you basically have to send it to the `master` branch. If the change can also be applied to the stable version, a core team member will apply the same change to the stable version so that the change is also included in the next revision update.

If the change is only applicable to the stable version and not to the `master` branch, please send it to the versioned branch. We basically only accept changes to the latest versioned branch (where the stable version is developed) unless the fix is critical.

If you want to make a new feature of the `master` branch available in the current stable version, please send a *backport PR* to the stable version (the latest `vN` branch). See the next section for details.

Note: a change that can be applied to both branches should be sent to the `master` branch. Each release of the stable version is also merged to the development version so that the change is also reflected to the next major version.

7.2.4 Feature Backport PRs

We basically do not backport any new features of the development version to the stable versions. If you desire to include the feature to the current stable version and you can work on the backport work, we welcome such a contribution. In such a case, you have to send a backport PR to the latest `vN` branch. **Note that we do not accept any feature backport PRs to older versions because we are not running quality assurance workflows (e.g. CI) for older versions so that we cannot ensure that the PR is correctly ported.**

There are some rules on sending a backport PR.

- Start the PR title from the prefix **[backport]**.
- Clarify the original PR number in the PR description (something like “This is a backport of #XXXX”).
- (optional) Write to the PR description the motivation of backporting the feature to the stable version.

Please follow these rules when you create a feature backport PR.

Note: PRs that do not include any changes/additions to APIs (e.g. bug fixes, documentation improvements) are usually backported by core dev members. It is also appreciated to make such a backport PR by any contributors, though, so that the overall development proceeds more smoothly!

7.3 Issues and Pull Requests

In this section, we explain how to file issues and send pull requests (PRs).

7.3.1 Issue/PR Labels

Issues and PRs are labeled by the following tags:

- **Bug:** bug reports (issues) and bug fixes (PRs)
- **Enhancement:** implementation improvements without breaking the interface
- **Feature:** feature requests (issues) and their implementations (PRs)
- **NoCompat:** disrupts backward compatibility
- **Test:** test fixes and updates
- **Document:** document fixes and improvements
- **Example:** fixes and improvements on the examples
- **Install:** fixes installation script

- **Contribution-Welcome:** issues that we request for contribution (only issues are categorized to this)
- **Other:** other issues and PRs

Multiple tags might be labeled to one issue/PR. **Note that revision releases cannot include PRs in Feature and NoCompat categories.**

7.3.2 How to File an Issue

On registering an issue, write precise explanations on how you want Chainer to be. Bug reports must include necessary and sufficient conditions to reproduce the bugs. Feature requests must include **what** you want to do (and **why** you want to do, if needed) with Chainer. You can contain your thoughts on **how** to realize it into the feature requests, though **what** part is most important for discussions.

Warning: If you have a question on usages of Chainer, it is highly recommended to send a post to [StackOverflow](#) or [Chainer User Group](#) instead of the issue tracker. The issue tracker is not a place to share knowledge on practices. We may suggest these places and immediately close how-to question issues.

7.3.3 How to Send a Pull Request

If you can write code to fix an issue, we encourage to send a PR.

First of all, before starting to write any code, do not forget to confirm the following points.

- Read through the [Coding Guidelines](#) and [Unit Testing](#).
- Check the appropriate branch that you should send the PR following [Git Branches](#). If you do not have any idea about selecting a branch, please choose the `master` branch.

In particular, **check the branch before writing any code**. The current source tree of the chosen branch is the starting point of your change.

After writing your code (**including unit tests and hopefully documentations!**), send a PR on GitHub. You have to write a precise explanation of **what** and **how** you fix; it is the first documentation of your code that developers read, which is a very important part of your PR.

Once you send a PR, it is automatically tested on [Travis CI](#) for Linux and Mac OS X, and on [AppVeyor](#) for Windows. Your PR needs to pass at least the test for Linux on Travis CI. After the automatic test passes, some of the core developers will start reviewing your code. Note that this automatic PR test only includes CPU tests.

Note: We are also running continuous integration with GPU tests for the `master` branch and the versioned branch of the latest major version. Since this service is currently running on our internal server, we do not use it for automatic PR tests to keep the server secure.

If you are planning to add a new feature or modify existing APIs, **it is recommended to open an issue and discuss the design first**. The design discussion needs lower cost for the core developers than code review. Following the consequences of the discussions, you can send a PR that is smoothly reviewed in a shorter time.

Even if your code is not complete, you can send a pull request as a *work-in-progress PR* by putting the `[WIP]` prefix to the PR title. If you write a precise explanation about the PR, core developers and other contributors can join the discussion about how to proceed the PR. WIP PR is also useful to have discussions based on a concrete code.

7.4 Coding Guidelines

Note: Coding guidelines are updated at v3.0. Those who have contributed to older versions should read the guidelines again.

We use [PEP 8](#) and a part of [OpenStack Style Guidelines](#) related to general coding style as our basic style guidelines.

To check your code, use `autopep8` and `flake8` command installed by `hacking` package:

```
$ pip install autopep8 hacking
$ autopep8 path/to/your/code.py
$ flake8 path/to/your/code.py
```

The `autopep8` supports automatically correct Python code to conform to the PEP 8 style guide:

```
$ autopep8 --in-place path/to/your/code.py
```

The `flake8` command lets you know the part of your code not obeying our style guidelines. Before sending a pull request, be sure to check that your code passes the `flake8` checking.

Note that `flake8` command is not perfect. It does not check some of the style guidelines. Here is a (not-complete) list of the rules that `flake8` cannot check.

- Relative imports are prohibited. [H304]
- Importing non-module symbols is prohibited.
- Import statements must be organized into three parts: standard libraries, third-party libraries, and internal imports. [H306]

In addition, we restrict the usage of *shortcut aliases* in any global-scope code. In particular, you cannot use shortcut aliases to designate a parent class in global-scope class definitions. When you want to make a class inheriting another class defined in another module, you have to spell out the full module name instead of importing a module that provides an alias.

For example, the following code is not allowed.

```
import chainer

class MyLink(chainer.Link): ...
```

Instead, import `chainer.link` and use that.

```
import chainer.link

class MyLink(chainer.link.Link): ...
```

If you feel the code too verbose, you can also use `from import` or `import as`.

```
from chainer import link

class MyLink(link.Link): ...
```

Note: From v3.0, we allow shortcut aliases used inside of functions and methods that are not called from any global scope code. For example, you can write `chainer.Variable` instead of `chainer.variable.Variable` inside of functions and methods. Use of such aliases is prohibited in the past for avoiding confusing errors related to cyclic dependencies; we relaxed the rule so that the library code looks similar to user code.

When you use such shortcut aliases, please be careful with cyclic imports. One of the typical pitfalls is a way to `import chainer.functions`. An import like `import chainer.functions as F` within modules under `chainer.functions` does not work. An import like `from chainer import functions` works well with Python 3, but does not with Python 2. We recommend you to use `import chainer.functions` and spell out like `chainer.functions.foo` in your methods.

Once you send a pull request, your coding style is automatically checked by [Travis-CI](#). The reviewing process starts after the check passes.

7.5 Unit Testing

Testing is one of the most important part of your code. You must write test cases and verify your implementation by following our testing guide.

Note that we are using `pytest` and `mock` package for testing, so install them before writing your code:

```
$ pip install pytest mock
```

7.5.1 How to Run Tests

You can run unit tests simply by running `python -m pytest` command at the repository root:

```
$ python -m pytest
```

or specify the test script that you want to run:

```
$ python -m pytest path/to/your/test.py
```

You can also run all unit tests under a specified directory:

```
$ python -m pytest tests/chainer_tests/<directory name>
```

It requires CUDA and cuDNN by default. In order to run unit tests that do not require CUDA and cuDNN, use `CHAINER_TEST_GPU_LIMIT=0` environment variable and `-m='not cudnn'` option:

```
$ export CHAINER_TEST_GPU_LIMIT=0
$ python -m pytest path/to/your/test.py -m='not cudnn'
```

Some GPU tests involve multiple GPUs. If you want to run GPU tests with insufficient number of GPUs, specify the number of available GPUs to `CHAINER_TEST_GPU_LIMIT`. For example, if you have only one GPU, launch `pytest` by the following command to skip multi-GPU tests:

```
$ export CHAINER_TEST_GPU_LIMIT=1
$ python -m pytest path/to/gpu/test.py
```

Some tests spend too much time. If you want to skip such tests, pass `-m='not slow'` option to the command:

```
$ python -m pytest path/to/your/test.py -m='not slow'
```

If you modify the code related to existing unit tests, you must run appropriate commands and confirm that the tests pass.

7.5.2 Test File and Directory Naming Conventions

Tests are put into the `tests/chainer_tests` directory. In order to enable test runner to find test scripts correctly, we are using special naming convention for the test subdirectories and the test scripts.

- The name of each subdirectory of `tests` must end with the `_tests` suffix.
- The name of each test script must start with the `test_` prefix.

When we write a test for a module, we use the appropriate path and file name for the test script whose correspondence to the tested module is clear. For example, if you want to write a test for a module `chainer.x.y.z`, the test script must be located at `tests/chainer_tests/x_tests/y_tests/test_z.py`.

7.5.3 How to Write Tests

There are many examples of unit tests under the `tests` directory, so reading some of them is a good and recommended way to learn how to write tests for Chainer. They simply use the `unittest` package of the standard library, while some tests are using utilities from `chainer.testing`.

In addition to the *Coding Guidelines* mentioned above, the following rules are applied to the test code:

- All test classes must inherit from `unittest.TestCase`.
- Use `unittest` features to write tests, except for the following cases:
 - Use `assert` statement instead of `self.assert*` methods (e.g., write `assert x == 1` instead of `self.assertEqual(x, 1)`).
 - Use `with pytest.raises(...):` instead of `with self.assertRaises(...):`.

Note: We are incrementally applying the above style. Some existing tests may be using the old style (`self.assertRaises`, etc.), but all newly written tests should follow the above style.

Even if your patch includes GPU-related code, your tests should not fail without GPU capability. Test functions that require CUDA must be tagged by `chainer.testing.attr.gpu` decorator:

```
import unittest
from chainer.testing import attr

class TestMyFunc(unittest.TestCase):
    ...

    @attr.gpu
    def test_my_gpu_func(self):
        ...
```

The functions tagged by the `gpu` decorator are skipped if `CHAINER_TEST_GPU_LIMIT=0` environment variable is set. We also have the `chainer.testing.attr.cudnn` decorator to let `pytest` know that the test depends on cuDNN. The test functions decorated by `cudnn` are skipped if `-m='not cudnn'` is given.

The test functions decorated by `gpu` must not depend on multiple GPUs. In order to write tests for multiple GPUs, use `chainer.testing.attr.multi_gpu()` decorator instead:

```
import unittest
from chainer.testing import attr

class TestMyFunc(unittest.TestCase):
```

(continues on next page)

(continued from previous page)

```
...

@attr.multi_gpu(2) # specify the number of required GPUs here
def test_my_two_gpu_func(self):
    ...
```

If your test requires too much time, add `chainer.testing.attr.slow` decorator. The test functions decorated by `slow` are skipped if `-m='not slow'` is given:

```
import unittest
from chainer.testing import attr

class TestMyFunc(unittest.TestCase):
    ...

    @attr.slow
    def test_my_slow_func(self):
        ...
```

Note: If you want to specify more than two attributes, use and operator like `-m='not cudnn and not slow'`. See detail in [the document of pytest](#).

Once you send a pull request, your code is automatically tested by [Travis-CI](#) except for tests annotated with “`gpu`”, “`multi_gpu`” and “`slow`”. Since Travis-CI does not support CUDA, we cannot check your CUDA-related code automatically. The reviewing process starts after the test passes. Note that reviewers will test your code without the option to check CUDA-related code.

Note: Some of numerically unstable tests might cause errors irrelevant to your changes. In such a case, we ignore the failures and go on to the review process, so do not worry about it!

7.6 Documentation

When adding a new feature to the framework, you also need to document it in the reference. For example, if you are adding a new function under `chainer.functions`, you need to add it to the [Functions](#) page.

Note: If you are unsure about how to fix the documentation, you can submit a pull request without doing so. Reviewers will help you fix the documentation appropriately.

The documentation source is stored under [docs directory](#) and written in [reStructuredText](#) format.

To build the documentation, you need to install [Sphinx](#):

```
$ pip install sphinx sphinx_rtd_theme
```

Then you can build the documentation in HTML format locally:

```
$ cd docs
$ make html
```


HTML files are generated under `build/html` directory. Open `index.html` with the browser and see if it is rendered as expected.

Note: Docstrings (documentation comments in the source code) are collected from the installed Chainer module. If you modified docstrings, make sure to install the module (e.g., using `pip install -e .`) before building the documentation.

8.1 It takes too long time to compile a computational graph. Can I skip it?

Chainer does not compile computational graphs, so you cannot skip it, or, I mean, you have already skipped it :).

It seems you have actually seen on-the-fly compilations of CUDA kernels. CuPy compiles kernels on demand to make kernels optimized to the number of dimensions and element types of input arguments. Pre-compilation is not available, because we have to compile an exponential number of kernels to support all CuPy functionalities. This restriction is unavoidable because Python cannot call CUDA/C++ template functions in generic way. Note that every framework using CUDA require compilation at some point; the difference between other statically-compiled frameworks (such as `cutorch`) and Chainer is whether a kernel is compiled at installation or at the first use.

These compilations should run only at the first use of the kernels. The compiled binaries are cached to the `$ (HOME) / .cupy/kernel_cache` directory by default. If you see that compilations run every time you run the same script, then the caching is failed. Please check that the directory is kept as is between multiple executions of the script. If your home directory is not suited to caching the kernels (e.g. in case that it uses NFS), change the kernel caching directory by setting the `CUPY_CACHE_DIR` environment variable to an appropriate path. See [CuPy Overview](#) for more details.

8.2 MNIST example does not converge in CPU mode on Mac OS X

Note: Mac OS X is not officially supported. Please use it at your own risk.

Many users have reported that MNIST example does not work correctly when using `vecLib` as NumPy backend on Mac OS X. `vecLib` is the default BLAS library installed on Mac OS X.

We recommend using other BLAS libraries such as [OpenBLAS](#).

To use an alternative BLAS library, it is necessary to reinstall NumPy. Here is an instruction to install NumPy with OpenBLAS using [Homebrew](#).

```
$ brew tap homebrew/science
$ brew install openblas
$ brew install numpy --with-openblas
```

If you want to install NumPy with pip, use [site.cfg](#) file.

For details of this problem, see [issue #704](#).

8.3 How do I fix InvalidType error?

Chainer raises an `InvalidType` exception when invalid inputs are given to [Functions](#). If you got `InvalidType`, generally you need to check if `dtype` and/or `shape` of inputs are valid for the function.

Here are some examples of `InvalidType` errors:

```
import chainer.functions as F
import numpy as np
```

```
x = np.arange(10) - 5
F.relu(x)
```

```
Traceback (most recent call last):
...
chainer.utils.type_check.InvalidType:
Invalid operation is performed in: ReLU (Forward)

Expect: in_types[0].dtype.kind == f
Actual: i != f
```

In this case, `kind` of `in_types[0]` (which means the first input to the function, `x`) is expected to be `f` (floating-point), whereas the input was `i` (signed integer). You need to cast the input appropriately before passing to the function (e.g., `x.astype(np.float32)`).

```
import chainer.functions as F
import numpy as np
```

```
x = np.ones((4, 4))
y = np.ones((3, 3))
F.concat([x, y])
```

```
Traceback (most recent call last):
...
chainer.utils.type_check.InvalidType:
Invalid operation is performed in: Concat (Forward)

Expect: in_types[0].shape[0] == in_types[1].shape[0]
Actual: 4 != 3
```

In this case, the function expects that `x.shape[0]` is equal to `y.shape[0]`, but actually it was 4 and 3, respectively.

See [Type Checks](#) for the detailed behavior of type checking system in Chainer.

8.4 How do I accelerate my model using iDeep on Intel CPU?

Follow these steps to utilize iDeep in your model.

8.4.1 Install iDeep

The following environments are recommended by iDeep.

- Ubuntu 14.04 / 16.04 LTS (64-bit) and CentOS 7 (64-bit)
- Python 2.7.5+, 3.5.2+, and 3.6.0+

On recommended systems, you can install iDeep wheel (binary distribution) by:

```
$ pip install 'ideep4py<2'
```

8.4.2 Enable iDeep Configuration

Currently iDeep is disabled by default because it is an experimental feature. You need to manually enable iDeep by changing `chainer.config.use_ideep` configuration to 'auto'. See [Configuring Chainer](#) for details.

The easiest way to change the configuration is to set environment variable as follows:

```
export CHAINER_USE_IDEEP="auto"
```

You can also use `chainer.using_config()` to change the configuration.

```
x = np.ones((3, 3), dtype='f')
with chainer.using_config('use_ideep', 'auto'):
    y = chainer.functions.relu(x)
print(type(y.data))
```

```
<class 'ideep4py.mdarray'>
```

8.4.3 Convert Your Model to iDeep

You need to call `model.to_intel64()` (in the same way you call `model.to_gpu()` to transfer your link to GPU) to convert the link to iDeep.

8.4.4 Run Your Model

Now your model is accelerated by iDeep!

Please note that not all functions and optimizers support iDeep acceleration. Also note that iDeep will not be used depending on the shape and data type of the input data.

8.5 My training process gets stuck when using MultiprocessIterator

When you are using OpenCV somewhere in your code and the `MultiprocessIterator` is used in the training code, the training loop may get stuck at some point. In such situation, there are several workarounds to prevent the process got stuck.

1. Set the environment variable as follows: `OMP_NUM_THREADS=1`
2. Add `cv2.setNumThreads(0)` right after `import cv2` in your training script.
3. Use *`MultithreadIterator`* instead of *`MultiprocessIterator`*.

This problem is originally reported here: [A training loop got stuck in a certain condition with multi-processing updater and opencv](#) for Chainer and the discussion on related problems is still going here: [OpenCV + Python multiprocessing breaks on OSX](#).

Performance Best Practices

This guide explains some tips and advice for maximizing the performance of Chainer.

9.1 Use the Latest Version

It is generally recommended to use the latest version of Chainer and its dependent libraries (CUDA, cuDNN, iDeep, etc.). Some of the new features and performance optimizations introduced in newer versions of dependent libraries may not be available in older versions of Chainer. Also, Chainer itself is incrementally being improved to provide better performance.

If you are using Chainer v4 or later, you can check the version configuration by:

```
chainer.print_runtime_info()
```

```
Chainer: 4.0.0
NumPy: 1.14.3
CuPy:
  CuPy Version      : 4.0.0
  CUDA Root         : /usr/local/cuda
  CUDA Build Version : 9000
  CUDA Driver Version : 9000
  CUDA Runtime Version : 9000
  cuDNN Build Version : 7100
  cuDNN Version      : 7100
  NCCL Build Version  : 2102
```

Generally, the Chainer team is maintaining the API between minor updates (e.g., v4.0 to v4.1) so that users can upgrade Chainer without modifying their code (see [API Compatibility Policy](#) for our policy). As for major updates, please refer to the [Upgrade Guide](#) to understand what should be done for migration.

9.2 Enable Hardware Accelerations

9.2.1 Using GPU

In most cases, running on GPU will give you better performance than on CPU. When using GPU, also make sure to install cuDNN, which is a library to accelerate deep neural network computations.

Note: You don't have to manually install cuDNN if you are using [CuPy wheels](#), which includes the latest version of cuDNN. Check the output of `chainer.print_runtime_info()`; if you see the cuDNN version number, it is installed properly and will be used by Chainer automatically.

Note: If you wish, you can manually disable use of cuDNN using `chainer.config.use_cudnn` configuration option. See [Configuring Chainer](#) for details.

9.2.2 Using CPU

If you are running Chainer on CPU, you can use [iDeep](#) to utilize vector instructions of CPU. See [Tips and FAQs](#) for steps to run your model with iDeep.

You can also improve performance by building NumPy linked to [Intel MKL](#). See [Numpy/Scipy with Intel® MKL and Intel® Compilers](#) for the detailed instructions.

Note: If you installed [numpy](#) package using Anaconda, you may already have MKL-linked NumPy. Check the output of `numpy.show_config()` to see what linear algebra library is linked.

Note: Use of iDeep and MKL-linked NumPy are orthogonal. You can use both of them at once to maximize the performance.

9.3 Migrate Data Preprocessing Code from NumPy to CuPy

If you are preprocessing your dataset or running data augmentation using NumPy, you may be able to use CuPy as a substitution to improve performance.

Note: It is **not always** efficient to use CuPy instead of NumPy, especially when the computation is not very heavy, or it cannot be done in batch.

9.4 Avoid Data Transfer

If you are using GPU, be aware of data transfer between CPU and GPU. For example, printing `chainer.Variable` on GPU (e.g., for debugging) will cause memory transfer from GPU to CPU, which will incur synchronization overhead.

You can use [NVIDIA Visual Profiler](#) to diagnose this kind of issue.

9.5 Optimize cuDNN Convolution

9.5.1 Workspace Size

Some convolution algorithms in cuDNN use additional GPU memory as a temporary buffer. This is called “workspace,” and users can adjust the upper limit of its size. By increasing the limit of workspace size, cuDNN may be able to use better (i.e., memory consuming but faster) algorithm.

The default size (in bytes) is:

```
>>> chainer.backends.cuda.get_max_workspace_size()
8388608
```

and can be adjusted using `chainer.backends.cuda.set_max_workspace_size()`.

Maximum required workspace size may vary depending on various conditions such as GPU hardware and batch size of inputs.

9.5.2 Auto-Tuner

Some convolution algorithms in cuDNN support the auto-tuner feature that finds the fastest convolution algorithm for given inputs. You can turn on this feature by setting `autotune` configuration to `True`.

See [Configuring Chainer](#) for detailed descriptions.

Note: Auto-tuner tries to find the best algorithm for every first observation of the input shape combination. Therefore, the first batch will become slower when auto-tuner is enabled. The result of auto-tuner is cached on memory so that it can be reused for data with the same input shape combination. In other words, algorithm selected in the first batch will be reused for the second and later batches, as long as the input shape combination is the same.

If you set `autotune` configuration to `False`, the default convolution algorithm will always be selected, regardless of the previous auto-tuner results.

Note: Auto-tuner always use the maximum workspace size.

9.6 Fine-Tune Configuration

There are some Chainer configuration values that affect performance. Although the default values work well in most cases, you can adjust the following configurations for better performance.

- `enable_backprop`

If you are running your model for inference (i.e., you don’t have to use back propagation because you are not training the model), you can set this configuration to `False` to improve performance and reduce memory consumption.

- `type_check`

By default, Chainer checks the integrity between input data and functions. This makes possible to display friendly message when, for example, data with invalid dtype or shape is given to a function. By setting this configuration to `False`, you can let Chainer skip such check to improve performance. It is recommended to turn off the check only for well-tested code and input data.

See *Configuring Chainer* for detailed descriptions.

9.7 Load Datasets Concurrently

If loading process of your dataset is I/O-bound or CPU-bound, consider using `chainer.iterators.MultithreadIterator` or `chainer.iterators.MultiprocessIterator` to load dataset concurrently using multiple threads or processes, instead of `chainer.iterators.SerialIterator` which works in a single thread in a single process.

9.8 Use Multiple GPUs

You can utilize multiple GPUs to make the training process faster.

For data parallelism, you can use `chainer.training.updaters.ParallelUpdater` or `chainer.training.updaters.MultiprocessParallelUpdater` instead of `chainer.training.updaters.StandardUpdater`. For model parallelism, you need to manually transfer each `chainer.Link` in your model to each device.

See *Using GPU(s) in Chainer* for the working examples of each case.

9.9 Use Multiple Nodes

You can scale-out the training process of your Chainer model to multiple-node cluster by using [ChainerMN](#), an additional package for Chainer which enables distributed deep learning. See [ChainerMN Official Documentation](#) for details.

This is a list of changes introduced in each release that users should be aware of when migrating from older versions. Most changes are carefully designed not to break existing code; however changes that may possibly break them are highlighted with a box.

10.1 Chainer v4

10.1.1 Introduction of Backend Namespace

We introduced `chainer.backends` subpackage for future support of various backend libraries other than NumPy and CuPy. By this change, `chainer.cuda` module is now moved to `chainer.backends.cuda`.

This does not break the existing code; you can safely continue to use `chainer.cuda` (e.g., from `chainer import cuda`) but it is now encouraged to use `from chainer.backends import cuda` instead.

10.1.2 Namespace Changes for Updaters

`chainer.training.StandardUpdater` and `chainer.training.ParallelUpdater` are now moved to `chainer.training.updaters.StandardUpdater` and `chainer.training.updaters.ParallelUpdater` respectively, to align with the namespace convention of other subpackages. See the discussion in #2982 for more details.

This change does not break the existing code; you can safely continue to use updater classes directly under `chainer.training` but it is now encouraged to use `chainer.training.updaters` instead.

10.1.3 Namespace Changes for Optimizer Hooks

Optimizer hook functions are moved from `chainer.optimizer.*` to `chainer.optimizer_hooks.*`. For example, `chainer.optimizer.WeightDecay` is now located `chainer.optimizer_hooks.WeightDecay`.

If the existing code is using hooks directly under `chainer.optimizer`, `DeprecationWarning` will be shown. You are now encouraged to use `chainer.optimizer_hooks` instead.

10.1.4 Prohibition of Mixed Use of Arrays on Different Devices in Function Arguments

Argument validation of functions is now stricened to check device consistency of argument variables to provide better error messages to users. Suppose the following code:

```
v1 = chainer.Variable(np.arange(10, dtype=np.float32))    # CPU
v2 = chainer.Variable(cupy.arange(10, dtype=cupy.float32)) # GPU

# The line below raises an exception, because arguments are on different device.
F.maximum(v1, v2)
```

Prior to v4, the above code raises an exception like `ValueError: object __array__ method not producing an array`, which was difficult to understand. In v4, the error message would become `TypeError: incompatible array types are mixed in the forward input (Maximum)`. This kind of error usually occurs by mistake (for example, not performing `to_gpu` for some variables).

Attention: As the argument validation is stricened, call of functions intentionally mixing NumPy/CuPy arrays in arguments will not work in Chainer v4. Please transfer all arrays to the same device before calling functions.

10.1.5 References to Function Nodes Not Retained in TimerHook and CupyMemoryProfilerHook

To reduce memory consumption, references to the function nodes will no longer be retained in the `chainer.function_hooks.CupyMemoryProfileHook` and `chainer.function_hooks.TimerHook`. See the discussion in #4300 for more details.

Attention: The existing code using function nodes retained in `call_history` attribute of these hooks will not work. The first element of `call_history` became the name of the function, instead of the function node instance itself. You can define your own function hook if you need to access the function node instances.

10.1.6 Update of Docker Images

Chainer official Docker images (see [Installation](#) for details) are now updated to use CUDA 8.0 and cuDNN 6.0. This change was introduced because CUDA 7.5 does not support NVIDIA Pascal GPUs.

To use these images, you may need to upgrade the NVIDIA driver on your host. See [Requirements of nvidia-docker](#) for details.

10.1.7 CuPy v4

Chainer v4 requires CuPy v4 if you need GPU support. Please see the [Upgrade Guide for CuPy v4](#) for details.

10.2 Chainer v3

10.2.1 Introduction of New-style Functions

This release introduces new-style functions (classes inheriting from `FunctionNode`) that support double backward (gradient of gradient). See the [Release Note for v3.0.0](#) for the usage of this feature.

Many of `Functions` are already migrated to new-style, although some of functions are still old-style (classes inheriting from `Function`). We are going to migrate more old-style functions to new-style in upcoming minor releases.

This does not break the existing code. Old-style functions (classes inheriting from `Function`) are still supported in v3 and future versions of Chainer.

If you are going to write new functions, it is encouraged to use `FunctionNode` to support double backward.

Attention: Users relying on undocumented function APIs (directly instantiating old-style classes) may experience an error like `TypeError: 'SomeFunction' object is not callable` after upgrading to v3. Please use the function APIs documented in `Functions`.

10.2.2 Changed Behavior of `matmul` Function

The behavior of `chainer.functions.matmul()` has been changed to behave like the corresponding NumPy function (`numpy.matmul()`). See the discussion in [#2426](#) for more details.

Attention: The existing code using `chainer.functions.matmul()` may require modification to work with Chainer v3.

Also note that `chainer.functions.batch_matmul()` is now deprecated by this change. You can rewrite it using `chainer.functions.matmul()`.

10.2.3 Removed `use_cudnn` Argument in `spatial_transformer_grid` and `spatial_transformer_sampler` Functions

`use_cudnn` argument has been removed from `chainer.functions.spatial_transformer_grid()` and `chainer.functions.spatial_transformer_sampler()`. See the discussion in [#2955](#) for more details.

Attention: The existing code using `use_cudnn` argument of `chainer.functions.spatial_transformer_grid()` and `chainer.functions.spatial_transformer_sampler()` require modification to work with Chainer v3. Please use the configuration context (e.g., with `chainer.config('use_cudnn', 'auto')`) to enable or disable use of cuDNN. See [Configuring Chainer](#) for details.

10.2.4 CuPy v2

Chainer v3 requires CuPy v2 if you need GPU support. Please see the [Upgrade Guide for CuPy v2](#) for details.

10.3 Chainer v2

See *Upgrade Guide from v1 to v2* for the changes introduced in Chainer v2.

10.3.1 Upgrade Guide from v1 to v2

This document provides detailed information of differences between Chainer v1 and v2. You will know by reading it which part of your code is required (or recommended) to be fixed when you upgrade Chainer from v1 to v2.

- *CuPy*
 - *CuPy has been separated from Chainer into a separate package*
- *Global configurations*
 - *Training mode is configured by a thread-local flag*
 - *Configurations are added and replace some of existing global flags*
- *Variable*
 - *Volatile flag is removed*
 - *Variable is not a part of a computational graph anymore*
 - *Parameter has to be an instance of Parameter class*
 - *Small changes to Variable*
- *Function*
 - *The force_tuple option of split_axis is True by default*
 - *Type check APIs are updated to enable lazy building of the error messages*
 - *Methods to release unneeded arrays are added*
- *Link/Chain/ChainList*
 - *wscale option is removed from links*
 - *bias option is removed from links*
 - *The bias vector is enabled by default in N-dimensional convolution links*
 - *init_weight function is removed*
 - *The order of arguments of GRU is changed*
 - *The default value of the forget bias for LSTM and StatelessLSTM is changed to 1*
 - *The interfaces of GRU and LSTM are aligned*
 - *Aliases of links in chainer.functions are removed*
 - *Parameter link is removed*
 - *New-style parameter registration APIs are added to Link*
 - *New-style child link registration APIs are added to Chain*
 - *The input-size placeholder of links are made optional*
- *Optimizer*

- *Deprecated methods of Optimizer are removed*
- *GradientMethod uses Link.cleargrads instead of Link.zerograde by default*
- *GradientMethod is redesigned to allow parameter-specific update rules*
- *Serializer*
 - *None is serializable*
- *Trainer and Extension*
 - *Updater and Evaluator pass raw data arrays to the loss function*
 - *trigger option is removed from snapshot and snapshot_object*
 - *Extension.invoke_before_training is removed*
 - *The dump_graph extension dumps the valid graph only at its first invocation*
- *Reporter*
 - *When a variable is reported, the variable is copied with the graph purged*
- *Other utilities*
 - *Some obsolete classes and functions are removed*

CuPy

CuPy has been separated from Chainer into a separate package

CuPy, which was originally a part of Chainer, has been separated into a different Python package since Chainer v2. It changes the way to set up Chainer with CUDA support. In particular, you have to separately install `cupy` package to enable CUDA support. See [Installation](#) for the recommended installation steps.

Fortunately, there is no need of updating your source code to catch up with this change.

Global configurations

Training mode is configured by a thread-local flag

In Chainer v2, the concept of *training mode* is added. It is represented by a thread-local flag `chainer.config.train`, which is a part of *the unified configuration*. When `chainer.config.train` is `True`, functions of Chainer run in the training mode, and otherwise they run in the test mode. For example, `BatchNormalization` and `dropout()` behave differently in each mode.

In Chainer v1, such a behavior was configured by the `train` or `test` argument of each function. **This train/test argument has been removed in Chainer v2.** If your code is using the `train` or `test` argument, you have to update it. In most cases, what you have to do is just removing the `train / test` argument from any function calls.

Example

Consider the following model definition and the code to call it in test mode written for Chainer v1.

```
# Chainer v1
import chainer.functions as F
```

(continues on next page)

(continued from previous page)

```

class MyModel(chainer.Link):
    ...

    def __call__(self, x, train=True):
        return f(F.dropout(x, train=train))

m = MyModel(...)
y = m(x, train=False)

```

In Chainer v2, it should be updated into the following code:

```

# Chainer v2
import chainer.functions as F

class MyModel(chainer.Link):
    ...

    def __call__(self, x):
        return f(F.dropout(x))

m = MyModel(...)
with chainer.config.using_config('train', False):
    y = m(x)

```

Configurations are added and replace some of existing global flags

There are many global settings moved to *the unified configuration* other than the training mode. Following is the complete list of the configuration entries that have corresponding features in Chainer v1.

chainer.config.cudnn_deterministic It is corresponding to the `deterministic` argument of some convolution functions in Chainer v1. **This argument has been removed since Chainer v2.** If you are using this argument, you have to use the `chainer.config.cudnn_deterministic` flag to change the behavior of the convolution functions.

chainer.config.debug It is corresponding to the debug mode in Chainer v1, which was configured by `set_debug()` and extracted by `is_debug()`. These functions are also available in Chainer v2, so you basically do not need to update the code related to the debug mode.

chainer.config.enable_backprop It is corresponding to the *backprop mode* in Chainer v1. The functions `no_backprop_mode()` and `force_backprop_mode()` are still available in Chainer v2, which automatically turns on/off the `enable_backprop` flag. One important difference from Chainer v1 is that **the volatile flag is removed from *Variable***. Therefore, there are more situations that you need to modify the `enable_backprop` flag.

chainer.config.keep_graph_on_report This flag configures whether or not to keep the computational graph alive for a reported variable. In Chainer v2, when a *Variable* object is reported by `report()`, a copy of the variable isolated from the computational graph is created and stored by default. Setting `True` to this flag, you can change this behavior and then the original *Variable* object is stored as is. See *When a variable is reported, the variable is copied with the graph purged* for the details.

chainer.config.train It is corresponding to the `train` or `test` argument of some functions in Chainer v1. **This argument has been removed since Chainer v2.** If you are using this argument, you have to use the `chainer.config.train` flag instead. See *Training mode is configured by a thread-local flag* for more details.

chainer.config.type_check It is corresponding to the `Function.type_check_enable` flag. If your code touches this flag, **you have to use `chainer.config.type_check` instead**. Note that the environment variable `CHAINER_TYPE_CHECK` is still available in Chainer v2, so if you are only using the environment variable, there is no need of updating your code.

chainer.config.use_cudnn It is corresponding to the `use_cudnn` argument of many functions that have cuDNN implementations. **This argument has been removed since Chainer v2**. If you are using this argument, you have to use the `chainer.config.use_cudnn` flag instead. *Note that this flag is ternary, not binary.* See [Configuring Chainer](#) for more details.

These configurations can be modified in two ways.

- Simply substituting a new value to an entry, like `chainer.config.train = False`.
- Using the `chainer.using_config` context manager. It can be used with the `with` statement of Python as follows:

```
with chainer.using_config('train', False):
    do something # this code runs with chainer.config.train == False
```

It recovers the original configuration after quitting the `with` block.

The `chainer.config` manages *the thread-local configuration*. You can also set the global configuration by modifying `chainer.global_config`. Note that the global configuration is used only if the entry of the thread-local configuration is not explicitly set up.

Variable

Volatile flag is removed

The `Variable.volatile` flag has been removed since Chainer v2.

Instead, the configuration `chainer.config.enable_backprop` can be used to enable/disable the automatic differentiation feature. If it is `True`, Chainer always creates a computational graph on the forward propagation, which corresponds to passing non-volatile variables in Chainer v1. Otherwise, Chainer does not create a graph, which corresponds to passing volatile variables in Chainer v1. The biggest difference is that `enable_backprop` is a thread-local flag, whereas `volatile` was a flag local to each [Variable](#) object. Note that `enable_backprop` flag has already existed in Chainer v1, which took effect only if all the inputs to the function have `volatile == 'auto'`.

The `chainer.config.enable_backprop` flag can be modified directly or by using `using_config()`. See [Configuring Chainer](#) for details. There is also a convenience function, `no_backprop_mode()`, to turn off the flag.

If you are using the `Variable.volatile` flag, you have to stop setting this flag (it will not take effect), and set the `enable_backprop` flag instead.

Example

Let `model` be your model, and consider the following code that calls it in volatile mode.

```
# Chainer v1
x_data = ... # ndarray
x = chainer.Variable(x_data, volatile=True)
y = model(x)
```

In Chainer v2, it should be updated as follows.

```
# Chainer v2
x_data = ... # ndarray
x = chainer.Variable(x_data)
with chainer.no_backprop_mode():
    y = model(x)
```

Variable is not a part of a computational graph anymore

The `Variable` class has been separated into two distinct classes, the `Variable` class and the `VariableNode` class, since Chainer v2. Every class: `Variable` object owns its own `VariableNode` object. A computational graph consists of `Function` objects and `VariableNode` objects. When one applies a `Function` to a `Variable`, the `VariableNode` object of the variable is extracted and set to one of the inputs of the function.

Note that the underlying data array of the variable is still held by the `Variable` object. It allows each `Function` implementation to release unneeded arrays from the computational graph, resulting in greatly reduced memory consumption.

This change does not affect most users' code. If you are directly traversing the computational graph by yourself or modifying the graph ad-hoc, you may have to update your code. In most cases, it is enough to just change `Variable` into `VariableNode` in the code traversing the computational graph.

Parameter has to be an instance of Parameter class

Chainer v2 has a subclass of `Variable` called `Parameter`. This class has an interface convenient on setting up a parameter variable registered to `Link`.

You basically do not need to update your code because `Link.add_param()` creates a `Parameter` object in Chainer v2. There is a new recommended way of registering parameters to a link in Chainer v2, though. [See here](#) for the recommended way of parameter registration.

Small changes to Variable

There are some changes on the interface and specification of methods.

- `len(variable)` returns the length of the first axis of the underlying array in Chainer v2. This is equivalent to `len(variable.data)`. It is different from the behavior of Chainer v1, in which `len` returned the total number of elements in the underlying array.
- `repr(variable)` returns a NumPy-like text representation of the underlying array in Chainer v2. In Chainer v1, it just returns a string that shows the name of the variable.

Function

The `force_tuple` option of `split_axis` is `True` by default

In Chainer v2, the `force_tuple` argument of `functions.split_axis()` is set to `True` by default. Therefore, it always returns a tuple regardless of the number of sections made after the split. It was `False` by default in Chainer v1.

Type check APIs are updated to enable lazy building of the error messages

In Chainer v2, the type check APIs are updated so that the overhead of checking types is greatly reduced. In order to achieve the overhead reduction, some APIs are changed.

If you have custom Function implementations that do type checking, you have to update your code. The following list shows which part has to be updated.

- Use `utils.type_check.eval()` instead of `Expr.eval`.
- Use `utils.type_check.make_variable()` to create a `utils.type_check.Variable` object instead of directly constructing it by yourself.
- Stop using `.name` attribute of any expression.

Background of this change: In Chainer v1, the type checking APIs build an abstract syntax tree (AST) based on each expression that tests some condition. The AST is used to emit a kind error message. However, building an AST requires constructions of many Python objects, which adds large Python overheads. In Chainer v2, the `Function.type_check_forward()` method is called once or twice. At the first call, the type checking APIs run in *light-weight mode*, where it does not build an AST and just checks the condition. The second call is made only if there is a test that fails, where it builds an AST. This change makes the ordinary path of running the type checking much faster, while keeping the kind error messages.

Methods to release unneeded arrays are added

As is written above, Chainer v2 introduced a new mechanism to reduce the memory consumption of each *Function* implementation. In many cases, a *Function* implementation does not need some input arrays in its backward computation. A new method called `Function.retain_inputs()` can be used to specify which input arrays are actually needed. This method must not be called from the outside of `Function.forward()`.

Example

For example, consider the following simple addition function.

```
class AddFunction(chainer.Function):
    def forward(self, inputs):
        return inputs[0] + inputs[1],

    def backward(self, inputs, grad_outputs):
        return grad_outputs[0], grad_outputs[0]
```

It can be seen that the backward computation of this function does not use any of the inputs. Then, specifying an empty tuple of indexes to `retain_inputs()` will reduce the memory overhead.

```
class AddFunction(chainer.Function):
    def forward(self, inputs):
        self.retain_inputs(()) # does not retain both inputs
        return inputs[0] + inputs[1],

    def backward(self, inputs, grad_outputs):
        return grad_outputs[0], grad_outputs[0]
```

In some cases, the function can (or have to) use the output arrays instead of the inputs in its backward computation. In Chainer v1, we have written code that store the output arrays to attributes of the *Function* object and reuse them in the `backward()` method. In Chainer v2, it is recommended to use `Function.retain_outputs()`

to declare which outputs are required in the backward computation. The retained output arrays can be accessed via `Function.output_data`.

Note: The existing `Function` implementations that store the output arrays to its attributes will run correctly in Chainer v2. There is no any memory overhead right now. It is recommended to use `retain_outputs()`, though, so that we can incorporate more memory optimization in the future.

Example

For example, consider the following simple implementation of the tanh function.

```
class TanhFunction(chainer.Function):
    def forward(self, inputs):
        xp = chainer.cuda.get_array_module(inputs[0])
        self.y = xp.tanh(inputs[0])
        return self.y,

    def backward(self, inputs, grad_outputs):
        one = self.y.dtype.type(1) # avoid type promotion
        return grad_outputs[0] * (one - self.y * self.y),
```

We can use `retain_outputs()` instead of preserving the output array by ourselves as follows.

```
class TanhFunction(chainer.Function):
    def forward(self, inputs):
        self.retain_outputs((0,))
        xp = chainer.cuda.get_array_module(inputs[0])
        return xp.tanh(inputs[0]),

    def backward(self, inputs, grad_outputs):
        y = self.output_data[0]
        one = y.dtype.type(1) # avoid type promotion
        return grad_outputs[0] * (one - y * y)
```

Link/Chain/ChainList

wscale option is removed from links

The `wscale` option has been removed from links since Chainer v2. **If you are using `wscale` option, you have to update your code.** The recommended way is to explicitly set the initializer.

Example

Consider the case of adding a `Linear` link with the weight initialized by 0.5x of the default initialization.

```
# Chainer v1
linear = chainer.links.Linear(10, 5, wscale=0.5)
```

Note that the default initializer of the weight matrix of `Linear` is a normal distribution of the standard deviation $1/\sqrt{fanin}$. Therefore, it can be fixed as follows.

```
# Chainer v2
linear = chainer.links.Linear(10, 5, initialW=chainer.initializers.Normal(0.5 / math.
↳sqrt(10)))
```

Or, by using the fact that `initializers.HeNormal` provides the initialization with a normal distribution of the standard deviation $scale * \sqrt{2/fanin}$, the following code is also equivalent to the original.

```
# Chainer v2, using HeNormal
linear = chainer.links.Linear(10, 5, initialW=chainer.initializers.HeNormal(0.5 /
↳math.sqrt(2)))
```

bias option is removed from links

In Chainer v2, the bias option is removed from the following links: `Linear`, `Convolution2D`, `Deconvolution2D`, and `DilatedConvolution2D`. The effect of this argument was duplicated with the `initial_bias` option. Use `initial_bias` instead.

The bias vector is enabled by default in N-dimensional convolution links

In Chainer v2, the bias parameter is enabled by default in `ConvolutionND` and `DeconvolutionND`. It was unintentionally disabled by default in Chainer v1.

If you are using `ConvolutionND` or `DeconvolutionND` without specifying the `initial_bias` argument, you have to fix your code. If you want to keep the old behavior (i.e., no bias vector is created by the link), pass `nobias=True` to the link at the construction. Otherwise it will automatically create a bias vector.

init_weight function is removed

The `chainer.initializers.init_weight` function that was used on weight initialization has been removed since Chainer v2.

You have to update your code if you are using `init_weight`. In most cases, the update is simple: pass an initializer to `Parameter`.

Example

Consider the following code that initializes a weight matrix randomly and a bias vector by zero.

```
# Chainer v1
class MyLink(chainer.Link):
    def __init__(self):
        super(MyLink, self).__init__(
            W=(10, 5),
            b=(5, ),
        )
        chainer.initializers.init_weight(self.W, chainer.initializers.Normal(0.05))
        self.b.data.fill(0)
    ...
```

This code should be fixed as follows (see the next topic for the use of `Parameter`).

```
# Chainer v2
class MyLink(chainer.Link):
    def __init__(self):
        super(MyLink, self).__init__()
        self.W = chainer.Parameter(chainer.initializers.Normal(0.05), (10, 5))
        self.b = chainer.Parameter(0, (5,))
    ...
```

The order of arguments of GRU is changed

In Chainer v2, the first two arguments of `GRU` is the input size and the output size. It was reversed in Chainer v1, causing an inconsistent interface compared to other links including `LSTM`. **If you are using `GRU`, you have to update your code.** The update is done by simply flipping the first two arguments.

Example

Consider the following code that creates a `GRU` link.

```
# Chainer v1
gru = chainer.links.GRU(20, 10)
```

It should be fixed into the following code.

```
# Chainer v2
gru = chainer.links.GRU(10, 20)
```

Note that if you were omitting the output size, the code works as is because `GRU` supports *the omitted input size*.

```
# Chainer v1/v2
gru = chainer.links.GRU(20)
```

The default value of the forget bias for LSTM and StatelessLSTM is changed to 1

In Chainer v2, the default forget bias value of `LSTM` and `StatelessLSTM` links is changed to 1. This change is based on [the paper reporting that using a large forget bias improves the training performance](#). The new behavior is also consistent with [the implementation of BasicLSTMCell in TensorFlow](#).

It will improve the most use cases of LSTMs, although this change would break the reproducibility of the existing experiments. **If you want to keep the same initialization procedure, you have to update your code.** The change is simple: pass `forget_bias_init=0` to `LSTM` and `StatelessLSTM`.

The interfaces of GRU and LSTM are aligned

In Chainer v1, `GRU` was *stateless*, as opposed to the current implementation. To align with the naming convention of LSTM links, we have changed the naming convention from Chainer v2 so that the shorthand name points the stateful links. **If you are using `StatelessGRU` for stateless version, whose implementation is identical to `chainer.links.GRU` in v1.**

Aliases of links in `chainer.functions` are removed

For the compatibility reason, there were some links that have aliases in the `chainer.functions` module. These aliases are removed in Chainer v2. Use `chainer.links` instead.

Parameter link is removed

The `chainer.links.Parameter` link is removed in Chainer v2. This link existed in Chainer v1 only for the backward compatibility. Use `chainer.Parameter` instead (for the new `Parameter` class, see *Parameter has to be an instance of Parameter class*).

New-style parameter registration APIs are added to Link

In Chainer v2, `Link.init_scope()` method returns a context manager that automatically registers a `Parameter` object to the link at setting it to an attribute. If you are using IDE like PyCharm, it is recommended to use this new-style parameter registration so that IDEs can easily detect the existence of the parameter as an attribute. It is also a good practice to use the new-style API even if you are not using IDEs, if you are planning to make the code public.

Note: The existing code that uses the conventional way of registering parameters are still valid.

Example

For example, the following link initialization code

```
# Chainer v1
class MyLink(chainer.Link):
    def __init__(self):
        super(MyLink, self).__init__(
            W=(10, 5),
            b=(5,),
        )
        chainer.initializers.Normal(0.05)(self.W.data)
        self.b.data.fill(0)
    ...
```

is recommended to be updated as follows.

```
# Chainer v2
class MyLink(chainer.Link):
    def __init__(self):
        super(MyLink, self).__init__()
        with self.init_scope():
            self.W = chainer.Parameter(chainer.initializers.Normal(0.05), (10, 5))
            self.b = chainer.Parameter(0, (5,)) # initialize by zero
    ...
```

Note: To keep a `Parameter` object as an attribute without registration, you can set the attribute without using the `self.init_scope():` block.

New-style child link registration APIs are added to Chain

Like *Parameter*, a *Link* object is also automatically registered to a *Chain* object by substitution to an attribute within a *init_scope()* scope. If you are using IDE like PyCharm, it is recommended to use the new-style child link registration so that IDEs can easily detect the existence of the child link as an attribute. It is also a good practice to use the new-style API even if you are not using IDEs, if you are planning to make the code public.

Note: The existing code that uses the conventional way of registering child links are still valid.

Example

For example, the following chain initialization code

```
# Chainer v1
class MyMLP(chainer.Chain):
    def __init__(self):
        super(MyMLP, self).__init__(
            layer1=L.Linear(None, 20),
            layer2=L.Linear(None, 30),
        )
    ...
```

is recommended to be updated as follows.

```
# Chainer v2
class MyMLP(chainer.Chain):
    def __init__(self):
        super(MyMLP, self).__init__()
        with self.init_scope():
            self.layer1 = L.Linear(20)
            self.layer2 = L.Linear(30)
```

Note that this example also demonstrates the use of new APIs with *the omitted input size*, explained below.

Note: To keep a *Link* object as an attribute without registration, you can set the attribute without using the `with self.init_scope():` block.

The input-size placeholder of links are made optional

In Chainer v2, the input size of many links, including *Linear* and *Convolution2D*, is made optional. In Chainer v1, we had to use `None` as the placeholder to specify that the input size should be determined at the first iteration. The placeholder can also be used in Chainer v2, although it is easier to just omit the input size.

See the previous item for the example of omitting the input size of *Linear*. The following links currently support the omitted input size.

- *Convolution2D*
- *Deconvolution2D*
- *DilatedConvolution2D*
- *Linear*

- `LSTM`
- `MLPConvolution2D`
- `StatelessLSTM`

Optimizer

Deprecated methods of Optimizer are removed

The following methods are removed from `Optimizer`. These methods have been already deprecated in the past versions. **If you are using these methods, you have to update your code.**

- `zero_grads`: use `Link.zerograds()` instead.
- `compute_grads_norm`: you can compute the gradient norm by iterating the list of parameters by `Link.params()`.
- `clip_grads`: use `GradientClipping` instead.
- `weight_decay`: use `WeightDecay` instead.
- `accumulate_grads`: use `Link.addgrads()` instead.

GradientMethod uses Link.cleargrads instead of Link.zerograds by default

In Chainer v2, `GradientMethod` clears the gradient before running backprop by `Link.cleargrads()`. It means that the gradient of each parameter is initialized by `None` instead of a zero array. Note that all the optimizer implementations provided by Chainer are subclasses of `GradientMethod`, and therefore this change affects all of them.

In most cases, you do not need to update your code. If your code relies on the zeroing initialization, you have to fix your code to explicitly initialize the gradient by zero, or to pass `False` to `GradientMethod.use_cleargrads()`.

GradientMethod is redesigned to allow parameter-specific update rules

In Chainer v2, the new class `UpdateRule` is used to define an update rule specific to each `Parameter` object. The `UpdateRule` is set to each `Parameter` object, and is used at each update step. This object implements an *update formula* using the data and gradient arrays.

Each `UpdateRule` object has `enabled` flag, which configures if the update rule should be applied to that parameter on update. By setting the flag to `False`, you can *freeze* the parameter. There is also a convenient method `Link.enable_update()` and `Link.disable_update()`, which configure the flag of each parameter under the link hierarchy. In other frameworks, a similar feature is called *layer freezing*. In Chainer v2, this is officially supported by these methods.

Each `UpdateRule` object can also hold its own hook functions similar to `Optimizer`. The built-in hook functions except for `GradientClipping` can also be used as a hook function of `UpdateRule`.

In most cases, you do not have to update your code because each optimizer automatically sets up an appropriate `UpdateRule` object to each parameter.

If you are using a custom gradient-based optimizer implementation, you need to update the implementation. The following list shows what you have to do.

- Write a subclass of `UpdateRule` that implements the update rule.

- Rewrite your `GradientMethod` implementation. The new implementation only has to set up the update rule for each parameter in the target link.

You can see live examples in [the optimizer implementations](#) provided by Chainer.

Serializer

None is serializable

In Chainer v2, all serializers start supporting `None` value to be serialized and deserialized. Users' code can rely on this feature, i.e., it can serialize and deserialize `None` value with any given serializer. This change only affects your code if it provides its own serializer implementations.

Trainer and Extension

Updater and Evaluator pass raw data arrays to the loss function

In Chainer v2, `Updater` and `Evaluator` pass raw data arrays to the loss function without wrapping them with `Variable`. You might need to update your code so that the loss function (in most cases, the model's `__call__`) accepts raw arrays.

Note that raw arrays can be directly passed to any `Function`; they are automatically wrapped by `Variable`. For example, if the input is directly passed to a `Function` object (or any function under `chainer.functions`), you do not need to update the code.

Example

Consider the following code that obtains the shape of the input via `Variable.data`.

```
# Chainer v1
class MyLink(chainer.Link):
    def __call__(self, x):
        shape = x.data.shape # valid if x is Variable, invalid if x is ndarray
        ...
```

It should be updated so that the link also accepts a raw array as the input. In this case, we have `Variable.shape` which is equivalent to `data.shape`, so you can simply write as follows.

```
# Chainer v2
class MyLink(chainer.Link):
    def __call__(self, x):
        shape = x.shape # valid regardless of x being Variable or ndarray
        ...
```

trigger option is removed from snapshot and snapshot_object

In Chainer v2, the `trigger` option is removed from the `snapshot()` and `snapshot_object()` extensions. The effect of the option was duplicated with the `trigger` option of `Trainer.extend`. **If you are passing the `trigger` argument to these extensions, you have to update your code.** The update can be done by passing the value to the corresponding `Trainer.extend`.

Example

Assume that `trainer` is an instance of `Trainer`, and consider that you were adding a `snapshot()` extension as follows.

```
# Chainer v1
trainer.extend(chainer.training.extensions.snapshot(trigger=(1000, 'iteration')))
```

It should be updated as follows (note that this code also works with Chainer v1).

```
# Chainer v1/v2
trainer.extend(chainer.training.extensions.snapshot(), trigger=(1000, 'iteration'))
```

Extension.invoke_before_training is removed

In Chainer v2, The attribute `invoke_before_training` of `Extension` is removed. Instead, the `Extension.initialize` method is added. This method is called by `Trainer.run` before entering the training loop.

In Chainer v1, the extension is just called before entering the training loop when `invoke_before_training` is `True`. **If you have a custom extension that has `invoke_before_training=True`, you have to update the code.** What you have to do is to remove the `invoke_before_training` flag and override `initialize()` method. If you are using the `make_extension()` decorator, you can set the `initialize` function by passing the `initializer` argument to `make_extension()`.

The dump_graph extension dumps the valid graph only at its first invocation

In Chainer v2, the `dump_graph()` extension dumps the valid computational graph only at its first invocation. **If you want to dump the graph more than once, you have to fix the code.** The easiest fix is setting the `chainer.config.keep_graph_on_report` flag to `True`. *Note that this fix will cancel the improvement on the memory consumption made in Chainer v2.* More memory-efficient fix is to dump the graph without using an extension, e.g. by customizing the loss function or the updater.

Here is the background of this change. In Chainer v2, *the Reporter copies reported variables with purging the computational graph by default*. On the other hand, the `dump_graph()` extension requires the computational graph reachable from the reported variable. In order to make the graph available, the `dump_graph()` extension turns on the `chainer.config.keep_graph_on_report` flag at its initializer (i.e., it turns on the graph before entering the training loop). Since we also wanted to achieve the memory efficiency, the `dump_graph()` extension **turns off the flag after dumping the graph at its first invocation** (strictly speaking, it recovers the original value). As a result, the computational graph is not available from the second invocation.

Since the `dump_graph()` recovers the original flag value at its invocation, you can keep the graph dumped more than once by changing the original flag value.

Reporter

When a variable is reported, the variable is copied with the graph purged

In Chainer v2, when a `Variable` object is reported using `report()` function (or directly using `Reporter`), a copy of the variable is made without preserving the computational graph. **If your code depends on the reachability of the computational graph from the reported variable, you have to update your code.** The easiest way to

update your code is setting `chainer.config.keep_graph_on_report` to `True`, then Chainer will keep the computational graph reachable from the reported variable.

The possible examples that are affected by this change are as follows (not exhaustive).

- A custom extension that runs backprop from a reported variable. It is definitely an example of assuming the reachability of the computational graph from the reported variable.
- An extension that visualizes the computational graph from a reported variable. If you are writing such an extension by yourself, you have to turn on the `keep_graph_on_report` flag. The `dump_graph()` extension is another example, for which see *the above item* for the details.

This change is made for the memory performance reason; with this change, the memory used by the computational graph for training is immediately released before invoking extensions. Therefore, *changing the behavior by overwriting `chainer.config.keep_graph_on_report` may increase the memory consumption*. It may cause an out-of-memory error if the computational graph of the loss function consumes almost all the memory available in your environment and there is an extension that uses a certain amount of memory (e.g. *Evaluator*).

Other utilities

Some obsolete classes and functions are removed

The following classes and functions are removed in Chainer v2.

- `chainer.Flag`
- `chainer.FunctionSet` (Use *Chain* or *ChainList* instead)
- `chainer.cuda.init` (It did nothing except for calling `check_cuda_available()`)
- `chainer.cuda.empty` (Use `cupy.empty()`)
- `chainer.cuda.empty_like` (Use `cupy.empty_like()`)
- `chainer.cuda.full` (Use `cupy.full()`)
- `chainer.cuda.full_like` (Use `cupy.full_like()`)
- `chainer.cuda.ones` (Use `cupy.ones()`)
- `chainer.cuda.ones_like` (Use `cupy.ones_like()`)
- `chainer.cuda.zeros` (Use `cupy.zeros()`)
- `chainer.cuda.zeros_like` (Use `cupy.zeros_like()`)

Comparison with Other Frameworks

11.1 A table for quick comparison

This table compares Chainer with other actively developed deep learning frameworks. Content is current as of July 2017.

		Chainer	Py-Torch	Tensor-Flow	Theano	Caffe	Caffe2	MXNet	Dynet	Pad-dle	Pad-dle	DL4J	CNTK	Neon	Knet	Dark-net	ThinC
Basics	Language	Python	Python	Python	Python	Python/MAT-LAB	C++/a-JIT	C++/Python	Python/Python	Python/Python	Python/C++	Brain-Script/Python	Python/C++	Python/C++	Julia	C	Python
	Approach	define-by-run	define-by-run	symbolic auto-grad	symbolic auto-grad	static	static/manual grads	symbolic auto-grad/manual grads/define-by-run ¹	define-by-run	symbolic auto-grad	static/manual grads/symbolic auto-grad ²	static/symbolic auto-grad	static/symbolic auto-grad ³	define-by-run	static	call-back-based define-by-run	
	CPU back-end package	NumPy	TH	Eigen	NumPy	TH	mshadow	Eigen		ND4J		NumPy	Julia			NumPy	
	GPU back-end package	CuPy	THC	Eigen	libg-puar-ray		THC	mshadow	Eigen		ND4J	neon	Kne-tAr-rays			CuPy	
	Primary sponsor	Pre-ferred Net-works	Face-book	Google	MILA	Face-book	Face-book	Ama-zon/Apache	CMU	Baidu	Sky-mind	Micro-soft	In-tel Ner-vana	Koç Uni-ver-sity	Joe Red-mon	Ex-plo-sion AI	
NNs	CNNs	full	full	full	full	full	full	full	partial	full	full	full	full	partial	full	none	
	RNNs	full	full	full	full	partial	full	full	full	full	full	full	partial	partial	partial	partial	
	Reverse-mode auto-grad	Y	Y	Y			torch-autograd	Y	Y	Y		Y	ngraph	Y		with closures	
	Forward-mode auto-grad			tensor-flow-forward-ad													
	Higher order grads	Y ⁴	Y	Y	Y									Y			
	Variable-length loops	na-tive	na-tive	while_loop	RNNs	na-tive	2017	na-tive	RNNs	na-tive	dynamic axis	na-tive	na-tive	na-tive	na-tive	na-tive	
	Dif-fer-ent ar-chi-tec-tures per	na-tive	na-tive	fold			torch-autograd	MinPy	na-tive					na-tive		na-tive	
800		Chapter 11. Comparison with Other Frameworks															

11.2 Benchmarks

Benchmarks for convolutional networks can be found at [convnet-benchmarks](#) while some NLP benchmarks are at [dynet-benchmark](#). Chainer wraps the latest available cuDNN kernels for CNNs and RNNs, so performance of most common networks that use these kernels is typically similar to that of other modern frameworks. As Chainer's define-by-run approach means the user's Python code is executed directly at runtime, particularly complex networks or those with very small tensor sizes may be slower than in static-graph frameworks.

¹ Define-by-run is in development as of June 2017 and tracked in [dmlc/mxnet#5705](#). It is also possible using the much slower MinPy extension.

² Symbolic autograd is in development as of June 2017 and tracked in [deeplearning4j/nd4j#1750](#).

³ Symbolic autograd is available only with ngraph backend (experimental).

⁴ Some functions do not support higher-order differentiation. See [chainer/chainer#4449](#).

⁵ Nervana provides kernels that are meant to compete with cuDNN.

⁶ Multiprocessing provides a significant performance improvement only for frameworks that use Python at runtime.

CHAPTER 12

License

Copyright (c) 2015 Preferred Infrastructure, Inc.

Copyright (c) 2015 Preferred Networks, Inc.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the “Software”), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED “AS IS”, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

CHAPTER 13

Indices and tables

- `genindex`
- `modindex`
- `search`

Bibliography

- [LeCun98] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11), 2278–2324, 1998.
- [Simonyan14] Simonyan, K. and Zisserman, A., Very Deep Convolutional Networks for Large-Scale Image Recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [He16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, Jian Sun. Deep Residual Learning for Image Recognition. *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770-778, 2016.
- [Graves2006] Alex Graves, Santiago Fernandez, Faustino Gomez, Jurgen Schmidhuber, [Connectionist Temporal Classification: Labelling Unsegmented Sequence Data with Recurrent Neural Networks](#)
- [Graves2012] Alex Graves, [Supervised Sequence Labelling with Recurrent Neural Networks](#)

C

- `chainer`, [724](#)
- `chainer.backends.cuda`, [727](#)
- `chainer.computational_graph`, [746](#)
- `chainer.dataset`, [689](#)
- `chainer.datasets`, [696](#)
- `chainer.exporters`, [749](#)
- `chainer.function_hooks`, [262](#)
- `chainer.functions`, [133](#)
- `chainer.initializers`, [645](#)
- `chainer.iterators`, [713](#)
- `chainer.links`, [272](#)
- `chainer.links.caffe`, [749](#)
- `chainer.serializers`, [717](#)
- `chainer.testing`, [756](#)
- `chainer.training`, [653](#)
- `chainer.utils`, [760](#)

Symbols

- `__abs__()` (chainer.Parameter method), 127
- `__abs__()` (chainer.Variable method), 119
- `__add__()` (chainer.Parameter method), 127
- `__add__()` (chainer.Sequential method), 612
- `__add__()` (chainer.Variable method), 120
- `__add__()` (chainer.utils.type_check.Expr method), 752
- `__bool__()` (chainer.Parameter method), 127
- `__bool__()` (chainer.Variable method), 119
- `__bool__()` (chainer.utils.type_check.Expr method), 752
- `__call__()` (chainer.AbstractSerializer method), 725
- `__call__()` (chainer.Chain method), 595
- `__call__()` (chainer.ChainList method), 600
- `__call__()` (chainer.Deserializer method), 725
- `__call__()` (chainer.Function method), 247
- `__call__()` (chainer.FunctionAdapter method), 251
- `__call__()` (chainer.FunctionNode method), 257
- `__call__()` (chainer.Initializer method), 645
- `__call__()` (chainer.Link method), 590
- `__call__()` (chainer.Sequential method), 606
- `__call__()` (chainer.Serializer method), 724
- `__call__()` (chainer.dataset.ConcatWithAsyncTransfer method), 694
- `__call__()` (chainer.initializers.Constant method), 646
- `__call__()` (chainer.initializers.GlorotNormal method), 649
- `__call__()` (chainer.initializers.GlorotUniform method), 652
- `__call__()` (chainer.initializers.HeNormal method), 649
- `__call__()` (chainer.initializers.HeUniform method), 652
- `__call__()` (chainer.initializers.Identity method), 646
- `__call__()` (chainer.initializers.LeCunNormal method), 648
- `__call__()` (chainer.initializers.LeCunUniform method), 651
- `__call__()` (chainer.initializers.NaN method), 647
- `__call__()` (chainer.initializers.Normal method), 648
- `__call__()` (chainer.initializers.One method), 647
- `__call__()` (chainer.initializers.Orthogonal method), 650
- `__call__()` (chainer.initializers.Uniform method), 651
- `__call__()` (chainer.initializers.Zero method), 647
- `__call__()` (chainer.links.BatchNormalization method), 475
- `__call__()` (chainer.links.BatchRenormalization method), 480
- `__call__()` (chainer.links.Bias method), 273
- `__call__()` (chainer.links.Bilinear method), 278
- `__call__()` (chainer.links.BinaryHierarchicalSoftmax method), 490
- `__call__()` (chainer.links.BlackOut method), 495
- `__call__()` (chainer.links.CRF1d method), 500
- `__call__()` (chainer.links.ChildSumTreeLSTM method), 284
- `__call__()` (chainer.links.Classifier method), 531
- `__call__()` (chainer.links.Convolution2D method), 290
- `__call__()` (chainer.links.ConvolutionND method), 295
- `__call__()` (chainer.links.Deconvolution2D method), 302
- `__call__()` (chainer.links.DeconvolutionND method), 307
- `__call__()` (chainer.links.DepthwiseConvolution2D method), 312
- `__call__()` (chainer.links.DilatedConvolution2D method), 318
- `__call__()` (chainer.links.EmbedID method), 323
- `__call__()` (chainer.links.GRU method), 328
- `__call__()` (chainer.links.GoogLeNet method), 544
- `__call__()` (chainer.links.Highway method), 333
- `__call__()` (chainer.links.Inception method), 338
- `__call__()` (chainer.links.InceptionBN method), 344
- `__call__()` (chainer.links.LSTM method), 361
- `__call__()` (chainer.links.LayerNormalization method), 485
- `__call__()` (chainer.links.Linear method), 350
- `__call__()` (chainer.links.LocalConvolution2D method), 355
- `__call__()` (chainer.links.MLPConvolution2D method), 366
- `__call__()` (chainer.links.Maxout method), 520
- `__call__()` (chainer.links.NStepBiGRU method), 377
- `__call__()` (chainer.links.NStepBiLSTM method), 383

`__call__()` (chainer.links.NStepBiRNReLU method), 389

`__call__()` (chainer.links.NStepBiRNNTanh method), 395

`__call__()` (chainer.links.NStepGRU method), 401

`__call__()` (chainer.links.NStepLSTM method), 407

`__call__()` (chainer.links.NStepRNReLU method), 413

`__call__()` (chainer.links.NStepRNNTanh method), 419

`__call__()` (chainer.links.NaryTreeLSTM method), 372

`__call__()` (chainer.links.NegativeSampling method), 525

`__call__()` (chainer.links.PReLU method), 510

`__call__()` (chainer.links.Parameter method), 424

`__call__()` (chainer.links.ResNet101Layers method), 564

`__call__()` (chainer.links.ResNet152Layers method), 571

`__call__()` (chainer.links.ResNet50Layers method), 558

`__call__()` (chainer.links.Scale method), 429

`__call__()` (chainer.links.SimplifiedDropconnect method), 505

`__call__()` (chainer.links.StatefulGRU method), 435

`__call__()` (chainer.links.StatefulMGU method), 446

`__call__()` (chainer.links.StatefulPeepholeLSTM method), 456

`__call__()` (chainer.links.StatefulZoneoutLSTM method), 461

`__call__()` (chainer.links.StatelessGRU method), 441

`__call__()` (chainer.links.StatelessLSTM method), 467

`__call__()` (chainer.links.StatelessMGU method), 451

`__call__()` (chainer.links.Swish method), 515

`__call__()` (chainer.links.TheanoFunction method), 578

`__call__()` (chainer.links.VGG16Layers method), 536

`__call__()` (chainer.links.caffe.CaffeFunction method), 583

`__call__()` (chainer.links.model.vision.resnet.ResNetLayers method), 551

`__call__()` (chainer.optimizer_hooks.GradientClipping method), 641

`__call__()` (chainer.optimizer_hooks.GradientHardClipping method), 642

`__call__()` (chainer.optimizer_hooks.GradientLARS method), 644

`__call__()` (chainer.optimizer_hooks.GradientNoise method), 643

`__call__()` (chainer.optimizer_hooks.Lasso method), 641

`__call__()` (chainer.optimizer_hooks.WeightDecay method), 640

`__call__()` (chainer.serializers.DictionarySerializer method), 718

`__call__()` (chainer.serializers.HDF5Deserializer method), 722

`__call__()` (chainer.serializers.HDF5Serializer method), 721

`__call__()` (chainer.serializers.NpzDeserializer method), 719

`__call__()` (chainer.training.Extension method), 663

`__call__()` (chainer.training.extensions.Evaluator method), 665

`__call__()` (chainer.training.extensions.ExponentialShift method), 673

`__call__()` (chainer.training.extensions.FailOnNonNumber method), 668

`__call__()` (chainer.training.extensions.LinearShift method), 674

`__call__()` (chainer.training.extensions.LogReport method), 678

`__call__()` (chainer.training.extensions.MicroAverage method), 667

`__call__()` (chainer.training.extensions.ParameterStatistics method), 671

`__call__()` (chainer.training.extensions.PlotReport method), 680

`__call__()` (chainer.training.extensions.PrintReport method), 675

`__call__()` (chainer.training.extensions.ProgressBar method), 677

`__call__()` (chainer.training.extensions.VariableStatisticsPlot method), 682

`__call__()` (chainer.training.triggers.BestValueTrigger method), 686

`__call__()` (chainer.training.triggers.EarlyStoppingTrigger method), 686

`__call__()` (chainer.training.triggers.IntervalTrigger method), 687

`__call__()` (chainer.training.triggers.ManualScheduleTrigger method), 688

`__call__()` (chainer.training.triggers.MaxValueTrigger method), 688

`__call__()` (chainer.training.triggers.MinValueTrigger method), 689

`__call__()` (chainer.training.triggers.TimeTrigger method), 689

`__call__()` (chainer.utils.type_check.Expr method), 752

`__copy__()` (chainer.Parameter method), 124

`__copy__()` (chainer.Variable method), 117

`__copy__()` (chainer.iterators.MultiprocessIterator method), 715

`__div__()` (chainer.Parameter method), 128

`__div__()` (chainer.Variable method), 120

`__enter__()` (chainer.DebugMode method), 745

`__enter__()` (chainer.FunctionHook method), 271

`__enter__()` (chainer.Reporter method), 734

`__enter__()` (chainer.dataset.Iterator method), 692

`__enter__()` (chainer.function_hooks.CUDAProfileHook method), 263

`__enter__()` (chainer.function_hooks.CupyMemoryProfileHook method), 264

`__enter__()` (chainer.function_hooks.PrintHook method), 266

`__enter__()` (chainer.function_hooks.TimerHook method), 268

[__enter__\(\) \(chainer.iterators.MultiprocessIterator method\), 715](#)
[__enter__\(\) \(chainer.iterators.MultithreadIterator method\), 716](#)
[__enter__\(\) \(chainer.iterators.SerialIterator method\), 714](#)
[__eq__\(\) \(chainer.Parameter method\), 127](#)
[__eq__\(\) \(chainer.Variable method\), 119](#)
[__eq__\(\) \(chainer.utils.type_check.Expr method\), 752](#)
[__exit__\(\) \(chainer.DebugMode method\), 745](#)
[__exit__\(\) \(chainer.FunctionHook method\), 271](#)
[__exit__\(\) \(chainer.Reporter method\), 734](#)
[__exit__\(\) \(chainer.dataset.Iterator method\), 692](#)
[__exit__\(\) \(chainer.function_hooks.CUDAProfileHook method\), 263](#)
[__exit__\(\) \(chainer.function_hooks.CupyMemoryProfileHook method\), 264](#)
[__exit__\(\) \(chainer.function_hooks.PrintHook method\), 266](#)
[__exit__\(\) \(chainer.function_hooks.TimerHook method\), 268](#)
[__exit__\(\) \(chainer.iterators.MultiprocessIterator method\), 715](#)
[__exit__\(\) \(chainer.iterators.MultithreadIterator method\), 716](#)
[__exit__\(\) \(chainer.iterators.SerialIterator method\), 714](#)
[__floordiv__\(\) \(chainer.Parameter method\), 128](#)
[__floordiv__\(\) \(chainer.Variable method\), 121](#)
[__floordiv__\(\) \(chainer.utils.type_check.Expr method\), 752](#)
[__ge__\(\) \(chainer.Parameter method\), 127](#)
[__ge__\(\) \(chainer.Variable method\), 119](#)
[__ge__\(\) \(chainer.utils.type_check.Expr method\), 752](#)
[__getitem__\(\) \(chainer.AbstractSerializer method\), 725](#)
[__getitem__\(\) \(chainer.Chain method\), 595](#)
[__getitem__\(\) \(chainer.ChainList method\), 600](#)
[__getitem__\(\) \(chainer.Deserializer method\), 726](#)
[__getitem__\(\) \(chainer.Parameter method\), 124](#)
[__getitem__\(\) \(chainer.Sequential method\), 607](#)
[__getitem__\(\) \(chainer.Serializer method\), 724](#)
[__getitem__\(\) \(chainer.Variable method\), 116](#)
[__getitem__\(\) \(chainer.dataset.DatasetMixin method\), 690](#)
[__getitem__\(\) \(chainer.datasets.ConcatenatedDataset method\), 698](#)
[__getitem__\(\) \(chainer.datasets.DictDataset method\), 697](#)
[__getitem__\(\) \(chainer.datasets.ImageDataset method\), 704](#)
[__getitem__\(\) \(chainer.datasets.LabeledImageDataset method\), 708](#)
[__getitem__\(\) \(chainer.datasets.MultiZippedImageDataset method\), 707](#)
[__getitem__\(\) \(chainer.datasets.SubDataset method\), 700](#)
[__getitem__\(\) \(chainer.datasets.TransformDataset method\), 703](#)
[__getitem__\(\) \(chainer.datasets.TupleDataset method\), 697](#)
[__getitem__\(\) \(chainer.datasets.ZippedImageDataset method\), 706](#)
[__getitem__\(\) \(chainer.links.ChildSumTreeLSTM method\), 284](#)
[__getitem__\(\) \(chainer.links.Classifier method\), 531](#)
[__getitem__\(\) \(chainer.links.GRU method\), 328](#)
[__getitem__\(\) \(chainer.links.GoogLeNet method\), 544](#)
[__getitem__\(\) \(chainer.links.Highway method\), 333](#)
[__getitem__\(\) \(chainer.links.Inception method\), 338](#)
[__getitem__\(\) \(chainer.links.InceptionBN method\), 344](#)
[__getitem__\(\) \(chainer.links.LSTM method\), 361](#)
[__getitem__\(\) \(chainer.links.MLPConvolution2D method\), 366](#)
[__getitem__\(\) \(chainer.links.Maxout method\), 520](#)
[__getitem__\(\) \(chainer.links.NStepBiGRU method\), 377](#)
[__getitem__\(\) \(chainer.links.NStepBiLSTM method\), 383](#)
[__getitem__\(\) \(chainer.links.NStepBiRNNReLU method\), 389](#)
[__getitem__\(\) \(chainer.links.NStepBiRNNTanh method\), 395](#)
[__getitem__\(\) \(chainer.links.NStepGRU method\), 401](#)
[__getitem__\(\) \(chainer.links.NStepLSTM method\), 407](#)
[__getitem__\(\) \(chainer.links.NStepRNNReLU method\), 413](#)
[__getitem__\(\) \(chainer.links.NStepRNNTanh method\), 419](#)
[__getitem__\(\) \(chainer.links.NaryTreeLSTM method\), 372](#)
[__getitem__\(\) \(chainer.links.ResNet101Layers method\), 564](#)
[__getitem__\(\) \(chainer.links.ResNet152Layers method\), 571](#)
[__getitem__\(\) \(chainer.links.ResNet50Layers method\), 558](#)
[__getitem__\(\) \(chainer.links.Scale method\), 429](#)
[__getitem__\(\) \(chainer.links.StatefulGRU method\), 435](#)
[__getitem__\(\) \(chainer.links.StatefulMGU method\), 446](#)
[__getitem__\(\) \(chainer.links.StatefulPeepholeLSTM method\), 456](#)
[__getitem__\(\) \(chainer.links.StatefulZoneoutLSTM method\), 461](#)
[__getitem__\(\) \(chainer.links.StatelessGRU method\), 441](#)
[__getitem__\(\) \(chainer.links.StatelessLSTM method\), 468](#)
[__getitem__\(\) \(chainer.links.StatelessMGU method\), 451](#)
[__getitem__\(\) \(chainer.links.VGG16Layers method\), 537](#)
[__getitem__\(\) \(chainer.links.caffe.CaffeFunction method\), 584](#)
[__getitem__\(\) \(chainer.links.model.vision.resnet.ResNetLayers method\), 551](#)
[__getitem__\(\) \(chainer.serializers.DictionarySerializer](#)

method), 719
__getitem__() (chainer.serializers.HDF5Deserializer method), 722
__getitem__() (chainer.serializers.HDF5Serializer method), 721
__getitem__() (chainer.serializers.NpzDeserializer method), 720
__getitem__() (chainer.utils.type_check.Expr method), 752
__gt__() (chainer.Parameter method), 127
__gt__() (chainer.Variable method), 119
__gt__() (chainer.utils.type_check.Expr method), 752
__iter__() (chainer.ChainList method), 600
__iter__() (chainer.Sequential method), 607
__iter__() (chainer.dataset.Iterator method), 692
__iter__() (chainer.iterators.MultiprocessIterator method), 715
__iter__() (chainer.iterators.MultithreadIterator method), 717
__iter__() (chainer.iterators.SerialIterator method), 714
__iter__() (chainer.links.MLPConvolution2D method), 366
__iter__() (chainer.links.NStepBiGRU method), 378
__iter__() (chainer.links.NStepBiLSTM method), 384
__iter__() (chainer.links.NStepBiRNReLU method), 389
__iter__() (chainer.links.NStepBiRNNTanh method), 395
__iter__() (chainer.links.NStepGRU method), 401
__iter__() (chainer.links.NStepLSTM method), 407
__iter__() (chainer.links.NStepRNReLU method), 413
__iter__() (chainer.links.NStepRNNTanh method), 419
__le__() (chainer.Parameter method), 127
__le__() (chainer.Variable method), 119
__le__() (chainer.utils.type_check.Expr method), 752
__len__() (chainer.ChainList method), 600
__len__() (chainer.Parameter method), 124
__len__() (chainer.Sequential method), 607
__len__() (chainer.Variable method), 117
__len__() (chainer.dataset.DatasetMixin method), 691
__len__() (chainer.datasets.ConcatenatedDataset method), 698
__len__() (chainer.datasets.DictDataset method), 697
__len__() (chainer.datasets.ImageDataset method), 705
__len__() (chainer.datasets.LabeledImageDataset method), 709
__len__() (chainer.datasets.MultiZippedImageDataset method), 707
__len__() (chainer.datasets.SubDataset method), 700
__len__() (chainer.datasets.TransformDataset method), 703
__len__() (chainer.datasets.TupleDataset method), 697
__len__() (chainer.datasets.ZippedImageDataset method), 706
__len__() (chainer.links.MLPConvolution2D method), 366
__len__() (chainer.links.NStepBiGRU method), 378
__len__() (chainer.links.NStepBiLSTM method), 383
__len__() (chainer.links.NStepBiRNReLU method), 389
__len__() (chainer.links.NStepBiRNNTanh method), 395
__len__() (chainer.links.NStepGRU method), 401
__len__() (chainer.links.NStepLSTM method), 407
__len__() (chainer.links.NStepRNReLU method), 413
__len__() (chainer.links.NStepRNNTanh method), 419
__lt__() (chainer.Parameter method), 127
__lt__() (chainer.Variable method), 119
__lt__() (chainer.utils.type_check.Expr method), 752
__matmul__() (chainer.Parameter method), 129
__matmul__() (chainer.Variable method), 121
__mul__() (chainer.Parameter method), 128
__mul__() (chainer.Variable method), 120
__mul__() (chainer.utils.type_check.Expr method), 752
__ne__() (chainer.Parameter method), 127
__ne__() (chainer.Variable method), 119
__ne__() (chainer.utils.type_check.Expr method), 752
__neg__() (chainer.Parameter method), 127
__neg__() (chainer.Variable method), 119
__neg__() (chainer.utils.type_check.Expr method), 752
__next__() (chainer.dataset.Iterator method), 692
__next__() (chainer.iterators.MultiprocessIterator method), 715
__next__() (chainer.iterators.MultithreadIterator method), 717
__next__() (chainer.iterators.SerialIterator method), 714
__nonzero__() (chainer.Parameter method), 127
__nonzero__() (chainer.Variable method), 119
__nonzero__() (chainer.utils.type_check.Expr method), 752
__pow__() (chainer.Parameter method), 129
__pow__() (chainer.Variable method), 121
__pow__() (chainer.utils.type_check.Expr method), 752
__radd__() (chainer.Parameter method), 127
__radd__() (chainer.Variable method), 120
__radd__() (chainer.utils.type_check.Expr method), 752
__rdiv__() (chainer.Parameter method), 128
__rdiv__() (chainer.Variable method), 120
__rfloordiv__() (chainer.Parameter method), 128
__rfloordiv__() (chainer.Variable method), 121
__rfloordiv__() (chainer.utils.type_check.Expr method), 752
__rmatmul__() (chainer.Parameter method), 129
__rmatmul__() (chainer.Variable method), 121
__rmul__() (chainer.Parameter method), 128
__rmul__() (chainer.Variable method), 120
__rmul__() (chainer.utils.type_check.Expr method), 752
__rpow__() (chainer.Parameter method), 129
__rpow__() (chainer.Variable method), 121

[__rsub__\(\)](#) (chainer.Parameter method), 128
[__rsub__\(\)](#) (chainer.Variable method), 120
[__rsub__\(\)](#) (chainer.utils.type_check.Expr method), 752
[__rtruediv__\(\)](#) (chainer.Parameter method), 128
[__rtruediv__\(\)](#) (chainer.Variable method), 121
[__rtruediv__\(\)](#) (chainer.utils.type_check.Expr method), 752
[__setitem__\(\)](#) (chainer.Sequential method), 607
[__sub__\(\)](#) (chainer.Parameter method), 128
[__sub__\(\)](#) (chainer.Variable method), 120
[__sub__\(\)](#) (chainer.utils.type_check.Expr method), 752
[__truediv__\(\)](#) (chainer.Parameter method), 128
[__truediv__\(\)](#) (chainer.Variable method), 120
[__truediv__\(\)](#) (chainer.utils.type_check.Expr method), 752

A

[absolute\(\)](#) (in module chainer.functions), 217
[absolute_error\(\)](#) (in module chainer.functions), 201
[AbstractSerializer](#) (class in chainer), 725
[accuracy\(\)](#) (in module chainer.functions), 198
[AdaDelta](#) (class in chainer.optimizers), 612
[AdaGrad](#) (class in chainer.optimizers), 615
[Adam](#) (class in chainer.optimizers), 617
[add\(\)](#) (chainer.DictSummary method), 737
[add\(\)](#) (chainer.Summary method), 736
[add\(\)](#) (in module chainer.functions), 133
[add_hook\(\)](#) (chainer.Function method), 247
[add_hook\(\)](#) (chainer.FunctionAdapter method), 251
[add_hook\(\)](#) (chainer.FunctionNode method), 257
[add_hook\(\)](#) (chainer.GradientMethod method), 638
[add_hook\(\)](#) (chainer.Optimizer method), 633
[add_hook\(\)](#) (chainer.optimizers.AdaDelta method), 613
[add_hook\(\)](#) (chainer.optimizers.AdaGrad method), 615
[add_hook\(\)](#) (chainer.optimizers.Adam method), 617
[add_hook\(\)](#) (chainer.optimizers.MomentumSGD method), 620
[add_hook\(\)](#) (chainer.optimizers.NesterovAG method), 622
[add_hook\(\)](#) (chainer.optimizers.RMSprop method), 624
[add_hook\(\)](#) (chainer.optimizers.RMSpropGraves method), 626
[add_hook\(\)](#) (chainer.optimizers.SGD method), 628
[add_hook\(\)](#) (chainer.optimizers.SMORMS3 method), 631
[add_hook\(\)](#) (chainer.UpdateRule method), 635
[add_link\(\)](#) (chainer.Chain method), 595
[add_link\(\)](#) (chainer.ChainList method), 600
[add_link\(\)](#) (chainer.links.caffe.CaffeFunction method), 584
[add_link\(\)](#) (chainer.links.ChildSumTreeLSTM method), 284
[add_link\(\)](#) (chainer.links.Classifier method), 531
[add_link\(\)](#) (chainer.links.GoogLeNet method), 544

[add_link\(\)](#) (chainer.links.GRU method), 328
[add_link\(\)](#) (chainer.links.Highway method), 333
[add_link\(\)](#) (chainer.links.Inception method), 338
[add_link\(\)](#) (chainer.links.InceptionBN method), 344
[add_link\(\)](#) (chainer.links.LSTM method), 361
[add_link\(\)](#) (chainer.links.Maxout method), 520
[add_link\(\)](#) (chainer.links.MLPConvolution2D method), 367
[add_link\(\)](#) (chainer.links.model.vision.resnet.ResNetLayers method), 551
[add_link\(\)](#) (chainer.links.NaryTreeLSTM method), 372
[add_link\(\)](#) (chainer.links.NStepBiGRU method), 378
[add_link\(\)](#) (chainer.links.NStepBiLSTM method), 384
[add_link\(\)](#) (chainer.links.NStepBiRNNTanh method), 389
[add_link\(\)](#) (chainer.links.NStepBiRNNTanh method), 395
[add_link\(\)](#) (chainer.links.NStepGRU method), 401
[add_link\(\)](#) (chainer.links.NStepLSTM method), 407
[add_link\(\)](#) (chainer.links.NStepRNNTanh method), 413
[add_link\(\)](#) (chainer.links.NStepRNNTanh method), 419
[add_link\(\)](#) (chainer.links.ResNet101Layers method), 564
[add_link\(\)](#) (chainer.links.ResNet152Layers method), 571
[add_link\(\)](#) (chainer.links.ResNet50Layers method), 558
[add_link\(\)](#) (chainer.links.Scale method), 429
[add_link\(\)](#) (chainer.links.StatefulGRU method), 435
[add_link\(\)](#) (chainer.links.StatefulMGU method), 446
[add_link\(\)](#) (chainer.links.StatefulPeepholeLSTM method), 457
[add_link\(\)](#) (chainer.links.StatefulZoneoutLSTM method), 461
[add_link\(\)](#) (chainer.links.StatelessGRU method), 441
[add_link\(\)](#) (chainer.links.StatelessLSTM method), 468
[add_link\(\)](#) (chainer.links.StatelessMGU method), 451
[add_link\(\)](#) (chainer.links.VGG16Layers method), 537
[add_link\(\)](#) (chainer.Sequential method), 607
[add_observer\(\)](#) (chainer.Reporter method), 734
[add_observers\(\)](#) (chainer.Reporter method), 734
[add_param\(\)](#) (chainer.Chain method), 595
[add_param\(\)](#) (chainer.ChainList method), 600
[add_param\(\)](#) (chainer.Link method), 590
[add_param\(\)](#) (chainer.links.BatchNormalization method), 475
[add_param\(\)](#) (chainer.links.BatchRenormalization method), 480
[add_param\(\)](#) (chainer.links.Bias method), 274
[add_param\(\)](#) (chainer.links.Bilinear method), 279
[add_param\(\)](#) (chainer.links.BinaryHierarchicalSoftmax method), 490
[add_param\(\)](#) (chainer.links.BlackOut method), 495
[add_param\(\)](#) (chainer.links.caffe.CaffeFunction method), 584
[add_param\(\)](#) (chainer.links.ChildSumTreeLSTM method), 284

`add_param()` (chainer.links.Classifier method), 531
`add_param()` (chainer.links.Convolution2D method), 290
`add_param()` (chainer.links.ConvolutionND method), 295
`add_param()` (chainer.links.CRF1d method), 500
`add_param()` (chainer.links.Deconvolution2D method), 302
`add_param()` (chainer.links.DeconvolutionND method), 307
`add_param()` (chainer.links.DepthwiseConvolution2D method), 312
`add_param()` (chainer.links.DilatedConvolution2D method), 318
`add_param()` (chainer.links.EmbedID method), 323
`add_param()` (chainer.links.GoogLeNet method), 544
`add_param()` (chainer.links.GRU method), 328
`add_param()` (chainer.links.Highway method), 333
`add_param()` (chainer.links.Inception method), 339
`add_param()` (chainer.links.InceptionBN method), 344
`add_param()` (chainer.links.LayerNormalization method), 485
`add_param()` (chainer.links.Linear method), 350
`add_param()` (chainer.links.LocalConvolution2D method), 355
`add_param()` (chainer.links.LSTM method), 361
`add_param()` (chainer.links.Maxout method), 520
`add_param()` (chainer.links.MLPConvolution2D method), 367
`add_param()` (chainer.links.model.vision.resnet.ResNetLayer method), 552
`add_param()` (chainer.links.NaryTreeLSTM method), 372
`add_param()` (chainer.links.NegativeSampling method), 525
`add_param()` (chainer.links.NStepBiGRU method), 378
`add_param()` (chainer.links.NStepBiLSTM method), 384
`add_param()` (chainer.links.NStepBiRNNTanh method), 389
`add_param()` (chainer.links.NStepBiRNNTanh method), 395
`add_param()` (chainer.links.NStepGRU method), 401
`add_param()` (chainer.links.NStepLSTM method), 407
`add_param()` (chainer.links.NStepRNNTanh method), 413
`add_param()` (chainer.links.NStepRNNTanh method), 419
`add_param()` (chainer.links.Parameter method), 424
`add_param()` (chainer.links.PReLU method), 510
`add_param()` (chainer.links.ResNet101Layers method), 565
`add_param()` (chainer.links.ResNet152Layers method), 571
`add_param()` (chainer.links.ResNet50Layers method), 558
`add_param()` (chainer.links.Scale method), 430
`add_param()` (chainer.links.SimplifiedDropconnect method), 505
`add_param()` (chainer.links.StatefulGRU method), 436
`add_param()` (chainer.links.StatefulMGU method), 446
`add_param()` (chainer.links.StatefulPeepholeLSTM method), 457
`add_param()` (chainer.links.StatefulZoneoutLSTM method), 462
`add_param()` (chainer.links.StatelessGRU method), 442
`add_param()` (chainer.links.StatelessLSTM method), 468
`add_param()` (chainer.links.StatelessMGU method), 451
`add_param()` (chainer.links.Swish method), 515
`add_param()` (chainer.links.TheanoFunction method), 578
`add_param()` (chainer.links.VGG16Layers method), 537
`add_param()` (chainer.Sequential method), 607
`add_persistent()` (chainer.Chain method), 596
`add_persistent()` (chainer.ChainList method), 601
`add_persistent()` (chainer.Link method), 590
`add_persistent()` (chainer.links.BatchNormalization method), 475
`add_persistent()` (chainer.links.BatchRenormalization method), 480
`add_persistent()` (chainer.links.Bias method), 274
`add_persistent()` (chainer.links.Bilinear method), 279
`add_persistent()` (chainer.links.BinaryHierarchicalSoftmax method), 491
`add_persistent()` (chainer.links.BlackOut method), 495
`add_persistent()` (chainer.links.caffe.CaffeFunction method), 584
`add_persistent()` (chainer.links.ChildSumTreeLSTM method), 285
`add_persistent()` (chainer.links.Classifier method), 532
`add_persistent()` (chainer.links.Convolution2D method), 291
`add_persistent()` (chainer.links.ConvolutionND method), 296
`add_persistent()` (chainer.links.CRF1d method), 500
`add_persistent()` (chainer.links.Deconvolution2D method), 302
`add_persistent()` (chainer.links.DeconvolutionND method), 307
`add_persistent()` (chainer.links.DepthwiseConvolution2D method), 312
`add_persistent()` (chainer.links.DilatedConvolution2D method), 318
`add_persistent()` (chainer.links.EmbedID method), 324
`add_persistent()` (chainer.links.GoogLeNet method), 545
`add_persistent()` (chainer.links.GRU method), 328
`add_persistent()` (chainer.links.Highway method), 334
`add_persistent()` (chainer.links.Inception method), 339
`add_persistent()` (chainer.links.InceptionBN method), 344
`add_persistent()` (chainer.links.LayerNormalization method), 486
`add_persistent()` (chainer.links.Linear method), 350

- `add_persistent()` (chainer.links.LocalConvolution2D method), 355
- `add_persistent()` (chainer.links.LSTM method), 361
- `add_persistent()` (chainer.links.Maxout method), 521
- `add_persistent()` (chainer.links.MLPConvolution2D method), 367
- `add_persistent()` (chainer.links.model.vision.resnet.ResNetLayers method), 552
- `add_persistent()` (chainer.links.NaryTreeLSTM method), 373
- `add_persistent()` (chainer.links.NegativeSampling method), 526
- `add_persistent()` (chainer.links.NStepBiGRU method), 378
- `add_persistent()` (chainer.links.NStepBiLSTM method), 384
- `add_persistent()` (chainer.links.NStepBiRNNReLU method), 390
- `add_persistent()` (chainer.links.NStepBiRNNTanh method), 396
- `add_persistent()` (chainer.links.NStepGRU method), 402
- `add_persistent()` (chainer.links.NStepLSTM method), 408
- `add_persistent()` (chainer.links.NStepRNNReLU method), 414
- `add_persistent()` (chainer.links.NStepRNNTanh method), 420
- `add_persistent()` (chainer.links.Parameter method), 425
- `add_persistent()` (chainer.links.PReLU method), 510
- `add_persistent()` (chainer.links.ResNet101Layers method), 565
- `add_persistent()` (chainer.links.ResNet152Layers method), 572
- `add_persistent()` (chainer.links.ResNet50Layers method), 559
- `add_persistent()` (chainer.links.Scale method), 430
- `add_persistent()` (chainer.links.SimplifiedDropconnect method), 505
- `add_persistent()` (chainer.links.StatefulGRU method), 436
- `add_persistent()` (chainer.links.StatefulMGU method), 447
- `add_persistent()` (chainer.links.StatefulPeepholeLSTM method), 457
- `add_persistent()` (chainer.links.StatefulZoneoutLSTM method), 462
- `add_persistent()` (chainer.links.StatelessGRU method), 442
- `add_persistent()` (chainer.links.StatelessLSTM method), 468
- `add_persistent()` (chainer.links.StatelessMGU method), 452
- `add_persistent()` (chainer.links.Swish method), 515
- `add_persistent()` (chainer.links.TheanoFunction method), 578
- `add_persistent()` (chainer.links.VGG16Layers method), 538
- `add_persistent()` (chainer.Sequential method), 607
- `added()` (chainer.function_hooks.CUDAProfileHook method), 263
- `added()` (chainer.function_hooks.CupyMemoryProfileHook method), 264
- `added()` (chainer.function_hooks.PrintHook method), 266
- `added()` (chainer.function_hooks.TimerHook method), 268
- `added()` (chainer.FunctionHook method), 271
- `addgrad()` (chainer.Parameter method), 124
- `addgrad()` (chainer.Variable method), 117
- `addgrads()` (chainer.Chain method), 596
- `addgrads()` (chainer.ChainList method), 601
- `addgrads()` (chainer.Link method), 590
- `addgrads()` (chainer.links.BatchNormalization method), 475
- `addgrads()` (chainer.links.BatchRenormalization method), 481
- `addgrads()` (chainer.links.Bias method), 274
- `addgrads()` (chainer.links.Bilinear method), 279
- `addgrads()` (chainer.links.BinaryHierarchicalSoftmax method), 491
- `addgrads()` (chainer.links.BlackOut method), 496
- `addgrads()` (chainer.links.caffe.CaffeFunction method), 585
- `addgrads()` (chainer.links.ChildSumTreeLSTM method), 285
- `addgrads()` (chainer.links.Classifier method), 532
- `addgrads()` (chainer.links.Convolution2D method), 291
- `addgrads()` (chainer.links.ConvolutionND method), 296
- `addgrads()` (chainer.links.CRF1d method), 500
- `addgrads()` (chainer.links.Deconvolution2D method), 302
- `addgrads()` (chainer.links.DeconvolutionND method), 308
- `addgrads()` (chainer.links.DepthwiseConvolution2D method), 313
- `addgrads()` (chainer.links.DilatedConvolution2D method), 319
- `addgrads()` (chainer.links.EmbedID method), 324
- `addgrads()` (chainer.links.GoogLeNet method), 545
- `addgrads()` (chainer.links.GRU method), 329
- `addgrads()` (chainer.links.Highway method), 334
- `addgrads()` (chainer.links.Inception method), 339
- `addgrads()` (chainer.links.InceptionBN method), 345
- `addgrads()` (chainer.links.LayerNormalization method), 486
- `addgrads()` (chainer.links.Linear method), 350
- `addgrads()` (chainer.links.LocalConvolution2D method), 355
- `addgrads()` (chainer.links.LSTM method), 362
- `addgrads()` (chainer.links.Maxout method), 521
- `addgrads()` (chainer.links.MLPConvolution2D method), 367

- `addgrads()` (chainer.links.model.vision.resnet.ResNetLayers method), 552
 - `addgrads()` (chainer.links.NaryTreeLSTM method), 373
 - `addgrads()` (chainer.links.NegativeSampling method), 526
 - `addgrads()` (chainer.links.NStepBiGRU method), 378
 - `addgrads()` (chainer.links.NStepBiLSTM method), 384
 - `addgrads()` (chainer.links.NStepBiRNReLU method), 390
 - `addgrads()` (chainer.links.NStepBiRNNTanh method), 396
 - `addgrads()` (chainer.links.NStepGRU method), 402
 - `addgrads()` (chainer.links.NStepLSTM method), 408
 - `addgrads()` (chainer.links.NStepRNReLU method), 414
 - `addgrads()` (chainer.links.NStepRNNTanh method), 420
 - `addgrads()` (chainer.links.Parameter method), 425
 - `addgrads()` (chainer.links.PReLU method), 510
 - `addgrads()` (chainer.links.ResNet101Layers method), 565
 - `addgrads()` (chainer.links.ResNet152Layers method), 572
 - `addgrads()` (chainer.links.ResNet50Layers method), 559
 - `addgrads()` (chainer.links.Scale method), 430
 - `addgrads()` (chainer.links.SimplifiedDropconnect method), 506
 - `addgrads()` (chainer.links.StatefulGRU method), 436
 - `addgrads()` (chainer.links.StatefulMGU method), 447
 - `addgrads()` (chainer.links.StatefulPeepholeLSTM method), 457
 - `addgrads()` (chainer.links.StatefulZoneoutLSTM method), 462
 - `addgrads()` (chainer.links.StatelessGRU method), 442
 - `addgrads()` (chainer.links.StatelessLSTM method), 469
 - `addgrads()` (chainer.links.StatelessMGU method), 452
 - `addgrads()` (chainer.links.Swish method), 516
 - `addgrads()` (chainer.links.TheanoFunction method), 579
 - `addgrads()` (chainer.links.VGG16Layers method), 538
 - `addgrads()` (chainer.Sequential method), 607
 - `alpha` (chainer.optimizers.Adam attribute), 619
 - `alpha` (chainer.optimizers.RMSprop attribute), 626
 - `alpha` (chainer.optimizers.RMSpropGraves attribute), 628
 - `amsgrad` (chainer.optimizers.Adam attribute), 619
 - `append()` (chainer.ChainList method), 601
 - `append()` (chainer.links.MLPConvolution2D method), 367
 - `append()` (chainer.links.NStepBiGRU method), 378
 - `append()` (chainer.links.NStepBiLSTM method), 384
 - `append()` (chainer.links.NStepBiRNReLU method), 390
 - `append()` (chainer.links.NStepBiRNNTanh method), 396
 - `append()` (chainer.links.NStepGRU method), 402
 - `append()` (chainer.links.NStepLSTM method), 408
 - `append()` (chainer.links.NStepRNReLU method), 414
 - `append()` (chainer.links.NStepRNNTanh method), 420
 - `append()` (chainer.Sequential method), 608
 - `apply()` (chainer.FunctionAdapter method), 251
 - `apply()` (chainer.FunctionNode method), 257
 - `arccos()` (in module chainer.functions), 218
 - `arcsin()` (in module chainer.functions), 218
 - `arctan()` (in module chainer.functions), 218
 - `arctan2()` (in module chainer.functions), 218
 - `argmax()` (chainer.links.CRF1d method), 500
 - `argmax()` (in module chainer.functions), 219
 - `argmax_crf1d()` (in module chainer.functions), 206
 - `argmin()` (in module chainer.functions), 219
 - `array` (chainer.Parameter attribute), 129
 - `array` (chainer.Variable attribute), 121
 - `as_variable()` (in module chainer), 123
 - `assert_allclose()` (in module chainer.testing), 756
 - `available()` (chainer.training.extensions.PlotReport static method), 680
 - `available()` (chainer.training.extensions.VariableStatisticsPlot static method), 682
 - `available()` (chainer.training.updaters.MultiprocessParallelUpdater static method), 661
 - `available_layers` (chainer.links.GoogLeNet attribute), 549
 - `available_layers` (chainer.links.model.vision.resnet.ResNetLayers attribute), 557
 - `available_layers` (chainer.links.ResNet101Layers attribute), 570
 - `available_layers` (chainer.links.ResNet152Layers attribute), 576
 - `available_layers` (chainer.links.ResNet50Layers attribute), 563
 - `available_layers` (chainer.links.VGG16Layers attribute), 542
 - `average()` (in module chainer.functions), 219
 - `average_pooling_2d()` (in module chainer.functions), 239
 - `average_pooling_nd()` (in module chainer.functions), 240
 - `avg_mean` (chainer.links.BatchNormalization attribute), 479
 - `avg_mean` (chainer.links.BatchRenormalization attribute), 484
 - `avg_var` (chainer.links.BatchNormalization attribute), 479
 - `avg_var` (chainer.links.BatchRenormalization attribute), 484
- ## B
- `backward()` (chainer.Function method), 248
 - `backward()` (chainer.FunctionAdapter method), 252
 - `backward()` (chainer.FunctionNode method), 257
 - `backward()` (chainer.Parameter method), 125
 - `backward()` (chainer.Variable method), 117
 - `backward_accumulate()` (chainer.FunctionAdapter method), 252
 - `backward_accumulate()` (chainer.FunctionNode method), 258
 - `backward_cpu()` (chainer.Function method), 248
 - `backward_gpu()` (chainer.Function method), 248
 - `backward_postprocess()` (chainer.function_hooks.CUDAProfileHook method), 263

- `backward_postprocess()` (chainer.function_hooks.CupyMemoryProfileHook method), 264
- `backward_postprocess()` (chainer.function_hooks.PrintHookcall_for_each_param (chainer.optimizer_hooks.GradientNoise attribute), 643
- `backward_postprocess()` (chainer.function_hooks.TimerHookcall_for_each_param (chainer.optimizer_hooks.Lasso attribute), 641
- `backward_postprocess()` (chainer.FunctionHook method), call_for_each_param (chainer.optimizer_hooks.WeightDecay attribute), 640
- `backward_preprocess()` (chainer.function_hooks.CUDAProfileHookcall_hooks() (chainer.GradientMethod method), 638
- `backward_preprocess()` (chainer.function_hooks.CupyMemoryProfileHookcall_hooks() (chainer.Optimizer method), 633
- `backward_preprocess()` (chainer.function_hooks.CupyMemoryProfileHookcall_hooks() (chainer.optimizers.AdaDelta method), 613
- `backward_preprocess()` (chainer.function_hooks.CupyMemoryProfileHookcall_hooks() (chainer.optimizers.AdaGrad method), 615
- `backward_preprocess()` (chainer.function_hooks.PrintHookcall_hooks() (chainer.optimizers.Adam method), 617
- `backward_preprocess()` (chainer.function_hooks.PrintHookcall_hooks() (chainer.optimizers.MomentumSGD method), 620
- `backward_preprocess()` (chainer.function_hooks.TimerHookcall_hooks() (chainer.optimizers.NesterovAG method), 622
- `backward_preprocess()` (chainer.FunctionHook method), call_hooks() (chainer.optimizers.RMSprop method), 624
- `backward_preprocess()` (chainer.FunctionHook method), call_hooks() (chainer.optimizers.RMSpropGraves method), 626
- `batch_det()` (in module chainer.functions), 222
- `batch_inv()` (in module chainer.functions), 219
- `batch_l2_norm_squared()` (in module chainer.functions), call_hooks() (chainer.optimizers.SGD method), 629
- `batch_l2_norm_squared()` (in module chainer.functions), call_hooks() (chainer.optimizers.SMORMS3 method), 631
- `batch_matmul()` (in module chainer.functions), 220
- `batch_normalization()` (in module chainer.functions), 236
- `batch_renormalization()` (in module chainer.functions), 237
- `BatchNormalization` (class in chainer.links), 473
- `BatchRenormalization` (class in chainer.links), 479
- `bernoulli_nll()` (in module chainer.functions), 201
- `BestValueTrigger` (class in chainer.training.triggers), 685
- `beta` (chainer.links.BatchNormalization attribute), 479
- `beta` (chainer.links.BatchRenormalization attribute), 484
- `beta1` (chainer.optimizers.Adam attribute), 619
- `beta2` (chainer.optimizers.Adam attribute), 619
- `Bias` (class in chainer.links), 273
- `bias()` (in module chainer.functions), 220
- `Bilinear` (class in chainer.links), 278
- `bilinear()` (in module chainer.functions), 174
- `binary_accuracy()` (in module chainer.functions), 198
- `BinaryHierarchicalSoftmax` (class in chainer.links), 489
- `black_out()` (in module chainer.functions), 202
- `BlackOut` (class in chainer.links), 495
- `broadcast()` (in module chainer.functions), 147
- `broadcast_to()` (in module chainer.functions), 148
- `build_computational_graph()` (in module chainer.computational_graph), 747
- C**
- `cache_or_load_file()` (in module chainer.dataset), 696
- `cached_download()` (in module chainer.dataset), 695
- `CaffeFunction` (class in chainer.links.caffe), 582
- `call_for_each_param` (chainer.optimizer_hooks.GradientHardClipping attribute), 642
- `call_for_each_param` (chainer.optimizer_hooks.GradientLARS attribute), 644
- `call_for_each_param` (chainer.optimizer_hooks.GradientNoise attribute), 643
- `call_for_each_param` (chainer.optimizer_hooks.Lasso attribute), 641
- `call_for_each_param` (chainer.optimizer_hooks.WeightDecay attribute), 640
- `call_hooks()` (chainer.GradientMethod method), 638
- `call_hooks()` (chainer.Optimizer method), 633
- `call_hooks()` (chainer.optimizers.AdaDelta method), 613
- `call_hooks()` (chainer.optimizers.AdaGrad method), 615
- `call_hooks()` (chainer.optimizers.Adam method), 617
- `call_hooks()` (chainer.optimizers.MomentumSGD method), 620
- `call_hooks()` (chainer.optimizers.NesterovAG method), 622
- `call_hooks()` (chainer.optimizers.RMSprop method), 624
- `call_hooks()` (chainer.optimizers.RMSpropGraves method), 626
- `call_hooks()` (chainer.optimizers.SGD method), 629
- `call_hooks()` (chainer.optimizers.SMORMS3 method), 631
- `cast()` (in module chainer.functions), 148
- `ceil()` (in module chainer.functions), 221
- `Chain` (class in chainer), 594
- `chainer` (module), 115, 588, 724
- `chainer.backends.cuda` (module), 727
- `chainer.computational_graph` (module), 746
- `chainer.dataset` (module), 689
- `chainer.datasets` (module), 696
- `chainer.exporters` (module), 749
- `chainer.function_hooks` (module), 262
- `chainer.functions` (module), 133
- `chainer.initializers` (module), 645
- `chainer.iterators` (module), 713
- `chainer.links` (module), 272
- `chainer.links.caffe` (module), 749
- `chainer.serializers` (module), 717
- `chainer.testing` (module), 756
- `chainer.training` (module), 653
- `chainer.utils` (module), 760
- `ChainList` (class in chainer), 600
- `check_backward()` (in module chainer.gradient_check), 754
- `check_type_forward()` (chainer.Function method), 249
- `check_type_forward()` (chainer.FunctionAdapter method), 253
- `check_type_forward()` (chainer.FunctionNode method), 258
- `children()` (chainer.Chain method), 596
- `children()` (chainer.ChainList method), 601
- `children()` (chainer.Link method), 590

- `children()` (chainer.links.BatchNormalization method), 476
- `children()` (chainer.links.BatchRenormalization method), 481
- `children()` (chainer.links.Bias method), 274
- `children()` (chainer.links.Bilinear method), 279
- `children()` (chainer.links.BinaryHierarchicalSoftmax method), 491
- `children()` (chainer.links.BlackOut method), 496
- `children()` (chainer.links.caffe.CaffeFunction method), 585
- `children()` (chainer.links.ChildSumTreeLSTM method), 285
- `children()` (chainer.links.Classifier method), 532
- `children()` (chainer.links.Convolution2D method), 291
- `children()` (chainer.links.ConvolutionND method), 296
- `children()` (chainer.links.CRF1d method), 501
- `children()` (chainer.links.Deconvolution2D method), 303
- `children()` (chainer.links.DeconvolutionND method), 308
- `children()` (chainer.links.DepthwiseConvolution2D method), 313
- `children()` (chainer.links.DilatedConvolution2D method), 319
- `children()` (chainer.links.EmbedID method), 324
- `children()` (chainer.links.GoogLeNet method), 545
- `children()` (chainer.links.GRU method), 329
- `children()` (chainer.links.Highway method), 334
- `children()` (chainer.links.Inception method), 339
- `children()` (chainer.links.InceptionBN method), 345
- `children()` (chainer.links.LayerNormalization method), 486
- `children()` (chainer.links.Linear method), 350
- `children()` (chainer.links.LocalConvolution2D method), 356
- `children()` (chainer.links.LSTM method), 362
- `children()` (chainer.links.Maxout method), 521
- `children()` (chainer.links.MLPConvolution2D method), 367
- `children()` (chainer.links.model.vision.resnet.ResNetLayers method), 552
- `children()` (chainer.links.NaryTreeLSTM method), 373
- `children()` (chainer.links.NegativeSampling method), 526
- `children()` (chainer.links.NStepBiGRU method), 379
- `children()` (chainer.links.NStepBiLSTM method), 384
- `children()` (chainer.links.NStepBiRNNReLU method), 390
- `children()` (chainer.links.NStepBiRNNTanh method), 396
- `children()` (chainer.links.NStepGRU method), 402
- `children()` (chainer.links.NStepLSTM method), 408
- `children()` (chainer.links.NStepRNNReLU method), 414
- `children()` (chainer.links.NStepRNNTanh method), 420
- `children()` (chainer.links.Parameter method), 425
- `children()` (chainer.links.PReLU method), 510
- `children()` (chainer.links.ResNet101Layers method), 566
- `children()` (chainer.links.ResNet152Layers method), 572
- `children()` (chainer.links.ResNet50Layers method), 559
- `children()` (chainer.links.Scale method), 430
- `children()` (chainer.links.SimplifiedDropconnect method), 506
- `children()` (chainer.links.StatefulGRU method), 436
- `children()` (chainer.links.StatefulMGU method), 447
- `children()` (chainer.links.StatefulPeepholeLSTM method), 458
- `children()` (chainer.links.StatefulZoneoutLSTM method), 463
- `children()` (chainer.links.StatelessGRU method), 442
- `children()` (chainer.links.StatelessLSTM method), 469
- `children()` (chainer.links.StatelessMGU method), 452
- `children()` (chainer.links.Swish method), 516
- `children()` (chainer.links.TheanoFunction method), 579
- `children()` (chainer.links.VGG16Layers method), 538
- `children()` (chainer.Sequential method), 608
- `ChildSumTreeLSTM` (class in chainer.links), 283
- `classification_summary()` (in module chainer.functions), 199
- `Classifier` (class in chainer.links), 530
- `clear()` (chainer.Sequential method), 608
- `clear_memo()` (in module chainer.backends.cuda), 731
- `cleargrad()` (chainer.Parameter method), 125
- `cleargrad()` (chainer.Variable method), 118
- `cleargrads()` (chainer.Chain method), 596
- `cleargrads()` (chainer.ChainList method), 601
- `cleargrads()` (chainer.Link method), 590
- `cleargrads()` (chainer.links.BatchNormalization method), 476
- `cleargrads()` (chainer.links.BatchRenormalization method), 481
- `cleargrads()` (chainer.links.Bias method), 274
- `cleargrads()` (chainer.links.Bilinear method), 279
- `cleargrads()` (chainer.links.BinaryHierarchicalSoftmax method), 491
- `cleargrads()` (chainer.links.BlackOut method), 496
- `cleargrads()` (chainer.links.caffe.CaffeFunction method), 585
- `cleargrads()` (chainer.links.ChildSumTreeLSTM method), 285
- `cleargrads()` (chainer.links.Classifier method), 532
- `cleargrads()` (chainer.links.Convolution2D method), 291
- `cleargrads()` (chainer.links.ConvolutionND method), 296
- `cleargrads()` (chainer.links.CRF1d method), 501
- `cleargrads()` (chainer.links.Deconvolution2D method), 303
- `cleargrads()` (chainer.links.DeconvolutionND method), 308
- `cleargrads()` (chainer.links.DepthwiseConvolution2D method), 313
- `cleargrads()` (chainer.links.DilatedConvolution2D method), 319

- cleargrads() (chainer.links.EmbedID method), 324
- cleargrads() (chainer.links.GoogLeNet method), 545
- cleargrads() (chainer.links.GRU method), 329
- cleargrads() (chainer.links.Highway method), 334
- cleargrads() (chainer.links.Inception method), 339
- cleargrads() (chainer.links.InceptionBN method), 345
- cleargrads() (chainer.links.LayerNormalization method), 486
- cleargrads() (chainer.links.Linear method), 350
- cleargrads() (chainer.links.LocalConvolution2D method), 356
- cleargrads() (chainer.links.LSTM method), 362
- cleargrads() (chainer.links.Maxout method), 521
- cleargrads() (chainer.links.MLPConvolution2D method), 368
- cleargrads() (chainer.links.model.vision.resnet.ResNetLayers connect_trainer() method), 552
- cleargrads() (chainer.links.NaryTreeLSTM method), 373
- cleargrads() (chainer.links.NegativeSampling method), 526
- cleargrads() (chainer.links.NStepBiGRU method), 379
- cleargrads() (chainer.links.NStepBiLSTM method), 385
- cleargrads() (chainer.links.NStepBiRNNReLU method), 390
- cleargrads() (chainer.links.NStepBiRNNTanh method), 396
- cleargrads() (chainer.links.NStepGRU method), 402
- cleargrads() (chainer.links.NStepLSTM method), 408
- cleargrads() (chainer.links.NStepRNNReLU method), 414
- cleargrads() (chainer.links.NStepRNNTanh method), 420
- cleargrads() (chainer.links.Parameter method), 425
- cleargrads() (chainer.links.PReLU method), 510
- cleargrads() (chainer.links.ResNet101Layers method), 566
- cleargrads() (chainer.links.ResNet152Layers method), 572
- cleargrads() (chainer.links.ResNet50Layers method), 559
- cleargrads() (chainer.links.Scale method), 430
- cleargrads() (chainer.links.SimplifiedDropconnect method), 506
- cleargrads() (chainer.links.StatefulGRU method), 437
- cleargrads() (chainer.links.StatefulMGU method), 447
- cleargrads() (chainer.links.StatefulPeepholeLSTM method), 458
- cleargrads() (chainer.links.StatefulZoneoutLSTM method), 463
- cleargrads() (chainer.links.StatelessGRU method), 442
- cleargrads() (chainer.links.StatelessLSTM method), 469
- cleargrads() (chainer.links.StatelessMGU method), 452
- cleargrads() (chainer.links.Swish method), 516
- cleargrads() (chainer.links.TheanoFunction method), 579
- cleargrads() (chainer.links.VGG16Layers method), 538
- cleargrads() (chainer.Sequential method), 608
- clip() (in module chainer.functions), 221
- clipped_relu() (in module chainer.functions), 134
- ComputationalGraph (class in chainer.computational_graph), 748
- compute_accuracy (chainer.links.Classifier attribute), 535
- compute_mean() (chainer.DictSummary method), 737
- compute_mean() (chainer.Summary method), 736
- concat() (in module chainer.functions), 149
- concat_examples() (in module chainer.dataset), 692
- ConcatenatedDataset (class in chainer.datasets), 698
- ConcatWithAsyncTransfer (class in chainer.dataset), 694
- config (in module chainer), 742
- connect_trainer() (chainer.training.Updater method), 656
- connect_trainer() (chainer.training.updaters.MultiprocessParallelUpdater method), 661
- connect_trainer() (chainer.training.updaters.ParallelUpdater method), 659
- connect_trainer() (chainer.training.updaters.StandardUpdater method), 658
- connectionist_temporal_classification() (in module chainer.functions), 203
- Constant (class in chainer.initializers), 646
- contrastive() (in module chainer.functions), 204
- convert_caffemodel_to_npz() (chainer.links.GoogLeNet class method), 545
- convert_caffemodel_to_npz() (chainer.links.model.vision.resnet.ResNetLayers class method), 552
- convert_caffemodel_to_npz() (chainer.links.ResNet101Layers class method), 566
- convert_caffemodel_to_npz() (chainer.links.ResNet152Layers class method), 572
- convert_caffemodel_to_npz() (chainer.links.ResNet50Layers class method), 559
- convert_caffemodel_to_npz() (chainer.links.VGG16Layers class method), 538
- Convolution2D (class in chainer.links), 288
- convolution_2d() (in module chainer.functions), 174
- convolution_nd() (in module chainer.functions), 176
- ConvolutionND (class in chainer.links), 295
- copy() (chainer.Chain method), 596
- copy() (chainer.ChainList method), 601
- copy() (chainer.Link method), 591
- copy() (chainer.links.BatchNormalization method), 476
- copy() (chainer.links.BatchRenormalization method), 481
- copy() (chainer.links.Bias method), 274
- copy() (chainer.links.Bilinear method), 280
- copy() (chainer.links.BinaryHierarchicalSoftmax method), 491
- copy() (chainer.links.BlackOut method), 496

`copy()` (chainer.links.caffe.CaffeFunction method), 585
`copy()` (chainer.links.ChildSumTreeLSTM method), 285
`copy()` (chainer.links.Classifier method), 532
`copy()` (chainer.links.Convolution2D method), 291
`copy()` (chainer.links.ConvolutionND method), 296
`copy()` (chainer.links.CRF1d method), 501
`copy()` (chainer.links.Deconvolution2D method), 303
`copy()` (chainer.links.DeconvolutionND method), 308
`copy()` (chainer.links.DepthwiseConvolution2D method), 313
`copy()` (chainer.links.DilatedConvolution2D method), 319
`copy()` (chainer.links.EmbedID method), 324
`copy()` (chainer.links.GoogLeNet method), 545
`copy()` (chainer.links.GRU method), 329
`copy()` (chainer.links.Highway method), 334
`copy()` (chainer.links.Inception method), 340
`copy()` (chainer.links.InceptionBN method), 345
`copy()` (chainer.links.LayerNormalization method), 486
`copy()` (chainer.links.Linear method), 351
`copy()` (chainer.links.LocalConvolution2D method), 356
`copy()` (chainer.links.LSTM method), 362
`copy()` (chainer.links.Maxout method), 521
`copy()` (chainer.links.MLPConvolution2D method), 368
`copy()` (chainer.links.model.vision.resnet.ResNetLayers method), 553
`copy()` (chainer.links.NaryTreeLSTM method), 373
`copy()` (chainer.links.NegativeSampling method), 526
`copy()` (chainer.links.NStepBiGRU method), 379
`copy()` (chainer.links.NStepBiLSTM method), 385
`copy()` (chainer.links.NStepBiRNReLU method), 390
`copy()` (chainer.links.NStepBiRNNTanh method), 396
`copy()` (chainer.links.NStepGRU method), 402
`copy()` (chainer.links.NStepLSTM method), 408
`copy()` (chainer.links.NStepRNReLU method), 414
`copy()` (chainer.links.NStepRNNTanh method), 420
`copy()` (chainer.links.Parameter method), 425
`copy()` (chainer.links.PReLU method), 511
`copy()` (chainer.links.ResNet101Layers method), 566
`copy()` (chainer.links.ResNet152Layers method), 572
`copy()` (chainer.links.ResNet50Layers method), 559
`copy()` (chainer.links.Scale method), 431
`copy()` (chainer.links.SimplifiedDropconnect method), 506
`copy()` (chainer.links.StatefulGRU method), 437
`copy()` (chainer.links.StatefulMGU method), 447
`copy()` (chainer.links.StatefulPeepholeLSTM method), 458
`copy()` (chainer.links.StatefulZoneoutLSTM method), 463
`copy()` (chainer.links.StatelessGRU method), 443
`copy()` (chainer.links.StatelessLSTM method), 469
`copy()` (chainer.links.StatelessMGU method), 452
`copy()` (chainer.links.Swish method), 516
`copy()` (chainer.links.TheanoFunction method), 579
`copy()` (chainer.links.VGG16Layers method), 538
`copy()` (chainer.Sequential method), 608
`copy()` (in module chainer.backends.cuda), 729
`copy()` (in module chainer.functions), 150
`copydata()` (chainer.Parameter method), 125
`copydata()` (chainer.Variable method), 118
`copyparams()` (chainer.Chain method), 597
`copyparams()` (chainer.ChainList method), 602
`copyparams()` (chainer.Link method), 591
`copyparams()` (chainer.links.BatchNormalization method), 476
`copyparams()` (chainer.links.BatchRenormalization method), 481
`copyparams()` (chainer.links.Bias method), 275
`copyparams()` (chainer.links.Bilinear method), 280
`copyparams()` (chainer.links.BinaryHierarchicalSoftmax method), 491
`copyparams()` (chainer.links.BlackOut method), 496
`copyparams()` (chainer.links.caffe.CaffeFunction method), 585
`copyparams()` (chainer.links.ChildSumTreeLSTM method), 285
`copyparams()` (chainer.links.Classifier method), 532
`copyparams()` (chainer.links.Convolution2D method), 292
`copyparams()` (chainer.links.ConvolutionND method), 297
`copyparams()` (chainer.links.CRF1d method), 501
`copyparams()` (chainer.links.Deconvolution2D method), 303
`copyparams()` (chainer.links.DeconvolutionND method), 308
`copyparams()` (chainer.links.DepthwiseConvolution2D method), 313
`copyparams()` (chainer.links.DilatedConvolution2D method), 319
`copyparams()` (chainer.links.EmbedID method), 324
`copyparams()` (chainer.links.GoogLeNet method), 546
`copyparams()` (chainer.links.GRU method), 329
`copyparams()` (chainer.links.Highway method), 335
`copyparams()` (chainer.links.Inception method), 340
`copyparams()` (chainer.links.InceptionBN method), 345
`copyparams()` (chainer.links.LayerNormalization method), 486
`copyparams()` (chainer.links.Linear method), 351
`copyparams()` (chainer.links.LocalConvolution2D method), 356
`copyparams()` (chainer.links.LSTM method), 362
`copyparams()` (chainer.links.Maxout method), 522
`copyparams()` (chainer.links.MLPConvolution2D method), 368
`copyparams()` (chainer.links.model.vision.resnet.ResNetLayers method), 553

- `copyparams()` (`chainer.links.NaryTreeLSTM` method), 373
- `copyparams()` (`chainer.links.NegativeSampling` method), 526
- `copyparams()` (`chainer.links.NStepBiGRU` method), 379
- `copyparams()` (`chainer.links.NStepBiLSTM` method), 385
- `copyparams()` (`chainer.links.NStepBiRNNReLU` method), 391
- `copyparams()` (`chainer.links.NStepBiRNNTanh` method), 397
- `copyparams()` (`chainer.links.NStepGRU` method), 403
- `copyparams()` (`chainer.links.NStepLSTM` method), 409
- `copyparams()` (`chainer.links.NStepRNNReLU` method), 415
- `copyparams()` (`chainer.links.NStepRNNTanh` method), 421
- `copyparams()` (`chainer.links.Parameter` method), 426
- `copyparams()` (`chainer.links.PReLU` method), 511
- `copyparams()` (`chainer.links.ResNet101Layers` method), 566
- `copyparams()` (`chainer.links.ResNet152Layers` method), 573
- `copyparams()` (`chainer.links.ResNet50Layers` method), 560
- `copyparams()` (`chainer.links.Scale` method), 431
- `copyparams()` (`chainer.links.SimplifiedDropconnect` method), 506
- `copyparams()` (`chainer.links.StatefulGRU` method), 437
- `copyparams()` (`chainer.links.StatefulMGU` method), 448
- `copyparams()` (`chainer.links.StatefulPeepholeLSTM` method), 458
- `copyparams()` (`chainer.links.StatefulZoneoutLSTM` method), 463
- `copyparams()` (`chainer.links.StatelessGRU` method), 443
- `copyparams()` (`chainer.links.StatelessLSTM` method), 469
- `copyparams()` (`chainer.links.StatelessMGU` method), 453
- `copyparams()` (`chainer.links.Swish` method), 516
- `copyparams()` (`chainer.links.TheanoFunction` method), 579
- `copyparams()` (`chainer.links.VGG16Layers` method), 538
- `copyparams()` (`chainer.Sequential` method), 608
- `cos()` (in module `chainer.functions`), 221
- `cosh()` (in module `chainer.functions`), 222
- `count()` (`chainer.Sequential` method), 608
- `count()` (`chainer.utils.type_check.TypeInfoTuple` method), 753
- `count_by_layer_type()` (`chainer.Sequential` method), 608
- `count_params()` (`chainer.Chain` method), 597
- `count_params()` (`chainer.ChainList` method), 602
- `count_params()` (`chainer.Link` method), 591
- `count_params()` (`chainer.links.BatchNormalization` method), 476
- `count_params()` (`chainer.links.BatchRenormalization` method), 481
- `count_params()` (`chainer.links.Bias` method), 275
- `count_params()` (`chainer.links.Bilinear` method), 280
- `count_params()` (`chainer.links.BinaryHierarchicalSoftmax` method), 491
- `count_params()` (`chainer.links.BlackOut` method), 496
- `count_params()` (`chainer.links.caffe.CaffeFunction` method), 585
- `count_params()` (`chainer.links.ChildSumTreeLSTM` method), 285
- `count_params()` (`chainer.links.Classifier` method), 533
- `count_params()` (`chainer.links.Convolution2D` method), 292
- `count_params()` (`chainer.links.ConvolutionND` method), 297
- `count_params()` (`chainer.links.CRF1d` method), 501
- `count_params()` (`chainer.links.Deconvolution2D` method), 303
- `count_params()` (`chainer.links.DeconvolutionND` method), 308
- `count_params()` (`chainer.links.DepthwiseConvolution2D` method), 313
- `count_params()` (`chainer.links.DilatedConvolution2D` method), 319
- `count_params()` (`chainer.links.EmbedID` method), 324
- `count_params()` (`chainer.links.GoogLeNet` method), 546
- `count_params()` (`chainer.links.GRU` method), 329
- `count_params()` (`chainer.links.Highway` method), 335
- `count_params()` (`chainer.links.Inception` method), 340
- `count_params()` (`chainer.links.InceptionBN` method), 345
- `count_params()` (`chainer.links.LayerNormalization` method), 487
- `count_params()` (`chainer.links.Linear` method), 351
- `count_params()` (`chainer.links.LocalConvolution2D` method), 356
- `count_params()` (`chainer.links.LSTM` method), 362
- `count_params()` (`chainer.links.Maxout` method), 522
- `count_params()` (`chainer.links.MLPConvolution2D` method), 368
- `count_params()` (`chainer.links.model.vision.resnet.ResNetLayers` method), 553
- `count_params()` (`chainer.links.NaryTreeLSTM` method), 374
- `count_params()` (`chainer.links.NegativeSampling` method), 527
- `count_params()` (`chainer.links.NStepBiGRU` method), 379
- `count_params()` (`chainer.links.NStepBiLSTM` method), 385
- `count_params()` (`chainer.links.NStepBiRNNReLU` method), 391
- `count_params()` (`chainer.links.NStepBiRNNTanh` method), 397
- `count_params()` (`chainer.links.NStepGRU` method), 403
- `count_params()` (`chainer.links.NStepLSTM` method), 409

- `count_params()` (chainer.links.NStepRNReLU method), 415
 - `count_params()` (chainer.links.NStepRNNTanh method), 421
 - `count_params()` (chainer.links.Parameter method), 426
 - `count_params()` (chainer.links.PReLU method), 511
 - `count_params()` (chainer.links.ResNet101Layers method), 566
 - `count_params()` (chainer.links.ResNet152Layers method), 573
 - `count_params()` (chainer.links.ResNet50Layers method), 560
 - `count_params()` (chainer.links.Scale method), 431
 - `count_params()` (chainer.links.SimplifiedDropconnect method), 506
 - `count_params()` (chainer.links.StatefulGRU method), 437
 - `count_params()` (chainer.links.StatefulMGU method), 448
 - `count_params()` (chainer.links.StatefulPeepholeLSTM method), 458
 - `count_params()` (chainer.links.StatefulZoneoutLSTM method), 463
 - `count_params()` (chainer.links.StatelessGRU method), 443
 - `count_params()` (chainer.links.StatelessLSTM method), 469
 - `count_params()` (chainer.links.StatelessMGU method), 453
 - `count_params()` (chainer.links.Swish method), 516
 - `count_params()` (chainer.links.TheanoFunction method), 579
 - `count_params()` (chainer.links.VGG16Layers method), 539
 - `count_params()` (chainer.Sequential method), 609
 - `create_huffman_tree()` (chainer.links.BinaryHierarchicalSoftmax static method), 492
 - `create_update_rule()` (chainer.GradientMethod method), 638
 - `create_update_rule()` (chainer.optimizers.AdaDelta method), 613
 - `create_update_rule()` (chainer.optimizers.AdaGrad method), 615
 - `create_update_rule()` (chainer.optimizers.Adam method), 618
 - `create_update_rule()` (chainer.optimizers.MomentumSGD method), 620
 - `create_update_rule()` (chainer.optimizers.NesterovAG method), 622
 - `create_update_rule()` (chainer.optimizers.RMSprop method), 624
 - `create_update_rule()` (chainer.optimizers.RMSpropGraves method), 627
 - `create_update_rule()` (chainer.optimizers.SGD method), 629
 - `create_update_rule()` (chainer.optimizers.SMORMS3 method), 631
 - `creator` (chainer.Parameter attribute), 129
 - `creator` (chainer.Variable attribute), 121
 - `creator` (chainer.variable.VariableNode attribute), 132
 - `creator_node` (chainer.Parameter attribute), 129
 - `creator_node` (chainer.Variable attribute), 122
 - `creator_node` (chainer.variable.VariableNode attribute), 132
 - `crelu()` (in module chainer.functions), 135
 - `CRF1d` (class in chainer.links), 499
 - `crf1d()` (in module chainer.functions), 205
 - `cross_covariance()` (in module chainer.functions), 207
 - `CUDAProfileHook` (class in chainer.function_hooks), 263
 - `cumsum()` (in module chainer.functions), 222
 - `CupyMemoryProfileHook` (class in chainer.function_hooks), 264
- ## D
- `data` (chainer.Parameter attribute), 130
 - `data` (chainer.Variable attribute), 122
 - `data` (chainer.variable.VariableNode attribute), 132
 - `DatasetMixin` (class in chainer.dataset), 690
 - `debug_print()` (chainer.Parameter method), 125
 - `debug_print()` (chainer.Variable method), 118
 - `DebugMode` (class in chainer), 745
 - `Deconvolution2D` (class in chainer.links), 300
 - `deconvolution_2d()` (in module chainer.functions), 178
 - `deconvolution_nd()` (in module chainer.functions), 180
 - `DeconvolutionND` (class in chainer.links), 306
 - `decov()` (in module chainer.functions), 207
 - `default_name` (chainer.training.Extension attribute), 663
 - `default_name` (chainer.training.extensions.Evaluator attribute), 666
 - `default_name` (chainer.training.extensions.ExponentialShift attribute), 673
 - `default_name` (chainer.training.extensions.FailOnNonNumber attribute), 669
 - `default_name` (chainer.training.extensions.LinearShift attribute), 675
 - `default_name` (chainer.training.extensions.LogReport attribute), 679
 - `default_name` (chainer.training.extensions.MicroAverage attribute), 668
 - `default_name` (chainer.training.extensions.ParameterStatistics attribute), 671
 - `default_name` (chainer.training.extensions.PlotReport attribute), 681
 - `default_name` (chainer.training.extensions.PrintReport attribute), 676
 - `default_name` (chainer.training.extensions.ProgressBar attribute), 677
 - `default_name` (chainer.training.extensions.VariableStatisticsPlot attribute), 682

- default_statistics (chainer.training.extensions.ParameterStatistics attribute), 671
- delete_hook() (chainer.Function method), 249
- delete_hook() (chainer.FunctionAdapter method), 253
- delete_hook() (chainer.FunctionNode method), 258
- deleted() (chainer.function_hooks.CUDAProfileHook method), 263
- deleted() (chainer.function_hooks.CupyMemoryProfileHook method), 265
- deleted() (chainer.function_hooks.PrintHook method), 267
- deleted() (chainer.function_hooks.TimerHook method), 269
- deleted() (chainer.FunctionHook method), 271
- depth2space() (in module chainer.functions), 150
- depthwise_convolution_2d() (in module chainer.functions), 182
- DepthwiseConvolution2D (class in chainer.links), 311
- Deserializer (class in chainer), 725
- det() (in module chainer.functions), 222
- DictDataset (class in chainer.datasets), 697
- DictionarySerializer (class in chainer.serializers), 718
- DictSummary (class in chainer), 736
- dilated_convolution_2d() (in module chainer.functions), 183
- DilatedConvolution2D (class in chainer.links), 316
- disable_update() (chainer.Chain method), 597
- disable_update() (chainer.ChainList method), 602
- disable_update() (chainer.Link method), 591
- disable_update() (chainer.links.BatchNormalization method), 476
- disable_update() (chainer.links.BatchRenormalization method), 482
- disable_update() (chainer.links.Bias method), 275
- disable_update() (chainer.links.Bilinear method), 280
- disable_update() (chainer.links.BinaryHierarchicalSoftmax method), 492
- disable_update() (chainer.links.BlackOut method), 497
- disable_update() (chainer.links.caffe.CaffeFunction method), 585
- disable_update() (chainer.links.ChildSumTreeLSTM method), 286
- disable_update() (chainer.links.Classifier method), 533
- disable_update() (chainer.links.Convolution2D method), 292
- disable_update() (chainer.links.ConvolutionND method), 297
- disable_update() (chainer.links.CRF1d method), 501
- disable_update() (chainer.links.Deconvolution2D method), 303
- disable_update() (chainer.links.DeconvolutionND method), 308
- disable_update() (chainer.links.DepthwiseConvolution2D method), 313
- disable_update() (chainer.links.DilatedConvolution2D method), 319
- disable_update() (chainer.links.EmbedID method), 325
- disable_update() (chainer.links.GoogLeNet method), 546
- disable_update() (chainer.links.GRU method), 329
- disable_update() (chainer.links.Highway method), 335
- disable_update() (chainer.links.Inception method), 340
- disable_update() (chainer.links.InceptionBN method), 346
- disable_update() (chainer.links.LayerNormalization method), 487
- disable_update() (chainer.links.Linear method), 351
- disable_update() (chainer.links.LocalConvolution2D method), 356
- disable_update() (chainer.links.LSTM method), 363
- disable_update() (chainer.links.Maxout method), 522
- disable_update() (chainer.links.MLPConvolution2D method), 368
- disable_update() (chainer.links.model.vision.resnet.ResNetLayers method), 553
- disable_update() (chainer.links.NaryTreeLSTM method), 374
- disable_update() (chainer.links.NegativeSampling method), 527
- disable_update() (chainer.links.NStepBiGRU method), 379
- disable_update() (chainer.links.NStepBiLSTM method), 385
- disable_update() (chainer.links.NStepBiRNNReLU method), 391
- disable_update() (chainer.links.NStepBiRNNTanh method), 397
- disable_update() (chainer.links.NStepGRU method), 403
- disable_update() (chainer.links.NStepLSTM method), 409
- disable_update() (chainer.links.NStepRNNReLU method), 415
- disable_update() (chainer.links.NStepRNNTanh method), 421
- disable_update() (chainer.links.Parameter method), 426
- disable_update() (chainer.links.PReLU method), 511
- disable_update() (chainer.links.ResNet101Layers method), 566
- disable_update() (chainer.links.ResNet152Layers method), 573
- disable_update() (chainer.links.ResNet50Layers method), 560
- disable_update() (chainer.links.Scale method), 431
- disable_update() (chainer.links.SimplifiedDropconnect method), 507
- disable_update() (chainer.links.StatefulGRU method), 437
- disable_update() (chainer.links.StatefulMGU method), 448

- `disable_update()` (chainer.links.StatefulPeepholeLSTM method), 458
 - `disable_update()` (chainer.links.StatefulZoneoutLSTM method), 463
 - `disable_update()` (chainer.links.StatelessGRU method), 443
 - `disable_update()` (chainer.links.StatelessLSTM method), 469
 - `disable_update()` (chainer.links.StatelessMGU method), 453
 - `disable_update()` (chainer.links.Swish method), 517
 - `disable_update()` (chainer.links.TheanoFunction method), 580
 - `disable_update()` (chainer.links.VGG16Layers method), 539
 - `disable_update()` (chainer.Sequential method), 609
 - `dropout()` (in module chainer.functions), 232
 - `dstack()` (in module chainer.functions), 151
 - `dtype` (chainer.Parameter attribute), 130
 - `dtype` (chainer.Variable attribute), 122
 - `dump()` (chainer.computational_graph.ComputationalGraph method), 748
 - `dump_graph()`, 58
 - `dump_graph()` (in module chainer.training.extensions), 683
- ## E
- `EarlyStoppingTrigger` (class in chainer.training.triggers), 686
 - `elapsed_time` (chainer.training.Trainer attribute), 655
 - `elementwise()` (in module chainer.backends.cuda), 731
 - `elu()` (in module chainer.functions), 135
 - `embed_id()` (in module chainer.functions), 184
 - `EmbedID` (class in chainer.links), 322
 - `enable_update()` (chainer.Chain method), 597
 - `enable_update()` (chainer.ChainList method), 602
 - `enable_update()` (chainer.Link method), 591
 - `enable_update()` (chainer.links.BatchNormalization method), 476
 - `enable_update()` (chainer.links.BatchRenormalization method), 482
 - `enable_update()` (chainer.links.Bias method), 275
 - `enable_update()` (chainer.links.Bilinear method), 280
 - `enable_update()` (chainer.links.BinaryHierarchicalSoftmax method), 492
 - `enable_update()` (chainer.links.BlackOut method), 497
 - `enable_update()` (chainer.links.caffe.CaffeFunction method), 586
 - `enable_update()` (chainer.links.ChildSumTreeLSTM method), 286
 - `enable_update()` (chainer.links.Classifier method), 533
 - `enable_update()` (chainer.links.Convolution2D method), 292
 - `enable_update()` (chainer.links.ConvolutionND method), 297
 - `enable_update()` (chainer.links.CRF1d method), 501
 - `enable_update()` (chainer.links.Deconvolution2D method), 303
 - `enable_update()` (chainer.links.DeconvolutionND method), 308
 - `enable_update()` (chainer.links.DepthwiseConvolution2D method), 314
 - `enable_update()` (chainer.links.DilatedConvolution2D method), 320
 - `enable_update()` (chainer.links.EmbedID method), 325
 - `enable_update()` (chainer.links.GoogLeNet method), 546
 - `enable_update()` (chainer.links.GRU method), 330
 - `enable_update()` (chainer.links.Highway method), 335
 - `enable_update()` (chainer.links.Inception method), 340
 - `enable_update()` (chainer.links.InceptionBN method), 346
 - `enable_update()` (chainer.links.LayerNormalization method), 487
 - `enable_update()` (chainer.links.Linear method), 351
 - `enable_update()` (chainer.links.LocalConvolution2D method), 356
 - `enable_update()` (chainer.links.LSTM method), 363
 - `enable_update()` (chainer.links.Maxout method), 522
 - `enable_update()` (chainer.links.MLPConvolution2D method), 368
 - `enable_update()` (chainer.links.model.vision.resnet.ResNetLayers method), 553
 - `enable_update()` (chainer.links.NaryTreeLSTM method), 374
 - `enable_update()` (chainer.links.NegativeSampling method), 527
 - `enable_update()` (chainer.links.NStepBiGRU method), 379
 - `enable_update()` (chainer.links.NStepBiLSTM method), 385
 - `enable_update()` (chainer.links.NStepBiRNNReLU method), 391
 - `enable_update()` (chainer.links.NStepBiRNNTanh method), 397
 - `enable_update()` (chainer.links.NStepGRU method), 403
 - `enable_update()` (chainer.links.NStepLSTM method), 409
 - `enable_update()` (chainer.links.NStepRNNReLU method), 415
 - `enable_update()` (chainer.links.NStepRNNTanh method), 421
 - `enable_update()` (chainer.links.Parameter method), 426
 - `enable_update()` (chainer.links.PReLU method), 511
 - `enable_update()` (chainer.links.ResNet101Layers method), 567
 - `enable_update()` (chainer.links.ResNet152Layers method), 573
 - `enable_update()` (chainer.links.ResNet50Layers method), 560

- `enable_update()` (chainer.links.Scale method), 431
 - `enable_update()` (chainer.links.SimplifiedDropconnect method), 507
 - `enable_update()` (chainer.links.StatefulGRU method), 437
 - `enable_update()` (chainer.links.StatefulMGU method), 448
 - `enable_update()` (chainer.links.StatefulPeepholeLSTM method), 458
 - `enable_update()` (chainer.links.StatefulZoneoutLSTM method), 463
 - `enable_update()` (chainer.links.StatelessGRU method), 443
 - `enable_update()` (chainer.links.StatelessLSTM method), 469
 - `enable_update()` (chainer.links.StatelessMGU method), 453
 - `enable_update()` (chainer.links.Swish method), 517
 - `enable_update()` (chainer.links.TheanoFunction method), 580
 - `enable_update()` (chainer.links.VGG16Layers method), 539
 - `enable_update()` (chainer.Sequential method), 609
 - `epoch` (chainer.GradientMethod attribute), 639
 - `epoch` (chainer.Optimizer attribute), 635
 - `epoch` (chainer.optimizers.AdaDelta attribute), 614
 - `epoch` (chainer.optimizers.AdaGrad attribute), 616
 - `epoch` (chainer.optimizers.Adam attribute), 619
 - `epoch` (chainer.optimizers.MomentumSGD attribute), 621
 - `epoch` (chainer.optimizers.NesterovAG attribute), 623
 - `epoch` (chainer.optimizers.RMSprop attribute), 626
 - `epoch` (chainer.optimizers.RMSpropGraves attribute), 628
 - `epoch` (chainer.optimizers.SGD attribute), 630
 - `epoch` (chainer.optimizers.SMORMS3 attribute), 632
 - `epoch` (chainer.training.updaters.MultiprocessParallelUpdater attribute), 662
 - `epoch` (chainer.training.updaters.ParallelUpdater attribute), 660
 - `epoch` (chainer.training.updaters.StandardUpdater attribute), 658
 - `epoch_detail` (chainer.iterators.MultiprocessIterator attribute), 716
 - `epoch_detail` (chainer.iterators.MultithreadIterator attribute), 717
 - `epoch_detail` (chainer.iterators.SerialIterator attribute), 714
 - `epoch_detail` (chainer.training.updaters.MultiprocessParallelUpdater attribute), 662
 - `epoch_detail` (chainer.training.updaters.ParallelUpdater attribute), 660
 - `epoch_detail` (chainer.training.updaters.StandardUpdater attribute), 658
 - `eps` (chainer.optimizers.AdaDelta attribute), 614
 - `eps` (chainer.optimizers.AdaGrad attribute), 616
 - `eps` (chainer.optimizers.Adam attribute), 619
 - `eps` (chainer.optimizers.RMSprop attribute), 626
 - `eps` (chainer.optimizers.RMSpropGraves attribute), 628
 - `eps` (chainer.optimizers.SMORMS3 attribute), 632
 - `erf()` (in module chainer.functions), 223
 - `erfc()` (in module chainer.functions), 223
 - `eta` (chainer.optimizers.Adam attribute), 619
 - `eval()` (chainer.utils.type_check.Expr method), 752
 - `evaluate()` (chainer.training.extensions.Evaluator method), 665
 - `Evaluator`, 58
 - `Evaluator` (class in chainer.training.extensions), 664
 - `exp()` (in module chainer.functions), 223
 - `expand_dims()` (in module chainer.functions), 152
 - `expect()` (in module chainer.utils.type_check), 753
 - `experimental()` (in module chainer.utils), 737
 - `expm1()` (in module chainer.functions), 223
 - `ExponentialShift` (class in chainer.training.extensions), 672
 - `export()` (in module chainer.exporters.caffe), 749
 - `Expr` (class in chainer.utils.type_check), 751
 - `extend()` (chainer.Sequential method), 609
 - `extend()` (chainer.training.Trainer method), 655
 - `Extension` (class in chainer.training), 662
 - `extract()` (chainer.links.GoogLeNet method), 546
 - `extract()` (chainer.links.model.vision.resnet.ResNetLayers method), 553
 - `extract()` (chainer.links.ResNet101Layers method), 567
 - `extract()` (chainer.links.ResNet152Layers method), 573
 - `extract()` (chainer.links.ResNet50Layers method), 560
 - `extract()` (chainer.links.VGG16Layers method), 539
- ## F
- `f1_score()` (in module chainer.functions), 200
 - `FailOnNonNumber` (class in chainer.training.extensions), 668
 - `fft()` (in module chainer.functions), 223
 - `fill_value` (chainer.initializers.Constant attribute), 646
 - `fill_value` (chainer.initializers.NaN attribute), 648
 - `fill_value` (chainer.initializers.One attribute), 647
 - `fill_value` (chainer.initializers.Zero attribute), 647
 - `finalize()` (chainer.dataset.Iterator method), 692
 - `finalize()` (chainer.iterators.MultiprocessIterator method), 715
 - `finalize()` (chainer.iterators.MultithreadIterator method), 717
 - `finalize()` (chainer.iterators.SerialIterator method), 714
 - `finalize()` (chainer.training.Extension method), 663
 - `finalize()` (chainer.training.extensions.Evaluator method), 666
 - `finalize()` (chainer.training.extensions.ExponentialShift method), 673
 - `finalize()` (chainer.training.extensions.FailOnNonNumber method), 669

- finalize() (chainer.training.extensions.LinearShift method), 674
- finalize() (chainer.training.extensions.LogReport method), 678
- finalize() (chainer.training.extensions.MicroAverage method), 668
- finalize() (chainer.training.extensions.ParameterStatistics method), 671
- finalize() (chainer.training.extensions.PlotReport method), 680
- finalize() (chainer.training.extensions.PrintReport method), 676
- finalize() (chainer.training.extensions.ProgressBar method), 677
- finalize() (chainer.training.extensions.VariableStatisticsPlot method), 682
- finalize() (chainer.training.Updater method), 656
- finalize() (chainer.training.updaters.MultiprocessParallelUpdater method), 661
- finalize() (chainer.training.updaters.ParallelUpdater method), 659
- finalize() (chainer.training.updaters.StandardUpdater method), 658
- fix() (in module chainer.functions), 224
- fixed_batch_normalization() (in module chainer.functions), 237
- fixed_batch_renormalization() (in module chainer.functions), 238
- flatten() (chainer.Sequential method), 609
- flatten() (in module chainer.functions), 153
- flip() (in module chainer.functions), 154
- flip1r() (in module chainer.functions), 154
- flipud() (in module chainer.functions), 154
- floor() (in module chainer.functions), 224
- fmod() (in module chainer.functions), 224
- force_backprop_mode() (in module chainer), 261
- forget() (in module chainer.functions), 246
- forward() (chainer.Function method), 249
- forward() (chainer.FunctionAdapter method), 253
- forward() (chainer.FunctionNode method), 258
- forward_cpu() (chainer.Function method), 249
- forward_cpu() (chainer.FunctionAdapter method), 253
- forward_cpu() (chainer.FunctionNode method), 259
- forward_gpu() (chainer.Function method), 249
- forward_gpu() (chainer.FunctionAdapter method), 253
- forward_gpu() (chainer.FunctionNode method), 259
- forward_postprocess() (chainer.function_hooks.CUDAProfileHook method), 263
- forward_postprocess() (chainer.function_hooks.CupyMemoryProfileHook method), 265
- forward_postprocess() (chainer.function_hooks.PrintHook method), 267
- forward_postprocess() (chainer.function_hooks.TimerHook method), 269
- forward_postprocess() (chainer.FunctionHook method), 271
- forward_preprocess() (chainer.function_hooks.CUDAProfileHook method), 263
- forward_preprocess() (chainer.function_hooks.CupyMemoryProfileHook method), 265
- forward_preprocess() (chainer.function_hooks.PrintHook method), 267
- forward_preprocess() (chainer.function_hooks.TimerHook method), 269
- forward_preprocess() (chainer.FunctionHook method), 272
- function (chainer.FunctionAdapter attribute), 255
- Function (class in chainer), 247
- FunctionAdapter (class in chainer), 251
- FunctionHook (class in chainer), 270
- FunctionNode (class in chainer), 255
- functions (chainer.links.GoogLeNet attribute), 549
- functions (chainer.links.model.vision.resnet.ResNetLayers attribute), 557
- functions (chainer.links.ResNet101Layers attribute), 570
- functions (chainer.links.ResNet152Layers attribute), 576
- functions (chainer.links.ResNet50Layers attribute), 563
- functions (chainer.links.VGG16Layers attribute), 542
- ## G
- gamma (chainer.links.BatchNormalization attribute), 479
- gamma (chainer.links.BatchRenormalization attribute), 484
- gaussian() (in module chainer.functions), 233
- gaussian_kl_divergence() (in module chainer.functions), 208
- gaussian_nll() (in module chainer.functions), 208
- generate_array() (in module chainer.initializers), 652
- get_all_iterators() (chainer.training.extensions.Evaluator method), 666
- get_all_optimizers() (chainer.training.Updater method), 656
- get_all_optimizers() (chainer.training.updaters.MultiprocessParallelUpdater method), 661
- get_all_optimizers() (chainer.training.updaters.ParallelUpdater method), 659
- get_all_optimizers() (chainer.training.updaters.StandardUpdater method), 658
- get_all_targets() (chainer.training.extensions.Evaluator method), 666
- get_array_module() (in module chainer.backends.cuda), 731
- get_cifar10() (in module chainer.datasets), 711
- get_cifar100() (in module chainer.datasets), 711
- get_conv_outsize() (in module chainer.utils), 726
- get_cross_validation_datasets() (in module chainer.datasets), 701

- [get_cross_validation_datasets_random\(\)](#) (in module `chainer.datasets`), 702
[get_current_reporter\(\)](#) (in module `chainer`), 735
[get_dataset_root\(\)](#) (in module `chainer.dataset`), 695
[get_deconv_outsize\(\)](#) (in module `chainer.utils`), 727
[get_device\(\)](#) (in module `chainer.backends.cuda`), 728
[get_device_from_array\(\)](#) (in module `chainer.backends.cuda`), 729
[get_device_from_id\(\)](#) (in module `chainer.backends.cuda`), 728
[get_dict\(\)](#) (`chainer.optimizer.Hyperparameter` method), 637
[get_example\(\)](#) (`chainer.dataset.DatasetMixin` method), 691
[get_example\(\)](#) (`chainer.datasets.ConcatenatedDataset` method), 699
[get_example\(\)](#) (`chainer.datasets.ImageDataset` method), 705
[get_example\(\)](#) (`chainer.datasets.LabeledImageDataset` method), 709
[get_example\(\)](#) (`chainer.datasets.MultiZippedImageDataset` method), 707
[get_example\(\)](#) (`chainer.datasets.SubDataset` method), 700
[get_example\(\)](#) (`chainer.datasets.TransformDataset` method), 703
[get_example\(\)](#) (`chainer.datasets.ZippedImageDataset` method), 706
[get_extension\(\)](#) (`chainer.training.Trainer` method), 655
[get_fashion_mnist\(\)](#) (in module `chainer.datasets`), 710
[get_item\(\)](#) (in module `chainer.functions`), 154
[get_iterator\(\)](#) (`chainer.training.extensions.Evaluator` method), 666
[get_iterator\(\)](#) (`chainer.training.updaters.MultiprocessParallelUpdater` method), 661
[get_iterator\(\)](#) (`chainer.training.updaters.ParallelUpdater` method), 660
[get_iterator\(\)](#) (`chainer.training.updaters.StandardUpdater` method), 658
[get_max_workspace_size\(\)](#) (in module `chainer.backends.cuda`), 732
[get_mnist\(\)](#) (in module `chainer.datasets`), 710
[get_optimizer\(\)](#) (`chainer.training.Updater` method), 656
[get_optimizer\(\)](#) (`chainer.training.updaters.MultiprocessParallelUpdater` method), 661
[get_optimizer\(\)](#) (`chainer.training.updaters.ParallelUpdater` method), 660
[get_optimizer\(\)](#) (`chainer.training.updaters.StandardUpdater` method), 658
[get_ptb_words\(\)](#) (in module `chainer.datasets`), 712
[get_ptb_words_vocabulary\(\)](#) (in module `chainer.datasets`), 712
[get_retained_inputs\(\)](#) (`chainer.FunctionAdapter` method), 254
[get_retained_inputs\(\)](#) (`chainer.FunctionNode` method), 259
[get_retained_outputs\(\)](#) (`chainer.FunctionAdapter` method), 254
[get_retained_outputs\(\)](#) (`chainer.FunctionNode` method), 259
[get_svhn\(\)](#) (in module `chainer.datasets`), 712
[get_target\(\)](#) (`chainer.training.extensions.Evaluator` method), 666
[get_training_length\(\)](#) (`chainer.training.triggers.EarlyStoppingTrigger` method), 687
[get_training_length\(\)](#) (`chainer.training.triggers.IntervalTrigger` method), 687
[get_trigger\(\)](#) (in module `chainer.training`), 685
[get_variable\(\)](#) (`chainer.variable.VariableNode` method), 131
[get_variable_or_none\(\)](#) (`chainer.variable.VariableNode` method), 131
[global_config](#) (in module `chainer`), 742
[GlobalConfig](#) (class in `chainer.configuration`), 742
[GlorotNormal](#) (class in `chainer.initializers`), 649
[GlorotUniform](#) (class in `chainer.initializers`), 651
[GoogLeNet](#) (class in `chainer.links`), 543
[grad](#) (`chainer.Parameter` attribute), 130
[grad](#) (`chainer.Variable` attribute), 122
[grad](#) (`chainer.variable.VariableNode` attribute), 132
[grad\(\)](#) (in module `chainer`), 261
[grad_var](#) (`chainer.Parameter` attribute), 130
[grad_var](#) (`chainer.Variable` attribute), 122
[grad_var](#) (`chainer.variable.VariableNode` attribute), 132
[GradientClipping](#) (class in `chainer.optimizer_hooks`), 641
[GradientHardClipping](#) (class in `chainer.optimizer_hooks`), 642
[GradientLARS](#) (class in `chainer.optimizer_hooks`), 643
[GradientMethod](#) (class in `chainer`), 637
[GradientNoise](#) (class in `chainer.optimizer_hooks`), 642
[GRU](#) (class in `chainer.links`), 327
[gumbel_softmax\(\)](#) (in module `chainer.functions`), 234
- ## H
- [hard_sigmoid\(\)](#) (in module `chainer.functions`), 136
[HDF5Deserializer](#) (class in `chainer.serializers`), 722
[HDF5Serializer](#) (class in `chainer.serializers`), 721
[HeNormal](#) (class in `chainer.initializers`), 649
[HeUniform](#) (class in `chainer.initializers`), 652
[Highway](#) (class in `chainer.links`), 332
[hinge\(\)](#) (in module `chainer.functions`), 209
[hstack\(\)](#) (in module `chainer.functions`), 155
[huber_loss\(\)](#) (in module `chainer.functions`), 210
[Hyperparameter](#) (class in `chainer.optimizer`), 637
- ## I
- [Identity](#) (class in `chainer.initializers`), 646
[identity\(\)](#) (in module `chainer.functions`), 224
[ifft\(\)](#) (in module `chainer.functions`), 224

- `ignore_label` (chainer.links.EmbedID attribute), 327
- `im2col()` (in module chainer.functions), 156
- `ImageDataset` (class in chainer.datasets), 704
- `Inception` (class in chainer.links), 338
- `InceptionBN` (class in chainer.links), 343
- `index()` (chainer.Sequential method), 609
- `index()` (chainer.utils.type_check.TypeInfoTuple method), 753
- `init_hx()` (chainer.links.NStepBiGRU method), 380
- `init_hx()` (chainer.links.NStepBiLSTM method), 385
- `init_hx()` (chainer.links.NStepBiRNNReLU method), 391
- `init_hx()` (chainer.links.NStepBiRNNTanh method), 397
- `init_hx()` (chainer.links.NStepGRU method), 403
- `init_hx()` (chainer.links.NStepLSTM method), 409
- `init_hx()` (chainer.links.NStepRNNReLU method), 415
- `init_hx()` (chainer.links.NStepRNNTanh method), 421
- `init_scope()` (chainer.Chain method), 597
- `init_scope()` (chainer.ChainList method), 602
- `init_scope()` (chainer.Link method), 591
- `init_scope()` (chainer.links.BatchNormalization method), 477
- `init_scope()` (chainer.links.BatchRenormalization method), 482
- `init_scope()` (chainer.links.Bias method), 275
- `init_scope()` (chainer.links.Bilinear method), 280
- `init_scope()` (chainer.links.BinaryHierarchicalSoftmax method), 492
- `init_scope()` (chainer.links.BlackOut method), 497
- `init_scope()` (chainer.links.caffe.CaffeFunction method), 586
- `init_scope()` (chainer.links.ChildSumTreeLSTM method), 286
- `init_scope()` (chainer.links.Classifier method), 533
- `init_scope()` (chainer.links.Convolution2D method), 292
- `init_scope()` (chainer.links.ConvolutionND method), 297
- `init_scope()` (chainer.links.CRF1d method), 501
- `init_scope()` (chainer.links.Deconvolution2D method), 303
- `init_scope()` (chainer.links.DeconvolutionND method), 309
- `init_scope()` (chainer.links.DepthwiseConvolution2D method), 314
- `init_scope()` (chainer.links.DilatedConvolution2D method), 320
- `init_scope()` (chainer.links.EmbedID method), 325
- `init_scope()` (chainer.links.GoogLeNet method), 547
- `init_scope()` (chainer.links.GRU method), 330
- `init_scope()` (chainer.links.Highway method), 335
- `init_scope()` (chainer.links.Inception method), 340
- `init_scope()` (chainer.links.InceptionBN method), 346
- `init_scope()` (chainer.links.LayerNormalization method), 487
- `init_scope()` (chainer.links.Linear method), 351
- `init_scope()` (chainer.links.LocalConvolution2D method), 357
- `init_scope()` (chainer.links.LSTM method), 363
- `init_scope()` (chainer.links.Maxout method), 522
- `init_scope()` (chainer.links.MLPConvolution2D method), 368
- `init_scope()` (chainer.links.model.vision.resnet.ResNetLayers method), 554
- `init_scope()` (chainer.links.NaryTreeLSTM method), 374
- `init_scope()` (chainer.links.NegativeSampling method), 527
- `init_scope()` (chainer.links.NStepBiGRU method), 380
- `init_scope()` (chainer.links.NStepBiLSTM method), 385
- `init_scope()` (chainer.links.NStepBiRNNReLU method), 391
- `init_scope()` (chainer.links.NStepBiRNNTanh method), 397
- `init_scope()` (chainer.links.NStepGRU method), 403
- `init_scope()` (chainer.links.NStepLSTM method), 409
- `init_scope()` (chainer.links.NStepRNNReLU method), 415
- `init_scope()` (chainer.links.NStepRNNTanh method), 421
- `init_scope()` (chainer.links.Parameter method), 426
- `init_scope()` (chainer.links.PReLU method), 511
- `init_scope()` (chainer.links.ResNet101Layers method), 567
- `init_scope()` (chainer.links.ResNet152Layers method), 574
- `init_scope()` (chainer.links.ResNet50Layers method), 561
- `init_scope()` (chainer.links.Scale method), 431
- `init_scope()` (chainer.links.SimplifiedDropconnect method), 507
- `init_scope()` (chainer.links.StatefulGRU method), 437
- `init_scope()` (chainer.links.StatefulMGU method), 448
- `init_scope()` (chainer.links.StatefulPeepholeLSTM method), 459
- `init_scope()` (chainer.links.StatefulZoneoutLSTM method), 463
- `init_scope()` (chainer.links.StatelessGRU method), 443
- `init_scope()` (chainer.links.StatelessLSTM method), 470
- `init_scope()` (chainer.links.StatelessMGU method), 453
- `init_scope()` (chainer.links.Swish method), 517
- `init_scope()` (chainer.links.TheanoFunction method), 580
- `init_scope()` (chainer.links.VGG16Layers method), 540
- `init_scope()` (chainer.Sequential method), 609
- `init_state()` (chainer.UpdateRule method), 636
- `initialize()` (chainer.Parameter method), 126
- `initialize()` (chainer.training.Extension method), 663
- `initialize()` (chainer.training.extensions.Evaluator method), 666
- `initialize()` (chainer.training.extensions.ExponentialShift method), 673
- `initialize()` (chainer.training.extensions.FailOnNonNumber method), 669

- [initialize\(\)](#) (chainer.training.extensions.LinearShift method), 674
[initialize\(\)](#) (chainer.training.extensions.LogReport method), 678
[initialize\(\)](#) (chainer.training.extensions.MicroAverage method), 668
[initialize\(\)](#) (chainer.training.extensions.ParameterStatistics method), 671
[initialize\(\)](#) (chainer.training.extensions.PlotReport method), 680
[initialize\(\)](#) (chainer.training.extensions.PrintReport method), 676
[initialize\(\)](#) (chainer.training.extensions.ProgressBar method), 677
[initialize\(\)](#) (chainer.training.extensions.VariableStatisticsPlot method), 682
[initializer](#) (chainer.Parameter attribute), 130
[Initializer](#) (class in chainer), 645
[inputs](#) (chainer.Function attribute), 250
[inputs](#) (chainer.FunctionAdapter attribute), 255
[inputs](#) (chainer.FunctionNode attribute), 260
[insert\(\)](#) (chainer.Sequential method), 610
[IntervalTrigger](#) (class in chainer.training.triggers), 687
[inv\(\)](#) (in module chainer.functions), 225
[is_debug\(\)](#) (in module chainer), 745
[is_new_epoch](#) (chainer.training.updaters.MultiprocessParallelUpdater attribute), 662
[is_new_epoch](#) (chainer.training.updaters.ParallelUpdater attribute), 660
[is_new_epoch](#) (chainer.training.updaters.StandardUpdater attribute), 658
[Iterator](#) (class in chainer.dataset), 691
- ## L
- [label](#) (chainer.Function attribute), 250
[label](#) (chainer.FunctionAdapter attribute), 255
[label](#) (chainer.FunctionNode attribute), 260
[label](#) (chainer.Parameter attribute), 130
[label](#) (chainer.Variable attribute), 122
[label](#) (chainer.variable.VariableNode attribute), 132
[LabeledImageDataset](#) (class in chainer.datasets), 708
[Lasso](#) (class in chainer.optimizer_hooks), 640
[layer_normalization\(\)](#) (in module chainer.functions), 238
[LayerNormalization](#) (class in chainer.links), 484
[lazy_grad_sum](#) (chainer.FunctionAdapter attribute), 255
[lazy_grad_sum](#) (chainer.FunctionNode attribute), 260
[leaky_relu\(\)](#) (in module chainer.functions), 136
[LeCunNormal](#) (class in chainer.initializers), 648
[LeCunUniform](#) (class in chainer.initializers), 651
[Linear](#) (class in chainer.links), 348
[linear\(\)](#) (in module chainer.functions), 185
[linear_interpolate\(\)](#) (in module chainer.functions), 225
[LinearShift](#) (class in chainer.training.extensions), 674
[Link](#) (class in chainer), 588
[links\(\)](#) (chainer.Chain method), 598
[links\(\)](#) (chainer.ChainList method), 602
[links\(\)](#) (chainer.Link method), 592
[links\(\)](#) (chainer.links.BatchNormalization method), 477
[links\(\)](#) (chainer.links.BatchRenormalization method), 482
[links\(\)](#) (chainer.links.Bias method), 276
[links\(\)](#) (chainer.links.Bilinear method), 281
[links\(\)](#) (chainer.links.BinaryHierarchicalSoftmax method), 492
[links\(\)](#) (chainer.links.BlackOut method), 497
[links\(\)](#) (chainer.links.caffe.CaffeFunction method), 586
[links\(\)](#) (chainer.links.ChildSumTreeLSTM method), 286
[links\(\)](#) (chainer.links.Classifier method), 533
[links\(\)](#) (chainer.links.Convolution2D method), 292
[links\(\)](#) (chainer.links.ConvolutionND method), 298
[links\(\)](#) (chainer.links.CRF1d method), 502
[links\(\)](#) (chainer.links.Deconvolution2D method), 304
[links\(\)](#) (chainer.links.DeconvolutionND method), 309
[links\(\)](#) (chainer.links.DepthwiseConvolution2D method), 314
[links\(\)](#) (chainer.links.DilatedConvolution2D method), 320
[links\(\)](#) (chainer.links.EmbedID method), 325
[links\(\)](#) (chainer.links.GoogLeNet method), 547
[links\(\)](#) (chainer.links.GRU method), 330
[links\(\)](#) (chainer.links.Highway method), 335
[links\(\)](#) (chainer.links.Inception method), 341
[links\(\)](#) (chainer.links.InceptionBN method), 346
[links\(\)](#) (chainer.links.LayerNormalization method), 487
[links\(\)](#) (chainer.links.Linear method), 352
[links\(\)](#) (chainer.links.LocalConvolution2D method), 357
[links\(\)](#) (chainer.links.LSTM method), 363
[links\(\)](#) (chainer.links.Maxout method), 522
[links\(\)](#) (chainer.links.MLPConvolution2D method), 369
[links\(\)](#) (chainer.links.model.vision.resnet.ResNetLayers method), 554
[links\(\)](#) (chainer.links.NaryTreeLSTM method), 374
[links\(\)](#) (chainer.links.NegativeSampling method), 527
[links\(\)](#) (chainer.links.NStepBiGRU method), 380
[links\(\)](#) (chainer.links.NStepBiLSTM method), 386
[links\(\)](#) (chainer.links.NStepBiRNNTanh method), 392
[links\(\)](#) (chainer.links.NStepBiRNNTanh method), 398
[links\(\)](#) (chainer.links.NStepGRU method), 404
[links\(\)](#) (chainer.links.NStepLSTM method), 410
[links\(\)](#) (chainer.links.NStepRNNTanh method), 416
[links\(\)](#) (chainer.links.NStepRNNTanh method), 422
[links\(\)](#) (chainer.links.Parameter method), 426
[links\(\)](#) (chainer.links.PReLU method), 512
[links\(\)](#) (chainer.links.ResNet101Layers method), 568
[links\(\)](#) (chainer.links.ResNet152Layers method), 574
[links\(\)](#) (chainer.links.ResNet50Layers method), 561
[links\(\)](#) (chainer.links.Scale method), 432
[links\(\)](#) (chainer.links.SimplifiedDropconnect method), 507

`links()` (chainer.links.StatefulGRU method), 438
`links()` (chainer.links.StatefulMGU method), 449
`links()` (chainer.links.StatefulPeepholeLSTM method), 459
`links()` (chainer.links.StatefulZoneoutLSTM method), 464
`links()` (chainer.links.StatelessGRU method), 444
`links()` (chainer.links.StatelessLSTM method), 470
`links()` (chainer.links.StatelessMGU method), 453
`links()` (chainer.links.Swish method), 517
`links()` (chainer.links.TheanoFunction method), 580
`links()` (chainer.links.VGG16Layers method), 540
`links()` (chainer.Sequential method), 610
`load()` (chainer.Deserializer method), 726
`load()` (chainer.serializers.HDF5Deserializer method), 723
`load()` (chainer.serializers.NpzDeserializer method), 720
`load_hdf5()` (in module chainer.serializers), 723
`load_npz()` (in module chainer.serializers), 720
`local_convolution_2d()` (in module chainer.functions), 186
`local_function_hooks` (chainer.Function attribute), 250
`local_function_hooks` (chainer.FunctionAdapter attribute), 255
`local_function_hooks` (chainer.FunctionNode attribute), 260
`local_response_normalization()` (in module chainer.functions), 238
`LocalConfig` (class in chainer.configuration), 743
`LocalConvolution2D` (class in chainer.links), 354
`log` (chainer.training.extensions.LogReport attribute), 679
`log()` (in module chainer.functions), 225
`log10()` (in module chainer.functions), 225
`log1p()` (in module chainer.functions), 226
`log2()` (in module chainer.functions), 226
`log_softmax()` (in module chainer.functions), 137
`LogReport`, 57
`LogReport` (class in chainer.training.extensions), 677
`logsumexp()` (in module chainer.functions), 226
`lr` (chainer.optimizers.AdaGrad attribute), 616
`lr` (chainer.optimizers.Adam attribute), 619
`lr` (chainer.optimizers.MomentumSGD attribute), 621
`lr` (chainer.optimizers.NesterovAG attribute), 623
`lr` (chainer.optimizers.RMSprop attribute), 626
`lr` (chainer.optimizers.RMSpropGraves attribute), 628
`lr` (chainer.optimizers.SGD attribute), 630
`lr` (chainer.optimizers.SMORMS3 attribute), 632
`LSTM` (class in chainer.links), 359
`lstm()` (in module chainer.functions), 138

M

`make_extension()` (in module chainer.training), 663
`make_statistics()` (chainer.DictSummary method), 737
`make_statistics()` (chainer.Summary method), 736

`ManualScheduleTrigger` (class in chainer.training.triggers), 687
`matmul()` (in module chainer.functions), 226
`max()` (in module chainer.functions), 227
`max_pooling_2d()` (in module chainer.functions), 240
`max_pooling_nd()` (in module chainer.functions), 241
`maximum()` (in module chainer.functions), 227
`Maxout` (class in chainer.links), 519
`maxout()` (in module chainer.functions), 139
`MaxValueTrigger` (class in chainer.training.triggers), 688
`mean()` (in module chainer.functions), 227
`mean_absolute_error()` (in module chainer.functions), 211
`mean_squared_error()` (in module chainer.functions), 211
`memoize()` (in module chainer.backends.cuda), 730
`MicroAverage` (class in chainer.training.extensions), 667
`min()` (in module chainer.functions), 228
`minimum()` (in module chainer.functions), 228
`MinValueTrigger` (class in chainer.training.triggers), 688
`MLPConvolution2D` (class in chainer.links), 365
`momentum` (chainer.optimizers.MomentumSGD attribute), 621
`momentum` (chainer.optimizers.NesterovAG attribute), 623
`momentum` (chainer.optimizers.RMSpropGraves attribute), 628
`MomentumSGD` (class in chainer.optimizers), 619
`MultiprocessIterator` (class in chainer.iterators), 715
`MultiprocessParallelUpdater` (class in chainer.training.updaters), 660
`MultithreadIterator` (class in chainer.iterators), 716
`MultiZippedImageDataset` (class in chainer.datasets), 706

N

`n_cells` (chainer.links.NStepBiGRU attribute), 382
`n_cells` (chainer.links.NStepBiLSTM attribute), 388
`n_cells` (chainer.links.NStepBiRNReLU attribute), 394
`n_cells` (chainer.links.NStepBiRNNTanh attribute), 400
`n_cells` (chainer.links.NStepGRU attribute), 406
`n_cells` (chainer.links.NStepLSTM attribute), 412
`n_cells` (chainer.links.NStepRNReLU attribute), 418
`n_cells` (chainer.links.NStepRNNTanh attribute), 424
`n_step_bigru()` (in module chainer.functions), 187
`n_step_bilstm()` (in module chainer.functions), 188
`n_step_birnn()` (in module chainer.functions), 191
`n_step_gru()` (in module chainer.functions), 193
`n_step_lstm()` (in module chainer.functions), 194
`n_step_rnn()` (in module chainer.functions), 196
`n_weights` (chainer.links.NStepBiGRU attribute), 382
`n_weights` (chainer.links.NStepBiLSTM attribute), 388
`n_weights` (chainer.links.NStepBiRNReLU attribute), 394
`n_weights` (chainer.links.NStepBiRNNTanh attribute), 400
`n_weights` (chainer.links.NStepGRU attribute), 406

- `n_weights` (chainer.links.NStepLSTM attribute), 412
- `n_weights` (chainer.links.NStepRNReLU attribute), 418
- `n_weights` (chainer.links.NStepRNNTanh attribute), 424
- `name` (chainer.function_hooks.CUDAProfileHook attribute), 264
- `name` (chainer.function_hooks.CupyMemoryProfileHook attribute), 266
- `name` (chainer.function_hooks.PrintHook attribute), 268
- `name` (chainer.function_hooks.TimerHook attribute), 269
- `name` (chainer.FunctionHook attribute), 272
- `name` (chainer.optimizer_hooks.GradientClipping attribute), 642
- `name` (chainer.optimizer_hooks.GradientHardClipping attribute), 642
- `name` (chainer.optimizer_hooks.GradientLARS attribute), 644
- `name` (chainer.optimizer_hooks.GradientNoise attribute), 643
- `name` (chainer.optimizer_hooks.Lasso attribute), 641
- `name` (chainer.optimizer_hooks.WeightDecay attribute), 640
- `name` (chainer.Parameter attribute), 130
- `name` (chainer.training.extensions.Evaluator attribute), 666
- `name` (chainer.Variable attribute), 122
- `namedlinks()` (chainer.Chain method), 598
- `namedlinks()` (chainer.ChainList method), 603
- `namedlinks()` (chainer.Link method), 592
- `namedlinks()` (chainer.links.BatchNormalization method), 477
- `namedlinks()` (chainer.links.BatchRenormalization method), 482
- `namedlinks()` (chainer.links.Bias method), 276
- `namedlinks()` (chainer.links.Bilinear method), 281
- `namedlinks()` (chainer.links.BinaryHierarchicalSoftmax method), 492
- `namedlinks()` (chainer.links.BlackOut method), 497
- `namedlinks()` (chainer.links.caffe.CaffeFunction method), 586
- `namedlinks()` (chainer.links.ChildSumTreeLSTM method), 286
- `namedlinks()` (chainer.links.Classifier method), 533
- `namedlinks()` (chainer.links.Convolution2D method), 292
- `namedlinks()` (chainer.links.ConvolutionND method), 298
- `namedlinks()` (chainer.links.CRF1d method), 502
- `namedlinks()` (chainer.links.Deconvolution2D method), 304
- `namedlinks()` (chainer.links.DeconvolutionND method), 309
- `namedlinks()` (chainer.links.DepthwiseConvolution2D method), 314
- `namedlinks()` (chainer.links.DilatedConvolution2D method), 320
- `namedlinks()` (chainer.links.EmbedID method), 325
- `namedlinks()` (chainer.links.GoogLeNet method), 547
- `namedlinks()` (chainer.links.GRU method), 330
- `namedlinks()` (chainer.links.Highway method), 335
- `namedlinks()` (chainer.links.Inception method), 341
- `namedlinks()` (chainer.links.InceptionBN method), 346
- `namedlinks()` (chainer.links.LayerNormalization method), 487
- `namedlinks()` (chainer.links.Linear method), 352
- `namedlinks()` (chainer.links.LocalConvolution2D method), 357
- `namedlinks()` (chainer.links.LSTM method), 363
- `namedlinks()` (chainer.links.Maxout method), 523
- `namedlinks()` (chainer.links.MLPConvolution2D method), 369
- `namedlinks()` (chainer.links.model.vision.resnet.ResNetLayers method), 554
- `namedlinks()` (chainer.links.NaryTreeLSTM method), 374
- `namedlinks()` (chainer.links.NegativeSampling method), 527
- `namedlinks()` (chainer.links.NStepBiGRU method), 380
- `namedlinks()` (chainer.links.NStepBiLSTM method), 386
- `namedlinks()` (chainer.links.NStepBiRNReLU method), 392
- `namedlinks()` (chainer.links.NStepBiRNNTanh method), 398
- `namedlinks()` (chainer.links.NStepGRU method), 404
- `namedlinks()` (chainer.links.NStepLSTM method), 410
- `namedlinks()` (chainer.links.NStepRNReLU method), 416
- `namedlinks()` (chainer.links.NStepRNNTanh method), 422
- `namedlinks()` (chainer.links.Parameter method), 427
- `namedlinks()` (chainer.links.PReLU method), 512
- `namedlinks()` (chainer.links.ResNet101Layers method), 568
- `namedlinks()` (chainer.links.ResNet152Layers method), 574
- `namedlinks()` (chainer.links.ResNet50Layers method), 561
- `namedlinks()` (chainer.links.Scale method), 432
- `namedlinks()` (chainer.links.SimplifiedDropconnect method), 507
- `namedlinks()` (chainer.links.StatefulGRU method), 438
- `namedlinks()` (chainer.links.StatefulMGU method), 449
- `namedlinks()` (chainer.links.StatefulPeepholeLSTM method), 459
- `namedlinks()` (chainer.links.StatefulZoneoutLSTM method), 464
- `namedlinks()` (chainer.links.StatelessGRU method), 444
- `namedlinks()` (chainer.links.StatelessLSTM method), 470
- `namedlinks()` (chainer.links.StatelessMGU method), 453
- `namedlinks()` (chainer.links.Swish method), 517

`namedlinks()` (chainer.links.TheanoFunction method), 580

`namedlinks()` (chainer.links.VGG16Layers method), 540

`namedlinks()` (chainer.Sequential method), 610

`namedparams()` (chainer.Chain method), 598

`namedparams()` (chainer.ChainList method), 603

`namedparams()` (chainer.Link method), 592

`namedparams()` (chainer.links.BatchNormalization method), 477

`namedparams()` (chainer.links.BatchRenormalization method), 482

`namedparams()` (chainer.links.Bias method), 276

`namedparams()` (chainer.links.Bilinear method), 281

`namedparams()` (chainer.links.BinaryHierarchicalSoftmax method), 493

`namedparams()` (chainer.links.BlackOut method), 497

`namedparams()` (chainer.links.caffe.CaffeFunction method), 586

`namedparams()` (chainer.links.ChildSumTreeLSTM method), 286

`namedparams()` (chainer.links.Classifier method), 533

`namedparams()` (chainer.links.Convolution2D method), 293

`namedparams()` (chainer.links.ConvolutionND method), 298

`namedparams()` (chainer.links.CRF1d method), 502

`namedparams()` (chainer.links.Deconvolution2D method), 304

`namedparams()` (chainer.links.DeconvolutionND method), 309

`namedparams()` (chainer.links.DepthwiseConvolution2D method), 314

`namedparams()` (chainer.links.DilatedConvolution2D method), 320

`namedparams()` (chainer.links.EmbedID method), 325

`namedparams()` (chainer.links.GoogLeNet method), 547

`namedparams()` (chainer.links.GRU method), 330

`namedparams()` (chainer.links.Highway method), 336

`namedparams()` (chainer.links.Inception method), 341

`namedparams()` (chainer.links.InceptionBN method), 346

`namedparams()` (chainer.links.LayerNormalization method), 487

`namedparams()` (chainer.links.Linear method), 352

`namedparams()` (chainer.links.LocalConvolution2D method), 357

`namedparams()` (chainer.links.LSTM method), 363

`namedparams()` (chainer.links.Maxout method), 523

`namedparams()` (chainer.links.MLPConvolution2D method), 369

`namedparams()` (chainer.links.model.vision.resnet.ResNetLayer method), 555

`namedparams()` (chainer.links.NaryTreeLSTM method), 374

`namedparams()` (chainer.links.NegativeSampling method), 527

`namedparams()` (chainer.links.NStepBiGRU method), 380

`namedparams()` (chainer.links.NStepBiLSTM method), 386

`namedparams()` (chainer.links.NStepBiRNNReLU method), 392

`namedparams()` (chainer.links.NStepBiRNNTanh method), 398

`namedparams()` (chainer.links.NStepGRU method), 404

`namedparams()` (chainer.links.NStepLSTM method), 410

`namedparams()` (chainer.links.NStepRNReLU method), 416

`namedparams()` (chainer.links.NStepRNNTanh method), 422

`namedparams()` (chainer.links.Parameter method), 427

`namedparams()` (chainer.links.PReLU method), 512

`namedparams()` (chainer.links.ResNet101Layers method), 568

`namedparams()` (chainer.links.ResNet152Layers method), 575

`namedparams()` (chainer.links.ResNet50Layers method), 561

`namedparams()` (chainer.links.Scale method), 432

`namedparams()` (chainer.links.SimplifiedDropconnect method), 507

`namedparams()` (chainer.links.StatefulGRU method), 438

`namedparams()` (chainer.links.StatefulMGU method), 449

`namedparams()` (chainer.links.StatefulPeepholeLSTM method), 459

`namedparams()` (chainer.links.StatefulZoneoutLSTM method), 464

`namedparams()` (chainer.links.StatelessGRU method), 444

`namedparams()` (chainer.links.StatelessLSTM method), 470

`namedparams()` (chainer.links.StatelessMGU method), 454

`namedparams()` (chainer.links.Swish method), 517

`namedparams()` (chainer.links.TheanoFunction method), 580

`namedparams()` (chainer.links.VGG16Layers method), 540

`namedparams()` (chainer.Sequential method), 610

`NaN` (class in chainer.initializers), 647

`NaryTreeLSTM` (class in chainer.links), 371

`ndim` (chainer.Parameter attribute), 130

`ndim` (chainer.Variable attribute), 122

`negative_sampling()` (in module chainer.functions), 211

`NegativeSampling` (class in chainer.links), 525

`NesterovAG` (class in chainer.optimizers), 622

`new_epoch()` (chainer.GradientMethod method), 638

`new_epoch()` (chainer.Optimizer method), 633

- `new_epoch()` (chainer.optimizers.AdaDelta method), 613
 - `new_epoch()` (chainer.optimizers.AdaGrad method), 615
 - `new_epoch()` (chainer.optimizers.Adam method), 618
 - `new_epoch()` (chainer.optimizers.MomentumSGD method), 620
 - `new_epoch()` (chainer.optimizers.NesterovAG method), 622
 - `new_epoch()` (chainer.optimizers.RMSprop method), 624
 - `new_epoch()` (chainer.optimizers.RMSpropGraves method), 627
 - `new_epoch()` (chainer.optimizers.SGD method), 629
 - `new_epoch()` (chainer.optimizers.SMORMS3 method), 631
 - `next()` (chainer.dataset.Iterator method), 692
 - `next()` (chainer.iterators.MultiprocessIterator method), 715
 - `next()` (chainer.iterators.MultithreadIterator method), 717
 - `next()` (chainer.iterators.SerialIterator method), 714
 - `no_backprop_mode()` (in module chainer), 261
 - `node` (chainer.Function attribute), 250
 - `node` (chainer.Parameter attribute), 130
 - `node` (chainer.Variable attribute), 122
 - `Normal` (class in chainer.initializers), 648
 - `normalize()` (in module chainer.functions), 239
 - `NpzDeserializer` (class in chainer.serializers), 719
 - `NStepBiGRU` (class in chainer.links), 376
 - `NStepBiLSTM` (class in chainer.links), 382
 - `NStepBiRNReLU` (class in chainer.links), 388
 - `NStepBiRNNTanh` (class in chainer.links), 394
 - `NStepGRU` (class in chainer.links), 400
 - `NStepLSTM` (class in chainer.links), 406
 - `NStepRNReLU` (class in chainer.links), 412
 - `NStepRNNTanh` (class in chainer.links), 418
 - `numerical_grad()` (in module chainer.gradient_check), 756
- ## O
- `observe_lr()` (in module chainer.training.extensions), 672
 - `observe_value()` (in module chainer.training.extensions), 672
 - `One` (class in chainer.initializers), 647
 - `Optimizer` (class in chainer), 633
 - `Orthogonal` (class in chainer.initializers), 650
 - `output_data` (chainer.Function attribute), 251
 - `output_data` (chainer.FunctionAdapter attribute), 255
 - `output_data` (chainer.FunctionNode attribute), 260
 - `outputs` (chainer.Function attribute), 251
 - `outputs` (chainer.FunctionAdapter attribute), 255
 - `outputs` (chainer.FunctionNode attribute), 260
- ## P
- `pad()` (in module chainer.functions), 157
 - `pad_sequence()` (in module chainer.functions), 158
 - `ParallelUpdater` (class in chainer.training.updaters), 659
 - `Parameter` (class in chainer), 123
 - `Parameter` (class in chainer.links), 424
 - `ParameterStatistics` (class in chainer.training.extensions), 670
 - `params()` (chainer.Chain method), 598
 - `params()` (chainer.ChainList method), 603
 - `params()` (chainer.Link method), 592
 - `params()` (chainer.links.BatchNormalization method), 477
 - `params()` (chainer.links.BatchRenormalization method), 482
 - `params()` (chainer.links.Bias method), 276
 - `params()` (chainer.links.Bilinear method), 281
 - `params()` (chainer.links.BinaryHierarchicalSoftmax method), 493
 - `params()` (chainer.links.BlackOut method), 497
 - `params()` (chainer.links.caffe.CaffeFunction method), 586
 - `params()` (chainer.links.ChildSumTreeLSTM method), 286
 - `params()` (chainer.links.Classifier method), 534
 - `params()` (chainer.links.Convolution2D method), 293
 - `params()` (chainer.links.ConvolutionND method), 298
 - `params()` (chainer.links.CRF1d method), 502
 - `params()` (chainer.links.Deconvolution2D method), 304
 - `params()` (chainer.links.DeconvolutionND method), 309
 - `params()` (chainer.links.DepthwiseConvolution2D method), 314
 - `params()` (chainer.links.DilatedConvolution2D method), 320
 - `params()` (chainer.links.EmbedID method), 325
 - `params()` (chainer.links.GoogLeNet method), 547
 - `params()` (chainer.links.GRU method), 330
 - `params()` (chainer.links.Highway method), 336
 - `params()` (chainer.links.Inception method), 341
 - `params()` (chainer.links.InceptionBN method), 346
 - `params()` (chainer.links.LayerNormalization method), 488
 - `params()` (chainer.links.Linear method), 352
 - `params()` (chainer.links.LocalConvolution2D method), 357
 - `params()` (chainer.links.LSTM method), 363
 - `params()` (chainer.links.Maxout method), 523
 - `params()` (chainer.links.MLPConvolution2D method), 369
 - `params()` (chainer.links.model.vision.resnet.ResNetLayers method), 555
 - `params()` (chainer.links.NaryTreeLSTM method), 375
 - `params()` (chainer.links.NegativeSampling method), 528
 - `params()` (chainer.links.NStepBiGRU method), 380
 - `params()` (chainer.links.NStepBiLSTM method), 386
 - `params()` (chainer.links.NStepBiRNReLU method), 392
 - `params()` (chainer.links.NStepBiRNNTanh method), 398
 - `params()` (chainer.links.NStepGRU method), 404
 - `params()` (chainer.links.NStepLSTM method), 410

- `params()` (chainer.links.NStepRNReLU method), 416
 - `params()` (chainer.links.NStepRNNTanh method), 422
 - `params()` (chainer.links.Parameter method), 427
 - `params()` (chainer.links.PReLU method), 512
 - `params()` (chainer.links.ResNet101Layers method), 568
 - `params()` (chainer.links.ResNet152Layers method), 575
 - `params()` (chainer.links.ResNet50Layers method), 561
 - `params()` (chainer.links.Scale method), 432
 - `params()` (chainer.links.SimplifiedDropconnect method), 507
 - `params()` (chainer.links.StatefulGRU method), 438
 - `params()` (chainer.links.StatefulMGU method), 449
 - `params()` (chainer.links.StatefulPeepholeLSTM method), 459
 - `params()` (chainer.links.StatefulZoneoutLSTM method), 464
 - `params()` (chainer.links.StatelessGRU method), 444
 - `params()` (chainer.links.StatelessLSTM method), 470
 - `params()` (chainer.links.StatelessMGU method), 454
 - `params()` (chainer.links.Swish method), 517
 - `params()` (chainer.links.TheanoFunction method), 580
 - `params()` (chainer.links.VGG16Layers method), 540
 - `params()` (chainer.Sequential method), 610
 - `parent` (chainer.optimizer.Hyperparameter attribute), 637
 - `permute()` (in module chainer.functions), 158
 - `PlotReport`, 58
 - `PlotReport` (class in chainer.training.extensions), 679
 - `pop()` (chainer.Sequential method), 610
 - `precision()` (in module chainer.functions), 200
 - `predict()` (chainer.links.GoogLeNet method), 548
 - `predict()` (chainer.links.model.vision.resnet.ResNetLayers method), 555
 - `predict()` (chainer.links.ResNet101Layers method), 568
 - `predict()` (chainer.links.ResNet152Layers method), 575
 - `predict()` (chainer.links.ResNet50Layers method), 562
 - `predict()` (chainer.links.VGG16Layers method), 540
 - `PReLU` (class in chainer.links), 509
 - `prelu()` (in module chainer.functions), 140
 - `prepare()` (in module chainer.links.model.vision.googlenet), 550
 - `prepare()` (in module chainer.links.model.vision.resnet), 577
 - `prepare()` (in module chainer.links.model.vision.vgg), 542
 - `previous_epoch_detail` (chainer.iterators.MultiprocessIterator attribute), 716
 - `previous_epoch_detail` (chainer.iterators.MultithreadIterator attribute), 717
 - `previous_epoch_detail` (chainer.iterators.SerialIterator attribute), 714
 - `previous_epoch_detail` (chainer.training.updaters.MultiprocessParallelUpdater attribute), 662
 - `previous_epoch_detail` (chainer.training.updaters.ParallelUpdater attribute), 660
 - `previous_epoch_detail` (chainer.training.updaters.StandardUpdater attribute), 658
 - `print_report()` (chainer.function_hooks.CupyMemoryProfileHook method), 265
 - `print_report()` (chainer.function_hooks.TimerHook method), 269
 - `PrintHook` (class in chainer.function_hooks), 266
 - `PrintReport`, 58
 - `PrintReport` (class in chainer.training.extensions), 675
 - `priority` (chainer.training.Extension attribute), 663
 - `priority` (chainer.training.extensions.Evaluator attribute), 666
 - `priority` (chainer.training.extensions.ExponentialShift attribute), 673
 - `priority` (chainer.training.extensions.FailOnNonNumber attribute), 669
 - `priority` (chainer.training.extensions.LinearShift attribute), 675
 - `priority` (chainer.training.extensions.LogReport attribute), 679
 - `priority` (chainer.training.extensions.MicroAverage attribute), 668
 - `priority` (chainer.training.extensions.ParameterStatistics attribute), 671
 - `priority` (chainer.training.extensions.PlotReport attribute), 681
 - `priority` (chainer.training.extensions.PrintReport attribute), 676
 - `priority` (chainer.training.extensions.ProgressBar attribute), 677
 - `priority` (chainer.training.extensions.VariableStatisticsPlot attribute), 683
 - `prod()` (in module chainer.functions), 228
 - `ProgressBar` (class in chainer.training.extensions), 676
- ## R
- `r2_score()` (in module chainer.functions), 200
 - `rank` (chainer.Function attribute), 251
 - `rank` (chainer.FunctionAdapter attribute), 255
 - `rank` (chainer.FunctionNode attribute), 260
 - `rank` (chainer.Parameter attribute), 130
 - `rank` (chainer.Variable attribute), 122
 - `rank` (chainer.variable.VariableNode attribute), 133
 - `reallocate_cleared_grads()` (chainer.GradientMethod method), 638
 - `reallocate_cleared_grads()` (chainer.optimizers.AdaDelta method), 613
 - `reallocate_cleared_grads()` (chainer.optimizers.AdaGrad method), 615
 - `reallocate_cleared_grads()` (chainer.optimizers.Adam method), 618
 - `reallocate_cleared_grads()` (chainer.optimizers.MomentumSGD method), 620

[reallocate_cleared_grads\(\)](#)
 (chainer.optimizers.NesterovAG method), [622](#)
[reallocate_cleared_grads\(\)](#) (chainer.optimizers.RMSprop method), [625](#)
[reallocate_cleared_grads\(\)](#)
 (chainer.optimizers.RMSpropGraves method), [627](#)
[reallocate_cleared_grads\(\)](#) (chainer.optimizers.SGD method), [629](#)
[reallocate_cleared_grads\(\)](#)
 (chainer.optimizers.SMORMS3 method), [631](#)
[recall\(\)](#) (in module chainer.functions), [200](#)
[reduce\(\)](#) (in module chainer.backends.cuda), [731](#)
[register_persistent\(\)](#) (chainer.Chain method), [598](#)
[register_persistent\(\)](#) (chainer.ChainList method), [603](#)
[register_persistent\(\)](#) (chainer.Link method), [592](#)
[register_persistent\(\)](#) (chainer.links.BatchNormalization method), [477](#)
[register_persistent\(\)](#) (chainer.links.BatchRenormalization method), [483](#)
[register_persistent\(\)](#) (chainer.links.Bias method), [276](#)
[register_persistent\(\)](#) (chainer.links.Bilinear method), [281](#)
[register_persistent\(\)](#) (chainer.links.BinaryHierarchicalSoftmax method), [493](#)
[register_persistent\(\)](#) (chainer.links.BlackOut method), [498](#)
[register_persistent\(\)](#) (chainer.links.caffe.CaffeFunction method), [586](#)
[register_persistent\(\)](#) (chainer.links.ChildSumTreeLSTM method), [287](#)
[register_persistent\(\)](#) (chainer.links.Classifier method), [534](#)
[register_persistent\(\)](#) (chainer.links.Convolution2D method), [293](#)
[register_persistent\(\)](#) (chainer.links.ConvolutionND method), [298](#)
[register_persistent\(\)](#) (chainer.links.CRF1d method), [502](#)
[register_persistent\(\)](#) (chainer.links.Deconvolution2D method), [304](#)
[register_persistent\(\)](#) (chainer.links.DeconvolutionND method), [309](#)
[register_persistent\(\)](#) (chainer.links.DepthwiseConvolution2D method), [314](#)
[register_persistent\(\)](#) (chainer.links.DilatedConvolution2D method), [320](#)
[register_persistent\(\)](#) (chainer.links.EmbedID method), [326](#)
[register_persistent\(\)](#) (chainer.links.GoogLeNet method), [548](#)
[register_persistent\(\)](#) (chainer.links.GRU method), [330](#)
[register_persistent\(\)](#) (chainer.links.Highway method), [336](#)
[register_persistent\(\)](#) (chainer.links.Inception method), [341](#)
[register_persistent\(\)](#) (chainer.links.InceptionBN method), [347](#)
[register_persistent\(\)](#) (chainer.links.LayerNormalization method), [488](#)
[register_persistent\(\)](#) (chainer.links.Linear method), [352](#)
[register_persistent\(\)](#) (chainer.links.LocalConvolution2D method), [357](#)
[register_persistent\(\)](#) (chainer.links.LSTM method), [364](#)
[register_persistent\(\)](#) (chainer.links.Maxout method), [523](#)
[register_persistent\(\)](#) (chainer.links.MLPConvolution2D method), [369](#)
[register_persistent\(\)](#) (chainer.links.model.vision.resnet.ResNetLayers method), [555](#)
[register_persistent\(\)](#) (chainer.links.NaryTreeLSTM method), [375](#)
[register_persistent\(\)](#) (chainer.links.NegativeSampling method), [528](#)
[register_persistent\(\)](#) (chainer.links.NStepBiGRU method), [380](#)
[register_persistent\(\)](#) (chainer.links.NStepBiLSTM method), [386](#)
[register_persistent\(\)](#) (chainer.links.NStepBiRNNTanh method), [392](#)
[register_persistent\(\)](#) (chainer.links.NStepBiRNNTanh method), [398](#)
[register_persistent\(\)](#) (chainer.links.NStepGRU method), [404](#)
[register_persistent\(\)](#) (chainer.links.NStepLSTM method), [410](#)
[register_persistent\(\)](#) (chainer.links.NStepRNNTanh method), [416](#)
[register_persistent\(\)](#) (chainer.links.NStepRNNTanh method), [422](#)
[register_persistent\(\)](#) (chainer.links.Parameter method), [427](#)
[register_persistent\(\)](#) (chainer.links.PReLU method), [512](#)
[register_persistent\(\)](#) (chainer.links.ResNet101Layers method), [568](#)
[register_persistent\(\)](#) (chainer.links.ResNet152Layers method), [575](#)
[register_persistent\(\)](#) (chainer.links.ResNet50Layers method), [562](#)
[register_persistent\(\)](#) (chainer.links.Scale method), [432](#)
[register_persistent\(\)](#) (chainer.links.SimplifiedDropconnect method), [507](#)
[register_persistent\(\)](#) (chainer.links.StatefulGRU method), [438](#)
[register_persistent\(\)](#) (chainer.links.StatefulMGU method), [449](#)
[register_persistent\(\)](#) (chainer.links.StatefulPeepholeLSTM method), [459](#)
[register_persistent\(\)](#) (chainer.links.StatefulZoneoutLSTM method), [464](#)

`register_persistent()` (chainer.links.StatelessGRU method), 444
`register_persistent()` (chainer.links.StatelessLSTM method), 470
`register_persistent()` (chainer.links.StatelessMGU method), 454
`register_persistent()` (chainer.links.Swish method), 518
`register_persistent()` (chainer.links.TheanoFunction method), 581
`register_persistent()` (chainer.links.VGG16Layers method), 541
`register_persistent()` (chainer.Sequential method), 610
`register_statistics()` (chainer.training.extensions.ParameterStatistics method), 671
`relu()` (in module chainer.functions), 141
`remove()` (chainer.Sequential method), 610
`remove_by_layer_type()` (chainer.Sequential method), 610
`remove_hook()` (chainer.GradientMethod method), 638
`remove_hook()` (chainer.Optimizer method), 634
`remove_hook()` (chainer.optimizers.AdaDelta method), 613
`remove_hook()` (chainer.optimizers.AdaGrad method), 615
`remove_hook()` (chainer.optimizers.Adam method), 618
`remove_hook()` (chainer.optimizers.MomentumSGD method), 620
`remove_hook()` (chainer.optimizers.NesterovAG method), 622
`remove_hook()` (chainer.optimizers.RMSprop method), 625
`remove_hook()` (chainer.optimizers.RMSpropGraves method), 627
`remove_hook()` (chainer.optimizers.SGD method), 629
`remove_hook()` (chainer.optimizers.SMORMS3 method), 631
`remove_hook()` (chainer.UpdateRule method), 636
`repeat` (chainer.iterators.MultithreadIterator attribute), 717
`repeat` (chainer.iterators.SerialIterator attribute), 714
`repeat()` (chainer.Chain method), 598
`repeat()` (chainer.ChainList method), 603
`repeat()` (chainer.Link method), 592
`repeat()` (chainer.links.BatchNormalization method), 478
`repeat()` (chainer.links.BatchRenormalization method), 483
`repeat()` (chainer.links.Bias method), 276
`repeat()` (chainer.links.Bilinear method), 281
`repeat()` (chainer.links.BinaryHierarchicalSoftmax method), 493
`repeat()` (chainer.links.BlackOut method), 498
`repeat()` (chainer.links.caffe.CaffeFunction method), 587
`repeat()` (chainer.links.ChildSumTreeLSTM method), 287
`repeat()` (chainer.links.Classifier method), 534
`repeat()` (chainer.links.Convolution2D method), 293
`repeat()` (chainer.links.ConvolutionND method), 298
`repeat()` (chainer.links.CRF1d method), 502
`repeat()` (chainer.links.Deconvolution2D method), 304
`repeat()` (chainer.links.DeconvolutionND method), 310
`repeat()` (chainer.links.DepthwiseConvolution2D method), 315
`repeat()` (chainer.links.DilatedConvolution2D method), 321
`repeat()` (chainer.links.EmbedID method), 326
`repeat()` (chainer.links.GoogLeNet method), 548
`repeat()` (chainer.links.GRU method), 331
`repeat()` (chainer.links.Highway method), 336
`repeat()` (chainer.links.Inception method), 341
`repeat()` (chainer.links.InceptionBN method), 347
`repeat()` (chainer.links.LayerNormalization method), 488
`repeat()` (chainer.links.Linear method), 352
`repeat()` (chainer.links.LocalConvolution2D method), 358
`repeat()` (chainer.links.LSTM method), 364
`repeat()` (chainer.links.Maxout method), 523
`repeat()` (chainer.links.MLPConvolution2D method), 369
`repeat()` (chainer.links.model.vision.resnet.ResNetLayers method), 555
`repeat()` (chainer.links.NaryTreeLSTM method), 375
`repeat()` (chainer.links.NegativeSampling method), 528
`repeat()` (chainer.links.NStepBiGRU method), 380
`repeat()` (chainer.links.NStepBiLSTM method), 386
`repeat()` (chainer.links.NStepBiRNNReLU method), 392
`repeat()` (chainer.links.NStepBiRNNTanh method), 398
`repeat()` (chainer.links.NStepGRU method), 404
`repeat()` (chainer.links.NStepLSTM method), 410
`repeat()` (chainer.links.NStepRNNReLU method), 416
`repeat()` (chainer.links.NStepRNNTanh method), 422
`repeat()` (chainer.links.Parameter method), 427
`repeat()` (chainer.links.PReLU method), 512
`repeat()` (chainer.links.ResNet101Layers method), 569
`repeat()` (chainer.links.ResNet152Layers method), 575
`repeat()` (chainer.links.ResNet50Layers method), 562
`repeat()` (chainer.links.Scale method), 432
`repeat()` (chainer.links.SimplifiedDropconnect method), 508
`repeat()` (chainer.links.StatefulGRU method), 438
`repeat()` (chainer.links.StatefulMGU method), 449
`repeat()` (chainer.links.StatefulPeepholeLSTM method), 460
`repeat()` (chainer.links.StatefulZoneoutLSTM method), 464
`repeat()` (chainer.links.StatelessGRU method), 444
`repeat()` (chainer.links.StatelessLSTM method), 471
`repeat()` (chainer.links.StatelessMGU method), 454
`repeat()` (chainer.links.Swish method), 518
`repeat()` (chainer.links.TheanoFunction method), 581
`repeat()` (chainer.links.VGG16Layers method), 541
`repeat()` (chainer.Sequential method), 611

- repeat() (in module chainer.functions), 159
 - report() (chainer.Reporter method), 734
 - report() (in module chainer), 735
 - report_key_template (chainer.training.extensions.ParameterStatistics attribute), 671
 - report_scope() (in module chainer), 736
 - Reporter (class in chainer), 733
 - requires_grad (chainer.Parameter attribute), 130
 - requires_grad (chainer.Variable attribute), 122
 - requires_grad (chainer.variable.VariableNode attribute), 133
 - reset() (chainer.iterators.MultiprocessIterator method), 716
 - reset() (chainer.iterators.MultithreadIterator method), 717
 - reset() (chainer.iterators.SerialIterator method), 714
 - reset_state() (chainer.links.GRU method), 331
 - reset_state() (chainer.links.LSTM method), 364
 - reset_state() (chainer.links.StatefulGRU method), 439
 - reset_state() (chainer.links.StatefulMGU method), 450
 - reset_state() (chainer.links.StatefulPeepholeLSTM method), 460
 - reset_state() (chainer.links.StatefulZoneoutLSTM method), 465
 - reshape() (chainer.Parameter method), 126
 - reshape() (chainer.Variable method), 118
 - reshape() (in module chainer.functions), 160
 - resize_images() (in module chainer.functions), 160
 - ResNet101Layers (class in chainer.links), 564
 - ResNet152Layers (class in chainer.links), 570
 - ResNet50Layers (class in chainer.links), 557
 - ResNetLayers (class in chainer.links.model.vision.resnet), 550
 - retain_data() (chainer.Parameter method), 126
 - retain_data() (chainer.Variable method), 118
 - retain_data() (chainer.variable.VariableNode method), 131
 - retain_inputs() (chainer.Function method), 250
 - retain_inputs() (chainer.FunctionAdapter method), 254
 - retain_inputs() (chainer.FunctionNode method), 259
 - retain_outputs() (chainer.Function method), 250
 - retain_outputs() (chainer.FunctionAdapter method), 254
 - retain_outputs() (chainer.FunctionNode method), 260
 - rho (chainer.optimizers.AdaDelta attribute), 614
 - RMSprop (class in chainer.optimizers), 624
 - RMSpropGraves (class in chainer.optimizers), 626
 - rnn() (chainer.links.NStepBiGRU method), 381
 - rnn() (chainer.links.NStepBiLSTM method), 387
 - rnn() (chainer.links.NStepBiRNNReLU method), 393
 - rnn() (chainer.links.NStepBiRNNTanh method), 399
 - rnn() (chainer.links.NStepGRU method), 405
 - rnn() (chainer.links.NStepLSTM method), 411
 - rnn() (chainer.links.NStepRNReLU method), 417
 - rnn() (chainer.links.NStepRNNTanh method), 423
 - roi_pooling_2d() (in module chainer.functions), 242
 - rollaxis() (in module chainer.functions), 161
 - rsqrt() (in module chainer.functions), 229
 - run() (chainer.training.Trainer method), 655
- ## S
- sample() (chainer.utils.WalkerAlias method), 732
 - sample_cpu() (chainer.utils.WalkerAlias method), 733
 - sample_data (chainer.links.BlackOut attribute), 499
 - sample_gpu() (chainer.utils.WalkerAlias method), 733
 - save() (chainer.Serializer method), 724
 - save() (chainer.serializers.DictionarySerializer method), 719
 - save() (chainer.serializers.HDF5Serializer method), 722
 - save_hdf5() (in module chainer.serializers), 723
 - save_npz() (in module chainer.serializers), 720
 - save_plot_using_module() (chainer.training.extensions.VariableStatisticsPlot method), 682
 - Scale (class in chainer.links), 429
 - scale() (in module chainer.functions), 229
 - scatter_add() (in module chainer.functions), 161
 - scope() (chainer.Reporter method), 734
 - select_item() (in module chainer.functions), 162
 - selu() (in module chainer.functions), 141
 - separate() (in module chainer.functions), 162
 - Sequential (class in chainer), 605
 - SerialIterator (class in chainer.iterators), 713
 - serialize() (chainer.Chain method), 599
 - serialize() (chainer.ChainList method), 604
 - serialize() (chainer.dataset.Iterator method), 692
 - serialize() (chainer.DictSummary method), 737
 - serialize() (chainer.GradientMethod method), 638
 - serialize() (chainer.iterators.MultiprocessIterator method), 716
 - serialize() (chainer.iterators.MultithreadIterator method), 717
 - serialize() (chainer.iterators.SerialIterator method), 714
 - serialize() (chainer.Link method), 593
 - serialize() (chainer.links.BatchNormalization method), 478
 - serialize() (chainer.links.BatchRenormalization method), 483
 - serialize() (chainer.links.Bias method), 277
 - serialize() (chainer.links.Bilinear method), 282
 - serialize() (chainer.links.BinaryHierarchicalSoftmax method), 494
 - serialize() (chainer.links.BlackOut method), 498
 - serialize() (chainer.links.caffe.CaffeFunction method), 587
 - serialize() (chainer.links.ChildSumTreeLSTM method), 287
 - serialize() (chainer.links.Classifier method), 535
 - serialize() (chainer.links.Convolution2D method), 294
 - serialize() (chainer.links.ConvolutionND method), 299

- `serialize()` (chainer.links.CRF1d method), 503
- `serialize()` (chainer.links.Deconvolution2D method), 305
- `serialize()` (chainer.links.DeconvolutionND method), 310
- `serialize()` (chainer.links.DepthwiseConvolution2D method), 315
- `serialize()` (chainer.links.DilatedConvolution2D method), 321
- `serialize()` (chainer.links.EmbedID method), 326
- `serialize()` (chainer.links.GoogLeNet method), 549
- `serialize()` (chainer.links.GRU method), 331
- `serialize()` (chainer.links.Highway method), 337
- `serialize()` (chainer.links.Inception method), 342
- `serialize()` (chainer.links.InceptionBN method), 347
- `serialize()` (chainer.links.LayerNormalization method), 489
- `serialize()` (chainer.links.Linear method), 353
- `serialize()` (chainer.links.LocalConvolution2D method), 358
- `serialize()` (chainer.links.LSTM method), 364
- `serialize()` (chainer.links.Maxout method), 524
- `serialize()` (chainer.links.MLPConvolution2D method), 370
- `serialize()` (chainer.links.model.vision.resnet.ResNetLayers method), 556
- `serialize()` (chainer.links.NaryTreeLSTM method), 376
- `serialize()` (chainer.links.NegativeSampling method), 529
- `serialize()` (chainer.links.NStepBiGRU method), 381
- `serialize()` (chainer.links.NStepBiLSTM method), 387
- `serialize()` (chainer.links.NStepBiRNReLU method), 393
- `serialize()` (chainer.links.NStepBiRNNTanh method), 399
- `serialize()` (chainer.links.NStepGRU method), 405
- `serialize()` (chainer.links.NStepLSTM method), 411
- `serialize()` (chainer.links.NStepRNReLU method), 417
- `serialize()` (chainer.links.NStepRNNTanh method), 423
- `serialize()` (chainer.links.Parameter method), 428
- `serialize()` (chainer.links.PReLU method), 513
- `serialize()` (chainer.links.ResNet101Layers method), 569
- `serialize()` (chainer.links.ResNet152Layers method), 576
- `serialize()` (chainer.links.ResNet50Layers method), 563
- `serialize()` (chainer.links.Scale method), 433
- `serialize()` (chainer.links.SimplifiedDropconnect method), 508
- `serialize()` (chainer.links.StatefulGRU method), 439
- `serialize()` (chainer.links.StatefulMGU method), 450
- `serialize()` (chainer.links.StatefulPeepholeLSTM method), 460
- `serialize()` (chainer.links.StatefulZoneoutLSTM method), 465
- `serialize()` (chainer.links.StatelessGRU method), 445
- `serialize()` (chainer.links.StatelessLSTM method), 471
- `serialize()` (chainer.links.StatelessMGU method), 455
- `serialize()` (chainer.links.Swish method), 518
- `serialize()` (chainer.links.TheanoFunction method), 581
- `serialize()` (chainer.links.VGG16Layers method), 542
- `serialize()` (chainer.Optimizer method), 634
- `serialize()` (chainer.optimizers.AdaDelta method), 613
- `serialize()` (chainer.optimizers.AdaGrad method), 616
- `serialize()` (chainer.optimizers.Adam method), 618
- `serialize()` (chainer.optimizers.MomentumSGD method), 620
- `serialize()` (chainer.optimizers.NesterovAG method), 623
- `serialize()` (chainer.optimizers.RMSprop method), 625
- `serialize()` (chainer.optimizers.RMSpropGraves method), 627
- `serialize()` (chainer.optimizers.SGD method), 629
- `serialize()` (chainer.optimizers.SMORMS3 method), 631
- `serialize()` (chainer.Sequential method), 611
- `serialize()` (chainer.Summary method), 736
- `serialize()` (chainer.training.Extension method), 663
- `serialize()` (chainer.training.extensions.Evaluator method), 666
- `serialize()` (chainer.training.extensions.ExponentialShift method), 673
- `serialize()` (chainer.training.extensions.FailOnNonNumber method), 669
- `serialize()` (chainer.training.extensions.LinearShift method), 674
- `serialize()` (chainer.training.extensions.LogReport method), 679
- `serialize()` (chainer.training.extensions.MicroAverage method), 668
- `serialize()` (chainer.training.extensions.ParameterStatistics method), 671
- `serialize()` (chainer.training.extensions.PlotReport method), 680
- `serialize()` (chainer.training.extensions.PrintReport method), 676
- `serialize()` (chainer.training.extensions.ProgressBar method), 677
- `serialize()` (chainer.training.extensions.VariableStatisticsPlot method), 682
- `serialize()` (chainer.training.Trainer method), 655
- `serialize()` (chainer.training.triggers.IntervalTrigger method), 687
- `serialize()` (chainer.training.triggers.ManualScheduleTrigger method), 688
- `serialize()` (chainer.training.triggers.TimeTrigger method), 689
- `serialize()` (chainer.training.Updater method), 656
- `serialize()` (chainer.training.updaters.MultiprocessParallelUpdater method), 661
- `serialize()` (chainer.training.updaters.ParallelUpdater method), 660
- `serialize()` (chainer.training.updaters.StandardUpdater method), 658
- `serialize()` (chainer.UpdateRule method), 636
- `Serializer` (class in chainer), 724

- set_creator() (chainer.Parameter method), 126
- set_creator() (chainer.Variable method), 118
- set_creator() (chainer.variable.VariableNode method), 131
- set_creator_node() (chainer.Parameter method), 126
- set_creator_node() (chainer.Variable method), 118
- set_creator_node() (chainer.variable.VariableNode method), 131
- set_dataset_root() (in module chainer.dataset), 695
- set_debug() (in module chainer), 745
- set_loss_scale() (chainer.GradientMethod method), 639
- set_loss_scale() (chainer.Optimizer method), 634
- set_loss_scale() (chainer.optimizers.AdaDelta method), 614
- set_loss_scale() (chainer.optimizers.AdaGrad method), 616
- set_loss_scale() (chainer.optimizers.Adam method), 618
- set_loss_scale() (chainer.optimizers.MomentumSGD method), 621
- set_loss_scale() (chainer.optimizers.NesterovAG method), 623
- set_loss_scale() (chainer.optimizers.RMSprop method), 625
- set_loss_scale() (chainer.optimizers.RMSpropGraves method), 627
- set_loss_scale() (chainer.optimizers.SGD method), 629
- set_loss_scale() (chainer.optimizers.SMORMS3 method), 631
- set_max_workspace_size() (in module chainer.backends.cuda), 732
- set_state() (chainer.links.GRU method), 331
- set_state() (chainer.links.LSTM method), 365
- set_state() (chainer.links.StatefulGRU method), 439
- set_state() (chainer.links.StatefulMGU method), 450
- set_state() (chainer.links.StatefulZoneoutLSTM method), 465
- setup() (chainer.GradientMethod method), 639
- setup() (chainer.Optimizer method), 634
- setup() (chainer.optimizers.AdaDelta method), 614
- setup() (chainer.optimizers.AdaGrad method), 616
- setup() (chainer.optimizers.Adam method), 618
- setup() (chainer.optimizers.MomentumSGD method), 621
- setup() (chainer.optimizers.NesterovAG method), 623
- setup() (chainer.optimizers.RMSprop method), 625
- setup() (chainer.optimizers.RMSpropGraves method), 627
- setup() (chainer.optimizers.SGD method), 629
- setup() (chainer.optimizers.SMORMS3 method), 632
- setup_workers() (chainer.training.updaters.MultiprocessParallelUpdater method), 478
- SGD (class in chainer.optimizers), 628
- shape (chainer.Parameter attribute), 130
- shape (chainer.Variable attribute), 122
- shift() (in module chainer.functions), 197
- show() (chainer.configuration.GlobalConfig method), 742
- show() (chainer.configuration.LocalConfig method), 743
- sigmoid() (in module chainer.functions), 141
- sigmoid_cross_entropy() (in module chainer.functions), 212
- sign() (in module chainer.functions), 230
- simplified_dropconnect() (in module chainer.functions), 234
- SimplifiedDropconnect (class in chainer.links), 504
- sin() (in module chainer.functions), 229
- sinh() (in module chainer.functions), 230
- size (chainer.Parameter attribute), 130
- size (chainer.utils.type_check.TypeInfo attribute), 753
- size (chainer.Variable attribute), 122
- size() (chainer.utils.type_check.TypeInfoTuple method), 753
- slstm() (in module chainer.functions), 142
- SMORMS3 (class in chainer.optimizers), 630
- snapshot(), 57
- snapshot() (in module chainer.training.extensions), 684
- snapshot_object(), 57
- snapshot_object() (in module chainer.training.extensions), 684
- softmax() (in module chainer.functions), 143
- softmax_cross_entropy() (in module chainer.functions), 213
- softplus() (in module chainer.functions), 144
- space2depth() (in module chainer.functions), 163
- spatial_pyramid_pooling_2d() (in module chainer.functions), 242
- spatial_transformer_grid() (in module chainer.functions), 164
- spatial_transformer_sampler() (in module chainer.functions), 165
- split_axis() (in module chainer.functions), 165
- split_dataset() (in module chainer.datasets), 700
- split_dataset_random() (in module chainer.datasets), 701
- sqrt() (in module chainer.functions), 230
- square() (in module chainer.functions), 230
- squared_difference() (in module chainer.functions), 231
- squared_error() (in module chainer.functions), 215
- squeeze() (in module chainer.functions), 166
- stack (chainer.Function attribute), 251
- stack (chainer.FunctionAdapter attribute), 255
- stack (chainer.FunctionNode attribute), 260
- stack() (in module chainer.functions), 167
- StandardUpdater (class in chainer.training.updaters), 657
- start_finetuning() (chainer.links.BatchNormalization method), 478
- start_finetuning() (chainer.links.BatchRenormalization method), 483
- state (chainer.UpdateRule attribute), 637
- StatefulGRU (class in chainer.links), 434

StatefulMGU (class in `chainer.links`), 446
StatefulPeepholeLSTM (class in `chainer.links`), 456
StatefulZoneoutLSTM (class in `chainer.links`), 461
StatelessGRU (class in `chainer.links`), 440
StatelessLSTM (class in `chainer.links`), 466
StatelessMGU (class in `chainer.links`), 451
SubDataset (class in `chainer.datasets`), 699
sum() (in module `chainer.functions`), 231
Summary (class in `chainer`), 736
summary() (`chainer.function_hooks.CupyMemoryProfileHook` method), 265
summary() (`chainer.function_hooks.TimerHook` method), 269
summary() (`chainer.Parameter` method), 126
summary() (`chainer.Variable` method), 118
swapaxes() (in module `chainer.functions`), 168
Swish (class in `chainer.links`), 514
swish() (in module `chainer.functions`), 144

T

t (`chainer.GradientMethod` attribute), 639
t (`chainer.Optimizer` attribute), 635
t (`chainer.optimizers.AdaDelta` attribute), 614
t (`chainer.optimizers.AdaGrad` attribute), 617
t (`chainer.optimizers.Adam` attribute), 619
t (`chainer.optimizers.MomentumSGD` attribute), 621
t (`chainer.optimizers.NesterovAG` attribute), 624
t (`chainer.optimizers.RMSprop` attribute), 626
t (`chainer.optimizers.RMSpropGraves` attribute), 628
t (`chainer.optimizers.SGD` attribute), 630
t (`chainer.optimizers.SMORMS3` attribute), 632
T (`chainer.Parameter` attribute), 129
T (`chainer.Variable` attribute), 121
tan() (in module `chainer.functions`), 232
tanh() (in module `chainer.functions`), 145
target (`chainer.GradientMethod` attribute), 639
target (`chainer.Optimizer` attribute), 635
target (`chainer.optimizers.AdaDelta` attribute), 614
target (`chainer.optimizers.AdaGrad` attribute), 617
target (`chainer.optimizers.Adam` attribute), 619
target (`chainer.optimizers.MomentumSGD` attribute), 621
target (`chainer.optimizers.NesterovAG` attribute), 624
target (`chainer.optimizers.RMSprop` attribute), 626
target (`chainer.optimizers.RMSpropGraves` attribute), 628
target (`chainer.optimizers.SGD` attribute), 630
target (`chainer.optimizers.SMORMS3` attribute), 632
tensordot() (in module `chainer.functions`), 232
TheanoFunction (class in `chainer.links`), 577
tile() (in module `chainer.functions`), 169
TimerHook (class in `chainer.function_hooks`), 268
TimeTrigger (class in `chainer.training.triggers`), 689
timing (`chainer.optimizer_hooks.GradientClipping` attribute), 642
timing (`chainer.optimizer_hooks.GradientHardClipping` attribute), 642
timing (`chainer.optimizer_hooks.GradientLARS` attribute), 644
timing (`chainer.optimizer_hooks.GradientNoise` attribute), 643
timing (`chainer.optimizer_hooks.Lasso` attribute), 641
timing (`chainer.optimizer_hooks.WeightDecay` attribute), 640
to_cpu() (`chainer.Chain` method), 599
to_cpu() (`chainer.ChainList` method), 604
to_cpu() (`chainer.Link` method), 593
to_cpu() (`chainer.links.BatchNormalization` method), 478
to_cpu() (`chainer.links.BatchRenormalization` method), 484
to_cpu() (`chainer.links.Bias` method), 277
to_cpu() (`chainer.links.Bilinear` method), 282
to_cpu() (`chainer.links.BinaryHierarchicalSoftmax` method), 494
to_cpu() (`chainer.links.BlackOut` method), 498
to_cpu() (`chainer.links.caffe.CaffeFunction` method), 587
to_cpu() (`chainer.links.ChildSumTreeLSTM` method), 288
to_cpu() (`chainer.links.Classifier` method), 535
to_cpu() (`chainer.links.Convolution2D` method), 294
to_cpu() (`chainer.links.ConvolutionND` method), 299
to_cpu() (`chainer.links.CRF1d` method), 503
to_cpu() (`chainer.links.Deconvolution2D` method), 305
to_cpu() (`chainer.links.DeconvolutionND` method), 310
to_cpu() (`chainer.links.DepthwiseConvolution2D` method), 315
to_cpu() (`chainer.links.DilatedConvolution2D` method), 321
to_cpu() (`chainer.links.EmbedID` method), 326
to_cpu() (`chainer.links.GoogLeNet` method), 549
to_cpu() (`chainer.links.GRU` method), 331
to_cpu() (`chainer.links.Highway` method), 337
to_cpu() (`chainer.links.Inception` method), 342
to_cpu() (`chainer.links.InceptionBN` method), 347
to_cpu() (`chainer.links.LayerNormalization` method), 489
to_cpu() (`chainer.links.Linear` method), 353
to_cpu() (`chainer.links.LocalConvolution2D` method), 358
to_cpu() (`chainer.links.LSTM` method), 365
to_cpu() (`chainer.links.Maxout` method), 524
to_cpu() (`chainer.links.MLPConvolution2D` method), 370
to_cpu() (`chainer.links.model.vision.resnet.ResNetLayers` method), 556
to_cpu() (`chainer.links.NaryTreeLSTM` method), 376
to_cpu() (`chainer.links.NegativeSampling` method), 529
to_cpu() (`chainer.links.NStepBiGRU` method), 381
to_cpu() (`chainer.links.NStepBiLSTM` method), 387
to_cpu() (`chainer.links.NStepBiRNNReLU` method), 393
to_cpu() (`chainer.links.NStepBiRNNTanh` method), 399

`to_cpu()` (chainer.links.NStepGRU method), 405
`to_cpu()` (chainer.links.NStepLSTM method), 411
`to_cpu()` (chainer.links.NStepRNNReLU method), 417
`to_cpu()` (chainer.links.NStepRNNTanh method), 423
`to_cpu()` (chainer.links.Parameter method), 428
`to_cpu()` (chainer.links.PReLU method), 513
`to_cpu()` (chainer.links.ResNet101Layers method), 569
`to_cpu()` (chainer.links.ResNet152Layers method), 576
`to_cpu()` (chainer.links.ResNet50Layers method), 563
`to_cpu()` (chainer.links.Scale method), 433
`to_cpu()` (chainer.links.SimplifiedDropconnect method), 508
`to_cpu()` (chainer.links.StatefulGRU method), 439
`to_cpu()` (chainer.links.StatefulMGU method), 450
`to_cpu()` (chainer.links.StatefulPeepholeLSTM method), 460
`to_cpu()` (chainer.links.StatefulZoneoutLSTM method), 465
`to_cpu()` (chainer.links.StatelessGRU method), 445
`to_cpu()` (chainer.links.StatelessLSTM method), 471
`to_cpu()` (chainer.links.StatelessMGU method), 455
`to_cpu()` (chainer.links.Swish method), 518
`to_cpu()` (chainer.links.TheanoFunction method), 581
`to_cpu()` (chainer.links.VGG16Layers method), 542
`to_cpu()` (chainer.Parameter method), 126
`to_cpu()` (chainer.Sequential method), 611
`to_cpu()` (chainer.utils.WalkerAlias method), 733
`to_cpu()` (chainer.Variable method), 118
`to_cpu()` (in module chainer.backends.cuda), 729
`to_device()` (in module chainer.dataset), 694
`to_gpu()` (chainer.Chain method), 599
`to_gpu()` (chainer.ChainList method), 604
`to_gpu()` (chainer.Link method), 593
`to_gpu()` (chainer.links.BatchNormalization method), 479
`to_gpu()` (chainer.links.BatchRenormalization method), 484
`to_gpu()` (chainer.links.Bias method), 277
`to_gpu()` (chainer.links.Bilinear method), 282
`to_gpu()` (chainer.links.BinaryHierarchicalSoftmax method), 494
`to_gpu()` (chainer.links.BlackOut method), 499
`to_gpu()` (chainer.links.caffe.CaffeFunction method), 588
`to_gpu()` (chainer.links.ChildSumTreeLSTM method), 288
`to_gpu()` (chainer.links.Classifier method), 535
`to_gpu()` (chainer.links.Convolution2D method), 294
`to_gpu()` (chainer.links.ConvolutionND method), 299
`to_gpu()` (chainer.links.CRF1d method), 503
`to_gpu()` (chainer.links.Deconvolution2D method), 305
`to_gpu()` (chainer.links.DeconvolutionND method), 310
`to_gpu()` (chainer.links.DepthwiseConvolution2D method), 316
`to_gpu()` (chainer.links.DilatedConvolution2D method), 322
`to_gpu()` (chainer.links.EmbedID method), 327
`to_gpu()` (chainer.links.GoogLeNet method), 549
`to_gpu()` (chainer.links.GRU method), 332
`to_gpu()` (chainer.links.Highway method), 337
`to_gpu()` (chainer.links.Inception method), 342
`to_gpu()` (chainer.links.InceptionBN method), 348
`to_gpu()` (chainer.links.LayerNormalization method), 489
`to_gpu()` (chainer.links.Linear method), 353
`to_gpu()` (chainer.links.LocalConvolution2D method), 358
`to_gpu()` (chainer.links.LSTM method), 365
`to_gpu()` (chainer.links.Maxout method), 524
`to_gpu()` (chainer.links.MLPConvolution2D method), 370
`to_gpu()` (chainer.links.model.vision.resnet.ResNetLayers method), 556
`to_gpu()` (chainer.links.NaryTreeLSTM method), 376
`to_gpu()` (chainer.links.NegativeSampling method), 529
`to_gpu()` (chainer.links.NStepBiGRU method), 381
`to_gpu()` (chainer.links.NStepBiLSTM method), 387
`to_gpu()` (chainer.links.NStepBiRNNReLU method), 393
`to_gpu()` (chainer.links.NStepBiRNNTanh method), 399
`to_gpu()` (chainer.links.NStepGRU method), 405
`to_gpu()` (chainer.links.NStepLSTM method), 411
`to_gpu()` (chainer.links.NStepRNNReLU method), 417
`to_gpu()` (chainer.links.NStepRNNTanh method), 423
`to_gpu()` (chainer.links.Parameter method), 428
`to_gpu()` (chainer.links.PReLU method), 513
`to_gpu()` (chainer.links.ResNet101Layers method), 569
`to_gpu()` (chainer.links.ResNet152Layers method), 576
`to_gpu()` (chainer.links.ResNet50Layers method), 563
`to_gpu()` (chainer.links.Scale method), 433
`to_gpu()` (chainer.links.SimplifiedDropconnect method), 509
`to_gpu()` (chainer.links.StatefulGRU method), 439
`to_gpu()` (chainer.links.StatefulMGU method), 450
`to_gpu()` (chainer.links.StatefulPeepholeLSTM method), 461
`to_gpu()` (chainer.links.StatefulZoneoutLSTM method), 466
`to_gpu()` (chainer.links.StatelessGRU method), 445
`to_gpu()` (chainer.links.StatelessLSTM method), 471
`to_gpu()` (chainer.links.StatelessMGU method), 455
`to_gpu()` (chainer.links.Swish method), 519
`to_gpu()` (chainer.links.TheanoFunction method), 582
`to_gpu()` (chainer.links.VGG16Layers method), 542
`to_gpu()` (chainer.Parameter method), 126
`to_gpu()` (chainer.Sequential method), 611
`to_gpu()` (chainer.utils.WalkerAlias method), 733
`to_gpu()` (chainer.Variable method), 118
`to_gpu()` (in module chainer.backends.cuda), 730
`to_intel64()` (chainer.Chain method), 599
`to_intel64()` (chainer.ChainList method), 604
`to_intel64()` (chainer.Link method), 593

- `to_intel64()` (chainer.links.BatchNormalization method), 479
- `to_intel64()` (chainer.links.BatchRenormalization method), 484
- `to_intel64()` (chainer.links.Bias method), 277
- `to_intel64()` (chainer.links.Bilinear method), 282
- `to_intel64()` (chainer.links.BinaryHierarchicalSoftmax method), 494
- `to_intel64()` (chainer.links.BlackOut method), 499
- `to_intel64()` (chainer.links.caffe.CaffeFunction method), 588
- `to_intel64()` (chainer.links.ChildSumTreeLSTM method), 288
- `to_intel64()` (chainer.links.Classifier method), 535
- `to_intel64()` (chainer.links.Convolution2D method), 294
- `to_intel64()` (chainer.links.ConvolutionND method), 299
- `to_intel64()` (chainer.links.CRF1d method), 504
- `to_intel64()` (chainer.links.Deconvolution2D method), 306
- `to_intel64()` (chainer.links.DeconvolutionND method), 311
- `to_intel64()` (chainer.links.DepthwiseConvolution2D method), 316
- `to_intel64()` (chainer.links.DilatedConvolution2D method), 322
- `to_intel64()` (chainer.links.EmbedID method), 327
- `to_intel64()` (chainer.links.GoogLeNet method), 549
- `to_intel64()` (chainer.links.GRU method), 332
- `to_intel64()` (chainer.links.Highway method), 337
- `to_intel64()` (chainer.links.Inception method), 342
- `to_intel64()` (chainer.links.InceptionBN method), 348
- `to_intel64()` (chainer.links.LayerNormalization method), 489
- `to_intel64()` (chainer.links.Linear method), 353
- `to_intel64()` (chainer.links.LocalConvolution2D method), 359
- `to_intel64()` (chainer.links.LSTM method), 365
- `to_intel64()` (chainer.links.Maxout method), 524
- `to_intel64()` (chainer.links.MLPConvolution2D method), 370
- `to_intel64()` (chainer.links.model.vision.resnet.ResNetLayers method), 556
- `to_intel64()` (chainer.links.NaryTreeLSTM method), 376
- `to_intel64()` (chainer.links.NegativeSampling method), 529
- `to_intel64()` (chainer.links.NStepBiGRU method), 382
- `to_intel64()` (chainer.links.NStepBiLSTM method), 388
- `to_intel64()` (chainer.links.NStepBiRNNTanh method), 393
- `to_intel64()` (chainer.links.NStepBiRNNTanh method), 399
- `to_intel64()` (chainer.links.NStepGRU method), 405
- `to_intel64()` (chainer.links.NStepLSTM method), 411
- `to_intel64()` (chainer.links.NStepRNNTanh method), 417
- `to_intel64()` (chainer.links.NStepRNNTanh method), 423
- `to_intel64()` (chainer.links.Parameter method), 428
- `to_intel64()` (chainer.links.PReLU method), 513
- `to_intel64()` (chainer.links.ResNet101Layers method), 570
- `to_intel64()` (chainer.links.ResNet152Layers method), 576
- `to_intel64()` (chainer.links.ResNet50Layers method), 563
- `to_intel64()` (chainer.links.Scale method), 433
- `to_intel64()` (chainer.links.SimplifiedDropconnect method), 509
- `to_intel64()` (chainer.links.StatefulGRU method), 440
- `to_intel64()` (chainer.links.StatefulMGU method), 450
- `to_intel64()` (chainer.links.StatefulPeepholeLSTM method), 461
- `to_intel64()` (chainer.links.StatefulZoneoutLSTM method), 466
- `to_intel64()` (chainer.links.StatelessGRU method), 445
- `to_intel64()` (chainer.links.StatelessLSTM method), 472
- `to_intel64()` (chainer.links.StatelessMGU method), 455
- `to_intel64()` (chainer.links.Swish method), 519
- `to_intel64()` (chainer.links.TheanoFunction method), 582
- `to_intel64()` (chainer.links.VGG16Layers method), 542
- `to_intel64()` (chainer.Parameter method), 126
- `to_intel64()` (chainer.Sequential method), 612
- `to_intel64()` (chainer.Variable method), 118
- `total_acquired_bytes()` (chainer.function_hooks.CupyMemoryProfileHook method), 265
- `total_time()` (chainer.function_hooks.TimerHook method), 269
- `total_used_bytes()` (chainer.function_hooks.CupyMemoryProfileHook method), 266
- `Trainer` (class in chainer.training), 653
- `TransformDataset` (class in chainer.datasets), 702
- `transpose()` (chainer.Parameter method), 126
- `transpose()` (chainer.Variable method), 119
- `transpose()` (in module chainer.functions), 170
- `transpose_sequence()` (in module chainer.functions), 171
- `tree_lstm()` (in module chainer.functions), 145
- `trigger` (chainer.training.Extension attribute), 663
- `trigger` (chainer.training.extensions.Evaluator attribute), 667
- `trigger` (chainer.training.extensions.ExponentialShift attribute), 673
- `trigger` (chainer.training.extensions.FailOnNonNumber attribute), 669
- `trigger` (chainer.training.extensions.LinearShift attribute), 675
- `trigger` (chainer.training.extensions.LogReport attribute), 679
- `trigger` (chainer.training.extensions.MicroAverage attribute), 668

- trigger (chainer.training.extensions.ParameterStatistics attribute), 672
 - trigger (chainer.training.extensions.PlotReport attribute), 681
 - trigger (chainer.training.extensions.PrintReport attribute), 676
 - trigger (chainer.training.extensions.ProgressBar attribute), 677
 - trigger (chainer.training.extensions.VariableStatisticsPlot attribute), 683
 - triplet() (in module chainer.functions), 215
 - TupleDataset (class in chainer.datasets), 697
 - TypeInfo (class in chainer.utils.type_check), 753
 - TypeInfoTuple (class in chainer.utils.type_check), 753
- ## U
- unary_math_function_unittest() (in module chainer.testing), 757
 - unchain() (chainer.Function method), 250
 - unchain() (chainer.FunctionAdapter method), 254
 - unchain() (chainer.FunctionNode method), 260
 - unchain() (chainer.Parameter method), 126
 - unchain() (chainer.Variable method), 119
 - unchain() (chainer.variable.VariableNode method), 132
 - unchain_backward() (chainer.Parameter method), 126
 - unchain_backward() (chainer.Variable method), 119
 - Uniform (class in chainer.initializers), 650
 - unpooling_2d() (in module chainer.functions), 243
 - unpooling_nd() (in module chainer.functions), 244
 - update() (chainer.GradientMethod method), 639
 - update() (chainer.Optimizer method), 634
 - update() (chainer.optimizers.AdaDelta method), 614
 - update() (chainer.optimizers.AdaGrad method), 616
 - update() (chainer.optimizers.Adam method), 618
 - update() (chainer.optimizers.MomentumSGD method), 621
 - update() (chainer.optimizers.NesterovAG method), 623
 - update() (chainer.optimizers.RMSprop method), 625
 - update() (chainer.optimizers.RMSpropGraves method), 627
 - update() (chainer.optimizers.SGD method), 630
 - update() (chainer.optimizers.SMORMS3 method), 632
 - update() (chainer.Parameter method), 127
 - update() (chainer.training.Updater method), 656
 - update() (chainer.training.updaters.MultiprocessParallelUpdater method), 661
 - update() (chainer.training.updaters.ParallelUpdater method), 660
 - update() (chainer.training.updaters.StandardUpdater method), 658
 - update() (chainer.UpdateRule method), 636
 - update_core() (chainer.training.updaters.MultiprocessParallelUpdater method), 662
 - update_core() (chainer.training.updaters.ParallelUpdater method), 660
 - update_core() (chainer.training.updaters.StandardUpdater method), 658
 - update_core() (chainer.UpdateRule method), 636
 - update_core_cpu() (chainer.UpdateRule method), 636
 - update_core_gpu() (chainer.UpdateRule method), 636
 - update_enabled (chainer.Chain attribute), 600
 - update_enabled (chainer.ChainList attribute), 604
 - update_enabled (chainer.Link attribute), 594
 - update_enabled (chainer.links.BatchNormalization attribute), 479
 - update_enabled (chainer.links.BatchRenormalization attribute), 484
 - update_enabled (chainer.links.Bias attribute), 278
 - update_enabled (chainer.links.Bilinear attribute), 283
 - update_enabled (chainer.links.BinaryHierarchicalSoftmax attribute), 494
 - update_enabled (chainer.links.BlackOut attribute), 499
 - update_enabled (chainer.links.caffe.CaffeFunction attribute), 588
 - update_enabled (chainer.links.ChildSumTreeLSTM attribute), 288
 - update_enabled (chainer.links.Classifier attribute), 535
 - update_enabled (chainer.links.Convolution2D attribute), 294
 - update_enabled (chainer.links.ConvolutionND attribute), 300
 - update_enabled (chainer.links.CRF1d attribute), 504
 - update_enabled (chainer.links.Deconvolution2D attribute), 306
 - update_enabled (chainer.links.DeconvolutionND attribute), 311
 - update_enabled (chainer.links.DepthwiseConvolution2D attribute), 316
 - update_enabled (chainer.links.DilatedConvolution2D attribute), 322
 - update_enabled (chainer.links.EmbedID attribute), 327
 - update_enabled (chainer.links.GoogLeNet attribute), 549
 - update_enabled (chainer.links.GRU attribute), 332
 - update_enabled (chainer.links.Highway attribute), 337
 - update_enabled (chainer.links.Inception attribute), 343
 - update_enabled (chainer.links.InceptionBN attribute), 348
 - update_enabled (chainer.links.LayerNormalization attribute), 489
 - update_enabled (chainer.links.Linear attribute), 354
 - update_enabled (chainer.links.LocalConvolution2D attribute), 359
 - update_enabled (chainer.links.LSTM attribute), 365
 - update_enabled (chainer.links.Maxout attribute), 524
 - update_enabled (chainer.links.MLPConvolution2D attribute), 371
 - update_enabled (chainer.links.model.vision.resnet.ResNetLayers

attribute), [557](#)
update_enabled (chainer.links.NaryTreeLSTM attribute), [376](#)
update_enabled (chainer.links.NegativeSampling attribute), [529](#)
update_enabled (chainer.links.NStepBiGRU attribute), [382](#)
update_enabled (chainer.links.NStepBiLSTM attribute), [388](#)
update_enabled (chainer.links.NStepBiRNNReLU attribute), [394](#)
update_enabled (chainer.links.NStepBiRNNTanh attribute), [400](#)
update_enabled (chainer.links.NStepGRU attribute), [406](#)
update_enabled (chainer.links.NStepLSTM attribute), [412](#)
update_enabled (chainer.links.NStepRNNReLU attribute), [418](#)
update_enabled (chainer.links.NStepRNNTanh attribute), [424](#)
update_enabled (chainer.links.Parameter attribute), [428](#)
update_enabled (chainer.links.PReLU attribute), [514](#)
update_enabled (chainer.links.ResNet101Layers attribute), [570](#)
update_enabled (chainer.links.ResNet152Layers attribute), [576](#)
update_enabled (chainer.links.ResNet50Layers attribute), [563](#)
update_enabled (chainer.links.Scale attribute), [434](#)
update_enabled (chainer.links.SimplifiedDropconnect attribute), [509](#)
update_enabled (chainer.links.StatefulGRU attribute), [440](#)
update_enabled (chainer.links.StatefulMGU attribute), [451](#)
update_enabled (chainer.links.StatefulPeepholeLSTM attribute), [461](#)
update_enabled (chainer.links.StatefulZoneoutLSTM attribute), [466](#)
update_enabled (chainer.links.StatelessGRU attribute), [446](#)
update_enabled (chainer.links.StatelessLSTM attribute), [472](#)
update_enabled (chainer.links.StatelessMGU attribute), [455](#)
update_enabled (chainer.links.Swish attribute), [519](#)
update_enabled (chainer.links.TheanoFunction attribute), [582](#)
update_enabled (chainer.links.VGG16Layers attribute), [542](#)
update_enabled (chainer.Sequential attribute), [612](#)
Updater (class in chainer.training), [656](#)
UpdateRule (class in chainer), [635](#)
upsampling_2d() (in module chainer.functions), [244](#)
use_bi_direction (chainer.links.NStepBiGRU attribute), [382](#)
use_bi_direction (chainer.links.NStepBiLSTM attribute), [388](#)
use_bi_direction (chainer.links.NStepBiRNNReLU attribute), [394](#)
use_bi_direction (chainer.links.NStepBiRNNTanh attribute), [400](#)
use_bi_direction (chainer.links.NStepGRU attribute), [406](#)
use_bi_direction (chainer.links.NStepLSTM attribute), [412](#)
use_bi_direction (chainer.links.NStepRNNReLU attribute), [418](#)
use_bi_direction (chainer.links.NStepRNNTanh attribute), [424](#)
use_cleargrads() (chainer.GradientMethod method), [639](#)
use_cleargrads() (chainer.optimizers.AdaDelta method), [614](#)
use_cleargrads() (chainer.optimizers.AdaGrad method), [616](#)
use_cleargrads() (chainer.optimizers.Adam method), [619](#)
use_cleargrads() (chainer.optimizers.MomentumSGD method), [621](#)
use_cleargrads() (chainer.optimizers.NesterovAG method), [623](#)
use_cleargrads() (chainer.optimizers.RMSprop method), [625](#)
use_cleargrads() (chainer.optimizers.RMSpropGraves method), [628](#)
use_cleargrads() (chainer.optimizers.SGD method), [630](#)
use_cleargrads() (chainer.optimizers.SMORMS3 method), [632](#)
use_fp32_update() (chainer.GradientMethod method), [639](#)
use_fp32_update() (chainer.optimizers.AdaDelta method), [614](#)
use_fp32_update() (chainer.optimizers.AdaGrad method), [616](#)
use_fp32_update() (chainer.optimizers.Adam method), [619](#)
use_fp32_update() (chainer.optimizers.MomentumSGD method), [621](#)
use_fp32_update() (chainer.optimizers.NesterovAG method), [623](#)
use_fp32_update() (chainer.optimizers.RMSprop method), [625](#)
use_fp32_update() (chainer.optimizers.RMSpropGraves method), [628](#)
use_fp32_update() (chainer.optimizers.SGD method), [630](#)
use_fp32_update() (chainer.optimizers.SMORMS3 method), [632](#)
use_fp32_update() (chainer.UpdateRule method), [636](#)
using_config() (in module chainer), [742](#)

V

Variable (class in chainer), 115
 VariableNode (class in chainer.variable), 130
 VariableStatisticsPlot (class
 chainer.training.extensions), 681
 VGG16Layers (class in chainer.links), 536
 vstack() (in module chainer.functions), 171

W

WalkerAlias (class in chainer.utils), 732
 weight_decay_rate (chainer.optimizers.Adam attribute),
 619
 WeightDecay (class in chainer.optimizer_hooks), 640
 where() (in module chainer.functions), 172
 within_init_scope (chainer.Chain attribute), 600
 within_init_scope (chainer.ChainList attribute), 604
 within_init_scope (chainer.Link attribute), 594
 within_init_scope (chainer.links.BatchNormalization at-
 tribute), 479
 within_init_scope (chainer.links.BatchRenormalization
 attribute), 484
 within_init_scope (chainer.links.Bias attribute), 278
 within_init_scope (chainer.links.Bilinear attribute), 283
 within_init_scope (chainer.links.BinaryHierarchicalSoftmax
 attribute), 494
 within_init_scope (chainer.links.BlackOut attribute), 499
 within_init_scope (chainer.links.caffe.CaffeFunction at-
 tribute), 588
 within_init_scope (chainer.links.ChildSumTreeLSTM at-
 tribute), 288
 within_init_scope (chainer.links.Classifier attribute), 535
 within_init_scope (chainer.links.Convolution2D at-
 tribute), 294
 within_init_scope (chainer.links.ConvolutionND at-
 tribute), 300
 within_init_scope (chainer.links.CRF1d attribute), 504
 within_init_scope (chainer.links.Deconvolution2D
 attribute), 306
 within_init_scope (chainer.links.DeconvolutionND at-
 tribute), 311
 within_init_scope (chainer.links.DepthwiseConvolution2D
 attribute), 316
 within_init_scope (chainer.links.DilatedConvolution2D
 attribute), 322
 within_init_scope (chainer.links.EmbedID attribute), 327
 within_init_scope (chainer.links.GoogLeNet attribute),
 549
 within_init_scope (chainer.links.GRU attribute), 332
 within_init_scope (chainer.links.Highway attribute), 337
 within_init_scope (chainer.links.Inception attribute), 343
 within_init_scope (chainer.links.InceptionBN attribute),
 348
 within_init_scope (chainer.links.LayerNormalization at-
 tribute), 489
 within_init_scope (chainer.links.Linear attribute), 354
 within_init_scope (chainer.links.LocalConvolution2D at-
 tribute), 359
 in within_init_scope (chainer.links.LSTM attribute), 365
 within_init_scope (chainer.links.Maxout attribute), 524
 within_init_scope (chainer.links.MLPConvolution2D at-
 tribute), 371
 within_init_scope (chainer.links.model.vision.resnet.ResNetLayers
 attribute), 557
 within_init_scope (chainer.links.NaryTreeLSTM at-
 tribute), 376
 within_init_scope (chainer.links.NegativeSampling at-
 tribute), 529
 within_init_scope (chainer.links.NStepBiGRU attribute),
 382
 within_init_scope (chainer.links.NStepBiLSTM at-
 tribute), 388
 within_init_scope (chainer.links.NStepBiRNNReLU at-
 tribute), 394
 within_init_scope (chainer.links.NStepBiRNNTanh at-
 tribute), 400
 within_init_scope (chainer.links.NStepGRU attribute),
 406
 within_init_scope (chainer.links.NStepLSTM attribute),
 412
 within_init_scope (chainer.links.NStepRNNReLU
 attribute), 418
 within_init_scope (chainer.links.NStepRNNTanh at-
 tribute), 424
 within_init_scope (chainer.links.Parameter attribute), 428
 within_init_scope (chainer.links.PReLU attribute), 514
 within_init_scope (chainer.links.ResNet101Layers
 attribute), 570
 within_init_scope (chainer.links.ResNet152Layers
 attribute), 576
 within_init_scope (chainer.links.ResNet50Layers at-
 tribute), 563
 within_init_scope (chainer.links.Scale attribute), 434
 within_init_scope (chainer.links.SimplifiedDropconnect
 attribute), 509
 within_init_scope (chainer.links.StatefulGRU attribute),
 440
 within_init_scope (chainer.links.StatefulMGU attribute),
 451
 within_init_scope (chainer.links.StatefulPeepholeLSTM
 attribute), 461
 within_init_scope (chainer.links.StatefulZoneoutLSTM
 attribute), 466
 within_init_scope (chainer.links.StatelessGRU attribute),
 446
 within_init_scope (chainer.links.StatelessLSTM at-
 tribute), 472
 within_init_scope (chainer.links.StatelessMGU attribute),
 455

`within_init_scope` (chainer.links.Swish attribute), 519
`within_init_scope` (chainer.links.TheanoFunction attribute), 582
`within_init_scope` (chainer.links.VGG16Layers attribute), 542
`within_init_scope` (chainer.Sequential attribute), 612

X

`xp` (chainer.Chain attribute), 600
`xp` (chainer.ChainList attribute), 605
`xp` (chainer.Link attribute), 594
`xp` (chainer.links.BatchNormalization attribute), 479
`xp` (chainer.links.BatchRenormalization attribute), 484
`xp` (chainer.links.Bias attribute), 278
`xp` (chainer.links.Bilinear attribute), 283
`xp` (chainer.links.BinaryHierarchicalSoftmax attribute), 494
`xp` (chainer.links.BlackOut attribute), 499
`xp` (chainer.links.caffe.CaffeFunction attribute), 588
`xp` (chainer.links.ChildSumTreeLSTM attribute), 288
`xp` (chainer.links.Classifier attribute), 535
`xp` (chainer.links.Convolution2D attribute), 294
`xp` (chainer.links.ConvolutionND attribute), 300
`xp` (chainer.links.CRF1d attribute), 504
`xp` (chainer.links.Deconvolution2D attribute), 306
`xp` (chainer.links.DeconvolutionND attribute), 311
`xp` (chainer.links.DepthwiseConvolution2D attribute), 316
`xp` (chainer.links.DilatedConvolution2D attribute), 322
`xp` (chainer.links.EmbedID attribute), 327
`xp` (chainer.links.GoogLeNet attribute), 549
`xp` (chainer.links.GRU attribute), 332
`xp` (chainer.links.Highway attribute), 337
`xp` (chainer.links.Inception attribute), 343
`xp` (chainer.links.InceptionBN attribute), 348
`xp` (chainer.links.LayerNormalization attribute), 489
`xp` (chainer.links.Linear attribute), 354
`xp` (chainer.links.LocalConvolution2D attribute), 359
`xp` (chainer.links.LSTM attribute), 365
`xp` (chainer.links.Maxout attribute), 524
`xp` (chainer.links.MLPConvolution2D attribute), 371
`xp` (chainer.links.model.vision.resnet.ResNetLayers attribute), 557
`xp` (chainer.links.NaryTreeLSTM attribute), 376
`xp` (chainer.links.NegativeSampling attribute), 529
`xp` (chainer.links.NStepBiGRU attribute), 382
`xp` (chainer.links.NStepBiLSTM attribute), 388
`xp` (chainer.links.NStepBiRNNReLU attribute), 394
`xp` (chainer.links.NStepBiRNNTanH attribute), 400
`xp` (chainer.links.NStepGRU attribute), 406
`xp` (chainer.links.NStepLSTM attribute), 412
`xp` (chainer.links.NStepRNNReLU attribute), 418
`xp` (chainer.links.NStepRNNTanh attribute), 424
`xp` (chainer.links.Parameter attribute), 428

`xp` (chainer.links.PReLU attribute), 514
`xp` (chainer.links.ResNet101Layers attribute), 570
`xp` (chainer.links.ResNet152Layers attribute), 577
`xp` (chainer.links.ResNet50Layers attribute), 563
`xp` (chainer.links.Scale attribute), 434
`xp` (chainer.links.SimplifiedDropconnect attribute), 509
`xp` (chainer.links.StatefulGRU attribute), 440
`xp` (chainer.links.StatefulMGU attribute), 451
`xp` (chainer.links.StatefulPeepholeLSTM attribute), 461
`xp` (chainer.links.StatefulZoneoutLSTM attribute), 466
`xp` (chainer.links.StatelessGRU attribute), 446
`xp` (chainer.links.StatelessLSTM attribute), 472
`xp` (chainer.links.StatelessMGU attribute), 455
`xp` (chainer.links.Swish attribute), 519
`xp` (chainer.links.TheanoFunction attribute), 582
`xp` (chainer.links.VGG16Layers attribute), 542
`xp` (chainer.Parameter attribute), 130
`xp` (chainer.Sequential attribute), 612
`xp` (chainer.Variable attribute), 122

Z

`Zero` (class in chainer.initializers), 646
`zero_grads()` (chainer.links.Bilinear method), 282
`zero_grads()` (chainer.Parameter method), 127
`zero_grads()` (chainer.Variable method), 119
`zero_grads()` (chainer.Chain method), 599
`zero_grads()` (chainer.ChainList method), 604
`zero_grads()` (chainer.Link method), 593
`zero_grads()` (chainer.links.BatchNormalization method), 479
`zero_grads()` (chainer.links.BatchRenormalization method), 484
`zero_grads()` (chainer.links.Bias method), 277
`zero_grads()` (chainer.links.Bilinear method), 282
`zero_grads()` (chainer.links.BinaryHierarchicalSoftmax method), 494
`zero_grads()` (chainer.links.BlackOut method), 499
`zero_grads()` (chainer.links.caffe.CaffeFunction method), 588
`zero_grads()` (chainer.links.ChildSumTreeLSTM method), 288
`zero_grads()` (chainer.links.Classifier method), 535
`zero_grads()` (chainer.links.Convolution2D method), 294
`zero_grads()` (chainer.links.ConvolutionND method), 299
`zero_grads()` (chainer.links.CRF1d method), 504
`zero_grads()` (chainer.links.Deconvolution2D method), 306
`zero_grads()` (chainer.links.DeconvolutionND method), 311
`zero_grads()` (chainer.links.DepthwiseConvolution2D method), 316
`zero_grads()` (chainer.links.DilatedConvolution2D method), 322
`zero_grads()` (chainer.links.EmbedID method), 327

[zerograds\(\) \(chainer.links.GoogLeNet method\), 549](#)
[zerograds\(\) \(chainer.links.GRU method\), 332](#)
[zerograds\(\) \(chainer.links.Highway method\), 337](#)
[zerograds\(\) \(chainer.links.Inception method\), 342](#)
[zerograds\(\) \(chainer.links.InceptionBN method\), 348](#)
[zerograds\(\) \(chainer.links.LayerNormalization method\), 489](#)
[zerograds\(\) \(chainer.links.Linear method\), 353](#)
[zerograds\(\) \(chainer.links.LocalConvolution2D method\), 359](#)
[zerograds\(\) \(chainer.links.LSTM method\), 365](#)
[zerograds\(\) \(chainer.links.Maxout method\), 524](#)
[zerograds\(\) \(chainer.links.MLPConvolution2D method\), 370](#)
[zerograds\(\) \(chainer.links.model.vision.resnet.ResNetLayers method\), 556](#)
[zerograds\(\) \(chainer.links.NaryTreeLSTM method\), 376](#)
[zerograds\(\) \(chainer.links.NegativeSampling method\), 529](#)
[zerograds\(\) \(chainer.links.NStepBiGRU method\), 382](#)
[zerograds\(\) \(chainer.links.NStepBiLSTM method\), 388](#)
[zerograds\(\) \(chainer.links.NStepBiRNNReLU method\), 393](#)
[zerograds\(\) \(chainer.links.NStepBiRNNTanh method\), 399](#)
[zerograds\(\) \(chainer.links.NStepGRU method\), 405](#)
[zerograds\(\) \(chainer.links.NStepLSTM method\), 411](#)
[zerograds\(\) \(chainer.links.NStepRNNReLU method\), 417](#)
[zerograds\(\) \(chainer.links.NStepRNNTanh method\), 423](#)
[zerograds\(\) \(chainer.links.Parameter method\), 428](#)
[zerograds\(\) \(chainer.links.PReLU method\), 513](#)
[zerograds\(\) \(chainer.links.ResNet101Layers method\), 570](#)
[zerograds\(\) \(chainer.links.ResNet152Layers method\), 576](#)
[zerograds\(\) \(chainer.links.ResNet50Layers method\), 563](#)
[zerograds\(\) \(chainer.links.Scale method\), 433](#)
[zerograds\(\) \(chainer.links.SimplifiedDropconnect method\), 509](#)
[zerograds\(\) \(chainer.links.StatefulGRU method\), 440](#)
[zerograds\(\) \(chainer.links.StatefulMGU method\), 450](#)
[zerograds\(\) \(chainer.links.StatefulPeepholeLSTM method\), 461](#)
[zerograds\(\) \(chainer.links.StatefulZoneoutLSTM method\), 466](#)
[zerograds\(\) \(chainer.links.StatelessGRU method\), 445](#)
[zerograds\(\) \(chainer.links.StatelessLSTM method\), 472](#)
[zerograds\(\) \(chainer.links.StatelessMGU method\), 455](#)
[zerograds\(\) \(chainer.links.Swish method\), 519](#)
[zerograds\(\) \(chainer.links.TheanoFunction method\), 582](#)
[zerograds\(\) \(chainer.links.VGG16Layers method\), 542](#)
[zerograds\(\) \(chainer.Sequential method\), 612](#)
[ZippedImageDataset \(class in chainer.datasets\), 705](#)
[zoneout\(\) \(in module chainer.functions\), 235](#)