
pytorch-pfn-extras

Preferred Networks, Inc.

Jan 19, 2024

CONTENTS

| | | |
|----------|------------------------------------|------------|
| 1 | User Guide | 3 |
| 1.1 | Quick Start | 3 |
| 1.2 | Trainer | 9 |
| 1.3 | Extensions | 15 |
| 1.4 | Utilities | 19 |
| 2 | API Reference | 37 |
| 2.1 | Package | 37 |
| 2.2 | Training Loop | 580 |
| 2.3 | Distributed Training | 582 |
| 2.4 | Check Pointing | 582 |
| 2.5 | Lazy Modules | 583 |
| 2.6 | ONNX | 583 |
| 2.7 | Datasets | 583 |
| 2.8 | Config | 584 |
| 2.9 | NumPy/CuPy Compatibility | 584 |
| | Python Module Index | 585 |
| | Index | 587 |

pytorch-pfn-extras (PPE) is a collection of supplementary components to accelerate research and development in PyTorch.

1.1 Quick Start

1.1.1 Quick Start

First, pytorch-pfn-extras organizes the training code implemented using PyTorch using the Trainer/Evaluator classes.

Next, it provides the following interfaces for training PyTorch models.

1. Addition of extensions for analysis and visualization
2. Runtime changes
3. Addition of custom training steps
4. Custom data handling

Step 1: Use Trainer

First, pass to the Trainer the Model and Optimizer you want to train.

Listing 1: quick_start_trainer.py

```
import pytorch_pfn_extras as ppe
import torch

class Model(torch.nn.Module):
    def __init__(self, *args, **kwargs) -> None:
        super().__init__(*args, **kwargs)
        self.linear = torch.nn.Linear(in_features=64, out_features=2)
        self.criterion = torch.nn.NLLLoss()

    def forward(self, x, target):
        y = self.linear.forward(x).log_softmax(dim=1)
        loss = self.criterion.forward(y, target)
        return {"loss": loss}

model = Model()
optimizer = torch.optim.SGD(model.parameters(), lr=0.01)
```

(continues on next page)

(continued from previous page)

```
device = (
    "cuda:0" # or any other PyTorch devices ('cpu', etc.) or PPE runtime names
)
epochs = 3
# Create a trainer with the defined model, optimizer, and other parameters
trainer = ppe.engine.create_trainer(
    models=model,
    optimizers=optimizer,
    max_epochs=epochs,
    evaluator=ppe.engine.create_evaluator(
        models=model,
        device=device,
    ),
    device=device,
)

# Send the model to device(GPU) for computation
ppe.to(model, device=device)

batch_size = 10
# Create 10 batches of random training data with dimension (batch_size x 64)
training_data = [
    {
        "x": torch.rand((batch_size, 64)),
        "target": torch.ones((batch_size,), dtype=torch.long),
    }
    for _ in range(10)
]
# Create 10 batches of random validation data with dimension (batch_size x 64)
validation_data = [
    {
        "x": torch.rand((batch_size, 64)),
        "target": torch.ones((batch_size,), dtype=torch.long),
    }
    for _ in range(10)
]

# Start the training and validation of the model
trainer.run(train_loader=training_data, val_loader=validation_data)

print("Finish training!")
```


Step 2: Get Log

Next, collect the logs of the training progress.

Listing 2: quick_start_log.py

```
import pytorch_pfn_extras as ppe
import torch

class Model(torch.nn.Module):
    def __init__(self, *args, **kwargs) -> None:
        super().__init__(*args, **kwargs)
        self.linear = torch.nn.Linear(in_features=64, out_features=2)
        self.criterion = torch.nn.NLLLoss()

    def forward(self, x, target):
        y = self.linear.forward(x).log_softmax(dim=1)
        loss = self.criterion.forward(y, target)
        return {"loss": loss}

model = Model()
optimizer = torch.optim.SGD(model.parameters(), lr=0.01)

device = "cuda:0"
epochs = 3
trainer = ppe.engine.create_trainer(
    models=model,
    optimizers=optimizer,
    max_epochs=epochs,
    evaluator=ppe.engine.create_evaluator(
        models=model,
        device=device,
        options={
            "eval_report_keys": [
                "loss"
            ], # Let the value of the loss be notified to the LogReport.
        },
    ),
    device=device,
    options={
        "train_report_keys": [
            "loss"
        ], # Let the value of the loss be notified to the LogReport.
    },
)

trainer.extend(
    ppe.training.extensions.LogReport()
) # It is an extension to collect parameters reported during training.

ppe.to(model, device=device)
```

(continues on next page)

(continued from previous page)

```

batch_size = 10
training_data = [
    {
        "x": torch.rand((batch_size, 64)),
        "target": torch.ones((batch_size,), dtype=torch.long),
    }
    for _ in range(10)
]
validation_data = [
    {
        "x": torch.rand((batch_size, 64)),
        "target": torch.ones((batch_size,), dtype=torch.long),
    }
    for _ in range(10)
]

trainer.run(train_loader=training_data, val_loader=validation_data)

print("Finish training!")

```

The logs of the collected learning progress are output to `./result/log`.

Step 3: Display of progress

Make it possible to check the progress of the learning.

Listing 3: quick_start_progress.py

```

import pytorch_pfn_extras as ppe
import torch

class Model(torch.nn.Module):
    def __init__(self, *args, **kwargs) -> None:
        super().__init__(*args, **kwargs)
        self.linear = torch.nn.Linear(in_features=64, out_features=2)
        self.criterion = torch.nn.NLLLoss()

    def forward(self, x, target):
        y = self.linear.forward(x).log_softmax(dim=1)
        loss = self.criterion.forward(y, target)
        return {"loss": loss}

model = Model()
optimizer = torch.optim.SGD(model.parameters(), lr=0.01)

device = "cuda:0"
epochs = 3
trainer = ppe.engine.create_trainer(
    models=model,

```

(continues on next page)

(continued from previous page)

```

optimizers=optimizer,
max_epochs=epochs,
evaluator=ppe.engine.create_evaluator(
    models=model,
    device=device,
    options={
        "eval_report_keys": ["loss"],
    },
),
device=device,
options={
    "train_report_keys": ["loss"],
},
)

trainer.extend(ppe.training.extensions.LogReport())
trainer.extend(ppe.training.extensions.ProgressBar())
trainer.extend(
    ppe.training.extensions.PrintReport( # Displays the collected logs interactively.
        [
            "epoch", # epoch, iteration, elapsed_time are automatically collected by
            ↪LogReport.
            "iteration",
            "elapsed_time",
            "train/loss", # The parameters specified by train_report_keys are collected
            ↪under keys with the 'train/' prefix.
            "val/loss", # The parameters specified by eval_report_keys are collected
            ↪under keys with the 'val/' prefix.
        ],
    )
)

ppe.to(model, device=device)

batch_size = 10
training_data = [
    {
        "x": torch.rand((batch_size, 64)),
        "target": torch.ones((batch_size,), dtype=torch.long),
    }
    for _ in range(10)
]
validation_data = [
    {
        "x": torch.rand((batch_size, 64)),
        "target": torch.ones((batch_size,), dtype=torch.long),
    }
    for _ in range(10)
]

trainer.run(train_loader=training_data, val_loader=validation_data)

```

(continues on next page)

(continued from previous page)

```
print("Finish training!")
```

Step 4: Save Model

Finally, save the trained model.

Listing 4: quick_start_save.py

```
import pytorch_pfn_extras as ppe
import torch

class Model(torch.nn.Module):
    def __init__(self, *args, **kwargs) -> None:
        super().__init__(*args, **kwargs)
        self.linear = torch.nn.Linear(in_features=64, out_features=2)
        self.criterion = torch.nn.NLLLoss()

    def forward(self, x, target):
        y = self.linear.forward(x).log_softmax(dim=1)
        loss = self.criterion.forward(y, target)
        return {"loss": loss}

model = Model()
optimizer = torch.optim.SGD(model.parameters(), lr=0.01)

device = "cuda:0"
epochs = 3
trainer = ppe.engine.create_trainer(
    models=model,
    optimizers=optimizer,
    max_epochs=epochs,
    evaluator=ppe.engine.create_evaluator(
        models=model,
        device=device,
        options={
            "eval_report_keys": ["loss"],
        },
    ),
    device=device,
    options={
        "train_report_keys": ["loss"],
    },
)

trainer.extend(ppe.training.extensions.LogReport())
trainer.extend(ppe.training.extensions.ProgressBar())
trainer.extend(
    ppe.training.extensions.PrintReport( # Displays the collected logs interactively.
```

(continues on next page)

(continued from previous page)

```

        [
            "epoch", # epoch, iteration, elapsed_time are automatically collected by
            ↪LogReport.
            "iteration",
            "elapsed_time",
            "train/loss", # The parameters specified by train_report_keys are collected
            ↪under keys with the 'train/' prefix.
            "val/loss", # The parameters specified by eval_report_keys are collected
            ↪under keys with the 'val/' prefix.
        ],
    )
)
trainer.extend(
    ppe.training.extensions.snapshot(target=model)
) # Save the model parameters after each epoch.

ppe.to(model, device=device)

batch_size = 10
training_data = [
    {
        "x": torch.rand((batch_size, 64)),
        "target": torch.ones((batch_size,), dtype=torch.long),
    }
    for _ in range(10)
]
validation_data = [
    {
        "x": torch.rand((batch_size, 64)),
        "target": torch.ones((batch_size,), dtype=torch.long),
    }
    for _ in range(10)
]

trainer.run(train_loader=training_data, val_loader=validation_data)

print("Finish training!")

```

The model parameters are stored with a file name that includes the time they were saved under `./result`.

Snapshots are generated using `state_dict()`. Please refer to the official PyTorch [docs](#) for how to load the model.

1.2 Trainer

1.2.1 Trainer and Evaluator

The Trainer and Evaluator provide the device-agnostic training framework for PyTorch. These APIs abstract the training process using different *runtimes*, handlers, and *logics*.

Concepts

- **Trainer** (`ppe.engine.create_trainer()`) abstracts the training loop, built on top of the `ExtensionsManager`.
- **Evaluator** (`ppe.engine.create_evaluator()`) abstracts the evaluation step and invoked from the Trainer (usually once in every epoch).
- **Runtime** (`ppe.runtime.BaseRuntime`) represents an environment used to execute models. Device-specific implementations will reside here. PPE provides the default Runtime that supports the PyTorch-native devices (`ppe.runtime.PyTorchRuntime`).
- **Handler** (`ppe.handler.Handler`) is a layer to support device-agnostic training. This is considered as a low-level API and in most cases users can just use the Handler provided by PPE.
- **Logic** (`ppe.handler.Logic`) is a set of callback functions that define the training logic (`optimizer.zero_grad()`, forward, backward, `optimizer.step()`). You can inherit the class and define your own training flow in case you need more complex training processes such as GAN.
- **Model** is a `torch.nn.Module` used for training and evaluation, whose inputs are dicts or keyword arguments and outputs of the forward pass is a dict.

Note that the default logic will perform backward in tensors returned by `model.forward` so you will need to perform the loss calculation inside the model itself.

Trainer at a glance

```
import torch
import torch.nn.functional as F

import pytorch_pfn_extras as ppe

class MyModel(torch.nn.Module):
    def __init__(self):
        super().__init__()
        self.w = torch.nn.LazyLinear(1)

    def forward(self, *, x, target):
        y = self.w(x)
        loss = F.nll_loss(y, target)
        prefix = 'train' if self.training else 'val'
        ppe.reporting.report({f'{prefix}/loss': loss.item()})
        return {'loss': loss}

model = MyModel()
optim = torch.optim.SGD(model.parameters(), lr=0.01)

extensions = [
    ppe.training.extensions.LogReport(),
    ppe.training.extensions.ProgressBar(),
    ppe.training.extensions.PrintReport(
        ['epoch', 'iteration', 'train/loss', 'val/loss']),
]
```

(continues on next page)

(continued from previous page)

```

device = 'cuda:0' # or any other PyTorch devices ('cpu', etc.) or PPE runtime names
epochs = 10
trainer = ppe.engine.create_trainer(
    model,
    optim,
    epochs,
    evaluator=ppe.engine.create_evaluator(
        model,
        device=device,
        progress_bar=True,
    ),
    device=device,
    extensions=extensions,
)

# Move the model to the device. This is almost equivalent to
# `model.to(device)`, but supports PPE runtimes as well as the PyTorch's
# built-in devices.
ppe.to(model, device)

# Using dummy data to illustrate the minimal working example.
# Notice that dict keys match with the kwargs of the forward method.
train_loader = torch.utils.data.DataLoader(
    [{'x': torch.rand(10, 64), 'target': torch.tensor([1])} for _ in range(1)],
    num_workers=8)
val_loader = torch.utils.data.DataLoader(
    [{'x': torch.rand(10, 64), 'target': torch.tensor([1])} for _ in range(1)],
    num_workers=8)

trainer.run(train_loader, val_loader)

```

Snapshot

To obtain and save the trained model for later use you can use the *Snapshot* extension, or directly invoke *state_dict* on the trainer itself.

Handler

The `ppe.handler.Handler` object is used to help the trainer and evaluator objects in the *Logic* and *Runtime* manipulation. This class should ideally never be overridden by the user if the desired functionality can be achieved through subclassing `BaseLogic` or `BaseRuntime`.

The handler object's main responsibility is to inspect all the submodules of a module to obtain the runtimes they have associated, and then execute their callbacks accordingly. In addition, it drives the actual model execution by using the user provided *Logic* object and deals with asynchronous execution in runtimes that provide support for it.

Runtime

By inheriting `ppe.runtime.BaseRuntime` and implementing your own runtime, you can use your non-standard devices with the training loop.

```
class MyRuntime(BaseRuntime):  
    ...  
  
# Register MyRuntime with device name "mydev"  
ppe.runtime.runtime_registry.register('mydev', MyRuntime)  
  
ppe.to(module_or_tensor, 'mydev')
```

See *Runtimes for Custom Devices* if you are interested in implementing your own runtime.

1.2.2 Logic for Custom Training and Evaluation

In the training and evaluation engines, `ppe.handler.BaseLogic` API is in charge of abstracting the algorithmic details of the training and evaluation loops.

Logic is an object that defines multiple callbacks used through the training and evaluation processes. With logic, we can implement training of complex models such as GANs.

Users wanting to define their own Logic for training can inherit from `ppe.handler.Logic` which implements the training and evaluation steps to train a single module.

Logic functions are not expected to be directly called by the user. They will be invoked by the Trainer and Evaluator engines.

Default Logic (`ppe.handler.Logic`)

PPE provides a default logic that performs the forward/backward/optimizer loop for a single model. This logic allows using some torch features such as AMP autocast and GradScaler and performs the backward pass on the outputs specified by the config option `backward_outputs`.

CodeBlock Logic (`ppe.handler.Logic`)

With the CodeBlock API, we provide a basic logic that uses it to perform the training of a single model. Similarly to the default logic AMP features are supported but by means of the Runtime. For more information check the codeblock documentation.

1.2.3 Runtimes for Custom Devices

Note: This documentation is intended for those implementing the own device backend for PPE training framework. Most users can just skip this chapter.

The `ppe.runtime.BaseRuntime` API is in charge of abstracting the device details and performing the movement of data and modules to the corresponding device.

A runtime is an object that defines multiple callbacks used through the training, evaluation, and regular model calls. With runtimes, we can implement training in devices other than `cpu` or `gpu` with minimal changes to the user code.

Users wanting to override only a few callbacks can inherit from `ppe.runtime.PyTorchRuntime` which implements the basic functionality for cpu and gpu devices.

Runtimes must be registered by calling the `ppe.runtime.runtime_registry.register(device_name, runtime_class)` function for them to be discoverable.

Use of `ppe.to` to transfer modules and batches to custom devices

If you have defined a new runtime for a custom device the `ppe.to` function allows moving a module or a tensor to the new device by invoking the `Runtime.move_tensor` and `Runtime.move_module` when needed.

The module will be tagged by adding a attribute named `_ppe_runtime` that holds the needed runtime. It is the responsibility of the user custom runtime to perform the actual movement to the device and apply all the transformations needed to a module so it can be correctly executed.

Usually, runtime writers will need to replace the given module forward function by a new one that performs the actual device execution.

```
class MyModule(torch.nn.Module):
    def __init__(self):
        super().__init__()
        self.layer = torch.nn.Linear(10, 10)

    def forward(x):
        return self.layer(x)

class MyMagicDeviceRuntime(ppe.runtime.BaseRuntime):
    def _device_forward(self, args):
        return run_batch_in_my_device(args)

    def move_module(self, module):
        # Registers a hook to initialize the module on the first batch
        # execution
        def hook(module, *args):
            module._ppe_runtime.initialize_module(module, args)

        self.hook = module.register_forward_pre_hook(hook)
        # Change the module forward to do the computation in the device
        module.forward = self._device_forward

    def initialize_module(self, module, loader_or_batch, optimizer=None):
        create_the_module_in_my_device(module, loader_or_batch, optimizer)

# Register the runtime class
ppe.runtime.runtime_registry.register('my_device', MyMagicDeviceRuntime)

# Create a regular module
module = MyModule()
# Move the module to the device
ppe.to(module, device='my_device')

for x in my_dataloader:
    # The first iteration will create the module in the device
    # and the next ones will directly execute the module in the device instead
```

(continues on next page)

```
# of executing the regular pytorch `forward` call.
y = model(x)
```

Please note that this is an oversimplified description and that developing a runtime that is 100% compatible with PyTorch requires to wrap the substitute forward function with `torch.autograd.Function` among several other concerns such as `state_dict` manipulation to ensure correctness.

Runtime Registry

When creating a new Runtime class for custom needs, they need to be registered in a global `runtime_registry` object as detailed above. This object is of the `_RuntimeRegistry` type and it maintains a map of strings and Runtime types. The keys are the devices passed to `ppe.to` and the types will be the type of the Runtime object that `ppe.to` will use to treat the module or tensor. Beware that users are not supposed to interact directly with this class, only with the `runtime_registry.register` to register new runtimes.

1.2.4 CodeBlocks for Abstracting Logic Steps

The `ppe.handler.CodeBlock` API provides a mean of abstracting the actions that are possible to be done in a model in a device agnostic way.

Currently there is support for two different actions using CodeBlock.

- **:function: `ppe.handler.update_parameters <pytorch_pfn_extras.handler.update_parameters>`**
takes a model, an optimizer and returns a CodeBlock object that performs the forward, backward and optimizer step at once.
- **:function: `ppe.handler.forward <pytorch_pfn_extras.handler.forward>`**
takes a model and returns a CodeBlock object that performs only the forward pass.

Executing CodeBlocks

For executing CodeBlock objects we need to add an **:method: `ppe.runtime.BaseRuntime.execute <pytorch_pfn_extras.runtime.BaseRuntime.execute>`** to the corresponding Runtime class. This method takes a CodeBlock and uses the information in the object to execute the CodeBlock in the device. Note that the **:method: `ppe.runtime.PyTorchRuntime.execute <pytorch_pfn_extras.runtime.PyTorchRuntime.execute>`** method provides support for using PyTorch AMP with autocast or gradient scaling if needed.

Moreover, you can execute CodeBlock objects outside the training API.

```
ppe.to(model, "cuda:0")
cblock = ppe.handler.update_parameters(model, optimizer)
outs = cblock(input_batch)
```

The only requirement is that the associated model has been assigned a device using `ppe.to`.

1.3 Extensions

1.3.1 Extensions Manager

Extensions Manager provides an interface to extend your training loop, by integrating it into your manual training loop or Ignite.

Extensions

See the [API Reference](#) for the list of built-in extensions.

How to use

Create an `ExtensionsManager` object and then wrap the iteration of your training loop inside the manager. `run_iteration()` context manager.

An example follows:

```
import pytorch_pfn_extras as ppe
from pytorch_pfn_extras.training import extensions

import time
import math

max_epoch = 10
iters_per_epoch = 938

# manager.extend(...) also works
my_extensions = [extensions.LogReport(),
                 extensions.ProgressBar(),
                 extensions.PrintReport(['epoch', 'iteration', 'sin', 'cos'])]

models = {}
optimizers = []
manager = ppe.training.ExtensionsManager(
    models, optimizers, max_epoch,
    extensions=my_extensions,
    iters_per_epoch=iters_per_epoch)

for epoch in range(max_epoch):
    for i in range(iters_per_epoch):
        with manager.run_iteration():
            ppe.reporting.report({
                'sin': math.sin(i * 2 * math.pi / iters_per_epoch),
                'cos': math.cos(i * 2 * math.pi / iters_per_epoch),
            })
            time.sleep(0.001)
```

In the examples folder there is a mnist using all the available extensions.

Usage with Ignite

Ignite is supported by using the `IgniteExtensionsManager` with the trainer as the first argument.

The user needs to define an ignite event to report the appropriated metrics for the extensions to use them.

```
manager = ppe.training.IgniteExtensionsManager(
    trainer, models, optimizers, epochs,
    extensions=my_extensions)

@trainer.on(Events.ITERATION_COMPLETED)
def report_loss(engine):
    ppe.reporting.report({'train/loss': engine.state.output})
```

Using Evaluators

Regular PyTorch

In order to report the results of the evaluation so they can be accessed by other extensions, an `Evaluation` extension needs to be created with the argument `eval_func` set to a function that gets the current data and target batches as parameters and reports the needed metrics. [Example](#)

The test function looks has the following signature

```
def test(args, model, device, data, target):
```

and is invoked once per batch in the validation dataloader. It is important to report the current validation loss or accuracy in order to the log report to see it.

```
def test(args, model, device, data, target):
    ...
    # Final result will be average of averages of the same size
    test_loss += F.nll_loss(output, target, reduction='mean').item()
    ppe.reporting.report({'val/loss': test_loss})
    pred = output.argmax(dim=1, keepdim=True)
    correct += pred.eq(target.view_as(pred)).sum().item()
    ppe.reporting.report({'val/acc': correct/len(data)})
```

Ignite

Just use the `IgniteEvaluator` extension with the ignite created evaluator as the first parameter and you are ready to go. [Example](#) The metrics defined when creating the evaluator with `create_supervised_evaluator` will be automatically reported

```
create_supervised_evaluator(model, metrics={'acc': Accuracy(), 'loss': Loss(F.nll_loss)}
↪, device)
```

Snapshots

It is possible to take snapshots by using the `snapshot` training extension just as in chainer.

Whenever the extension is triggered, it saves the status of the optimizer, model and extensions to the output folder in the same way as chainer. To load the snapshot and continue the training call `torch.load` and use the `ExtensionsManager.load_state_dict` to resume the training. The snapshots can be used outside the `pytorch-pfn-extras` module just by accessing the models, or optimizers fields of the loaded state.

Extensions execution order

The supported extensions honours the chainer priorities for execution. However, when using Ignite. Chainer extensions are executed after any user-defined ignite events. The idea is to use ignite events to report the metrics of the model, and after this, Chainer extensions will be executed in the chainer defined order.

If you want to execute an event-handler in between chainer extensions, create a Chainer-like extension and access the ignite engine on the `.engine` attribute of the manager object passed as a parameter when your extension is called.

1.3.2 Creating Extensions

It is possible to create an extension just by passing a function which receives the manager object as an argument to the manager extend call

```
def my_extension(manager):
    print('Epoch-Iteration: {}-{}'.format(manager.epoch, manager.iteration))

manager.extend(my_extension, trigger=(1, 'iteration'))
```

It is also possible to create extensions using the `ppe.training.extension.make_extension` decorator to add a specific trigger, default_name, priority. In addition, initializer, finalizer and on_error functions can be specified as well.

```
@ppe.training.extension.make_extension(finalizer=lambda: print('done'))
def my_extension(manager):
    print('Epoch-Iteration: {}-{}'.format(manager.epoch, manager.iteration))
```

Finally, it is possible to create an extension by subclassing the `ppe.training.extensions.Extension` class as shown below.

```
import pytorch_pfn_extras as ppe

class MyExtension(ppe.training.extension.Extension):
    def __init__(self, args):
        self.args = args

    def initialize(self, manager):
        """
        Automatically called before training. Optional.
        """
        pass

    def __call__(self, manager):
        """
```

(continues on next page)

```

        Called when the associated trigger is fired.
        """
        print('Epoch-Iteration: {}-{}'.format(manager.epoch, manager.iteration))

    def state_dict(self):
        """
        Used to serialize the state. Optional.
        """
        return {'args': self.args}

    def load_state_dict(self, state):
        """
        Used to deserialize the state. Optional.
        """
        self.args = state['args']

```

1.3.3 Reporting

`reporting.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.

```

import pytorch_pfn_extras as ppe

reporter = ppe.reporting.Reporter()
observer = object()
reporter.add_observer('my_observer', observer)
observation = {}

with reporter.scope(observation):
    reporter.report({'x': 1}, observer)

print(observation)
# outputs: {'my_observer/x': 1}

```

There is also a global API to add values:

```

import pytorch_pfn_extras as ppe

reporter = ppe.reporting.Reporter()
observer = object()
reporter.add_observer('my_observer', observer)

observation = {}
with reporter:
    with ppe.reporting.report_scope(observation):
        ppe.reporting.report({'x': 1}, observer)

print(observation)
# outputs: {'my_observer/x': 1}

```

The most important application of Reporter is to report observed values from different parts of the model in the training and validation procedures. `ExtensionsManager` objects hold their own `Reporter` object with the parameters of the target module registered as observers. `report()` can be used inside the modules to report the observed values (e.g., training loss, accuracy, activation statistics, etc.).

1.3.4 Distributed Snapshot

To take snapshots when using `torch.distributed` the only needed step is to provide the `saver_rank` keyword argument to the regular snapshot extension.

```
# saver_rank is the MPI rank which will write the actual snapshot.
snapshot = extensions.snapshot(saver_rank=saver_rank)
```

To resume the training, snapshots are loaded in every worker by using the `ExtensionsManager.load_state_dict` method, or the `extensions.snapshot.autoload` keyword argument.

1.4 Utilities

1.4.1 Lazy Modules

Lazy modules can automatically infer shapes of parameters based on the shape of the data given to the first forward invocation.

Following modules are provided:

- `ppe.nn.LazyBatchNorm1d`, `ppe.nn.LazyBatchNorm2d`, `ppe.nn.LazyBatchNorm3d`
 - Module that behaves as `torch.nn.BatchNorm[123]d` but `num_features` can be set to `None`.
 - These modules are now included as a part of PyTorch 1.9 release (`torch.nn.LazyBatchNormXd`, `pull-request`).

The following modules are now considered deprecated as now included as a part of PyTorch 1.8 release:

- `ppe.nn.LazyLinear`
 - Module that behaves as `torch.nn.Linear` but `in_features` can be set to `None`.
 - PyTorch-native implementation: (`torch.nn.LazyLinear`, `pull-request`)
- `ppe.nn.LazyConv1d`, `ppe.nn.LazyConv2d`, `ppe.nn.LazyConv3d`
 - Module that behaves as `torch.nn.Conv[123]d` but `in_channels` can be set to `None`.
 - PyTorch-native implementation: (`torch.nn.LazyConvXd`, `pull-request`)

Now that all lazy modules are merged to the upstream, we encourage you to migrate to PyTorch's lazy modules. We will keep these implementations only for backward compatibility.

Note that you need to run a “dummy” forward to initialize lazy parameters. See the example below:

```
import torch
import torch.nn.functional as F

import pytorch_pfn_extras as ppe

class Net(torch.nn.Module):
```

(continues on next page)

(continued from previous page)

```

def __init__(self):
    super().__init__()
    self.conv1 = ppe.nn.LazyConv2d(None, 20, 5, 1)
    self.conv2 = ppe.nn.LazyConv2d(None, 50, 5, 1)
    self.fc1 = ppe.nn.LazyLinear(None, 500)
    self.fc2 = ppe.nn.LazyLinear(None, 10)

def forward(self, x):
    x = F.relu(self.conv1(x))
    x = F.max_pool2d(x, 2, 2)
    x = F.relu(self.conv2(x))
    x = F.max_pool2d(x, 2, 2)
    x = x.flatten(start_dim=1)
    x = F.relu(self.fc1(x))
    x = self.fc2(x)
    return F.log_softmax(x, dim=1)

```

```
model = Net()
```

```
# Initialize lazy parameters.
```

```
dummy_input = ...
```

```
model(dummy_input)
```

```
# Pass parameters to the optimizer.
```

```
optimizer = torch.optim.SGD(
    model.parameters(), lr=args.lr, momentum=args.momentum)
```

```
# Run training loop.
```

```
# ...
```

You need to run a dummy forward before passing parameters to optimizers; otherwise optimizers cannot refer to lazily-initialized parameters. You will get a warning if you pass uninitialized lazy parameters to optimizers:

```

>>> model = ppe.nn.LazyLinear(None, 10)
>>> optimizer = torch.optim.SGD(model.parameters(), lr=0.1)
/.../pytorch-pfn-extras/pytorch_pfn_extras/nn/modules/lazy.py:127: UserWarning:
  Use of uninitialized lazy parameter in Optimizer has been detected.
  Maybe you forgot to run forward before passing `module.parameters()` to the
  ↪optimizer?

```

- *Config*
 - *Basic*
 - *Substitution*
 - * *Callable Substitution*
 - * *Substitution by Path*
 - * *Substitution by Attribute*
 - * *Default Value by Path Substitution*
 - * *Ignore Substitution*

* *Lazy Evaluation*

1.4.2 Config

Basic

```
from pytorch_pfn_extras.config import Config
import yaml
pre_eval_config = yaml.load('''
foo:
  bar: 'bar_value'
  ls:
    - 'first'
    - key0: 'value0'
      key1: 'value1'
baz: 'baz_value'
''')
config = Config(pre_eval_config)
```

Accessing config values:

```
print(config['/foo/ls/0'])
# 'first'
print(config['/foo/ls/1/key0'])
# 'value0'
print(config['/foo/ls'])
# ['first', {'key0': 'value0', 'key1': 'value1'}]
print(config['/baz'])
# 'baz_value'
```

Substitution

Callable Substitution

You could replace a value as the return value of a callable.

- `types` is an additional input to `Config`. `types` is a mapping from a callable's name to the actual callable.
- A sub-dictionary containing the key `type` invokes callable substitution.

```
pre_eval_config = yaml.load('''
name:
  type: concat
  x0: 'First'
  x1: 'Last'
''')

types = {
    'concat': lambda x0, x1: x0 + ' ' + x1
}

config = Config(pre_eval_config, types)
```

(continues on next page)

(continued from previous page)

```
# the value returned by
# concat(x0='First', x1='Last')
print(config['/name'])
# 'First Last'
```

Nested

```
pre_eval_config = yaml.load('''
name:
  type: concat
  x0: 'First'
  x1:
    type: concat
    x0: 'Middle'
    x1: 'Last'
''')
types = {
  'concat': lambda x0, x1: x0 + ' ' + x1
}
config = Config(pre_eval_config, types)
print(config['/name'])
# First Middle Last
```

Class

```
pre_eval_config = yaml.load('''
dataset:
  type: Dataset
  n_class: 10
''')

class Dataset(object):

    def __init__(self, n_class):
        self.n_class = n_class

types = {
  'Dataset': Dataset,
}

config = Config(pre_eval_config, types)
print(isinstance(config['/dataset'], Dataset))
# True
```

Substitution by Path

Absolute

@/absolute/path is replaced by the value at /absolute/path.

```
pre_eval_config = yaml.load('''
foo: 'FOO'
boo:
  baz: '@/foo'
''')
config = Config(pre_eval_config)
print(config['/boo/baz'])
# FOO
```

Relative

Relative path is also possible using @relative/path.

```
pre_eval_config = yaml.load('''
foo: 'FOO'
boo:
  baz: '@../foo'
''')
config = Config(pre_eval_config)
print(config['/boo/baz'])
# FOO
```

Substitution by Attribute

@/path/to/obj.attr_name is replaced by:

1. Use substitution by path to get an object at /path/to/obj.
2. Replace the config value by getattr(obj, attr_name), where obj is obtained at step 1.

```
pre_eval_config = yaml.load('''
dataset:
  type: Dataset
  n_class: 10
n_data: '@/dataset.n_data'
''')

class Dataset(object):

    def __init__(self, n_class):
        self.n_class = n_class
        self.n_data = 4

types = {
    'Dataset': Dataset,
```

(continues on next page)

(continued from previous page)

```

}

config = Config(pre_eval_config, types)
print(config['/n_data'])
# 4

```

Default Value by Path Substitution

`customize_type` is a decorator that sets default argument values by path substitution.

```

from pytorch_pfn_extras.config import customize_type

pre_eval_config = yaml.load('''
dataset:
  type: Dataset
  n_class: 5
''')

# If n_class is not passed, the value would be config['/n_class'].
# Both absolute and relative paths are allowed.
@customize_type(n_class='/n_class')
class Dataset(object):

    def __init__(self, n_class):
        self.n_class = n_class

types = {
    'Dataset': Dataset,
}

config = Config(pre_eval_config, types)
print(config['/dataset'].n_class)
# 5

```

Ignore Substitution

Access using `config['!/path']` instead of `config['/path']`.

```

pre_eval_config = yaml.load('''
name:
  type: concat
  x0: 'First'
  x1: 'Last'
''')

types = {
    'concat': lambda x0, x1: x0 + ' ' + x1
}

```

(continues on next page)

(continued from previous page)

```

config = Config(pre_eval_config, types)
print(config['!/name'])
# {'type': 'concat', 'x0': 'First', 'x1': 'Last'}

```

Lazy Evaluation

Callable substitution is lazily executed. This means that callables that are not dependent on the accessed value do not get executed.

```

pre_eval_config = yaml.load('''
foo:
  - type: f0
  - '@/bar'
bar:
  type: f1
baz:
  type: f2
''')

def f0():
    print('f0 called')
    return 'f0_return'

def f1():
    print('f1 called')
    return 'f1_return'

def f2():
    print('f2 called')
    return 'f2_return'

types = {
    'f0': f0,
    'f1': f1,
    'f2': f2,
}

config = Config(pre_eval_config, types)
config['/foo']  # f2 does not get called
# f0 called
# f1 called

```

1.4.3 pytorch_pfn_extras.onnx

Extensions to `torch.onnx.export`.

Installation

```
pip3 install "pytorch-pfn-extras[onnx]"
```

Or

1. Install pytorch-pfn-extras normally
2. Install onnx with `pip install onnx==1.7.0`

API

`pytorch_pfn_extras.onnx.export_testcase`

Instead of specifying file name in `torch.onnx.export`, `pytorch_pfn_extra.onnx.export_testcase` specifies directory to output ONNX model and test case in/out.

```
import torch
import torch.nn as nn
model = nn.Sequential(nn.Linear(5, 10, bias=False))
x = torch.zeros((2, 5))

import pytorch_pfn_extras.onnx as tou
tou.export_testcase(model, x, '/path/to/output')
```

Directory structure with following will be generated to `/path/to/output`:

```
$ tree /path/to/output
/path/to/output
├── meta.json
├── model.onnx
├── test_data_set_0
│   ├── input_0.pb
│   └── output_0.pb
```

- This directory structure format is inspired by ONNX official test data set: (Example: [node](#)). PyTorch's ONNX tests use this format too. (Reference: `export_onnx_tests_generator.py`)
 - There are scripts in `chainer-compiler/utils` to run inference in major runtime with the directory structure. For example to inference with ONNXRuntime, run `$ python run_onnx_onnxruntime.py /path/to/output` to use `input_N.pb` as input and compare numerically with its output `output_N.pb` (N is the index of test case).
- By default `meta.json` is generated too to track git infos, date times, etc. Add `metadata=False` argument to suppress this.

out_grad option

If `out_grad=True` is specified gradient will be dumped too, which is useful for debugging backward. `gradient_N.pb` and `gradient_input_N.pb` would be dumped to test case directory with in/out data. `gradient_input_N.pb` is the initial value of backward, and it's default value is ones tensor with same shape of output. Use `out_grad` to specify custom initial value (`torch.Tensor` type) for it.

```
model = nn.Sequential(nn.Linear(5, 10, bias=False))
x = torch.zeros((2, 5))

import pytorch_pfn_extras.onnx as tou
tou.export_testcase(model, x, '/path/to/output', out_grad=True)
```

```
$ tree /path/to/output
/path/to/output
├── meta.json
├── model.onnx
└── test_data_set_0
    ├── gradient_0.pb
    ├── gradient_input_0.pb
    ├── input_0.pb
    └── output_0.pb
```

model_overwrite option

Use `model_overwrite` option to create multiple data set like following:

```
import pytorch_pfn_extras.onnx as tou
tou.export_testcase(model, x1, '/path/to/output')
tou.export_testcase(model, x2, '/path/to/output', model_overwrite=False)
```

Following is the generated test cases of the above. `test_data_set_0` is the input `x1` and is its output, `test_data_set_1` is the input `x2` and its output.

```
$ tree /path/to/output
├── meta.json
├── model.onnx
├── test_data_set_0
│   ├── input_0.pb
│   └── output_0.pb
└── test_data_set_1
    ├── input_0.pb
    └── output_0.pb
```

strip_large_tensor_data option

This option strips large tensor in dumped files which is useful to reduce file size in usage such as benchmarking. Not only `model.onnx`, in/out, gradient data would be affected too. `large_tensor_threshold` could be used to specify threshold of large tensor size.

```
import torchvision
model = torchvision.models.resnet50(pretrained=True)
x = torch.zeros((1, 3, 224, 224))

import pytorch_pfn_extras.onnx as tou
tou.export_testcase(model, x, '/path/to/output')
tou.export_testcase(model, x, '/path/to/output2', strip_large_tensor_data=True)
```

```
$ ls -lh /path/to/output/model.onnx
-rwxrwxrwx 1 user user 98M Jun 24 23:34 /path/to/output/model.onnx
$ ls -lh /path/to/output2/model.onnx
-rwxrwxrwx 1 user user 64K Jun 24 23:34 /path/to/output2/model.onnx
```

This feature could be called from CLI:

```
$ python -m pytorch_pfn_extras.onnx.strip_large_tensor resnet50.onnx --out_onnx_path_
↳ resnet50_slim.onnx
$ ls -lh
-rwxrwxrwx 1 user user 98M Jun 30 09:13 resnet50.onnx
-rwxrwxrwx 1 user user 64K Jun 30 09:16 resnet50_slim.onnx
```

See `$ python -m pytorch_pfn_extras.onnx.strip_large_tensor -h` for help

Notes:

If an ONNX runtime does not support `no_raw_data` tensor, `unstrip_tensor.py` will resolve. See `$ python -m pytorch_pfn_extras.onnx.unstrip_tensor -h` for help

pytorch_pfn_extras.onnx.export

Function with same interface like `torch.onnx.export`. Unlike `torch.onnx.export`, you can use annotation feature (described below), `strip_large_tensor_data` options, or other `torch.onnx` extensions.

- `strip_large_tensor_data`: Same as `export_testcase`. Useful reducing file sizes.
- `return_output`: Returns output value of model execution. Note: Most output type would be `torch.Tensor`(not `onnx.TensorProto`)

```
model = nn.Sequential(nn.Linear(5, 10, bias=False))
x = torch.zeros((2, 5))

import io, onnx
bytesio = io.BytesIO()
pytorch_pfn_extras.onnx.export(model, x, bytesio)
onnx_proto = onnx.load(io.BytesIO(bytesio.getvalue()))
```


annotate

Feature to add custom ONNX attribute to specified `nn.Module`.

Notes:

- Annotated ONNX would be invalid ONNX format that doesn't pass check of `onnx.checker.check_model`.
- Only valid with `pytorch_pfn_extras.onnx.export_testcase` or `pytorch_pfn_extras.onnx.export`.
- **Only** the first ONNX node of modules like `nn.Linear`, `nn.GroupNorm`, etc. with multiple ONNX node would be annotated
 - For example `nn.Linear` with bias is split to `MatMul` -> `Add` graph. Only `MatMul` would be annotated. This is same in `apply_annotation` (described later) too.
- Use `apply_annotation` instead when the annotation target isn't `nn.Module`.

```
import pytorch_pfn_extras.onnx as tou

class Net(nn.Module):
    def __init__(self):
        super(Net, self).__init__()

        self.conv = nn.Conv2d(6, 9, 3)
        self.conv2 = nn.Conv2d(9, 12, 3)
        self.linear = nn.Linear(28, 20)
        self.linear2 = nn.Linear(20, 15)

    def forward(self, x):
        h = self.conv(x)
        with tou.annotate(key='value'):
            h = self.conv2(h)
            h = self.linear(h)
        h = self.linear2(h)
        return h

model = Net()
x = torch.randn((1, 6, 32, 32))
tou.export_testcase(model, x, '/path/to/output')
onnx_proto = onnx.load(os.path.join('/path/to/output', 'model.onnx'))
print(onnx.helper.printable_graph(onnx_proto.graph))
```

```
graph torch-jit-export (
  %input.1[FLOAT, 1x6x32x32]
) initializers (
  %17[FLOAT, 28x20]
  %18[FLOAT, 20x15]
  %conv.bias[FLOAT, 9]
  %conv.weight[FLOAT, 9x6x3x3]
  %conv2.bias[FLOAT, 12]
  %conv2.weight[FLOAT, 12x9x3x3]
  %linear.bias[FLOAT, 20]
  %linear2.bias[FLOAT, 15]
) {
```

(continues on next page)

(continued from previous page)

```

    %9 = Conv[dilations = [1, 1], group = 1, kernel_shape = [3, 3], pads = [0, 0, 0, 0],
    ↳ strides = [1, 1]](%input.1, %conv.weight, %conv.bias)
    %10 = Conv[dilations = [1, 1], group = 1, kernel_shape = [3, 3], key = 'value', pads =
    ↳ [0, 0, 0, 0], strides = [1, 1]](%9, %conv2.weight, %conv2.bias)
    %12 = MatMul[key = 'value'](%10, %17)
    %13 = Add(%12, %linear.bias)
    %15 = MatMul(%13, %18)
    %16 = Add(%15, %linear2.bias)
    return %16
}

```

In above example %10 = Conv and %12 = MatMul has key='value' attribute annotated.

apply_annotation

This annotates function call instead of annotating it with with.

The annotate target is nn.Module, so torch.nn.functional couldn't be annotated

```

import torch.nn.functional as F
import pytorch_pfn_extras.onnx as tou

class Net(nn.Module):
    def __init__(self):
        super(Net, self).__init__()

        self.conv = nn.Conv2d(6, 9, 3)
        self.conv2 = nn.Conv2d(9, 12, 3)
        self.linear = nn.Linear(28, 20)
        self.linear2 = nn.Linear(20, 15)

    def forward(self, x):
        h = self.conv(x)
        with tou.annotate(key='value'):
            h = self.conv2(h)
            h = F.relu(h)
            h = self.linear(h)
        h = self.linear2(h)
        return h

model = Net()
x = torch.randn((1, 6, 32, 32))
tou.export_testcase(model, x, '/path/to/output')
onnx_proto = onnx.load(os.path.join('/path/to/output', 'model.onnx'))
print(onnx.helper.printable_graph(onnx_proto.graph))

```

```

graph torch-jit-export (
    %input.1[FLOAT, 1x6x32x32]
) initializers (
    %18[FLOAT, 28x20]
    %19[FLOAT, 20x15]

```

(continues on next page)

(continued from previous page)

```

%conv.bias[FLOAT, 9]
%conv.weight[FLOAT, 9x6x3x3]
%conv2.bias[FLOAT, 12]
%conv2.weight[FLOAT, 12x9x3x3]
%linear.bias[FLOAT, 20]
%linear2.bias[FLOAT, 15]
) {
  %9 = Conv[dilations = [1, 1], group = 1, kernel_shape = [3, 3], pads = [0, 0, 0, 0],
→ strides = [1, 1]](%input.1, %conv.weight, %conv.bias)
  %10 = Conv[dilations = [1, 1], group = 1, kernel_shape = [3, 3], key = 'value', pads =
→ [0, 0, 0, 0], strides = [1, 1]](%9, %conv2.weight, %conv2.bias)
  %11 = Relu(%10)
  %13 = MatMul[key = 'value'](%11, %18)
  %14 = Add(%13, %linear.bias)
  %16 = MatMul(%14, %19)
  %17 = Add(%16, %linear2.bias)
  return %17
}

```

%10 = Conv and %13 = MatMul has key='value' attribute but %11 = Relu hasn't. By using apply_annotation all node in the function is annotated.

```

import pytorch_pfn_extras.onnx as tou

class Net(nn.Module):
    def __init__(self):
        super(Net, self).__init__()

        self.conv = nn.Conv2d(6, 9, 3)
        self.conv2 = nn.Conv2d(9, 12, 3)
        self.linear = nn.Linear(28, 20)
        self.linear2 = nn.Linear(20, 15)

    def forward(self, x):
        h = self.conv(x)
        def _f(x):
            h = self.conv2(x)
            h = F.relu(h)
            h = self.linear(h)
            return h
        h = tou.apply_annotation(_f, h, key='value')
        h = self.linear2(h)
        return h

model = Net()
x = torch.randn((1, 6, 32, 32))
tou.export_testcase(model, x, '/path/to/outout')
onnx_proto = onnx.load(os.path.join('/path/to/output', 'model.onnx'))
print(onnx.helper.printable_graph(onnx_proto.graph))

```

```

graph torch-jit-export (
  %input.1[FLOAT, 1x6x32x32]

```

(continues on next page)

(continued from previous page)

```

) initializers (
    %18[FLOAT, 28x20]
    %19[FLOAT, 20x15]
    %conv.bias[FLOAT, 9]
    %conv.weight[FLOAT, 9x6x3x3]
    %conv2.bias[FLOAT, 12]
    %conv2.weight[FLOAT, 12x9x3x3]
    %linear.bias[FLOAT, 20]
    %linear2.bias[FLOAT, 15]
) {
    %9 = Conv[dilations = [1, 1], group = 1, kernel_shape = [3, 3], pads = [0, 0, 0, 0],
    ↳ strides = [1, 1]](%input.1, %conv.weight, %conv.bias)
    %10 = Conv[dilations = [1, 1], group = 1, kernel_shape = [3, 3], key = 'value', pads =
    ↳ [0, 0, 0, 0], strides = [1, 1]](%9, %conv2.weight, %conv2.bias)
    %11 = Relu[key = 'value'](%10)
    %13 = MatMul[key = 'value'](%11, %18)
    %14 = Add(%13, %linear.bias)
    %16 = MatMul(%14, %19)
    %17 = Add(%16, %linear2.bias)
    return %17
}

```

Now %11 = Relu is annotated with key='value' attribute too.

scoped_anchor

This annotates scope's beginning and end of one or modules by adding Anchor node. Node would be named Anchor_N_start or Anchor_N_end (N is a index) and with op_type Identity.

- Adding custom parameter would add ONNX attribute and this will generate invalid ONNX in checker.
- Use this with `pytorch_pfn_extras.onnx.export_testcase` or `pytorch_pfn_extras.onnx.export`.
- When scope has multiple input/output only first input/output will get Anchor node added.
- N of node name is the index of pair beginning/end Anchor node like `Anchor_0_start`, `Anchor_0_end`.

```

import pytorch_pfn_extras.onnx as tou

class Net(nn.Module):
    def __init__(self):
        super(Net, self).__init__()

        self.conv = nn.Conv2d(6, 9, 3)
        self.conv2 = nn.Conv2d(9, 12, 3)
        self.linear = nn.Linear(28, 20)
        self.linear2 = nn.Linear(20, 15)

    def forward(self, x):
        h = self.conv(x)
        with tou.scoped_anchor(key='value'):
            h = self.conv2(h)
            h = self.linear(h)

```

(continues on next page)

(continued from previous page)

```

        h = self.linear2(h)
        return h

    def forward(self, x):
        with annotate(key='value'):
            return self.add(x)

model = Net()
x = torch.randn((1, 6, 32, 32))
out_dir = tou.export_testcase(model, x, '/path/to/output')
onnx_proto = onnx.load(os.path.join('/path/to/output', 'model.onnx'))
print(onnx.helper.printable_graph(onnx_proto.graph))

```

```

graph torch-jit-export (
    %input.1[FLOAT, 1x6x32x32]
) initializers (
    %23[FLOAT, 28x20]
    %24[FLOAT, 20x15]
    %conv.bias[FLOAT, 9]
    %conv.weight[FLOAT, 9x6x3x3]
    %conv2.bias[FLOAT, 12]
    %conv2.weight[FLOAT, 12x9x3x3]
    %linear.bias[FLOAT, 20]
    %linear2.bias[FLOAT, 15]
) {
    %9 = Conv[dilations = [1, 1], group = 1, kernel_shape = [3, 3], pads = [0, 0, 0, 0],
    ↳ strides = [1, 1]](%input.1, %conv.weight, %conv.bias)
    %11 = Identity[key = 'value'](%9)
    %12 = Conv[dilations = [1, 1], group = 1, kernel_shape = [3, 3], pads = [0, 0, 0, 0],
    ↳ strides = [1, 1]](%11, %conv2.weight, %conv2.bias)
    %16 = MatMul(%12, %23)
    %17 = Add(%16, %linear.bias)
    %19 = Identity[key = 'value'](%17)
    %21 = MatMul(%19, %24)
    %22 = Add(%21, %linear2.bias)
    return %22
}

```

%11 = Identity (node name = Anchor_0_start) and %19 = Identity (node name = Anchor_0_end) is added. key='value' is added as ONNX attribute.

non-nn.Module

The target of scope is only nn.Module. You can add adding sub nn.Module instead, if scope bound doesn't match nn.Module.

```

import pytorch_pfn_extras.onnx as tou

class Net(nn.Module):
    def __init__(self):
        super(Net, self).__init__()

```

(continues on next page)

(continued from previous page)

```

class _Net(nn.Module):
    def forward(self, x):
        return x + torch.ones((1,))
    self.add = _Net()

def forward(self, x):
    with tou.scoped_anchor(key='value'):
        return self.add(x)

model = Net()
x = torch.randn((1, 6, 32, 32))
out_dir = tou.export_testcase(model, x, '/path/to/output')
onnx_proto = onnx.load(os.path.join('/path/to/output', 'model.onnx'))
print(onnx.helper.printable_graph(onnx_proto.graph))

```

```

graph torch-jit-export (
    %x.1[FLOAT, 1x6x32x32]
) {
    %2 = Identity[key = 'value'](%x.1)
    %3 = Constant[value = <Tensor>]()
    %4 = Add(%2, %3)
    %6 = Identity[key = 'value'](%4)
    return %6
}

```

Or you can use `anchor` (described below) instead.

anchor (Future work)

Inserts `Anchor` node per each arbitrarily position of `nn.Module`. Node name would be `Anchor` and `op_type` would be `Identity`.

- Note: adding extra parameter would make extended ONNX format because it would be attribute.
- Please use it with `pytorch_pfn_extras.onnx.export_testcase` or `pytorch_pfn_extras.onnx.export`.

1.4.4 CUDA (CuPy Interoperability)

- `pytorch_pfn_extras.cuda.stream(stream)`
 - Context-manager that selects a given stream. This context manager also changes the CuPy's default stream if CuPy is available. When CuPy is not available, the functionality is the same as the PyTorch's counterpart, `torch.cuda.stream()`.
- `pytorch_pfn_extras.cuda.use_torch_mempool_in_cupy()`
 - Use PyTorch's memory pool in CuPy. If you want to use PyTorch's memory pool and non-default CUDA streams, streams must be created and managed using PyTorch (using `torch.cuda.Stream()` and `pytorch_pfn_extras.cuda.stream(stream)`). This feature requires CuPy v8.0+ and PyTorch v1.5+.
- `pytorch_pfn_extras.cuda.use_default_mempool_in_cupy()`
 - Use CuPy's default memory pool in CuPy.

- `pytorch_pfn_extras.from_ndarray(ndarray)`
 - Creates a Tensor from NumPy/CuPy ndarray.
- `pytorch_pfn_extras.as_ndarray(tensor)`
 - Creates a NumPy/CuPy ndarray from Tensor.
- `pytorch_pfn_extras.get_xp(tensor_device_or_ndarray)`
 - Returns numpy or cupy module for the given object.
- `pytorch_pfn_extras.as_numpy_dtype(torch_dtype)`
 - Returns NumPy dtype for the given torch dtype.
- `pytorch_pfn_extras.from_numpy_dtype(numpy_dtype)`
 - Returns torch dtype for the given NumPy dtype.

API REFERENCE

- `genindex`

2.1 Package

pytorch_pfn_extras

2.1.1 `pytorch_pfn_extras`

Functions

| | |
|--|---|
| <i>pytorch_pfn_extras.as_ndarray</i> (tensor) | Creates a <i>numpy.ndarray</i> or <i>cupy.ndarray</i> from <i>torch.Tensor</i> . |
| <i>pytorch_pfn_extras.as_numpy_dtype</i> (torch_dtype) | Returns NumPy dtype for the given PyTorch dtype. |
| <i>pytorch_pfn_extras.compile</i> (module[, ...]) | Compiles a module and an optimizer in a single graph using the provided backend. |
| <i>pytorch_pfn_extras.from_ndarray</i> (ndarray) | Creates a <i>torch.Tensor</i> from a <i>numpy.ndarray</i> or <i>cupy.ndarray</i> . |
| <i>pytorch_pfn_extras.from_numpy_dtype</i> (numpy_dtype) | Returns PyTorch dtype for the given NumPy dtype. |
| <i>pytorch_pfn_extras.get_xp</i> (obj) | Returns a module of ndarray implementation (<i>numpy</i> or <i>cupy</i>) for the given <i>obj</i> . |
| <i>pytorch_pfn_extras.map</i> (func, iterable[, ...]) | |
| <i>pytorch_pfn_extras.requires</i> (version[, package]) | |
| <i>pytorch_pfn_extras.to</i> (module_or_tensor, ...) | A function to transfer the given object to the given device. |

pytorch_pfn_extras.as_ndarray

`pytorch_pfn_extras.as_ndarray(tensor)`

Creates a `numpy.ndarray` or `cupy.ndarray` from `torch.Tensor`.

This method returns a tensor as a NumPy or CuPy ndarray depending on where the given *tensor* resides in. The *tensor* and the returned *ndarray* share the same underlying storage. Changes to the tensor will be reflected in the *ndarray* and vice versa. Note that changes made to *ndarray* cannot be tracked in the computational graph.

Parameters

tensor (*Tensor*) –

Return type

Any

pytorch_pfn_extras.as_numpy_dtype

`pytorch_pfn_extras.as_numpy_dtype(torch_dtype)`

Returns NumPy dtype for the given PyTorch dtype.

Parameters

torch_dtype (*dtype*) – PyTorch's dtype object.

Returns

NumPy type object.

Return type

Any

pytorch_pfn_extras.compile

`pytorch_pfn_extras.compile(module, optimizer=None, backend=None, *, generate_backward=True, decompositions=None)`

Compiles a module and an optimizer in a single graph using the provided backend.

Note: The backend object needs to be a callable accepting a `torch.fx.GraphModule` and a list of `torch.Tensor` and return a `Callable` as specified by <https://pytorch.org/docs/2.0/dynamo/custom-backends.html#custom-backends>

Note: Modules that are split in multiple graphs are not supported. `torch.compile` is called with the `fullgraph=True` argument.

Parameters

- **module** (*Module*) – `torch.nn.Module` to be compiled
- **optimizer** (*Optional[Optimizer]*) – Optimizer object associated to the module. It will be traced and its operations included in the module graph. Some dry run operations may be performed to fully initialize the optimizer status.
- **backend** (*optional*) – Object to process the graph and compile it for custom devices, will use PyTorch dynamo by default if not specified.
- **generate_backward** (*bool*) – Add the backward pass to the graph. Default is True.

- **decompositions** (*optional*) – Custom mapping for decompose a torch op into simple ops. Default is `None` and resorts to `torch._decomp.core_aten_decompositions()`

Return type*Callable[[...], Any]***pytorch_pfn_extras.from_ndarray**`pytorch_pfn_extras.from_ndarray(ndarray)`

Creates a `torch.Tensor` from a `numpy.ndarray` or `cupy.ndarray`.

Unlike `torch.from_numpy`, this method may make a copy when needed, e.g. when the given `ndarray` contains the negative strides which is not supported by PyTorch.

Parameters**ndarray** (*Any*) –**Return type***Tensor***pytorch_pfn_extras.from_numpy_dtype**`pytorch_pfn_extras.from_numpy_dtype(numpy_dtype)`

Returns PyTorch dtype for the given NumPy dtype.

Parameters**numpy_dtype** (*Any*) – NumPy's dtype object.**Returns**

PyTorch type object.

Return type*dtype***pytorch_pfn_extras.get_xp**`pytorch_pfn_extras.get_xp(obj)`

Returns a module of ndarray implementation (`numpy` or `cupy`) for the given `obj`.

The `obj` can be `torch.Tensor`, `torch.device` or NumPy/CuPy `ndarray`.

Parameters**obj** (*Union[Any, Tensor]*) –**Return type***Any*

pytorch_pfn_extras.map

`pytorch_pfn_extras.map(func, iterable, out_keys=None, device='cpu')`

Parameters

- **func** (*Callable*[[*Any*], *Any*]) –
- **iterable** (*Sequence*[*Any*]) –
- **out_keys** (*Optional*[*Set*[*str*]]) –
- **device** (*Any*) –

Return type

Sequence[*Any*]

pytorch_pfn_extras.requires

`pytorch_pfn_extras.requires(version, package='torch')`

Parameters

- **version** (*str*) –
- **package** (*str*) –

Return type

bool

pytorch_pfn_extras.to

`pytorch_pfn_extras.to(module_or_tensor, device, *, options=None, runtime_class=None, config=None)`

A function to transfer the given object to the given device.

If PyTorch's device type is given as the device argument, the behavior of this function is equivalent to `module_or_tensor.to(module_or_tensor, device)`.

Otherwise, this function uses the **Runtime** mechanism. This function looks for the Runtime for the device from the RuntimeRegistry and delegates the actual transfer operation to it.

See also the documentation of `ppe.runtime.BaseRuntime` for details.

Parameters

- **module_or_tensor** (*torch.nn.Module* or *torch.Tensor*) – An object to be transferred.
- **device** (*torch.device* or *str*) – The device that the input object is transferred to.
- **options** (*dict*, *optional*) – An options of dictionary type that is passed to `runtime_class.__init__` as an argument.
- **runtime_class** (*Optional*[*Type*[*BaseRuntime*]]) – A runtime class inherited from *BaseRuntime* class. If *None*, a runtime class is automatically selected based on the device argument from the runtime registry.
- **config** (*dict*, *optional*) – DEPRECATED. Use *options*.

Returns

A *torch.Tensor* with the specified device.

Return type
ModuleOrTensor

Modules

`pytorch_pfn_extras.config`

`pytorch_pfn_extras.config_types`

`pytorch_pfn_extras.cuda`

`pytorch_pfn_extras.dataloaders`

`pytorch_pfn_extras.dataset`

`pytorch_pfn_extras.distributed`

`pytorch_pfn_extras.engine`

`pytorch_pfn_extras.handler`

`pytorch_pfn_extras.logging`

`pytorch_pfn_extras.nn`

`pytorch_pfn_extras.onnx`

`pytorch_pfn_extras.ops`

`pytorch_pfn_extras.profiler`

`pytorch_pfn_extras.reporting`

`pytorch_pfn_extras.runtime`

`pytorch_pfn_extras.testing`

`pytorch_pfn_extras.torchscript`

`pytorch_pfn_extras.training`

`pytorch_pfn_extras.utils`

`pytorch_pfn_extras.writing`

pytorch_pfn_extras.config

Functions

*pytorch_pfn_extras.config.
customize_type(...)*

pytorch_pfn_extras.config.customize_type

pytorch_pfn_extras.config.customize_type(**default_kwargs)

Parameters

default_kwargs (Any) –

Return type

Callable[[Callable[[...], Any]], Callable[[...], Any]]

Classes

pytorch_pfn_extras.config.Config(config[,
types])

pytorch_pfn_extras.config.Config

class pytorch_pfn_extras.config.**Config**(config, types=None)

Bases: object

Methods

__init__(config[, types])

load_path(path, *[, loader, types])

update_via_args(args)

Parameters

- **config** (Any) –
- **types** (Optional[Mapping[str, Callable[[...], Any]]]) –

__init__(config, types=None)

Parameters

- **config** (Any) –
- **types** (Optional[Mapping[str, Callable[[...], Any]]]) –

Return type

None

classmethod `load_path(path, *, loader=None, types=None)`**Parameters**

- **path** (*str*) –
- **loader** (*Optional[Callable[[str], Any]]*) –
- **types** (*Optional[Mapping[str, Callable[[...], Any]]]*) –

Return type[Config](#)**update_via_args(args)****Parameters****args** (*Sequence[Tuple[str, Any]]*) –**Return type**

None

pytorch_pfn_extras.config_types**Functions**

```
pytorch_pfn_extras.config_types.  
load_path_with_optuna_types(...)  
pytorch_pfn_extras.config_types.  
optuna_types(trial)
```

pytorch_pfn_extras.config_types.load_path_with_optuna_types

```
pytorch_pfn_extras.config_types.load_path_with_optuna_types(path, trial, loader=None,  
                                                             types=None)
```

Parameters

- **path** (*str*) –
- **trial** (*optuna.trial.Trial*) –
- **loader** (*Optional[Callable[[str], Any]]*) –
- **types** (*Optional[Dict[str, Callable[[...], Any]]]*) –

Return type[Config](#)

pytorch_pfn_extras.config_types.optuna_types

`pytorch_pfn_extras.config_types.optuna_types(trial)`

Parameters

trial (*optuna.trial.Trial*) –

Return type

Dict[str, Any]

pytorch_pfn_extras.cuda

Functions

| | |
|--|--|
| <code>pytorch_pfn_extras.cuda.stream(stream)</code> | Context-manager that selects a given stream. |
| <code>pytorch_pfn_extras.cuda.use_default_mempool_in_cupy()</code> | Use the default memory pool in CuPy. |
| <code>pytorch_pfn_extras.cuda.use_torch_mempool_in_cupy()</code> | Use the PyTorch memory pool in CuPy. |

pytorch_pfn_extras.cuda.stream

`pytorch_pfn_extras.cuda.stream(stream)`

Context-manager that selects a given stream.

This context manager also changes the CuPy's default stream if CuPy is available. When CuPy is not available, the functionality is the same as the PyTorch's counterpart, `torch.cuda.stream()`.

Parameters

stream (*Optional[Stream]*) –

Return type

Generator[None, None, None]

pytorch_pfn_extras.cuda.use_default_mempool_in_cupy

`pytorch_pfn_extras.cuda.use_default_mempool_in_cupy()`

Use the default memory pool in CuPy.

Return type

None

pytorch_pfn_extras.cuda.use_torch_mempool_in_cupy

`pytorch_pfn_extras.cuda.use_torch_mempool_in_cupy()`

Use the PyTorch memory pool in CuPy.

If you want to use PyTorch's memory pool and non-default CUDA streams, streams must be created and managed using PyTorch (using `torch.cuda.Stream()` and `pytorch_pfn_extras.cuda.stream(stream)`).

Return type

None

pytorch_pfn_extras.dataloaders

Classes

| | |
|--|--------------|
| <code>pytorch_pfn_extras.dataloaders. DataLoader(dataset)</code> | Data loader. |
|--|--------------|

pytorch_pfn_extras.dataloaders.DataLoader

```
class pytorch_pfn_extras.dataloaders.DataLoader(dataset, batch_size=1, shuffle=None, sampler=None,
                                                batch_sampler=None, num_workers=0,
                                                collate_fn=None, pin_memory=False,
                                                drop_last=False, timeout=0, worker_init_fn=None,
                                                multiprocessing_context=None, generator=None, *,
                                                prefetch_factor=None, persistent_workers=False,
                                                pin_memory_device="")
```

Bases: `Generic[T_co]`

Data loader. Combines a dataset and a sampler, and provides an iterable over the given dataset.

The `DataLoader` supports both map-style and iterable-style datasets with single- or multi-process loading, customizing loading order and optional automatic batching (collation) and memory pinning.

See `torch.utils.data` documentation page for more details.

Parameters

- **dataset** (`Dataset`) – dataset from which to load the data.
- **batch_size** (`int`, *optional*) – how many samples per batch to load (default: 1).
- **shuffle** (`bool`, *optional*) – set to `True` to have the data reshuffled at every epoch (default: `False`).
- **sampler** (`Sampler` or `Iterable`, *optional*) – defines the strategy to draw samples from the dataset. Can be any `Iterable` with `__len__` implemented. If specified, `shuffle` must not be specified.
- **batch_sampler** (`Sampler` or `Iterable`, *optional*) – like `sampler`, but returns a batch of indices at a time. Mutually exclusive with `batch_size`, `shuffle`, `sampler`, and `drop_last`.
- **num_workers** (`int`, *optional*) – how many subprocesses to use for data loading. `0` means that the data will be loaded in the main process. (default: `0`)

- **collate_fn** (*Callable, optional*) – merges a list of samples to form a mini-batch of Tensor(s). Used when using batched loading from a map-style dataset.
- **pin_memory** (*bool, optional*) – If True, the data loader will copy Tensors into device/CUDA pinned memory before returning them. If your data elements are a custom type, or your collate_fn returns a batch that is a custom type, see the example below.
- **drop_last** (*bool, optional*) – set to True to drop the last incomplete batch, if the dataset size is not divisible by the batch size. If False and the size of dataset is not divisible by the batch size, then the last batch will be smaller. (default: False)
- **timeout** (*numeric, optional*) – if positive, the timeout value for collecting a batch from workers. Should always be non-negative. (default: 0)
- **worker_init_fn** (*Callable, optional*) – If not None, this will be called on each worker subprocess with the worker id (an int in [0, num_workers - 1]) as input, after seeding and before data loading. (default: None)
- **multiprocessing_context** (*str or multiprocessing.context.BaseContext, optional*) – If None, the default [multiprocessing context](#) of your operating system will be used. (default: None)
- **generator** (*torch.Generator, optional*) – If not None, this RNG will be used by RandomSampler to generate random indexes and multiprocessing to generate base_seed for workers. (default: None)
- **prefetch_factor** (*int, optional, keyword-only arg*) – Number of batches loaded in advance by each worker. 2 means there will be a total of 2 * num_workers batches prefetched across all workers. (default value depends on the set value for num_workers. If value of num_workers=0 default is None. Otherwise, if value of num_workers > 0 default is 2).
- **persistent_workers** (*bool, optional*) – If True, the data loader will not shut down the worker processes after a dataset has been consumed once. This allows to maintain the workers Dataset instances alive. (default: False)
- **pin_memory_device** (*str, optional*) – the device to [pin_memory](#) to if pin_memory is True.

Warning: If the spawn start method is used, `worker_init_fn` cannot be an unpicklable object, e.g., a lambda function. See [multiprocessing-best-practices](#) on more details related to multiprocessing in PyTorch.

Warning: `len(dataloader)` heuristic is based on the length of the sampler used. When `dataset` is an `IterableDataset`, it instead returns an estimate based on `len(dataset) / batch_size`, with proper rounding depending on `drop_last`, regardless of multi-process loading configurations. This represents the best guess PyTorch can make because PyTorch trusts user `dataset` code in correctly handling multi-process loading to avoid duplicate data.

However, if sharding results in multiple workers having incomplete last batches, this estimate can still be inaccurate, because (1) an otherwise complete batch can be broken into multiple ones and (2) more than one batch worth of samples can be dropped when `drop_last` is set. Unfortunately, PyTorch can not detect such cases in general.

See [`Dataset Types`_](#) for more details on these two types of datasets and how `IterableDataset` interacts with [`Multi-process data loading`_](#).

Warning: See reproducibility, and dataloader-workers-random-seed, and data-loading-randomness notes for random seed related questions.

Methods

`__init__(dataset[, batch_size, shuffle, ...])`

`check_worker_number_rationality()`

Attributes

`multiprocessing_context`

`dataset`

`batch_size`

`num_workers`

`pin_memory`

`drop_last`

`timeout`

`sampler`

`pin_memory_device`

`prefetch_factor`

`__init__(dataset, batch_size=1, shuffle=None, sampler=None, batch_sampler=None, num_workers=0, collate_fn=None, pin_memory=False, drop_last=False, timeout=0, worker_init_fn=None, multiprocessing_context=None, generator=None, *, prefetch_factor=None, persistent_workers=False, pin_memory_device="")`

Parameters

- **dataset** (`Dataset[T_co]`) –
- **batch_size** (`Optional[int]`) –
- **shuffle** (`Optional[bool]`) –
- **sampler** (`Optional[Union[Sampler, Iterable]]`) –
- **batch_sampler** (`Optional[Union[Sampler[List], Iterable[List]]]`) –
- **num_workers** (`int`) –

- `collate_fn` (*Optional*[*Callable*[[*List*[*T*]], *Any*]]) –
- `pin_memory` (*bool*) –
- `drop_last` (*bool*) –
- `timeout` (*float*) –
- `worker_init_fn` (*Optional*[*Callable*[[*int*], *None*]]) –
- `prefetch_factor` (*Optional*[*int*]) –
- `persistent_workers` (*bool*) –
- `pin_memory_device` (*str*) –

`batch_size`: *Optional*[*int*]

`check_worker_number_rationality`()

`dataset`: *Dataset*[*T_co*]

`drop_last`: *bool*

`property multiprocessing_context`

`num_workers`: *int*

`pin_memory`: *bool*

`pin_memory_device`: *str*

`prefetch_factor`: *Optional*[*int*]

`sampler`: *Union*[*Sampler*, *Iterable*]

`timeout`: *float*

Modules

pytorch_pfn_extras.dataloaders.data_loader

pytorch_pfn_extras.dataloaders.utils

pytorch_pfn_extras.dataloaders.data_loader

Functions

| | |
|---|--|
| <i>pytorch_pfn_extras.dataloaders.data_loader.default_collate</i> (batch) | Function that takes in a batch of data and puts the elements within the batch into a tensor with an additional outer dimension - batch size. |
| <i>pytorch_pfn_extras.dataloaders.data_loader.default_convert</i> (data) | Function that converts each NumPy array element into a <code>torch.Tensor</code> . |
| <i>pytorch_pfn_extras.dataloaders.data_loader.get_worker_info</i> () | Returns the information about the current DataLoader iterator worker process. |

pytorch_pfn_extras.dataloaders.data_loader.default_collate

`pytorch_pfn_extras.dataloaders.data_loader.default_collate(batch)`

Function that takes in a batch of data and puts the elements within the batch into a tensor with an additional outer dimension - batch size. The exact output type can be a `torch.Tensor`, a *Sequence* of `torch.Tensor`, a *Collection* of `torch.Tensor`, or left unchanged, depending on the input type. This is used as the default function for collation when *batch_size* or *batch_sampler* is defined in `DataLoader`.

Here is the general input type (based on the type of the element within the batch) to output type mapping:

- `torch.Tensor` -> `torch.Tensor` (with an added outer dimension batch size)
- NumPy Arrays -> `torch.Tensor`
- *float* -> `torch.Tensor`
- *int* -> `torch.Tensor`
- *str* -> *str* (unchanged)
- *bytes* -> *bytes* (unchanged)
- *Mapping*[*K*, *V*_{*i*}] -> *Mapping*[*K*, `default_collate`([*V*₁, *V*₂, ...])]
- *NamedTuple*[*V*₁_{*i*}, *V*₂_{*i*}, ...] -> *NamedTuple*[`default_collate`([*V*₁₁, *V*₁₂, ...]), `default_collate`([*V*₂₁, *V*₂₂, ...]), ...]
- *Sequence*[*V*₁_{*i*}, *V*₂_{*i*}, ...] -> *Sequence*[`default_collate`([*V*₁₁, *V*₁₂, ...]), `default_collate`([*V*₂₁, *V*₂₂, ...]), ...]

Parameters

batch – a single batch to be collated

Examples

```
>>> # xdoctest: +SKIP
>>> # Example with a batch of `int`s:
>>> default_collate([0, 1, 2, 3])
tensor([0, 1, 2, 3])
>>> # Example with a batch of `str`s:
>>> default_collate(['a', 'b', 'c'])
['a', 'b', 'c']
>>> # Example with `Map` inside the batch:
>>> default_collate([{'A': 0, 'B': 1}, {'A': 100, 'B': 100}])
{'A': tensor([ 0, 100]), 'B': tensor([ 1, 100])}
>>> # Example with `NamedTuple` inside the batch:
>>> Point = namedtuple('Point', ['x', 'y'])
>>> default_collate([Point(0, 0), Point(1, 1)])
Point(x=tensor([0, 1]), y=tensor([0, 1]))
>>> # Example with `Tuple` inside the batch:
>>> default_collate([(0, 1), (2, 3)])
[tensor([0, 2]), tensor([1, 3])]
>>> # Example with `List` inside the batch:
>>> default_collate([[0, 1], [2, 3]])
[tensor([0, 2]), tensor([1, 3])]
>>> # Two options to extend `default_collate` to handle specific type
>>> # Option 1: Write custom collate function and invoke `default_collate`
```

(continues on next page)

(continued from previous page)

```

>>> def custom_collate(batch):
...     elem = batch[0]
...     if isinstance(elem, CustomType): # Some custom condition
...         return ...
...     else: # Fall back to `default_collate`
...         return default_collate(batch)
>>> # Option 2: In-place modify `default_collate_fn_map`
>>> def collate_customtype_fn(batch, *, collate_fn_map=None):
...     return ...
>>> default_collate_fn_map.update(CustoType, collate_customtype_fn)
>>> default_collate(batch) # Handle `CustomType` automatically

```

pytorch_pfn_extras.dataloaders.data_loader.default_convert

pytorch_pfn_extras.dataloaders.data_loader.default_convert(*data*)

Function that converts each NumPy array element into a `torch.Tensor`. If the input is a *Sequence*, *Collection*, or *Mapping*, it tries to convert each element inside to a `torch.Tensor`. If the input is not an NumPy array, it is left unchanged. This is used as the default function for collation when both *batch_sampler* and *batch_size* are NOT defined in `DataLoader`.

The general input type to output type mapping is similar to that of `default_collate()`. See the description there for more details.

Parameters

data – a single data point to be converted

Examples

```

>>> # xdoctest: +SKIP
>>> # Example with `int`
>>> default_convert(0)
0
>>> # Example with NumPy array
>>> default_convert(np.array([0, 1]))
tensor([0, 1])
>>> # Example with NamedTuple
>>> Point = namedtuple('Point', ['x', 'y'])
>>> default_convert(Point(0, 0))
Point(x=0, y=0)
>>> default_convert(Point(np.array(0), np.array(0)))
Point(x=tensor(0), y=tensor(0))
>>> # Example with List
>>> default_convert([np.array([0, 1]), np.array([2, 3])])
[tensor([0, 1]), tensor([2, 3])]

```

pytorch_pfn_extras.dataloaders.data_loader.get_worker_info`pytorch_pfn_extras.dataloaders.data_loader.get_worker_info()`

Returns the information about the current DataLoader iterator worker process.

When called in a worker, this returns an object guaranteed to have the following attributes:

- `id`: the current worker id.
- `num_workers`: the total number of workers.
- `seed`: the random seed set for the current worker. This value is determined by main process RNG and the worker id. See `DataLoader`'s documentation for more details.
- `dataset`: the copy of the dataset object in **this** process. Note that this will be a different object in a different process than the one in the main process.

When called in the main process, this returns `None`.

Note: When used in a `worker_init_fn` passed over to `DataLoader`, this method can be useful to set up each worker process differently, for instance, using `worker_id` to configure the `dataset` object to only read a specific fraction of a sharded dataset, or use `seed` to seed other libraries used in dataset code.

Return type*Optional[WorkerInfo]***Classes**

`pytorch_pfn_extras.dataloaders.data_loader.DataLoader`. Data loader.
`DataLoader(dataset)`

pytorch_pfn_extras.dataloaders.data_loader.DataLoader

```
class pytorch_pfn_extras.dataloaders.data_loader.DataLoader(dataset, batch_size=1, shuffle=None,
    sampler=None, batch_sampler=None,
    num_workers=0, collate_fn=None,
    pin_memory=False, drop_last=False,
    timeout=0, worker_init_fn=None,
    multiprocessing_context=None,
    generator=None, *,
    prefetch_factor=None,
    persistent_workers=False,
    pin_memory_device="")
```

Bases: `Generic[T_co]`

Data loader. Combines a dataset and a sampler, and provides an iterable over the given dataset.

The `DataLoader` supports both map-style and iterable-style datasets with single- or multi-process loading, customizing loading order and optional automatic batching (collation) and memory pinning.

See `torch.utils.data` documentation page for more details.

Parameters

- **dataset** (*Dataset*) – dataset from which to load the data.
- **batch_size** (*int, optional*) – how many samples per batch to load (default: 1).
- **shuffle** (*bool, optional*) – set to True to have the data reshuffled at every epoch (default: False).
- **sampler** (*Sampler or Iterable, optional*) – defines the strategy to draw samples from the dataset. Can be any *Iterable* with `__len__` implemented. If specified, `shuffle` must not be specified.
- **batch_sampler** (*Sampler or Iterable, optional*) – like `sampler`, but returns a batch of indices at a time. Mutually exclusive with `batch_size`, `shuffle`, `sampler`, and `drop_last`.
- **num_workers** (*int, optional*) – how many subprocesses to use for data loading. 0 means that the data will be loaded in the main process. (default: 0)
- **collate_fn** (*Callable, optional*) – merges a list of samples to form a mini-batch of Tensor(s). Used when using batched loading from a map-style dataset.
- **pin_memory** (*bool, optional*) – If True, the data loader will copy Tensors into device/CUDA pinned memory before returning them. If your data elements are a custom type, or your `collate_fn` returns a batch that is a custom type, see the example below.
- **drop_last** (*bool, optional*) – set to True to drop the last incomplete batch, if the dataset size is not divisible by the batch size. If False and the size of dataset is not divisible by the batch size, then the last batch will be smaller. (default: False)
- **timeout** (*numeric, optional*) – if positive, the timeout value for collecting a batch from workers. Should always be non-negative. (default: 0)
- **worker_init_fn** (*Callable, optional*) – If not None, this will be called on each worker subprocess with the worker id (an int in `[0, num_workers - 1]`) as input, after seeding and before data loading. (default: None)
- **multiprocessing_context** (*str or multiprocessing.context.BaseContext, optional*) – If None, the default `multiprocessing context` of your operating system will be used. (default: None)
- **generator** (*torch.Generator, optional*) – If not None, this RNG will be used by RandomSampler to generate random indexes and multiprocessing to generate `base_seed` for workers. (default: None)
- **prefetch_factor** (*int, optional, keyword-only arg*) – Number of batches loaded in advance by each worker. 2 means there will be a total of `2 * num_workers` batches prefetched across all workers. (default value depends on the set value for `num_workers`. If value of `num_workers=0` default is None. Otherwise, if value of `num_workers > 0` default is 2).
- **persistent_workers** (*bool, optional*) – If True, the data loader will not shut down the worker processes after a dataset has been consumed once. This allows to maintain the workers *Dataset* instances alive. (default: False)
- **pin_memory_device** (*str, optional*) – the device to `pin_memory` to if `pin_memory` is True.

Warning: If the `spawn` start method is used, `worker_init_fn` cannot be an unpicklable object, e.g., a lambda function. See `multiprocessing-best-practices` on more details related to multiprocessing in PyTorch.

Warning: `len(data_loader)` heuristic is based on the length of the sampler used. When `dataset` is an `IterableDataset`, it instead returns an estimate based on `len(dataset) / batch_size`, with proper rounding depending on `drop_last`, regardless of multi-process loading configurations. This represents the best guess PyTorch can make because PyTorch trusts user `dataset` code in correctly handling multi-process loading to avoid duplicate data.

However, if sharding results in multiple workers having incomplete last batches, this estimate can still be inaccurate, because (1) an otherwise complete batch can be broken into multiple ones and (2) more than one batch worth of samples can be dropped when `drop_last` is set. Unfortunately, PyTorch can not detect such cases in general.

See `'Dataset Types'` for more details on these two types of datasets and how `IterableDataset` interacts with `'Multi-process data loading'`.

Warning: See reproducibility, and `data_loader_workers_random_seed`, and `data_loading_randomness` notes for random seed related questions.

Methods

`__init__(dataset[, batch_size, shuffle, ...])`

`check_worker_number_rationality()`

Attributes

`multiprocessing_context`

`dataset`

`batch_size`

`num_workers`

`pin_memory`

`drop_last`

`timeout`

`sampler`

`pin_memory_device`

`prefetch_factor`

```
__init__(dataset, batch_size=1, shuffle=None, sampler=None, batch_sampler=None, num_workers=0,
          collate_fn=None, pin_memory=False, drop_last=False, timeout=0, worker_init_fn=None,
          multiprocessing_context=None, generator=None, *, prefetch_factor=None,
          persistent_workers=False, pin_memory_device="")
```

Parameters

- **dataset** (`Dataset[T_co]`) –
- **batch_size** (`Optional[int]`) –
- **shuffle** (`Optional[bool]`) –
- **sampler** (`Optional[Union[Sampler, Iterable]]`) –
- **batch_sampler** (`Optional[Union[Sampler[List], Iterable[List]]]`) –
- **num_workers** (`int`) –
- **collate_fn** (`Optional[Callable[[List[T]], Any]]`) –
- **pin_memory** (`bool`) –
- **drop_last** (`bool`) –
- **timeout** (`float`) –
- **worker_init_fn** (`Optional[Callable[[int], None]]`) –
- **prefetch_factor** (`Optional[int]`) –
- **persistent_workers** (`bool`) –
- **pin_memory_device** (`str`) –

`batch_size:` `Optional[int]`

`check_worker_number_rationality()`

`dataset:` `Dataset[T_co]`

`drop_last:` `bool`

`property multiprocessing_context`

`num_workers:` `int`

`pin_memory:` `bool`

`pin_memory_device:` `str`

`prefetch_factor:` `Optional[int]`

`sampler:` `Union[Sampler, Iterable]`

`timeout:` `float`

pytorch_pfn_extras.dataloaders.utils

Classes

| | |
|--|---|
| <code>pytorch_pfn_extras.dataloaders.utils.CollateAsDict(names)</code> | Creates a collate function that converts inputs to a dict of tensors. |
|--|---|

pytorch_pfn_extras.dataloaders.utils.CollateAsDict

class pytorch_pfn_extras.dataloaders.utils.CollateAsDict(*names*, *collate_fn*=<function default_collate>)

Bases: object

Creates a collate function that converts inputs to a dict of tensors.

An instantiated callable object can be feeded to `torch.utils.data.DataLoader` as a `collate_fn` option.

Parameters

- **names** (*list of str*) – Names of keys of output dict.
- **collate_fn** (*function*) – A function preprocesses inputs.

Methods

`__init__(names[, collate_fn])`

`__call__(*args, **kwargs)`

Converts inputs the dataset generated to a dictionary of tensors.

Returns (dict of Tensor):

A dictionary with keys that specified as names option, and values as input tensors.

Parameters

- **args** (*Any*) –
- **kwargs** (*Any*) –

Return type

`Dict[str, Any]`

`__init__(names, collate_fn=<function default_collate>)`

Parameters

- **names** (*Sequence[str]*) –
- **collate_fn** (*Callable[[...], Any]*) –

Return type

`None`

pytorch_pfn_extras.dataset**Classes**

| | |
|--|---|
| <code>pytorch_pfn_extras.dataset.SharedDataset(sm_size)</code> | Dataset that caches the load samples in shared memory |
| <code>pytorch_pfn_extras.dataset.TabularDataset(...)</code> | An abstract class that represents tabular dataset. |

pytorch_pfn_extras.dataset.SharedDataset

class `pytorch_pfn_extras.dataset.SharedDataset`(*sm_size*, *cache_type*=<class 'pytorch_pfn_extras.dataset.shared_dataset.InfiniteCache'>)

Bases: `Dataset`

Dataset that caches the load samples in shared memory

Args

Methods

| |
|--|
| <code>__init__(sm_size[, cache_type])</code> |
| <code>cache_item(idx, x)</code> |
| <code>is_cached(idx)</code> |

| |
|--|
| <code>__init__(sm_size, cache_type=<class 'pytorch_pfn_extras.dataset.shared_dataset.InfiniteCache'>)</code> |
| <code>cache_item(idx, x)</code> |
| <code>is_cached(idx)</code> |

pytorch_pfn_extras.dataset.TabularDataset

class `pytorch_pfn_extras.dataset.TabularDataset`(*args, **kws)

Bases: `Dataset`

An abstract class that represents tabular dataset.

This class represents a tabular dataset. In a tabular dataset, all examples have the same number of elements. For example, all examples of the dataset below have three elements (`a[i]`, `b[i]`, and `c[i]`).

| | a | b | c |
|---|-------------------|-------------------|-------------------|
| 0 | <code>a[0]</code> | <code>b[0]</code> | <code>c[0]</code> |
| 1 | <code>a[1]</code> | <code>b[1]</code> | <code>c[1]</code> |
| 2 | <code>a[2]</code> | <code>b[2]</code> | <code>c[2]</code> |
| 3 | <code>a[3]</code> | <code>b[3]</code> | <code>c[3]</code> |

Since an example can be represented by both tuple and dict ((a[i], b[i], c[i]) and {'a': a[i], 'b': b[i], 'c': c[i]}), this class uses *mode* to indicate which representation will be used. If there is only one column, an example also can be represented by a value (a[i]). In this case, *mode* is None.

An inheritance should implement `__len__()`, *keys*, *mode* and *get_examples()*.

```
>>> import numpy as np
>>>
>>> from pytorch_pfn_extras import dataset
>>>
>>> class MyDataset(dataset.TabularDataset):
...     def __len__(self):
...         return 4
...
...     @property
...     def keys(self):
...         return ('a', 'b', 'c')
...
...     @property
...     def mode(self):
...         return tuple
...
...     def get_examples(self, indices, key_indices):
...         data = np.arange(12).reshape((4, 3))
...         if indices is not None:
...             data = data[indices]
...         if key_indices is not None:
...             data = data[:, list(key_indices)]
...         return tuple(data.transpose())
...
>>> dataset = MyDataset()
>>> len(dataset)
4
>>> dataset.keys
('a', 'b', 'c')
>>> dataset.astuple()[0]
(0, 1, 2)
>>> sorted(dataset.asdict()[0].items())
[('a', 0), ('b', 1), ('c', 2)]
>>>
>>> view = dataset.slice[[3, 2], ('c', 0)]
>>> len(view)
2
>>> view.keys
('c', 'a')
>>> view.astuple()[1]
(8, 6)
>>> sorted(view.asdict()[1].items())
[('a', 6), ('c', 8)]
```

Methods

| | |
|---|--|
| <code>__init__()</code> | |
| <code>asdict()</code> | Return a view with dict mode. |
| <code>astuple()</code> | Return a view with tuple mode. |
| <code>concat(*datasets)</code> | Stack datasets along rows. |
| <code>convert(data)</code> | Convert fetched data. |
| <code>fetch()</code> | Fetch data. |
| <code>get_example(i)</code> | |
| <code>get_examples(indices, key_indices)</code> | Return a part of data. |
| <code>join(*datasets)</code> | Stack datasets along columns. |
| <code>transform(keys, transform)</code> | Apply a transform to each example. |
| <code>transform_batch(keys, transform_batch)</code> | Apply a transform to examples. |
| <code>with_converter(converter)</code> | Override the behaviour of <code>convert()</code> . |

Attributes

| | |
|--------------------|-------------------------|
| <code>keys</code> | Names of columns. |
| <code>mode</code> | Mode of representation. |
| <code>slice</code> | Get a slice of dataset. |

`asdict()`

Return a view with dict mode.

Returns

A view whose `mode` is dict.

`astuple()`

Return a view with tuple mode.

Returns

A view whose `mode` is tuple.

`concat(*datasets)`

Stack datasets along rows.

Parameters

datasets (iterable of `TabularDataset`) – Datasets to be concatenated. All datasets must have the same `keys`.

Returns

A concatenated dataset.

`convert(data)`

Convert fetched data.

This method takes data fetched by `fetch()` and pre-process them before passing them to models. The default behaviour is converting each column into an ndarray. This behaviour can be overridden by `with_converter()`. If the dataset is constructed by `concat()` or `join()`, the converter of the first dataset is used.

Parameters

data (`tuple` or `dict`) – Data from `fetch()`.

Returns

A tuple or dict. Each value is an ndarray.

fetch()

Fetch data.

This method fetches all data of the dataset/view. Note that this method returns a column-major data (i.e. $([a[0], \dots, a[3]], \dots, [c[0], \dots, c[3]])$, $\{ 'a': [a[0], \dots, a[3]], \dots, 'c': [c[0], \dots, c[3]] \}$, or $[a[0], \dots, a[3]]$).

Returns

If *mode* is `tuple`, this method returns a tuple of lists/arrays. If *mode* is `dict`, this method returns a dict of lists/arrays.

get_example(*i*)**get_examples(*indices*, *key_indices*)**

Return a part of data.

Parameters

- **indices** (*list of ints or slice*) – Indices of requested rows. If this argument is `None`, it indicates all rows.
- **key_indices** (*tuple of ints*) – Indices of requested columns. If this argument is `None`, it indicates all columns.

Returns

tuple of lists/arrays

join(datasets*)**

Stack datasets along columns.

Args: datasets (iterable of *TabularDataset*):

Datasets to be concatenated. All datasets must have the same length

Returns

A joined dataset.

property keys

Names of columns.

A tuple of strings that indicate the names of columns.

property mode

Mode of representation.

This indicates the type of value returned by *fetch()* and *__getitem__()*. `tuple`, `dict`, and `None` are supported.

property slice

Get a slice of dataset.

Parameters

- **indices** (*list/array of ints/bools or slice*) – Requested rows.
- **keys** (*tuple of ints/strs or int or str*) – Requested columns.

Returns

A view of specified range.

transform(*keys*, *transform*)

Apply a transform to each example.

The transformations are a list where each element is a tuple that holds the transformation signature and a callable that is the transformation itself.

The transformation signature is a tuple of 2 elements with the first one being the keys of the dataset that are taken as inputs. And the last one the outputs it produces for the transformation *keys* argument.

When multiple transformations are specified, the outputs must be disjoint or *ValueError* will be risen.

Parameters

- **keys** (*tuple of strs*) – The keys of transformed examples.
- **transform** (*list of tuples*) – A list where each element specifies a transformation with a tuple with the transformation signature and a callable that takes an example and returns transformed example. *mode* of transformed dataset is determined by the transformed examples.

Returns

A transformed dataset.

transform_batch(*keys*, *transform_batch*)

Apply a transform to examples.

The transformations are a list where each element is a tuple that holds the transformation signature and a callable that is the transformation itself.

The transformation signature is a tuple of 2 elements with the first one being the keys of the dataset that are taken as inputs. And the last one the outputs it produces for the transformation *keys* argument.

When multiple transformations are specified, the outputs must be disjoint or *ValueError* will be risen.

Parameters

- **keys** (*tuple of strs*) – The keys of transformed examples.
- **transform_batch** (*list of tuples*) – A list where each element specifies a transformation with a tuple with the transformation signature and a callable that takes a batch of examples and returns a batch of transformed examples. *mode* of transformed dataset is determined by the transformed examples.

Returns

A transformed dataset.

with_converter(*converter*)

Override the behaviour of *convert()*.

This method overrides *convert()*.

Parameters

converter (*callable*) – A new converter.

Returns

A dataset with the new converter.

Exceptions

*pytorch_pfn_extras.dataset.
ItemNotFoundException*

pytorch_pfn_extras.dataset.ItemNotFoundException

exception pytorch_pfn_extras.dataset.ItemNotFoundException

Modules

pytorch_pfn_extras.dataset.shared_dataset

pytorch_pfn_extras.dataset.tabular

pytorch_pfn_extras.dataset.shared_dataset

Classes

*pytorch_pfn_extras.dataset.shared_dataset.
Cache()*

*pytorch_pfn_extras.dataset.shared_dataset.
InfiniteCache(sm_size)*

pytorch_pfn_extras.dataset.shared_dataset. Dataset that caches the load samples in shared memory
SharedDataset(sm_size)

pytorch_pfn_extras.dataset.shared_dataset.Cache

class pytorch_pfn_extras.dataset.shared_dataset.Cache

Bases: object

Methods

__init__()

add_to_cache(idx, x)

get_value(idx)

is_cached(idx)

add_to_cache(*idx, x*)

`get_value(idx)`

`is_cached(idx)`

pytorch_pfn_extras.dataset.shared_dataset.InfiniteCache

class pytorch_pfn_extras.dataset.shared_dataset.**InfiniteCache**(*sm_size*)

Bases: [Cache](#)

Methods

`__init__(sm_size)`

`add_to_cache(idx, x)`

`get_value(idx)`

`is_cached(idx)`

`__init__(sm_size)`

`add_to_cache(idx, x)`

`get_value(idx)`

`is_cached(idx)`

pytorch_pfn_extras.dataset.shared_dataset.SharedDataset

class pytorch_pfn_extras.dataset.shared_dataset.**SharedDataset**(*sm_size*, *cache_type*=<class 'pytorch_pfn_extras.dataset.shared_dataset.InfiniteCache'>)

Bases: [Dataset](#)

Dataset that caches the load samples in shared memory

Args

Methods

`__init__(sm_size[, cache_type])`

`cache_item(idx, x)`

`is_cached(idx)`

`__init__(sm_size, cache_type=<class 'pytorch_pfn_extras.dataset.shared_dataset.InfiniteCache'>)`

```
cache_item(idx, x)
```

```
is_cached(idx)
```

Exceptions

```
pytorch_pfn_extras.dataset.shared_dataset.  
ItemNotFoundException
```

pytorch_pfn_extras.dataset.shared_dataset.ItemNotFoundException

```
exception pytorch_pfn_extras.dataset.shared_dataset.ItemNotFoundException
```

pytorch_pfn_extras.dataset.tabular

Functions

```
pytorch_pfn_extras.dataset.tabular.  
from_data(data, *)
```

Create a TabularDataset from lists/arrays/callables.

pytorch_pfn_extras.dataset.tabular.from_data

```
pytorch_pfn_extras.dataset.tabular.from_data(data, *, size=None)
```

Create a TabularDataset from lists/arrays/callables.

```
>>> from pytorch_pfn_extras.dataset import tabular
>>>
>>> dataset = tabular.from_data([0, 1, 2])
>>> dataset[0]
0
>>> dataset = tabular.from_data([0, 1, 2], [3, 4, 5])
>>> dataset[0]
(0, 3)
>>> dataset = tabular.from_data([('a', [0, 1, 2]), ('b', [3, 4, 5])])
>>> dataset.keys
('a', 'b')
>>> dataset[0]
(0, 3)
>>> dataset = tabular.from_data({'a': [0, 1, 2], 'b': [3, 4, 5]})
>>> sorted(dataset[0].items())
[('a', 0), ('b', 3)]
>>> dataset = tabular.from_data('a', lambda i: i * i, size=10)
>>> dataset[5]
25
```

Parameters

- **data** (*list, array, tuple, or dict*) – Data in following format.

- *list/array*
 - *(str, list/array/callable)*
 - *((str, ...), callable)*
 - *((list/array)/(str, list/array/callable) /((key, ...), callable), ...)*
 - *{str: (list/array/callable)/(str, ...): callable, ...}*
- **size** (*int*) – The length of the dataset. This argument is required when no lists/arrays exist in data.

Returns

A *TabularDataset*.

Classes

| | |
|---|---|
| <i>pytorch_pfn_extras.dataset.tabular. DelegateDataset(dataset)</i> | A helper class to implement a <i>TabularDataset</i> . |
|---|---|

pytorch_pfn_extras.dataset.tabular.DelegateDataset

class `pytorch_pfn_extras.dataset.tabular.DelegateDataset(dataset)`

Bases: *TabularDataset*

A helper class to implement a *TabularDataset*.

This class wraps an instance of *TabularDataset* and provides methods of *TabularDataset*. This class is useful to create a custom dataset class by inheriting it.

```
>>> import numpy as np
>>>
>>> from pytorch_pfn_extras.dataset import tabular
>>>
>>> class MyDataset(tabular.DelegateDataset):
...     def __init__(self):
...         super().__init__(tabular.from_data((
...             ('a', np.arange(10)),
...             ('b', self.get_b),
...             ('c', [3, 1, 4, 5, 9, 2, 6, 8, 7, 0]),
...             (('d', 'e'), self.get_de))))
...
...     def get_b(self, i):
...         return 'b[{}]'.format(i)
...
...     def get_de(self, i):
...         return {'d': 'd[{}]'.format(i), 'e': 'e[{}]'.format(i)}
...
>>> dataset = MyDataset()
>>> len(dataset)
10
>>> dataset.keys
```

(continues on next page)

(continued from previous page)

```
( 'a', 'b', 'c', 'd', 'e' )
>>> dataset[0]
(0, 'b[0]', 3, 'd[0]', 'e[0]')
```

Parameters

dataset (`pytorch_pfn_extras.dataset.TabularDataset`) – An underlying dataset.

Methods

| | |
|---|--|
| <code>__init__(dataset)</code> | |
| <code>asdict()</code> | Return a view with dict mode. |
| <code>astuple()</code> | Return a view with tuple mode. |
| <code>concat(*datasets)</code> | Stack datasets along rows. |
| <code>convert(data)</code> | Convert fetched data. |
| <code>fetch()</code> | Fetch data. |
| <code>get_example(i)</code> | |
| <code>get_examples(indices, key_indices)</code> | Return a part of data. |
| <code>join(*datasets)</code> | Stack datasets along columns. |
| <code>transform(keys, transform)</code> | Apply a transform to each example. |
| <code>transform_batch(keys, transform_batch)</code> | Apply a transform to examples. |
| <code>with_converter(converter)</code> | Override the behaviour of <code>convert()</code> . |

Attributes

| | |
|--------------------|-------------------------|
| <code>keys</code> | Names of columns. |
| <code>mode</code> | Mode of representation. |
| <code>slice</code> | Get a slice of dataset. |

`__init__(dataset)`

`get_examples(indices, key_indices)`

Return a part of data.

Parameters

- **indices** (*list of ints or slice*) – Indices of requested rows. If this argument is `None`, it indicates all rows.
- **key_indices** (*tuple of ints*) – Indices of requested columns. If this argument is `None`, it indicates all columns.

Returns

tuple of lists/arrays

property keys

Names of columns.

A tuple of strings that indicate the names of columns.

property mode

Mode of representation.

This indicates the type of value returned by `fetch()` and `__getitem__()`. tuple, dict, and None are supported.

Modules

`pytorch_pfn_extras.dataset.tabular.
delegate_dataset`

`pytorch_pfn_extras.dataset.tabular.
from_data(data, *)` Create a TabularDataset from lists/arrays/callables.

`pytorch_pfn_extras.dataset.tabular.
tabular_dataset`

pytorch_pfn_extras.dataset.tabular.delegate_dataset**Classes**

`pytorch_pfn_extras.dataset.tabular.
delegate_dataset.DelegateDataset(dataset)` A helper class to implement a TabularDataset.

pytorch_pfn_extras.dataset.tabular.delegate_dataset.DelegateDataset**class** `pytorch_pfn_extras.dataset.tabular.delegate_dataset.DelegateDataset(dataset)`

Bases: *TabularDataset*

A helper class to implement a TabularDataset.

This class wraps an instance of *TabularDataset* and provides methods of *TabularDataset*. This class is useful to create a custom dataset class by inheriting it.

```
>>> import numpy as np
>>>
>>> from pytorch_pfn_extras.dataset import tabular
>>>
>>> class MyDataset(tabular.DelegateDataset):
...     def __init__(self):
...         super().__init__(tabular.from_data((
...             ('a', np.arange(10)),
...             ('b', self.get_b),
...             ('c', [3, 1, 4, 5, 9, 2, 6, 8, 7, 0]),
...             (('d', 'e'), self.get_de))))
...
...     def get_b(self, i):
...         return 'b[{}]'.format(i)
...
...     def get_de(self, i):
```

(continues on next page)

(continued from previous page)

```

...         return {'d': 'd[{}]'.format(i), 'e': 'e[{}]'.format(i)}
...
>>> dataset = MyDataset()
>>> len(dataset)
10
>>> dataset.keys
('a', 'b', 'c', 'd', 'e')
>>> dataset[0]
(0, 'b[0]', 3, 'd[0]', 'e[0]')

```

Parameters

dataset (`pytorch_pfn_extras.dataset.TabularDataset`) – An underlying dataset.

Methods

| | |
|---|--|
| <code>__init__(dataset)</code> | |
| <code>asdict()</code> | Return a view with dict mode. |
| <code>astuple()</code> | Return a view with tuple mode. |
| <code>concat(*datasets)</code> | Stack datasets along rows. |
| <code>convert(data)</code> | Convert fetched data. |
| <code>fetch()</code> | Fetch data. |
| <code>get_example(i)</code> | |
| <code>get_examples(indices, key_indices)</code> | Return a part of data. |
| <code>join(*datasets)</code> | Stack datasets along columns. |
| <code>transform(keys, transform)</code> | Apply a transform to each example. |
| <code>transform_batch(keys, transform_batch)</code> | Apply a transform to examples. |
| <code>with_converter(converter)</code> | Override the behaviour of <code>convert()</code> . |

Attributes

| | |
|--------------------|-------------------------|
| <code>keys</code> | Names of columns. |
| <code>mode</code> | Mode of representation. |
| <code>slice</code> | Get a slice of dataset. |

`__init__(dataset)`

`get_examples(indices, key_indices)`

Return a part of data.

Parameters

- **indices** (*list of ints or slice*) – Indices of requested rows. If this argument is None, it indicates all rows.
- **key_indices** (*tuple of ints*) – Indices of requested columns. If this argument is None, it indicates all columns.

Returns

tuple of lists/arrays

property keys

Names of columns.

A tuple of strings that indicate the names of columns.

property mode

Mode of representation.

This indicates the type of value returned by `fetch()` and `__getitem__()`. `tuple`, `dict`, and `None` are supported.

pytorch_pfn_extras.dataset.tabular.tabular_dataset**Classes**

| | |
|---|--|
| <code>pytorch_pfn_extras.dataset.tabular.tabular_dataset.Dataset(...)</code> | An abstract class representing a <i>Dataset</i> . |
| <code>pytorch_pfn_extras.dataset.tabular.tabular_dataset.TabularDataset(...)</code> | An abstract class that represents tabular dataset. |

pytorch_pfn_extras.dataset.tabular.tabular_dataset.Dataset

class `pytorch_pfn_extras.dataset.tabular.tabular_dataset.Dataset(*args, **kwargs)`

Bases: `Generic[T_co]`

An abstract class representing a *Dataset*.

All datasets that represent a map from keys to data samples should subclass it. All subclasses should overwrite `__getitem__()`, supporting fetching a data sample for a given key. Subclasses could also optionally overwrite `__len__()`, which is expected to return the size of the dataset by many `Sampler` implementations and the default options of `DataLoader`. Subclasses could also optionally implement `__getitems__()`, for speedup batched samples loading. This method accepts list of indices of samples of batch and returns list of samples.

Note: `DataLoader` by default constructs a index sampler that yields integral indices. To make it work with a map-style dataset with non-integral indices/keys, a custom sampler must be provided.

Methods

`__init__()`

pytorch_pfn_extras.dataset.tabular.tabular_dataset.TabularDataset

class pytorch_pfn_extras.dataset.tabular.tabular_dataset.TabularDataset(*args, **kwargs)

Bases: *Dataset*

An abstract class that represents tabular dataset.

This class represents a tabular dataset. In a tabular dataset, all examples have the same number of elements. For example, all examples of the dataset below have three elements (a[i], b[i], and c[i]).

| | a | b | c |
|---|------|------|------|
| 0 | a[0] | b[0] | c[0] |
| 1 | a[1] | b[1] | c[1] |
| 2 | a[2] | b[2] | c[2] |
| 3 | a[3] | b[3] | c[3] |

Since an example can be represented by both tuple (a[i], b[i], c[i]) and {'a': a[i], 'b': b[i], 'c': c[i]}, this class uses *mode* to indicate which representation will be used. If there is only one column, an example also can be represented by a value (a[i]). In this case, *mode* is None.

An inheritance should implement `__len__()`, *keys*, *mode* and *get_examples()*.

```
>>> import numpy as np
>>>
>>> from pytorch_pfn_extras import dataset
>>>
>>> class MyDataset(dataset.TabularDataset):
...     def __len__(self):
...         return 4
...
...     @property
...     def keys(self):
...         return ('a', 'b', 'c')
...
...     @property
...     def mode(self):
...         return tuple
...
...     def get_examples(self, indices, key_indices):
...         data = np.arange(12).reshape((4, 3))
...         if indices is not None:
...             data = data[indices]
...         if key_indices is not None:
...             data = data[:, list(key_indices)]
...         return tuple(data.transpose())
...
>>> dataset = MyDataset()
>>> len(dataset)
4
>>> dataset.keys
('a', 'b', 'c')
>>> dataset.astuple()[0]
(0, 1, 2)
```

(continues on next page)

(continued from previous page)

```

>>> sorted(dataset.asdict()[0].items())
[('a', 0), ('b', 1), ('c', 2)]
>>>
>>> view = dataset.slice[[3, 2], ('c', 0)]
>>> len(view)
2
>>> view.keys
('c', 'a')
>>> view.astuple()[1]
(8, 6)
>>> sorted(view.asdict()[1].items())
[('a', 6), ('c', 8)]

```

Methods

| | |
|---|--|
| <code>__init__()</code> | |
| <code>asdict()</code> | Return a view with dict mode. |
| <code>astuple()</code> | Return a view with tuple mode. |
| <code>concat(*datasets)</code> | Stack datasets along rows. |
| <code>convert(data)</code> | Convert fetched data. |
| <code>fetch()</code> | Fetch data. |
| <code>get_example(i)</code> | |
| <code>get_examples(indices, key_indices)</code> | Return a part of data. |
| <code>join(*datasets)</code> | Stack datasets along columns. |
| <code>transform(keys, transform)</code> | Apply a transform to each example. |
| <code>transform_batch(keys, transform_batch)</code> | Apply a transform to examples. |
| <code>with_converter(converter)</code> | Override the behaviour of <code>convert()</code> . |

Attributes

| | |
|--------------------|-------------------------|
| <code>keys</code> | Names of columns. |
| <code>mode</code> | Mode of representation. |
| <code>slice</code> | Get a slice of dataset. |

`asdict()`

Return a view with dict mode.

Returns

A view whose `mode` is dict.

`astuple()`

Return a view with tuple mode.

Returns

A view whose `mode` is tuple.

concat(*datasets)

Stack datasets along rows.

Parameters

datasets (iterable of *TabularDataset*) – Datasets to be concatenated. All datasets must have the same *keys*.

Returns

A concatenated dataset.

convert(data)

Convert fetched data.

This method takes data fetched by *fetch()* and pre-process them before passing them to models. The default behaviour is converting each column into an ndarray. This behaviour can be overridden by *with_converter()*. If the dataset is constructed by *concat()* or *join()*, the converter of the first dataset is used.

Parameters

data (*tuple* or *dict*) – Data from *fetch()*.

Returns

A tuple or dict. Each value is an ndarray.

fetch()

Fetch data.

This method fetches all data of the dataset/view. Note that this method returns a column-major data (i.e. ([a[0], ..., a[3]], ..., [c[0], ... c[3]]), {'a': [a[0], ..., a[3]], ..., 'c': [c[0], ..., c[3]]}, or [a[0], ..., a[3]]).

Returns

If *mode* is *tuple*, this method returns a tuple of lists/arrays. If *mode* is *dict*, this method returns a dict of lists/arrays.

get_example(i)**get_examples(indices, key_indices)**

Return a part of data.

Parameters

- **indices** (*list of ints* or *slice*) – Indices of requested rows. If this argument is *None*, it indicates all rows.
- **key_indices** (*tuple of ints*) – Indices of requested columns. If this argument is *None*, it indicates all columns.

Returns

tuple of lists/arrays

join(*datasets)

Stack datasets along columns.

Args: **datasets** (iterable of *TabularDataset*):

Datasets to be concatenated. All datasets must have the same length

Returns

A joined dataset.

property keys

Names of columns.

A tuple of strings that indicate the names of columns.

property mode

Mode of representation.

This indicates the type of value returned by `fetch()` and `__getitem__()`. `tuple`, `dict`, and `None` are supported.

property slice

Get a slice of dataset.

Parameters

- **indices** (*list/array of ints/bools or slice*) – Requested rows.
- **keys** (*tuple of ints/strs or int or str*) – Requested columns.

Returns

A view of specified range.

transform(*keys, transform*)

Apply a transform to each example.

The transformations are a list where each element is a tuple that holds the transformation signature and a callable that is the transformation itself.

The transformation signature is a tuple of 2 elements with the first one being the keys of the dataset that are taken as inputs. And the last one the outputs it produces for the transformation *keys* argument.

When multiple transformations are specified, the outputs must be disjoint or *ValueError* will be risen.

Parameters

- **keys** (*tuple of strs*) – The keys of transformed examples.
- **transform** (*list of tuples*) – A list where each element specifies a transformation with a tuple with the transformation signature and a callable that takes an example and returns transformed example. *mode* of transformed dataset is determined by the transformed examples.

Returns

A transformed dataset.

transform_batch(*keys, transform_batch*)

Apply a transform to examples.

The transformations are a list where each element is a tuple that holds the transformation signature and a callable that is the transformation itself.

The transformation signature is a tuple of 2 elements with the first one being the keys of the dataset that are taken as inputs. And the last one the outputs it produces for the transformation *keys* argument.

When multiple transformations are specified, the outputs must be disjoint or *ValueError* will be risen.

Parameters

- **keys** (*tuple of strs*) – The keys of transformed examples.
- **transform_batch** (*list of tuples*) – A list where each element specifies a transformation with a tuple with the transformation signature and a callable that takes a batch of

examples and returns a batch of transformed examples. *mode* of transformed dataset is determined by the transformed examples.

Returns

A transformed dataset.

with_converter(*converter*)

Override the behaviour of *convert()*.

This method overrides *convert()*.

Parameters

converter (*callable*) – A new converter.

Returns

A dataset with the new converter.

pytorch_pfn_extras.distributed

Functions

| | |
|--|--|
| <i>pytorch_pfn_extras.distributed.create_distributed_subset_indices(...)</i> | Returns a indices of a dataset to be used for the current process. |
| <i>pytorch_pfn_extras.distributed.initialize_ompi_environment(*)</i> | Initialize <i>torch.distributed</i> environments with values taken from OpenMPI. |

pytorch_pfn_extras.distributed.create_distributed_subset_indices

`pytorch_pfn_extras.distributed.create_distributed_subset_indices(num_total_samples,
num_replicas=None,
rank=None, shuffle=True,
seed=None)`

Returns a indices of a dataset to be used for the current process.

Parameters

- **num_total_samples** (*int*) – The size of the dataset.
- **num_replicas** (*Optional[int]*) – Number of processes participating in the training. By default, `torch.distributed.get_world_size()` is used.
- **rank** (*Optional[int]*) – Rank of the current process within *num_replicas*. By default, `torch.distributed.get_rank()` is used.
- **shuffle** (*bool*) – If True (default), shuffle the indices.
- **seed** (*Optional[int]*) – Random seed used to shuffle.

Return type

List[int]

pytorch_pfn_extras.distributed.initialize_ompi_environment

```
pytorch_pfn_extras.distributed.initialize_ompi_environment(*, backend='gloo', init_method='tcp',
                                                         world_size=1, rank=0, local_rank=0,
                                                         addr='localhost', port='1234',
                                                         timeout=1800)
```

Initialize *torch.distributed* environments with values taken from OpenMPI.

Parameters

- **backend** (*str*) – The backend to be used, only "gloo" and "nccl" are supported. Defaults to "gloo".
- **init_method** (*str*) – Initialization method used by torch, only "tcp" and "env" are supported. Defaults to "tcp".
- **world_size** (*int*) – The total world size to be used in case it is not specified in MPI env vars. Defaults to 1.
- **rank** (*int*) – The process rank to be used in case it is not specified in MPI env vars. Defaults to 0.
- **local_rank** (*int*) – The process local rank to be used in case it is not specified in MPI env vars. Defaults to 0.
- **addr** (*str*) – The address of the master process of *torch.distributed*. Defaults to "localhost"
- **port** (*str*) – The port of the master process of *torch.distributed*. Defaults to "1234"
- **timeout** (*int*) – Timeout seconds for *torch.distributed* collective communication. Defaults to 1800.

Return type

Tuple[int, int, int]

Classes

| | |
|--|---|
| <code>pytorch_pfn_extras.distributed. DistributedValidationSampler(dataset)</code> | Distributed sampler without duplication |
|--|---|

pytorch_pfn_extras.distributed.DistributedValidationSampler

```
class pytorch_pfn_extras.distributed.DistributedValidationSampler(dataset, num_replicas=None,
                                                                rank=None, shuffle=True,
                                                                seed=0)
```

Bases: `Sampler`

Distributed sampler without duplication

This sampler splits the input dataset to each worker process in distributed setup without allowing repetition. It is for evaluation purpose such as `DistributedEvaluator`. This does not guarantee each worker to get the same number of samples, so for training do not use this sampler (use `PyTorch DistributedSampler` instead).

Methods

`__init__(dataset[, num_replicas, rank, ...])`

`__init__(dataset, num_replicas=None, rank=None, shuffle=True, seed=0)`

Parameters

- **dataset** (*Sized*) –
- **num_replicas** (*Optional[int]*) –
- **rank** (*Optional[int]*) –
- **shuffle** (*bool*) –
- **seed** (*int*) –

Return type

None

pytorch_pfn_extras.engine

Functions

| | |
|--|------------------------------|
| <code>pytorch_pfn_extras.engine.cast(typ, val)</code> | Cast a value to a type. |
| <code>pytorch_pfn_extras.engine.create_evaluator(...)</code> | Creates an evaluator object. |
| <code>pytorch_pfn_extras.engine.create_trainer(...)</code> | Creates a trainer object. |
| <code>pytorch_pfn_extras.engine.default_transform_model(n, x)</code> | |
| <code>pytorch_pfn_extras.engine.filter_state_objects(args)</code> | |
| <code>pytorch_pfn_extras.engine.filter_state_objects_dict(args)</code> | |

pytorch_pfn_extras.engine.cast

`pytorch_pfn_extras.engine.cast(typ, val)`

Cast a value to a type.

This returns the value unchanged. To the type checker this signals that the return value has the designated type, but at runtime we intentionally don't check anything (we want this to be as fast as possible).

pytorch_pfn_extras.engine.create_evaluator

```
pytorch_pfn_extras.engine.create_evaluator(models, *, progress_bar=False, device='cpu',
                                           metrics=None, logic=None, handler_class=None,
                                           options=None, runtime_options=None, profile=None,
                                           distributed=False)
```

Creates an evaluator object. The return value of this function is expected to be fed to *ppe.engine.create_trainer* as an argument.

Parameters

- **models** (*Union[Module, Mapping[str, Module]]*) – Map of string to `torch.nn.Module` or an actual `Module`. In most cases, this argument is the same as the *model* argument of *ppe.engine.create_trainer*.
- **progress_bar** (*bool*) – If *True*, a progress bar is enabled in evaluation.
- **device** (*str or torch.device*) – Device name used for selecting a corresponding runtime class.
- **metrics** (*list of metrics*) – List of metrics, which computes various quantities and update output for the reporting.
- **logic** (*Optional[Logic]*) – A logic object. If *None* is given, an logic object is instantiated from the default logic class.
- **handler_class** (*Optional[Type[Handler]]*) – A handler class that instantiates a handler object. If *None* is given, *ppe.handler.Handler* is used as a default handler class.
- **options** (*Optional[Dict[str, Any]]*) – Options that are set to the handler and logic object. See the documentation of *ppe.handler.Handler* and *ppe.handler.Logic* for details.
- **runtime_options** (*Optional[Mapping[str, Any]]*) – Options that are set to the runtime object. See the documentation of *ppe.handler.Handler* for details.
- **profile** (*Optional[profile]*) – A *torch.profiler.profile* object to collect the performance metrics.
- **distributed** (*bool*) – Flag to determine whether to create a distributed-enabled evaluator. If set to *True*, the created evaluator will support distributed execution.

Return type

Evaluator

pytorch_pfn_extras.engine.create_trainer

```
pytorch_pfn_extras.engine.create_trainer(models, optimizers, max_epochs, *, extensions=None,
                                          out_dir='result', stop_trigger=None, writer=None,
                                          evaluator=None, device='cpu', logic=None,
                                          transform_model=<function default_transform_model>,
                                          handler_class=None, options=None, runtime_options=None,
                                          profile=None, **kwargs)
```

Creates a trainer object.

Parameters

- **models** (*Union[Module, Mapping[str, Module]]*) – Map of string to `Module` or an actual `Module`.

- **optimizers** (*Union[Optimizer, Mapping[str, Optimizer]]*) – Map of string to Optimizer or an actual Optimizer.
- **max_epochs** (*int*) – Number of epochs in the whole training loop. Ignored if *stop_trigger* is passed as a kwarg.
- **extensions** (*Optional[Sequence[Union[extension.ExtensionLike, extension.ExtensionEntry]]]*) – List of extensions to be registered to the trainer.
- **out_dir** (*str*) – Output directory (default: *result*).
- **stop_trigger** (*trigger, optional*) – Trigger that can be consulted to determine whether training has concluded. The default is an interval trigger set to *max_epochs*.
- **writer** (*Optional[writing.Writer]*) – Writer that can be used by extensions to write data to custom filesystems.
- **evaluator** (*Optional[Union[Evaluator, Tuple[Evaluator, TriggerLike], Mapping[str, Union[Evaluator, Tuple[Evaluator, TriggerLike]]]]]*) – Evaluator that is used in evaluation phase. If *None* is given, the evaluation is skipped. Evaluators can be created with *pytorch_pfn_extras.engine.create_evaluator()*.
- **device** (*str or torch.device*) – Device name used for selecting a corresponding runtime class.
- **logic** (*Optional[BaseLogic]*) – A logic object. If *None* is given, an logic object is instantiated from the default logic class.
- **transform_model** (*Callable[[str, Module], Module]*) – A function to transform a model structure, often used to unwrap the a module from DDP module.
- **handler_class** (*Optional[Type[Handler]]*) – A handler class that instantiates a handler object. If *None* is given, *ppe.handler.Handler* is used as a default handler class.
- **options** (*Optional[Dict[str, Any]]*) – Options that are set to the handler and logic object. See the documentation of *ppe.handler.Handler* and *ppe.handler.Logic* for details.
- **runtime_options** (*Optional[Mapping[str, Any]]*) – Options that are set to the runtime object. See the documentation of *ppe.runtime.PyTorchRuntime* for details.
- **profile** (*Optional[profile]*) – A *torch.profiler.profile* object to collect the performance metrics.
- **kwargs** (*Any*) –

Return type

Trainer

pytorch_pfn_extras.engine.default_transform_model

`pytorch_pfn_extras.engine.default_transform_model(n, x)`

Parameters

- **n** (*str*) –
- **x** (*Module*) –

Return type

Module

pytorch_pfn_extras.engine.filter_state_objects

`pytorch_pfn_extras.engine.filter_state_objects(args, key_name="")`

Parameters

- **args** (*Any*) –
- **key_name** (*str*) –

Return type

List[Tuple[str, StateObjectProtocol]]

pytorch_pfn_extras.engine.filter_state_objects_dict

`pytorch_pfn_extras.engine.filter_state_objects_dict(args, key_name='option')`

Parameters

- **args** (*Dict[str, Any]*) –
- **key_name** (*str*) –

Return type

List[Tuple[str, StateObjectProtocol]]

Classes

`pytorch_pfn_extras.engine.`

`DistributedEvaluator(...)`

`pytorch_pfn_extras.engine.Evaluator(handler,`
`...)`

`pytorch_pfn_extras.engine.`

`NamedTuple(typename)`

Typed version of namedtuple.

`pytorch_pfn_extras.engine.`

`StateObjectProtocol(...)`

`pytorch_pfn_extras.engine.Trainer(handler, ...)`

pytorch_pfn_extras.engine.DistributedEvaluator

class `pytorch_pfn_extras.engine.DistributedEvaluator`(*handler, models, *, progress_bar=False,*
metrics=None, profile=None)

Bases: `Evaluator`

Methods

`__init__`(handler, models, *[, progress_bar, ...])

`run`(loader, *[, eval_len])

Executes the evaluation loop.

Parameters

- **handler** (`BaseHandler`) –
- **models** (`Union[Module, Mapping[str, Module]]`) –
- **progress_bar** (`bool`) –
- **metrics** (`Optional[Sequence[MetricType]]`) –
- **profile** (`Optional[profile]`) –

`__init__`(handler, models, *, progress_bar=False, metrics=None, profile=None)

Parameters

- **handler** (`BaseHandler`) –
- **models** (`Union[Module, Mapping[str, Module]]`) –
- **progress_bar** (`bool`) –
- **metrics** (`Optional[Sequence[MetricType]]`) –
- **profile** (`Optional[profile]`) –

pytorch_pfn_extras.engine.Evaluator

class `pytorch_pfn_extras.engine.Evaluator`(handler, models, *, progress_bar=False, metrics=None, profile=None)

Bases: `object`

Methods

`__init__`(handler, models, *[, progress_bar, ...])

`run`(loader, *[, eval_len])

Executes the evaluation loop.

Parameters

- **handler** (`BaseHandler`) –
- **models** (`Union[Module, Mapping[str, Module]]`) –
- **progress_bar** (`bool`) –
- **metrics** (`Optional[Sequence[MetricType]]`) –
- **profile** (`Optional[profile]`) –

```
__init__(handler, models, *, progress_bar=False, metrics=None, profile=None)
```

Parameters

- **handler** (`BaseHandler`) –
- **models** (`Union[Module, Mapping[str, Module]]`) –
- **progress_bar** (`bool`) –
- **metrics** (`Optional[Sequence[MetricType]]`) –
- **profile** (`Optional[profile]`) –

```
run(loader, *, eval_len=None)
```

Executes the evaluation loop.

Parameters

- **loader** (`torch.utils.data.DataLoader`) – A data loader for evaluation.
- **eval_len** (`int, optional`) – The number of iterations per one evaluation epoch.

Return type

None

pytorch_pfn_extras.engine.NamedTuple

```
class pytorch_pfn_extras.engine.NamedTuple(typename, fields=None, /, **kwargs)
```

Bases: object

Typed version of namedtuple.

Usage in Python versions >= 3.6:

```
class Employee(NamedTuple):  
    name: str  
    id: int
```

This is equivalent to:

```
Employee = collections.namedtuple('Employee', ['name', 'id'])
```

The resulting class has an extra `__annotations__` attribute, giving a dict that maps field names to types. (The field names are also in the `_fields` attribute, which is part of the namedtuple API.) Alternative equivalent keyword syntax is also accepted:

```
Employee = NamedTuple('Employee', name=str, id=int)
```

In Python versions <= 3.5 use:

```
Employee = NamedTuple('Employee', [('name', str), ('id', int)])
```

Methods

`__init__()`

pytorch_pfn_extras.engine.StateObjectProtocol**class** pytorch_pfn_extras.engine.StateObjectProtocol(*args, **kwargs)

Bases: Protocol

Methods

`__init__(*args, **kwargs)`

`load_state_dict(state_dict)`

`state_dict()`

`__init__(*args, **kwargs)``load_state_dict(state_dict)`**Parameters****state_dict** (*Dict[str, Any]*) –**Return type**

None

`state_dict()`**Return type***Dict[str, Any]***pytorch_pfn_extras.engine.Trainer****class** pytorch_pfn_extras.engine.Trainer(handler, *, evaluator, models, profile=None, **kwargs)

Bases: object

Methods

`__init__(handler, *, evaluator, models[, ...])`

`extend(extension[, name, trigger, priority, ...])`

`get_optimizer(name)`

`is_epoch_last_iter(idx)`

`load_state_dict(to_load)`

`run(train_loader[, val_loader, train_len, ...])` Executes the training loop.

`set_optimizer(name, optimizer)`

`state_dict()`

Attributes

`epoch`

`epoch_detail`

`evaluator`

`is_before_training`

`iteration`

`manager`

`models`

`optimizers`

`stop_trigger`

Parameters

- **handler** (`handler_module.BaseHandler`) –
- **evaluator** (`Optional[Union[Evaluator, Tuple[Evaluator, Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]], Mapping[str, Union[Evaluator, Tuple[Evaluator, Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]]]]]`) –
- **models** (`Union[Module, Mapping[str, Module]]`) –
- **profile** (`Optional[profile]`) –

- **kwargs** (*Any*) –

__init__ (*handler*, *, *evaluator*, *models*, *profile=None*, ***kwargs*)

Parameters

- **handler** (*handler_module.BaseHandler*) –
- **evaluator** (*Optional[Union[Evaluator, Tuple[Evaluator, Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]], Mapping[str, Union[Evaluator, Tuple[Evaluator, Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]]]]]*) –
- **models** (*Union[Module, Mapping[str, Module]]*) –
- **profile** (*Optional[profile]*) –
- **kwargs** (*Any*) –

property epoch: *int*

property epoch_detail: *float*

property evaluator: *Optional[Evaluator]*

extend (*extension*, *name=None*, *trigger=None*, *priority=None*, *, *call_before_training=False*, ***kwargs*)

Parameters

- **extension** (*Union[extension.ExtensionLike, ExtensionEntry]*) –
- **name** (*Optional[str]*) –
- **trigger** (*TriggerLike*) –
- **priority** (*Optional[int]*) –
- **call_before_training** (*bool*) –
- **kwargs** (*Any*) –

Return type

None

get_optimizer (*name*)

Parameters

name (*str*) –

Return type

Optimizer

property is_before_training: *bool*

is_epoch_last_iter (*idx*)

Parameters

idx (*int*) –

Return type

bool

property iteration: *int*

load_state_dict(*to_load*)

Parameters

to_load (*Dict[str, Any]*) –

Return type

None

property manager: `ExtensionsManager`

property models: `Mapping[str, Module]`

property optimizers: `Mapping[str, Optimizer]`

run(*train_loader*, *val_loader=None*, *, *train_len=None*, *eval_len=None*)

Executes the training loop.

Parameters

- **train_loader** (*torch.utils.data.DataLoader*) – A data loader for training.
- **val_loader** (*torch.utils.data.DataLoader*, *optional*) – A data loader passed to `Evaluator.run()`.
- **train_len** (*int*, *optional*) – The number of iterations per one training epoch. The default value is inferred from the size of training data loader.
- **eval_len** (*int*, *optional*) – The number of iterations per one evaluation epoch, passed to `Evaluator.run()`

Return type

None

See also:

- [`pytorch-pfn-extras.training._evaluator.Evaluator\(\)`](#)

set_optimizer(*name*, *optimizer*)

Parameters

- **name** (*str*) –
- **optimizer** (*Optimizer*) –

Return type

None

state_dict()

Return type

Dict[str, Any]

property stop_trigger: [`Trigger`](#)

pytorch_pfn_extras.handler**Functions**

| | |
|--|---|
| <code>pytorch_pfn_extras.handler.forward(block)</code> | Returns a <code>CodeBlock</code> that performs the forward pass for the given <code>torch.nn.Module</code> or another <code>CodeBlock</code> . |
| <code>pytorch_pfn_extras.handler.torch_autocast(...)</code> | |
| <code>pytorch_pfn_extras.handler.update_parameters(...)</code> | Returns a <code>CodeBlock</code> that performs the forward, backward passes and applies the optimizer step for the given <code>torch.nn.Module</code> or another <code>CodeBlock</code> . |

pytorch_pfn_extras.handler.forward

`pytorch_pfn_extras.handler.forward(block)`

Returns a `CodeBlock` that performs the forward pass for the given `torch.nn.Module` or another `CodeBlock`.

Parameters

block (*Callable*) – `torch.nn.Module` or `CodeBlock` to update the parameters.

Return type

`CodeBlock`

Returns: A `CodeBlock` object.

pytorch_pfn_extras.handler.torch_autocast

`pytorch_pfn_extras.handler.torch_autocast(enabled=True)`

Parameters

enabled (*bool*) –

Return type

`Generator[None, None, None]`

pytorch_pfn_extras.handler.update_parameters

`pytorch_pfn_extras.handler.update_parameters(block, optimizers, backprop_from=None, backprop_to=None)`

Returns a `CodeBlock` that performs the forward, backward passes and applies the optimizer step for the given `torch.nn.Module` or another `CodeBlock`.

Parameters

- **block** (*Callable*) – `torch.nn.Module` or `CodeBlock` to update the parameters.
- **optimizers** (*List[Optimizer]*) – The list of `Optimizer` that will be used for parameter update.
- **backprop_from** (*Optional[str]*) – Select a single output from the block execution to perform the gradient calculation.
- **backprop_to** (*Optional[Set[str]]*) – Name of the values where backpropagation will be stopped.

Return type

CodeBlock

Returns: A CodeBlock object.

Classes

| | |
|---|--|
| <code>pytorch_pfn_extras.handler.BaseHandler(...)</code> | Base class of Handler. |
| <code>pytorch_pfn_extras.handler.BaseLogic([options])</code> | |
| <code>pytorch_pfn_extras.handler.ClousureLogic([...])</code> | A set of methods that defines the training logic. |
| <code>pytorch_pfn_extras.handler.CodeBlock(func, ...)</code> | Class that is used to specify and apply actions to a callable. |
| <code>pytorch_pfn_extras.handler.CodeBlockLogic([...])</code> | A set of methods that defines the training logic. |
| <code>pytorch_pfn_extras.handler.Handler(logic, ...)</code> | A set of callback functions to perform device-specific operations. |
| <code>pytorch_pfn_extras.handler.Logic([...])</code> | A set of methods that defines the training logic. |

pytorch_pfn_extras.handler.BaseHandler**class** `pytorch_pfn_extras.handler.BaseHandler(logic, options, *args, **kwargs)`

Bases: object

Base class of Handler.

Parameters

- **logic** (`Logic`) – A logic.
- **options** (`Dict[str, Any]`) –
- **args** (`Any`) –
- **kwargs** (`Any`) –

Methods

| | |
|---|--|
| <code>__init__(logic, options, *args, **kwargs)</code> | Base class of Handler. |
| <code>consume_options(options)</code> | A method to update options of Handler. |
| <code>eval_loop_begin(evaluator)</code> | A method called before each evaluation step. |
| <code>eval_loop_end(evaluator)</code> | A method called after running all steps of the evaluation. |
| <code>eval_post_step(evaluator, batch_idx, batch, ...)</code> | A method called after each evaluation step. |
| <code>eval_setup(evaluator, loader)</code> | A method called only once when starting a training run. |
| <code>eval_step(evaluator, batch_idx, batch, ...)</code> | Evaluation iteration. |
| <code>train_cleanup(trainer)</code> | A method called only once when completing a training run. |
| <code>train_epoch_begin(trainer, loader)</code> | A method called when starting a new epoch. |
| <code>train_epoch_end(trainer)</code> | A method called when finishing an epoch. |
| <code>train_post_step(trainer, batch_idx, batch, ...)</code> | A method called after each training step. |
| <code>train_setup(trainer, loader)</code> | A method called only once when starting a training run. |
| <code>train_step(trainer, batch_idx, batch, ...)</code> | A training step. |
| <code>train_validation_begin(trainer, evaluator)</code> | A method called when starting a validation. |
| <code>train_validation_end(trainer, evaluator)</code> | A method called after validation. |

`__init__(logic, options, *args, **kwargs)`

Base class of Handler.

Parameters

- **logic** (`Logic`) – A logic.
- **options** (`Dict[str, Any]`) –
- **args** (`Any`) –
- **kwargs** (`Any`) –

Return type

None

`consume_options(options)`

A method to update options of Handler.

Note that the given dict will be modified.

Parameters

- **options** (`dict`) – Option key-values to be set.

Return type

None

`eval_loop_begin(evaluator)`

A method called before each evaluation step.

Parameters

- **evaluator** (`Evaluator`) – The evaluator.

Return type

None

eval_loop_end(*evaluator*)

A method called after running all steps of the evaluation.

Parameters

evaluator ([Evaluator](#)) –

Return type

None

eval_post_step(*evaluator, batch_idx, batch, outputs*)

A method called after each evaluation step.

Parameters

- **evaluator** ([Evaluator](#)) –
- **batch_idx** (*int*) –
- **batch** (*Any*) –
- **outputs** (*Any*) –

Return type

None

eval_setup(*evaluator, loader*)

A method called only once when starting a training run. When evaluator is not given, this method is not called.

Parameters

- **evaluator** ([Evaluator](#)) –
- **loader** (*Iterable[Any]*) –

Return type

None

eval_step(*evaluator, batch_idx, batch, complete_fn*)

Evaluation iteration.

Parameters

- **evaluator** ([Evaluator](#)) –
- **batch_idx** (*int*) –
- **batch** (*Any*) –
- **complete_fn** (*Callable[[int, Any], None]*) –

Return type

None

train_cleanup(*trainer*)

A method called only once when completing a training run.

Parameters

trainer ([Trainer](#)) –

Return type

None

train_epoch_begin(*trainer*, *loader*)

A method called when starting a new epoch.

Parameters

- **trainer** ([Trainer](#)) –
- **loader** ([Iterable\[Any\]](#)) –

Return type

None

train_epoch_end(*trainer*)

A method called when finishing an epoch.

Parameters

- **trainer** ([Trainer](#)) –

Return type

None

train_post_step(*trainer*, *batch_idx*, *batch*, *outputs*)

A method called after each training step.

Parameters

- **trainer** ([Trainer](#)) –
- **batch_idx** ([int](#)) –
- **batch** ([Any](#)) –
- **outputs** ([Any](#)) –

Return type

None

train_setup(*trainer*, *loader*)

A method called only once when starting a training run.

Parameters

- **trainer** ([Trainer](#)) –
- **loader** ([Iterable\[Any\]](#)) –

Return type

None

train_step(*trainer*, *batch_idx*, *batch*, *complete_fn*)

A training step.

Parameters

- **trainer** ([Trainer](#)) –
- **batch_idx** ([int](#)) –
- **batch** ([Any](#)) –
- **complete_fn** ([Callable\[\[int, Any\], None\]](#)) –

Return type

None

train_validation_begin(*trainer, evaluator*)

A method called when starting a validation.

Parameters

- **trainer** ([Trainer](#)) –
- **evaluator** ([Evaluator](#)) –

Return type

None

train_validation_end(*trainer, evaluator*)

A method called after validation.

Parameters

- **trainer** ([Trainer](#)) – The trainer that calls this method.
- **evaluator** ([Evaluator](#)) – The evaluator used for validation.

Return type

None

pytorch_pfn_extras.handler.BaseLogic**class** `pytorch_pfn_extras.handler.BaseLogic`(*options=None*)Bases: `object`**Methods**

| | |
|--|--|
| <code>__init__</code> ([options]) | |
| <code>consume_options</code> (options) | A method to update options of Logic. |
| <code>eval_step</code> (models, batch_idx, batch) | A method for an evaluation step. |
| <code>train_epoch_begin</code> (models, epoch, loader) | A method called when starting a new epoch of training. |
| <code>train_epoch_end</code> (models, epoch) | A method called when completing an epoch of training. |
| <code>train_step</code> (models, optimizers, batch_idx, batch) | A method invokes the models forward and backward passes. |
| <code>train_step_optimizers</code> (models, optimizers, ...) | A method in charge of stepping the provided optimizers. |
| <code>train_validation_begin</code> (models) | A method called when starting a validation. |
| <code>train_validation_end</code> (models) | A method called when the validation completes. |

Parameters**options** (*Optional*[*Dict*[*str*, *Any*]]) –`__init__`(*options=None*)**Parameters****options** (*Optional*[*Dict*[*str*, *Any*]]) –

consume_options(*options*)

A method to update options of Logic.

Note that the given dict will be modified.

Parameters

options (*dict*) – Option key-values to be set.

Return type

None

eval_step(*models, batch_idx, batch*)

A method for an evaluation step.

Parameters

- **models** (*dict of torch.nn.Module*) – The models.
- **batch_idx** (*int*) – Number of steps already finished.
- **batch** (*torch.Tensor, list of torch.Tensor, dict of torch.Tensor*) – Input tensors feeded to the model of the current step.

Return type

Any

train_epoch_begin(*models, epoch, loader*)

A method called when starting a new epoch of training.

Parameters

- **epoch** (*int*) – Number of epochs already finished.
- **models** (*dict of torch.nn.Module*) – The models.
- **loader** (*torch.utils.data.DataLoader*) – The data loader.

Return type

None

train_epoch_end(*models, epoch*)

A method called when completing an epoch of training.

Parameters

- **epoch** (*int*) – Number of epochs already finished.
- **models** (*dict of torch.nn.Module*) – The models.

Return type

None

train_step(*models, optimizers, batch_idx, batch*)

A method invokes the models forward and backward passes.

Optimizing is left to *train_step_optimizers* since maybe the user would like to aggregate the gradients of several iterations.

Parameters

- **models** (*dict of torch.nn.Module*) – The models.
- **optimizers** (*dict of torch.optim.Optimizer*) – The optimizers.
- **batch_idx** (*int*) – Number of training steps already finished.

- **batch** (*torch.Tensor, list of torch.Tensor, dict of torch.Tensor*) – Input tensors feeded to the model of the current step.

Return type

Any

train_step_optimizers(*models, optimizers, batch_idx*)

A method in charge of stepping the provided optimizers.

Parameters

- **optimizers** (*dict of torch.optim.Optimizer*) – The optimizers.
- **batch_idx** (*int*) – Number of steps already finished.
- **models** (*Mapping[str, Module]*) –

Return type

None

train_validation_begin(*models*)

A method called when starting a validation.

Parameters

models (*dict of torch.nn.Module*) – The models.

Return type

None

train_validation_end(*models*)

A method called when the validation completes.

Parameters

models (*dict of torch.nn.Module*) – The models.

Return type

None

pytorch_pfn_extras.handler.ClousureLogic

class pytorch_pfn_extras.handler.ClousureLogic(*model_name='main', options=None*)

Bases: *Logic*

A set of methods that defines the training logic.

Parameters

- **model_name** (*str*) – Name of the model. Default is 'main'.
- **options** (*dict, optional*) – The configuration options.
 - **'backward_outputs'** (*list of str*):
A list of names of outputs that require computation of the gradient.
 - **'autocast'** (*bool or dict*):
If `True`, `torch.autocast` is enabled, using `{"enabled": True, "device_type": "cuda"}` as autocast options. The default is `False` which corresponds to the following options `{"enabled": False, "device_type": "cuda"}`. If dict, options are passed to `torch.autocast`.
 - **'grad_scaler'** (*torch.cuda.amp.GradScaler*):
A gradient scaler that outputs are applied to.

Methods

| | |
|---|---|
| <code>__init__([model_name, options])</code> | A set of methods that defines the training logic. |
| <code>consume_options(options)</code> | A method to update options of Logic. |
| <code>eval_step(models, batch_idx, batch)</code> | A method for an evaluation step. |
| <code>train_epoch_begin(models, epoch, loader)</code> | A method called when starting a new epoch of training. |
| <code>train_epoch_end(models, epoch)</code> | A method called when completing an epoch of training. |
| <code>train_step(models, optimizers, batch_idx, batch)</code> | A method invokes the model forward and backward passes and performs an optimization step. |
| <code>train_step_optimizers(models, optimizers, ...)</code> | In clousure mode, the stepping of the optimizer cannot be changed. |
| <code>train_validation_begin(models)</code> | A method called when starting a validation. |
| <code>train_validation_end(models)</code> | A method called when the validation completes. |

`consume_options(options)`

A method to update options of Logic.

Note that the given dict will be modified.

Parameters

options (*dict*) – Option key-values to be set.

Return type

None

`train_step(models, optimizers, batch_idx, batch)`

A method invokes the model forward and backward passes and performs an optimization step.

Parameters

- **models** (*dict of torch.nn.Module*) – The models.
- **optimizers** (*dict of torch.optim.Optimizer*) – The optimizers.
- **batch_idx** (*int*) – Number of training steps already finished.
- **batch** (*torch.Tensor, list of torch.Tensor, dict of torch.Tensor*) – Input tensors feeded to the model of the current step.

Return type

Any

`train_step_optimizers(models, optimizers, batch_idx)`

In clousure mode, the stepping of the optimizer cannot be changed.

If you want to change the stepping of the optimizer, please use the normal Logic class.

Parameters

- **optimizers** (*dict of torch.optim.Optimizer*) – The optimizers.
- **batch_idx** (*int*) – Number of steps already finished.
- **models** (*Mapping[str, Module]*) –

Return type

None

pytorch_pfn_extras.handler.CodeBlock

```
class pytorch_pfn_extras.handler.CodeBlock(func, optimizers, backprop, backprop_from, backprop_to,
                                             state, runtime)
```

Bases: object

Class that is used to specify and apply actions to a callable.

CodeBlocks are used in Logic classes to write device agnostic codes, as the device runtime is in charge of doing the execution of the module with the actions requested from the codeblock

Parameters

- **func** (*Callable*) – The function to be operated according to the specified options.
- **optimizer** – The Optimizer that will be used for parameter update.
- **backprop** (*bool*) – Flag to specify if gradients are to be calculated.
- **backprop_from** (*Optional[str]*) – Select a single output from the block execution to perform the gradient calculation.
- **backprop_to** (*Optional[Set[str]]*) – Name of the values where backpropagation will be stopped.
- **state** (*Dict[str, Any]*) – Data that can be used during the CodeBlock execution.
- **optimizers** (*List[Optimizer]*) –
- **runtime** (*Any*) –

Methods

```
__init__(func, optimizers, backprop, ...)
```

```
load_state_dict(state)
```

```
state_dict()
```

Attributes

```
func
```

```
optimizers
```

```
backprop
```

```
backprop_from
```

```
backprop_to
```

```
state
```

```
runtime
```

__call__(*inputs*)

Call self as a function.

Parameters

inputs (*Any*) –

Return type

Any

__init__(*func, optimizers, backprop, backprop_from, backprop_to, state, runtime*)

Parameters

- **func** (*Callable*) –
- **optimizers** (*List[Optimizer]*) –
- **backprop** (*bool*) –
- **backprop_from** (*Optional[str]*) –
- **backprop_to** (*Optional[Set[str]]*) –
- **state** (*Dict[str, Any]*) –
- **runtime** (*Any*) –

Return type

None

backprop: *bool*

backprop_from: *Optional[str]*

backprop_to: *Optional[Set[str]]*

func: *Callable*

load_state_dict(*state*)

Parameters

state (*Dict[str, Any]*) –

Return type

None

optimizers: *List[Optimizer]*

runtime: *Any*

state: *Dict[str, Any]*

state_dict()

Return type

Dict[str, Any]

pytorch_pfn_extras.handler.CodeBlockLogic

class pytorch_pfn_extras.handler.CodeBlockLogic(*model_name='main', options=None*)

Bases: *BaseLogic*

A set of methods that defines the training logic.

Parameters

- **model_name** (*str*) – Name of the model. Default is 'main'.
- **options** (*dict, optional*) – The configuration options.
- **'backward_outputs'** (list of *str*):
A list of names of outputs that require computation of the gradient.

Methods

| | |
|---|---|
| <code>__init__([model_name, options])</code> | A set of methods that defines the training logic. |
| <code>consume_options(options)</code> | A method to update options of Logic. |
| <code>eval_step(models, batch_idx, batch)</code> | A method for an evaluation step. |
| <code>train_epoch_begin(models, epoch, loader)</code> | A method called when starting a new epoch of training. |
| <code>train_epoch_end(models, epoch)</code> | A method called when completing an epoch of training. |
| <code>train_step(models, optimizers, batch_idx, batch)</code> | A method invokes the model forward and backward passes. |
| <code>train_step_optimizers(models, optimizers, ...)</code> | A method in charge of stepping the provided optimizers. |
| <code>train_validation_begin(models)</code> | A method called when starting a validation. |
| <code>train_validation_end(models)</code> | A method called when the validation completes. |

`__init__` (*model_name='main', options=None*)

A set of methods that defines the training logic.

Parameters

- **model_name** (*str*) – Name of the model. Default is 'main'.
- **options** (*dict, optional*) – The configuration options.
- **'backward_outputs'** (list of *str*):
A list of names of outputs that require computation of the gradient.

Return type

None

consume_options (*options*)

A method to update options of Logic.

Note that the given dict will be modified.

Parameters

options (*dict*) – Option key-values to be set.

Return type

None

eval_step(*models*, *batch_idx*, *batch*)

A method for an evaluation step.

Parameters

- **models** (*dict of torch.nn.Module*) – The models.
- **batch_idx** (*int*) – Number of steps already finished.
- **batch** (*torch.Tensor, list of torch.Tensor, dict of torch.Tensor*) – Input tensors feeded to the model of the current step.

Return type

Any

train_epoch_begin(*models*, *epoch*, *loader*)

A method called when starting a new epoch of training.

Parameters

- **epoch** (*int*) – Number of epochs already finished.
- **models** (*dict of torch.nn.Module*) – The models.
- **loader** (*torch.utils.data.DataLoader*) – The data loader.

Return type

None

train_epoch_end(*models*, *epoch*)

A method called when completing an epoch of training.

Parameters

- **epoch** (*int*) – Number of epochs already finished.
- **models** (*dict of torch.nn.Module*) – The models.

Return type

None

train_step(*models*, *optimizers*, *batch_idx*, *batch*)

A method invokes the model forward and backward passes.

Optimizing is left to *train_step_optimizers* since maybe the user would like to aggregate the gradients of several iterations.

Parameters

- **models** (*dict of torch.nn.Module*) – The models.
- **optimizers** (*dict of torch.optim.Optimizer*) – The optimizers.
- **batch_idx** (*int*) – Number of training steps already finished.
- **batch** (*torch.Tensor, list of torch.Tensor, dict of torch.Tensor*) – Input tensors feeded to the model of the current step.

Return type

Any

train_validation_begin(*models*)

A method called when starting a validation.

Parameters

- **models** (*dict of torch.nn.Module*) – The models.

Return type

None

train_validation_end(*models*)

A method called when the validation completes.

Parameters**models** (*dict of torch.nn.Module*) – The models.**Return type**

None

pytorch_pfn_extras.handler.Handler**class** pytorch_pfn_extras.handler.**Handler**(*logic, entry_runtime, options*)Bases: *BaseHandler*

A set of callback functions to perform device-specific operations.

Parameters

- **logic** (*Logic*) – A logic.
- **entry_runtime** (*BaseRuntime*) – A runtime object.
- **options** (*dict*) – The configuration options.
 - **'eval_report_keys'** (list of str):
A list of names of outputs that are given as inputs of `reporting.report` after each evaluation step. Default is an empty list.
 - **'train_report_keys'** (list of str):
A list of names of outputs that are given as inputs of `reporting.report` after each training step. Default is an empty list.

Methods

| | |
|---|--|
| <code>__init__</code> (<i>logic, entry_runtime, options</i>) | A set of callback functions to perform device-specific operations. |
| <code>consume_options</code> (<i>options</i>) | A method to update options of Handler. |
| <code>eval_loop_begin</code> (<i>evaluator</i>) | A method called before each evaluation step. |
| <code>eval_loop_end</code> (<i>evaluator</i>) | A method called after running all steps of the evaluation. |
| <code>eval_post_step</code> (<i>evaluator, batch_idx, batch, ...</i>) | A method called after each evaluation step. |
| <code>eval_setup</code> (<i>evaluator, loader</i>) | Called only once when starting a training run. |
| <code>eval_step</code> (<i>evaluator, batch_idx, batch, ...</i>) | Evaluation iteration. |
| <code>train_cleanup</code> (<i>trainer</i>) | A method called only once when compleing a training run. |
| <code>train_epoch_begin</code> (<i>trainer, loader</i>) | A method called when starting a new epoch. |
| <code>train_epoch_end</code> (<i>trainer</i>) | A method called when finishing an epoch. |
| <code>train_post_step</code> (<i>trainer, batch_idx, batch, ...</i>) | A method called after each training step. |
| <code>train_setup</code> (<i>trainer, loader</i>) | A method called only once when starting a training run. |
| <code>train_step</code> (<i>trainer, batch_idx, batch, ...</i>) | A training step. |
| <code>train_validation_begin</code> (<i>trainer, evaluator</i>) | A method called when starting a validation. |
| <code>train_validation_end</code> (<i>trainer, evaluator</i>) | A method called after validation. |

__init__(*logic, entry_runtime, options*)

A set of callback functions to perform device-specific operations.

Parameters

- **logic** (*Logic*) – A logic.
- **entry_runtime** (*BaseRuntime*) – A runtime object.
- **options** (*dict*) – The configuration options.
 - **'eval_report_keys'** (list of str):
A list of names of outputs that are given as inputs of `reporting.report` after each evaluation step. Default is an empty list.
 - **'train_report_keys'** (list of str):
A list of names of outputs that are given as inputs of `reporting.report` after each training step. Default is an empty list.

Return type

None

consume_options(*options*)

A method to update options of Handler.

Note that the given dict will be modified.

Parameters

options (*dict*) – Option key-values to be set.

Return type

None

eval_loop_end(*evaluator*)

A method called after running all steps of the evaluation.

Parameters

evaluator (*Evaluator*) – The evaluator.

Return type

None

eval_post_step(*evaluator, batch_idx, batch, outputs*)

A method called after each evaluation step.

Parameters

- **evaluator** (*Evaluator*) – The evaluator.
- **batch_idx** (*int*) – Number of iterations already finished.
- **batch** (*dict of torch.Tensor*) – Input tensors of this batch.
- **complete_fn** (*callable*) – A callback function called after training step.
- **outputs** (*Any*) –

Return type

None

eval_setup(*evaluator, loader*)

Called only once when starting a training run. When evaluator is not given, this method is not called.

Parameters

- **evaluator** ([Evaluator](#)) – The evaluator.
- **loader** (`torch.utils.data.DataLoader`) – The data loader.

Return type

None

eval_step(*evaluator, batch_idx, batch, complete_fn*)

Evaluation iteration.

Parameters

- **evaluator** ([Evaluator](#)) – The evaluator.
- **batch_idx** (*int*) – Number of iterations already finished.
- **batch** (*dict of torch.Tensor*) – Input tensors of this batch.
- **complete_fn** (*callable*) – A callback function called after training step.

Return type

None

train_cleanup(*trainer*)

A method called only once when compleing a training run.

Parameters

- **trainer** ([Trainer](#)) – The trainer that calls this method.
- **loader** (`torch.utils.data.DataLoader`) – The data loader.

Return type

None

train_epoch_begin(*trainer, loader*)

A method called when starting a new epoch.

Parameters

- **trainer** ([Trainer](#)) – The trainer that calls this method.
- **loader** (`torch.utils.data.DataLoader`) – The data loader.

Return type

None

train_epoch_end(*trainer*)

A method called when finishing an epoch.

Parameters

- **trainer** ([Trainer](#)) – The trainer that calls this method.

Return type

None

train_post_step(*trainer, batch_idx, batch, outputs*)

A method called after each training step.

Parameters

- **trainer** ([Trainer](#)) – The trainer that calls this method.
- **batch_idx** (*int*) – Number of iterations
- **batch** (*dict of torch.Tensor*) – Input tensors of this batch.

- **outputs** (*dict of torch.Tensor*) – Output tensors of this batch.

Return type

None

train_setup(*trainer, loader*)

A method called only once when starting a training run.

Parameters

- **trainer** ([Trainer](#)) – The trainer that calls this method.
- **loader** (*torch.utils.data.DataLoader*) – The data loader.

Return type

None

train_step(*trainer, batch_idx, batch, complete_fn*)

A training step.

Parameters

- **trainer** ([Trainer](#)) – A trainer.
- **batch_idx** (*int*) – Number of iterations already finished.
- **batch** (*dict of torch.Tensor*) – Input tensors of this batch.
- **complete_fn** (*callable*) – A callback function called after training step.

Return type

None

train_validation_begin(*trainer, evaluator*)

A method called when starting a validation.

Parameters

- **evaluator** ([Evaluator](#)) – An evaluator.
- **trainer** ([Trainer](#)) –

Return type

None

train_validation_end(*trainer, evaluator*)

A method called after validation.

Parameters

- **trainer** ([Trainer](#)) – The trainer that calls this method.
- **evaluator** ([Evaluator](#)) – The evaluator used for validation.

Return type

None

pytorch_pfn_extras.handler.Logic

class pytorch_pfn_extras.handler.Logic(*model_name='main', options=None*)

Bases: [BaseLogic](#)

A set of methods that defines the training logic.

Parameters

- **model_name** (*str*) – Name of the model. Default is 'main'.
- **options** (*dict, optional*) – The configuration options.
 - **'backward_outputs'** (list of *str*):
A list of names of outputs that require computation of the gradient.
 - **'autocast'** (bool or dict):
If True, torch.autocast is enabled, using {"enabled": True, "device_type": "cuda"} as autocast options. The default is False which corresponds to the following options {"enabled": False, "device_type": "cuda"}. If dict, options are passed to torch.autocast.
 - **'grad_scaler'** (torch.cuda.amp.GradScaler):
A gradient scaler that outputs are applied to.

Methods

| | |
|--|---|
| __init__ ([<i>model_name</i> , <i>options</i>]) | A set of methods that defines the training logic. |
| consume_options (<i>options</i>) | A method to update options of Logic. |
| eval_step (<i>models</i> , <i>batch_idx</i> , <i>batch</i>) | A method for an evaluation step. |
| train_epoch_begin (<i>models</i> , <i>epoch</i> , <i>loader</i>) | A method called when starting a new epoch of training. |
| train_epoch_end (<i>models</i> , <i>epoch</i>) | A method called when completing an epoch of training. |
| train_step (<i>models</i> , <i>optimizers</i> , <i>batch_idx</i> , <i>batch</i>) | A method invokes the model forward and backward passes. |
| train_step_optimizers (<i>models</i> , <i>optimizers</i> , ...) | A method in charge of stepping the provided optimizers. |
| train_validation_begin (<i>models</i>) | A method called when starting a validation. |
| train_validation_end (<i>models</i>) | A method called when the validation completes. |

__init__ (*model_name='main', options=None*)

A set of methods that defines the training logic.

Parameters

- **model_name** (*str*) – Name of the model. Default is 'main'.
- **options** (*dict, optional*) – The configuration options.
 - **'backward_outputs'** (list of *str*):
A list of names of outputs that require computation of the gradient.
 - **'autocast'** (bool or dict):
If True, torch.autocast is enabled, using {"enabled": True, "device_type": "cuda"} as autocast options. The default is False

which corresponds to the following options {"enabled": False, "device_type": "cuda"}. If dict, options are passed to torch.autocast.

– **'grad_scaler' (torch.cuda.amp.GradScaler):**

A gradient scaler that outputs are applied to.

Return type

None

consume_options(options)

A method to update options of Logic.

Note that the given dict will be modified.

Parameters

options (*dict*) – Option key-values to be set.

Return type

None

eval_step(models, batch_idx, batch)

A method for an evaluation step.

Parameters

- **models** (*dict of torch.nn.Module*) – The models.
- **batch_idx** (*int*) – Number of steps already finished.
- **batch** (*torch.Tensor, list of torch.Tensor, dict of torch.Tensor*) – Input tensors feeded to the model of the current step.

Return type

Any

train_epoch_begin(models, epoch, loader)

A method called when starting a new epoch of training.

Parameters

- **epoch** (*int*) – Number of epochs already finished.
- **models** (*dict of torch.nn.Module*) – The models.
- **loader** (*torch.utils.data.DataLoader*) – The data loader.

Return type

None

train_epoch_end(models, epoch)

A method called when completing an epoch of training.

Parameters

- **epoch** (*int*) – Number of epochs already finished.
- **models** (*dict of torch.nn.Module*) – The models.

Return type

None

train_step(*models*, *optimizers*, *batch_idx*, *batch*)

A method invokes the model forward and backward passes.

Optimizing is left to *train_step_optimizers* since maybe the user would like to aggregate the gradients of several iterations.

Parameters

- **models** (*dict of torch.nn.Module*) – The models.
- **optimizers** (*dict of torch.optim.Optimizer*) – The optimizers.
- **batch_idx** (*int*) – Number of training steps already finished.
- **batch** (*torch.Tensor, list of torch.Tensor, dict of torch.Tensor*) – Input tensors feeded to the model of the current step.

Return type

Any

train_step_optimizers(*models*, *optimizers*, *batch_idx*)

A method in charge of stepping the provided optimizers.

Also a grad scaler will be used if defined.

Parameters

- **optimizers** (*dict of torch.optim.Optimizer*) – The optimizers.
- **batch_idx** (*int*) – Number of steps already finished.
- **models** (*Mapping[str, Module]*) –

Return type

None

train_validation_begin(*models*)

A method called when starting a validation.

Parameters

models (*dict of torch.nn.Module*) – The models.

Return type

None

train_validation_end(*models*)

A method called when the validation completes.

Parameters

models (*dict of torch.nn.Module*) – The models.

Return type

None

pytorch_pfn_extras.logging

Functions

| | |
|--|--|
| <code>pytorch_pfn_extras.logging.get_logger(name)</code> | Returns a child logger to be used by applications. |
|--|--|

pytorch_pfn_extras.logging.get_logger

`pytorch_pfn_extras.logging.get_logger(name)`

Returns a child logger to be used by applications.

Parameters

name (*str*) – Name used to register and retrieve the logger object.

Returns

A logging.Logger object used to log in the application code.

Return type

Logger

pytorch_pfn_extras.nn

Functions

| | |
|--|--|
| <code>pytorch_pfn_extras.nn.ensure(tensor[, ...])</code> | Checks the shape and type of a tensor. |
|--|--|

pytorch_pfn_extras.nn.ensure

`pytorch_pfn_extras.nn.ensure(tensor, shape=None, dtype=None, broadcastable=False, can_cast=False)`

Checks the shape and type of a tensor.

Parameters

- **shape** (*Optional[Tuple[Optional[int], ...]]*) – Tuple with the desired shape. If the input tensor shape is not compatible, *ValueError* will be raised. If *None* is set as a dimension value, that dimension will be ignored.
- **dtype** (*Optional[dtype]*) – Checks if the *dtype* of the input tensor matches the provided one.
- **broadcastable** (*bool*) – Check if the shapes are compatible using broadcasting rules.
- **can_cast** (*bool*) – Check if the input tensor can be casted to the provided type.
- **tensor** (*Tensor*) –

Return type

None

Classes

| | |
|--|--|
| <code>pytorch_pfn_extras.nn.Ensure(*[, shape, ...])</code> | Module to check the shape of a tensor. |
| <code>pytorch_pfn_extras.nn.ExtendedSequential()</code> | Sequential module with extended features from chainer. |
| <code>pytorch_pfn_extras.nn.LazyBatchNorm1d(...)</code> | BatchNorm1d module with lazy weight initialization. |
| <code>pytorch_pfn_extras.nn.LazyBatchNorm2d(...)</code> | BatchNorm2d module with lazy weight initialization. |
| <code>pytorch_pfn_extras.nn.LazyBatchNorm3d(...)</code> | BatchNorm3d module with lazy weight initialization. |
| <code>pytorch_pfn_extras.nn.LazyConv1d(...)</code> | Conv1d module with lazy weight initialization. |
| <code>pytorch_pfn_extras.nn.LazyConv2d(...)</code> | Conv2d module with lazy weight initialization. |
| <code>pytorch_pfn_extras.nn.LazyConv3d(...)</code> | Conv3d module with lazy weight initialization. |
| <code>pytorch_pfn_extras.nn.LazyLinear(...)</code> | Linear module with lazy weight initialization. |

`pytorch_pfn_extras.nn.Ensure`

class `pytorch_pfn_extras.nn.Ensure(*, shape=None, dtype=None, broadcastable=False, can_cast=False)`

Bases: Module

Module to check the shape of a tensor.

Parameters

- **shape** (*Optional[Tuple[Optional[int], ...]]*) – Tuple with the desired shape. If the input tensor shape is not compatible, *ValueError* will be raised. If *None* is set as a dimension value, that dimension will be ignored.
- **dtype** (*Optional[dtype]*) – Checks if the *dtype* of the input tensor matches the provided one.
- **broadcastable** (*bool*) – Check if the shapes are compatible using broadcasting rules.
- **can_cast** (*bool*) – Check if the input tensor can be casted to the provided type.

Initializes internal Module state, shared by both `nn.Module` and `ScriptModule`.

Methods

| | |
|--|--|
| <code>__init__(*[, shape, dtype, broadcastable, ...])</code> | Initializes internal Module state, shared by both <code>nn.Module</code> and <code>ScriptModule</code> . |
| <code>add_module(name, module)</code> | Adds a child module to the current module. |
| <code>apply(fn)</code> | Applies <code>fn</code> recursively to every submodule (as returned by <code>.children()</code>) as well as self. |
| <code>bfloat16()</code> | Casts all floating point parameters and buffers to <code>bfloat16</code> datatype. |
| <code>buffers([recurse])</code> | Returns an iterator over module buffers. |
| <code>children()</code> | Returns an iterator over immediate children modules. |
| <code>compile(*args, **kwargs)</code> | Compile this Module's forward using <code>torch.compile()</code> . |
| <code>cpu()</code> | Moves all model parameters and buffers to the CPU. |
| <code>cuda([device])</code> | Moves all model parameters and buffers to the GPU. |
| <code>double()</code> | Casts all floating point parameters and buffers to <code>double</code> datatype. |
| <code>eval()</code> | Sets the module in evaluation mode. |

continues on next page

Table 1 – continued from previous page

| | |
|---|---|
| <code>extra_repr()</code> | Set the extra representation of the module |
| <code>float()</code> | Casts all floating point parameters and buffers to <code>float</code> datatype. |
| <code>forward(input)</code> | Defines the computation performed at every call. |
| <code>get_buffer(target)</code> | Returns the buffer given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_extra_state()</code> | Returns any extra state to include in the module's <code>state_dict</code> . |
| <code>get_parameter(target)</code> | Returns the parameter given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_submodule(target)</code> | Returns the submodule given by <code>target</code> if it exists, otherwise throws an error. |
| <code>half()</code> | Casts all floating point parameters and buffers to <code>half</code> datatype. |
| <code>ipu([device])</code> | Moves all model parameters and buffers to the IPU. |
| <code>load_state_dict(state_dict[, strict, assign])</code> | Copies parameters and buffers from <code>state_dict</code> into this module and its descendants. |
| <code>modules()</code> | Returns an iterator over all modules in the network. |
| <code>named_buffers([prefix, recurse, ...])</code> | Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself. |
| <code>named_children()</code> | Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself. |
| <code>named_modules([memo, prefix, remove_duplicate])</code> | Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself. |
| <code>named_parameters([prefix, recurse, ...])</code> | Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself. |
| <code>parameters([recurse])</code> | Returns an iterator over module parameters. |
| <code>register_backward_hook(hook)</code> | Registers a backward hook on the module. |
| <code>register_buffer(name, tensor[, persistent])</code> | Adds a buffer to the module. |
| <code>register_forward_hook(hook, *[, prepend, ...])</code> | Registers a forward hook on the module. |
| <code>register_forward_pre_hook(hook, *[, ...])</code> | Registers a forward pre-hook on the module. |
| <code>register_full_backward_hook(hook[, prepend])</code> | Registers a backward hook on the module. |
| <code>register_full_backward_pre_hook(hook[, prepend])</code> | Registers a backward pre-hook on the module. |
| <code>register_load_state_dict_post_hook(hook)</code> | Registers a post hook to be run after module's <code>load_state_dict</code> is called. |
| <code>register_module(name, module)</code> | Alias for <code>add_module()</code> . |
| <code>register_parameter(name, param)</code> | Adds a parameter to the module. |
| <code>register_state_dict_pre_hook(hook)</code> | These hooks will be called with arguments: <code>self</code> , <code>prefix</code> , and <code>keep_vars</code> before calling <code>state_dict</code> on <code>self</code> . |
| <code>requires_grad_([requires_grad])</code> | Change if autograd should record operations on parameters in this module. |
| <code>set_extra_state(state)</code> | This function is called from <code>load_state_dict()</code> to handle any extra state found within the <code>state_dict</code> . |
| <code>share_memory()</code> | See <code>torch.Tensor.share_memory_()</code> |

continues on next page

Table 1 – continued from previous page

| | |
|--|---|
| <code>state_dict(*args[, destination, prefix, ...])</code> | Returns a dictionary containing references to the whole state of the module. |
| <code>to(*args, **kwargs)</code> | Moves and/or casts the parameters and buffers. |
| <code>to_empty(*, device[, recurse])</code> | Moves the parameters and buffers to the specified device without copying storage. |
| <code>train([mode])</code> | Sets the module in training mode. |
| <code>type(dst_type)</code> | Casts all parameters and buffers to <code>dst_type</code> . |
| <code>xpu([device])</code> | Moves all model parameters and buffers to the XPU. |
| <code>zero_grad([set_to_none])</code> | Resets gradients of all model parameters. |

Attributes

| | |
|------------------------------|--|
| <code>T_destination</code> | alias of <code>TypeVar('T_destination', bound=Dict[str, Any])</code> |
| <code>call_super_init</code> | |
| <code>dump_patches</code> | |

`__init__`(*, *shape=None, dtype=None, broadcastable=False, can_cast=False*)

Initializes internal Module state, shared by both `nn.Module` and `ScriptModule`.

Parameters

- **shape** (*Optional[Tuple[Optional[int], ...]]*) –
- **dtype** (*Optional[dtype]*) –
- **broadcastable** (*bool*) –
- **can_cast** (*bool*) –

forward(*input*)

Defines the computation performed at every call.

Should be overridden by all subclasses.

Note: Although the recipe for forward pass needs to be defined within this function, one should call the Module instance afterwards instead of this since the former takes care of running the registered hooks while the latter silently ignores them.

Parameters

input (*Tensor*) –

Return type

Tensor

training: `bool`

pytorch_pfn_extras.nn.ExtendedSequential

class pytorch_pfn_extras.nn.**ExtendedSequential**(*args: Module)

class pytorch_pfn_extras.nn.**ExtendedSequential**(arg: OrderedDict[str, Module])

Bases: Sequential

Sequential module with extended features from chainer.

Initializes internal Module state, shared by both nn.Module and ScriptModule.

Methods

| | |
|--|--|
| <code>__init__(*args)</code> | Initializes internal Module state, shared by both nn.Module and ScriptModule. |
| <code>add_module(name, module)</code> | Adds a child module to the current module. |
| <code>append(module)</code> | Appends a given module to the end. |
| <code>apply(fn)</code> | Applies <code>fn</code> recursively to every submodule (as returned by <code>.children()</code>) as well as self. |
| <code>bfloat16()</code> | Casts all floating point parameters and buffers to <code>bfloat16</code> datatype. |
| <code>buffers([recurse])</code> | Returns an iterator over module buffers. |
| <code>children()</code> | Returns an iterator over immediate children modules. |
| <code>compile(*args, **kwargs)</code> | Compile this Module's forward using <code>torch.compile()</code> . |
| <code>cpu()</code> | Moves all model parameters and buffers to the CPU. |
| <code>cuda([device])</code> | Moves all model parameters and buffers to the GPU. |
| <code>double()</code> | Casts all floating point parameters and buffers to <code>double</code> datatype. |
| <code>eval()</code> | Sets the module in evaluation mode. |
| <code>extend(sequential)</code> | |
| <code>extra_repr()</code> | Set the extra representation of the module |
| <code>float()</code> | Casts all floating point parameters and buffers to <code>float</code> datatype. |
| <code>forward(input)</code> | Defines the computation performed at every call. |
| <code>get_buffer(target)</code> | Returns the buffer given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_extra_state()</code> | Returns any extra state to include in the module's <code>state_dict</code> . |
| <code>get_parameter(target)</code> | Returns the parameter given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_submodule(target)</code> | Returns the submodule given by <code>target</code> if it exists, otherwise throws an error. |
| <code>half()</code> | Casts all floating point parameters and buffers to <code>half</code> datatype. |
| <code>insert(index, module)</code> | |
| <code>ipu([device])</code> | Moves all model parameters and buffers to the IPU. |
| <code>load_state_dict(state_dict[, strict, assign])</code> | Copies parameters and buffers from <code>state_dict</code> into this module and its descendants. |
| <code>modules()</code> | Returns an iterator over all modules in the network. |

continues on next page

Table 2 – continued from previous page

| | |
|---|---|
| <code>named_buffers([prefix, recurse, ...])</code> | Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself. |
| <code>named_children()</code> | Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself. |
| <code>named_modules([memo, prefix, remove_duplicate])</code> | Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself. |
| <code>named_parameters([prefix, recurse, ...])</code> | Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself. |
| <code>parameters([recurse])</code> | Returns an iterator over module parameters. |
| <code>pop(key)</code> | |
| <code>register_backward_hook(hook)</code> | Registers a backward hook on the module. |
| <code>register_buffer(name, tensor[, persistent])</code> | Adds a buffer to the module. |
| <code>register_forward_hook(hook, *[, prepend, ...])</code> | Registers a forward hook on the module. |
| <code>register_forward_pre_hook(hook, *[, ...])</code> | Registers a forward pre-hook on the module. |
| <code>register_full_backward_hook(hook[, prepend])</code> | Registers a backward hook on the module. |
| <code>register_full_backward_pre_hook(hook[, prepend])</code> | Registers a backward pre-hook on the module. |
| <code>register_load_state_dict_post_hook(hook)</code> | Registers a post hook to be run after module's <code>load_state_dict</code> is called. |
| <code>register_module(name, module)</code> | Alias for <code>add_module()</code> . |
| <code>register_parameter(name, param)</code> | Adds a parameter to the module. |
| <code>register_state_dict_pre_hook(hook)</code> | These hooks will be called with arguments: <code>self</code> , <code>prefix</code> , and <code>keep_vars</code> before calling <code>state_dict</code> on <code>self</code> . |
| <code>repeat(n_repeat[, mode])</code> | Repeats this Sequential multiple times. |
| <code>requires_grad_([requires_grad])</code> | Change if autograd should record operations on parameters in this module. |
| <code>set_extra_state(state)</code> | This function is called from <code>load_state_dict()</code> to handle any extra state found within the <code>state_dict</code> . |
| <code>share_memory()</code> | See <code>torch.Tensor.share_memory_()</code> |
| <code>state_dict(*args[, destination, prefix, ...])</code> | Returns a dictionary containing references to the whole state of the module. |
| <code>to(*args, **kwargs)</code> | Moves and/or casts the parameters and buffers. |
| <code>to_empty(*, device[, recurse])</code> | Moves the parameters and buffers to the specified device without copying storage. |
| <code>train([mode])</code> | Sets the module in training mode. |
| <code>type(dst_type)</code> | Casts all parameters and buffers to <code>dst_type</code> . |
| <code>xpu([device])</code> | Moves all model parameters and buffers to the XPU. |
| <code>zero_grad([set_to_none])</code> | Resets gradients of all model parameters. |

Attributes

| | |
|------------------------------|--|
| <code>T_destination</code> | alias of <code>TypeVar('T_destination', bound=Dict[str, Any])</code> |
| <code>call_super_init</code> | |
| <code>dump_patches</code> | |

repeat(*n_repeat*, *mode*='init')

Repeats this Sequential multiple times.

This method returns a `Sequential` object which has original *Sequential* multiple times repeatedly. The *mode* argument means how to copy this sequential to repeat.

The functions is supposed to behave the same way as *repeat* in *chainer*.

When the mode is set to *init*, the default value, modules will be copied and reinitialized by calling `reset_parameters` (or `_reset_parameters`) method.

To repeat user-defined modules, which have parameters or buffers, with *mode*='init' in this `Sequential`, you need to implement the `reset_parameters` or `_reset_parameters` method to the module to reinitialize parameters and (if necessary) buffers; otherwise the initialization cannot be performed and a warning message will be shown.

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.

Return type

`ExtendedSequential`

pytorch_pfn_extras.nn.LazyBatchNorm1d

class `pytorch_pfn_extras.nn.LazyBatchNorm1d`(*num_features*, *args, **kwargs)

Bases: `_LazyBatchNorm`, `BatchNorm1d`

`BatchNorm1d` module with lazy weight initialization.

When *num_features* is `None`, it is determined at the first time of the forward step.

Initializes internal Module state, shared by both `nn.Module` and `ScriptModule`.

Methods

| | |
|--|---|
| <code>__init__(num_features, *args, **kwargs)</code> | Initializes internal Module state, shared by both <code>nn.Module</code> and <code>ScriptModule</code> . |
| <code>add_module(name, module)</code> | Adds a child module to the current module. |
| <code>apply(fn)</code> | Applies <code>fn</code> recursively to every submodule (as returned by <code>.children()</code>) as well as self. |
| <code>bfloat16()</code> | Casts all floating point parameters and buffers to <code>bfloat16</code> datatype. |
| <code>buffers([recurse])</code> | Returns an iterator over module buffers. |
| <code>children()</code> | Returns an iterator over immediate children modules. |
| <code>compile(*args, **kwargs)</code> | Compile this Module's forward using <code>torch.compile()</code> . |
| <code>cpu()</code> | Moves all model parameters and buffers to the CPU. |
| <code>cuda([device])</code> | Moves all model parameters and buffers to the GPU. |
| <code>double()</code> | Casts all floating point parameters and buffers to <code>double</code> datatype. |
| <code>eval()</code> | Sets the module in evaluation mode. |
| <code>extra_repr()</code> | Set the extra representation of the module |
| <code>float()</code> | Casts all floating point parameters and buffers to <code>float</code> datatype. |
| <code>forward(input)</code> | Defines the computation performed at every call. |
| <code>get_buffer(target)</code> | Returns the buffer given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_extra_state()</code> | Returns any extra state to include in the module's <code>state_dict</code> . |
| <code>get_parameter(target)</code> | Returns the parameter given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_submodule(target)</code> | Returns the submodule given by <code>target</code> if it exists, otherwise throws an error. |
| <code>half()</code> | Casts all floating point parameters and buffers to <code>half</code> datatype. |
| <code>ipu([device])</code> | Moves all model parameters and buffers to the IPU. |
| <code>load_state_dict(state_dict[, strict, assign])</code> | Copies parameters and buffers from <code>state_dict</code> into this module and its descendants. |
| <code>modules()</code> | Returns an iterator over all modules in the network. |
| <code>named_buffers([prefix, recurse, ...])</code> | Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself. |
| <code>named_children()</code> | Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself. |
| <code>named_modules([memo, prefix, remove_duplicate])</code> | Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself. |
| <code>named_parameters([prefix, recurse, ...])</code> | Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself. |
| <code>parameters([recurse])</code> | Returns an iterator over module parameters. |
| <code>register_backward_hook(hook)</code> | Registers a backward hook on the module. |
| <code>register_buffer(name, tensor[, persistent])</code> | Adds a buffer to the module. |
| <code>register_forward_hook(hook, *[, prepend, ...])</code> | Registers a forward hook on the module. |

continues on next page

Table 3 – continued from previous page

| | |
|---|---|
| <code>register_forward_pre_hook(hook, *[, ...])</code> | Registers a forward pre-hook on the module. |
| <code>register_full_backward_hook(hook[, prepend])</code> | Registers a backward hook on the module. |
| <code>register_full_backward_pre_hook(hook[, prepend])</code> | Registers a backward pre-hook on the module. |
| <code>register_load_state_dict_post_hook(hook)</code> | Registers a post hook to be run after module's <code>load_state_dict</code> is called. |
| <code>register_module(name, module)</code> | Alias for <code>add_module()</code> . |
| <code>register_parameter(name, param)</code> | Adds a parameter to the module. |
| <code>register_state_dict_pre_hook(hook)</code> | These hooks will be called with arguments: <code>self</code> , <code>prefix</code> , and <code>keep_vars</code> before calling <code>state_dict</code> on <code>self</code> . |
| <code>requires_grad_([requires_grad])</code> | Change if autograd should record operations on parameters in this module. |
| <code>reset_parameters()</code> | |
| <code>reset_running_stats()</code> | |
| <code>set_extra_state(state)</code> | This function is called from <code>load_state_dict()</code> to handle any extra state found within the <i>state_dict</i> . |
| <code>share_memory()</code> | See <code>torch.Tensor.share_memory_()</code> |
| <code>state_dict(*args, **kwargs)</code> | Returns a dictionary containing a whole state of the module. |
| <code>to(*args, **kwargs)</code> | Moves and/or casts the parameters and buffers. |
| <code>to_empty(*, device[, recurse])</code> | Moves the parameters and buffers to the specified device without copying storage. |
| <code>train([mode])</code> | Sets the module in training mode. |
| <code>type(dst_type)</code> | Casts all parameters and buffers to <code>dst_type</code> . |
| <code>xpu([device])</code> | Moves all model parameters and buffers to the XPU. |
| <code>zero_grad([set_to_none])</code> | Resets gradients of all model parameters. |

Attributes

| | |
|---|--|
| <code>T_destination</code> | alias of <code>TypeVar('T_destination', bound=Dict[str, Any])</code> |
| <code>call_super_init</code> | |
| <code>dump_patches</code> | |
| <code>lazy_buffer_names</code> | |
| <code>lazy_parameter_names</code> | |
| <code>lazy_parameters_determined</code> | Returns if all lazy parameters are determined. |

Parameters

- `num_features` (*int*) –
- `args` (*Any*) –
- `kwargs` (*Any*) –

running_mean: Any

running_var: Any

pytorch_pfn_extras.nn.LazyBatchNorm2d

class pytorch_pfn_extras.nn.LazyBatchNorm2d(*num_features*, *args, **kwargs)

Bases: _LazyBatchNorm, BatchNorm2d

BatchNorm2d module with lazy weight initialization.

When *num_features* is None, it is determined at the first time of the forward step.

Initializes internal Module state, shared by both nn.Module and ScriptModule.

Methods

| | |
|--|--|
| <code>__init__(num_features, *args, **kwargs)</code> | Initializes internal Module state, shared by both nn.Module and ScriptModule. |
| <code>add_module(name, module)</code> | Adds a child module to the current module. |
| <code>apply(fn)</code> | Applies <i>fn</i> recursively to every submodule (as returned by <code>.children()</code>) as well as self. |
| <code>bfloat16()</code> | Casts all floating point parameters and buffers to bfloat16 datatype. |
| <code>buffers([recurse])</code> | Returns an iterator over module buffers. |
| <code>children()</code> | Returns an iterator over immediate children modules. |
| <code>compile(*args, **kwargs)</code> | Compile this Module's forward using <code>torch.compile()</code> . |
| <code>cpu()</code> | Moves all model parameters and buffers to the CPU. |
| <code>cuda([device])</code> | Moves all model parameters and buffers to the GPU. |
| <code>double()</code> | Casts all floating point parameters and buffers to double datatype. |
| <code>eval()</code> | Sets the module in evaluation mode. |
| <code>extra_repr()</code> | Set the extra representation of the module |
| <code>float()</code> | Casts all floating point parameters and buffers to float datatype. |
| <code>forward(input)</code> | Defines the computation performed at every call. |
| <code>get_buffer(target)</code> | Returns the buffer given by <i>target</i> if it exists, otherwise throws an error. |
| <code>get_extra_state()</code> | Returns any extra state to include in the module's <i>state_dict</i> . |
| <code>get_parameter(target)</code> | Returns the parameter given by <i>target</i> if it exists, otherwise throws an error. |
| <code>get_submodule(target)</code> | Returns the submodule given by <i>target</i> if it exists, otherwise throws an error. |
| <code>half()</code> | Casts all floating point parameters and buffers to half datatype. |
| <code>ipu([device])</code> | Moves all model parameters and buffers to the IPU. |
| <code>load_state_dict(state_dict[, strict, assign])</code> | Copies parameters and buffers from <i>state_dict</i> into this module and its descendants. |
| <code>modules()</code> | Returns an iterator over all modules in the network. |

continues on next page

Table 4 – continued from previous page

| | |
|---|---|
| <code>named_buffers([prefix, recurse, ...])</code> | Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself. |
| <code>named_children()</code> | Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself. |
| <code>named_modules([memo, prefix, remove_duplicate])</code> | Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself. |
| <code>named_parameters([prefix, recurse, ...])</code> | Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself. |
| <code>parameters([recurse])</code> | Returns an iterator over module parameters. |
| <code>register_backward_hook(hook)</code> | Registers a backward hook on the module. |
| <code>register_buffer(name, tensor[, persistent])</code> | Adds a buffer to the module. |
| <code>register_forward_hook(hook, *[, prepend, ...])</code> | Registers a forward hook on the module. |
| <code>register_forward_pre_hook(hook, *[, ...])</code> | Registers a forward pre-hook on the module. |
| <code>register_full_backward_hook(hook[, prepend])</code> | Registers a backward hook on the module. |
| <code>register_full_backward_pre_hook(hook[, prepend])</code> | Registers a backward pre-hook on the module. |
| <code>register_load_state_dict_post_hook(hook)</code> | Registers a post hook to be run after module's <code>load_state_dict</code> is called. |
| <code>register_module(name, module)</code> | Alias for <code>add_module()</code> . |
| <code>register_parameter(name, param)</code> | Adds a parameter to the module. |
| <code>register_state_dict_pre_hook(hook)</code> | These hooks will be called with arguments: <code>self</code> , <code>prefix</code> , and <code>keep_vars</code> before calling <code>state_dict</code> on <code>self</code> . |
| <code>requires_grad_([requires_grad])</code> | Change if autograd should record operations on parameters in this module. |
| <code>reset_parameters()</code> | |
| <code>reset_running_stats()</code> | |
| <code>set_extra_state(state)</code> | This function is called from <code>load_state_dict()</code> to handle any extra state found within the <code>state_dict</code> . |
| <code>share_memory()</code> | See <code>torch.Tensor.share_memory_()</code> |
| <code>state_dict(*args, **kwargs)</code> | Returns a dictionary containing a whole state of the module. |
| <code>to(*args, **kwargs)</code> | Moves and/or casts the parameters and buffers. |
| <code>to_empty(*, device[, recurse])</code> | Moves the parameters and buffers to the specified device without copying storage. |
| <code>train([mode])</code> | Sets the module in training mode. |
| <code>type(dst_type)</code> | Casts all parameters and buffers to <code>dst_type</code> . |
| <code>xpu([device])</code> | Moves all model parameters and buffers to the XPU. |
| <code>zero_grad([set_to_none])</code> | Resets gradients of all model parameters. |

Attributes

| | |
|---|--|
| <code>T_destination</code> | alias of <code>TypeVar('T_destination', bound=Dict[str, Any])</code> |
| <code>call_super_init</code> | |
| <code>dump_patches</code> | |
| <code>lazy_buffer_names</code> | |
| <code>lazy_parameter_names</code> | |
| <code>lazy_parameters_determined</code> | Returns if all lazy parameters are determined. |

Parameters

- **num_features** (*int*) –
- **args** (*Any*) –
- **kwargs** (*Any*) –

running_mean: *Any*

running_var: *Any*

pytorch_pfn_extras.nn.LazyBatchNorm3d

class `pytorch_pfn_extras.nn.LazyBatchNorm3d`(*num_features*, **args*, ***kwargs*)

Bases: `_LazyBatchNorm`, `BatchNorm3d`

BatchNorm3d module with lazy weight initialization.

When `num_features` is `None`, it is determined at the first time of the forward step.

Initializes internal Module state, shared by both `nn.Module` and `ScriptModule`.

Methods

| | |
|--|--|
| <code>__init__(num_features, *args, **kwargs)</code> | Initializes internal Module state, shared by both <code>nn.Module</code> and <code>ScriptModule</code> . |
| <code>add_module(name, module)</code> | Adds a child module to the current module. |
| <code>apply(fn)</code> | Applies <code>fn</code> recursively to every submodule (as returned by <code>.children()</code>) as well as self. |
| <code>bfloat16()</code> | Casts all floating point parameters and buffers to <code>bfloat16</code> datatype. |
| <code>buffers([recurse])</code> | Returns an iterator over module buffers. |
| <code>children()</code> | Returns an iterator over immediate children modules. |
| <code>compile(*args, **kwargs)</code> | Compile this Module's forward using <code>torch.compile()</code> . |
| <code>cpu()</code> | Moves all model parameters and buffers to the CPU. |
| <code>cuda([device])</code> | Moves all model parameters and buffers to the GPU. |

continues on next page

Table 5 – continued from previous page

| | |
|---|---|
| <code>double()</code> | Casts all floating point parameters and buffers to <code>double</code> datatype. |
| <code>eval()</code> | Sets the module in evaluation mode. |
| <code>extra_repr()</code> | Set the extra representation of the module |
| <code>float()</code> | Casts all floating point parameters and buffers to <code>float</code> datatype. |
| <code>forward(input)</code> | Defines the computation performed at every call. |
| <code>get_buffer(target)</code> | Returns the buffer given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_extra_state()</code> | Returns any extra state to include in the module's <code>state_dict</code> . |
| <code>get_parameter(target)</code> | Returns the parameter given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_submodule(target)</code> | Returns the submodule given by <code>target</code> if it exists, otherwise throws an error. |
| <code>half()</code> | Casts all floating point parameters and buffers to <code>half</code> datatype. |
| <code>ipu([device])</code> | Moves all model parameters and buffers to the IPU. |
| <code>load_state_dict(state_dict[, strict, assign])</code> | Copies parameters and buffers from <code>state_dict</code> into this module and its descendants. |
| <code>modules()</code> | Returns an iterator over all modules in the network. |
| <code>named_buffers([prefix, recurse, ...])</code> | Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself. |
| <code>named_children()</code> | Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself. |
| <code>named_modules([memo, prefix, remove_duplicate])</code> | Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself. |
| <code>named_parameters([prefix, recurse, ...])</code> | Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself. |
| <code>parameters([recurse])</code> | Returns an iterator over module parameters. |
| <code>register_backward_hook(hook)</code> | Registers a backward hook on the module. |
| <code>register_buffer(name, tensor[, persistent])</code> | Adds a buffer to the module. |
| <code>register_forward_hook(hook, *[, prepend, ...])</code> | Registers a forward hook on the module. |
| <code>register_forward_pre_hook(hook, *[, ...])</code> | Registers a forward pre-hook on the module. |
| <code>register_full_backward_hook(hook[, prepend])</code> | Registers a backward hook on the module. |
| <code>register_full_backward_pre_hook(hook[, prepend])</code> | Registers a backward pre-hook on the module. |
| <code>register_load_state_dict_post_hook(hook)</code> | Registers a post hook to be run after module's <code>load_state_dict</code> is called. |
| <code>register_module(name, module)</code> | Alias for <code>add_module()</code> . |
| <code>register_parameter(name, param)</code> | Adds a parameter to the module. |
| <code>register_state_dict_pre_hook(hook)</code> | These hooks will be called with arguments: <code>self</code> , <code>prefix</code> , and <code>keep_vars</code> before calling <code>state_dict</code> on <code>self</code> . |
| <code>requires_grad_([requires_grad])</code> | Change if autograd should record operations on parameters in this module. |

continues on next page

Table 5 – continued from previous page

| | |
|---|--|
| <code>reset_parameters()</code> | |
| <code>reset_running_stats()</code> | |
| <code>set_extra_state(state)</code> | This function is called from <code>load_state_dict()</code> to handle any extra state found within the <i>state_dict</i> . |
| <code>share_memory()</code> | See <code>torch.Tensor.share_memory_()</code> |
| <code>state_dict(*args, **kwargs)</code> | Returns a dictionary containing a whole state of the module. |
| <code>to(*args, **kwargs)</code> | Moves and/or casts the parameters and buffers. |
| <code>to_empty(*, device[, recurse])</code> | Moves the parameters and buffers to the specified device without copying storage. |
| <code>train([mode])</code> | Sets the module in training mode. |
| <code>type(dst_type)</code> | Casts all parameters and buffers to <i>dst_type</i> . |
| <code>xpu([device])</code> | Moves all model parameters and buffers to the XPU. |
| <code>zero_grad([set_to_none])</code> | Resets gradients of all model parameters. |

Attributes

| | |
|---|--|
| <code>T_destination</code> | alias of <code>TypeVar('T_destination', bound=Dict[str, Any])</code> |
| <code>call_super_init</code> | |
| <code>dump_patches</code> | |
| <code>lazy_buffer_names</code> | |
| <code>lazy_parameter_names</code> | |
| <code>lazy_parameters_determined</code> | Returns if all lazy parameters are determined. |

Parameters

- `num_features` (*int*) –
- `args` (*Any*) –
- `kwargs` (*Any*) –

`running_mean:` *Any*

`running_var:` *Any*

pytorch_pfn_extras.nn.LazyConv1d

class pytorch_pfn_extras.nn.LazyConv1d(*in_channels*, *args, **kwargs)

Bases: `_LazyConvNd`, `Conv1d`

Conv1d module with lazy weight initialization.

When *in_channels* is `None`, it is determined at the first time of the forward step.

Initializes internal Module state, shared by both `nn.Module` and `ScriptModule`.

Methods

| | |
|--|---|
| <code>__init__(in_channels, *args, **kwargs)</code> | Initializes internal Module state, shared by both <code>nn.Module</code> and <code>ScriptModule</code> . |
| <code>add_module(name, module)</code> | Adds a child module to the current module. |
| <code>apply(fn)</code> | Applies <code>fn</code> recursively to every submodule (as returned by <code>.children()</code>) as well as self. |
| <code>bfloat16()</code> | Casts all floating point parameters and buffers to <code>bfloat16</code> datatype. |
| <code>buffers([recurse])</code> | Returns an iterator over module buffers. |
| <code>children()</code> | Returns an iterator over immediate children modules. |
| <code>compile(*args, **kwargs)</code> | Compile this Module's forward using <code>torch.compile()</code> . |
| <code>cpu()</code> | Moves all model parameters and buffers to the CPU. |
| <code>cuda([device])</code> | Moves all model parameters and buffers to the GPU. |
| <code>double()</code> | Casts all floating point parameters and buffers to <code>double</code> datatype. |
| <code>eval()</code> | Sets the module in evaluation mode. |
| <code>extra_repr()</code> | Set the extra representation of the module |
| <code>float()</code> | Casts all floating point parameters and buffers to <code>float</code> datatype. |
| <code>forward(input)</code> | Defines the computation performed at every call. |
| <code>get_buffer(target)</code> | Returns the buffer given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_extra_state()</code> | Returns any extra state to include in the module's <code>state_dict</code> . |
| <code>get_parameter(target)</code> | Returns the parameter given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_submodule(target)</code> | Returns the submodule given by <code>target</code> if it exists, otherwise throws an error. |
| <code>half()</code> | Casts all floating point parameters and buffers to <code>half</code> datatype. |
| <code>ipu([device])</code> | Moves all model parameters and buffers to the IPU. |
| <code>load_state_dict(state_dict[, strict, assign])</code> | Copies parameters and buffers from <code>state_dict</code> into this module and its descendants. |
| <code>modules()</code> | Returns an iterator over all modules in the network. |
| <code>named_buffers([prefix, recurse, ...])</code> | Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself. |
| <code>named_children()</code> | Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself. |

continues on next page

Table 6 – continued from previous page

| | |
|---|---|
| <code>named_modules([memo, prefix, remove_duplicate])</code> | Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself. |
| <code>named_parameters([prefix, recurse, ...])</code> | Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself. |
| <code>parameters([recurse])</code> | Returns an iterator over module parameters. |
| <code>register_backward_hook(hook)</code> | Registers a backward hook on the module. |
| <code>register_buffer(name, tensor[, persistent])</code> | Adds a buffer to the module. |
| <code>register_forward_hook(hook, *[, prepend, ...])</code> | Registers a forward hook on the module. |
| <code>register_forward_pre_hook(hook, *[, ...])</code> | Registers a forward pre-hook on the module. |
| <code>register_full_backward_hook(hook[, prepend])</code> | Registers a backward hook on the module. |
| <code>register_full_backward_pre_hook(hook[, prepend])</code> | Registers a backward pre-hook on the module. |
| <code>register_load_state_dict_post_hook(hook)</code> | Registers a post hook to be run after module's <code>load_state_dict</code> is called. |
| <code>register_module(name, module)</code> | Alias for <code>add_module()</code> . |
| <code>register_parameter(name, param)</code> | Adds a parameter to the module. |
| <code>register_state_dict_pre_hook(hook)</code> | These hooks will be called with arguments: <code>self</code> , <code>prefix</code> , and <code>keep_vars</code> before calling <code>state_dict</code> on <code>self</code> . |
| <code>requires_grad_([requires_grad])</code> | Change if autograd should record operations on parameters in this module. |
| <code>reset_parameters()</code> | |
| <code>set_extra_state(state)</code> | This function is called from <code>load_state_dict()</code> to handle any extra state found within the <code>state_dict</code> . |
| <code>share_memory()</code> | See <code>torch.Tensor.share_memory_()</code> |
| <code>state_dict(*args, **kwargs)</code> | Returns a dictionary containing a whole state of the module. |
| <code>to(*args, **kwargs)</code> | Moves and/or casts the parameters and buffers. |
| <code>to_empty(*, device[, recurse])</code> | Moves the parameters and buffers to the specified device without copying storage. |
| <code>train([mode])</code> | Sets the module in training mode. |
| <code>type(dst_type)</code> | Casts all parameters and buffers to <code>dst_type</code> . |
| <code>xpu([device])</code> | Moves all model parameters and buffers to the XPU. |
| <code>zero_grad([set_to_none])</code> | Resets gradients of all model parameters. |

Attributes

| | |
|---|--|
| <code>T_destination</code> | alias of <code>TypeVar('T_destination', bound=Dict[str, Any])</code> |
| <code>call_super_init</code> | |
| <code>dump_patches</code> | |
| <code>lazy_buffer_names</code> | |
| <code>lazy_parameter_names</code> | |
| <code>lazy_parameters_determined</code> | Returns if all lazy parameters are determined. |

Parameters

- **`in_channels`** (*int*) –
- **`args`** (*Any*) –
- **`kwargs`** (*Any*) –

pytorch_pfn_extras.nn.LazyConv2d

class `pytorch_pfn_extras.nn.LazyConv2d`(*in_channels*, **args*, ***kwargs*)

Bases: `_LazyConvNd`, `Conv2d`

Conv2d module with lazy weight initialization.

When `in_channels` is `None`, it is determined at the first time of the forward step.

Initializes internal Module state, shared by both `nn.Module` and `ScriptModule`.

Methods

| | |
|---|--|
| <code>__init__(in_channels, *args, **kwargs)</code> | Initializes internal Module state, shared by both <code>nn.Module</code> and <code>ScriptModule</code> . |
| <code>add_module(name, module)</code> | Adds a child module to the current module. |
| <code>apply(fn)</code> | Applies <code>fn</code> recursively to every submodule (as returned by <code>.children()</code>) as well as self. |
| <code>bfloat16()</code> | Casts all floating point parameters and buffers to <code>bfloat16</code> datatype. |
| <code>buffers([recurse])</code> | Returns an iterator over module buffers. |
| <code>children()</code> | Returns an iterator over immediate children modules. |
| <code>compile(*args, **kwargs)</code> | Compile this Module's forward using <code>torch.compile()</code> . |
| <code>cpu()</code> | Moves all model parameters and buffers to the CPU. |
| <code>cuda([device])</code> | Moves all model parameters and buffers to the GPU. |
| <code>double()</code> | Casts all floating point parameters and buffers to <code>double</code> datatype. |
| <code>eval()</code> | Sets the module in evaluation mode. |
| <code>extra_repr()</code> | Set the extra representation of the module |

continues on next page

Table 7 – continued from previous page

| | |
|---|---|
| <code>float()</code> | Casts all floating point parameters and buffers to <code>float</code> datatype. |
| <code>forward(input)</code> | Defines the computation performed at every call. |
| <code>get_buffer(target)</code> | Returns the buffer given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_extra_state()</code> | Returns any extra state to include in the module's <code>state_dict</code> . |
| <code>get_parameter(target)</code> | Returns the parameter given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_submodule(target)</code> | Returns the submodule given by <code>target</code> if it exists, otherwise throws an error. |
| <code>half()</code> | Casts all floating point parameters and buffers to <code>half</code> datatype. |
| <code>ipu([device])</code> | Moves all model parameters and buffers to the IPU. |
| <code>load_state_dict(state_dict[, strict, assign])</code> | Copies parameters and buffers from <code>state_dict</code> into this module and its descendants. |
| <code>modules()</code> | Returns an iterator over all modules in the network. |
| <code>named_buffers([prefix, recurse, ...])</code> | Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself. |
| <code>named_children()</code> | Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself. |
| <code>named_modules([memo, prefix, remove_duplicate])</code> | Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself. |
| <code>named_parameters([prefix, recurse, ...])</code> | Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself. |
| <code>parameters([recurse])</code> | Returns an iterator over module parameters. |
| <code>register_backward_hook(hook)</code> | Registers a backward hook on the module. |
| <code>register_buffer(name, tensor[, persistent])</code> | Adds a buffer to the module. |
| <code>register_forward_hook(hook, *[, prepend, ...])</code> | Registers a forward hook on the module. |
| <code>register_forward_pre_hook(hook, *[, ...])</code> | Registers a forward pre-hook on the module. |
| <code>register_full_backward_hook(hook[, prepend])</code> | Registers a backward hook on the module. |
| <code>register_full_backward_pre_hook(hook[, prepend])</code> | Registers a backward pre-hook on the module. |
| <code>register_load_state_dict_post_hook(hook)</code> | Registers a post hook to be run after module's <code>load_state_dict</code> is called. |
| <code>register_module(name, module)</code> | Alias for <code>add_module()</code> . |
| <code>register_parameter(name, param)</code> | Adds a parameter to the module. |
| <code>register_state_dict_pre_hook(hook)</code> | These hooks will be called with arguments: <code>self</code> , <code>prefix</code> , and <code>keep_vars</code> before calling <code>state_dict</code> on <code>self</code> . |
| <code>requires_grad_([requires_grad])</code> | Change if autograd should record operations on parameters in this module. |
| <code>reset_parameters()</code> | |
| <code>set_extra_state(state)</code> | This function is called from <code>load_state_dict()</code> to handle any extra state found within the <code>state_dict</code> . |
| <code>share_memory()</code> | See <code>torch.Tensor.share_memory_()</code> |

continues on next page

Table 7 – continued from previous page

| | |
|---|---|
| <code>state_dict(*args, **kwargs)</code> | Returns a dictionary containing a whole state of the module. |
| <code>to(*args, **kwargs)</code> | Moves and/or casts the parameters and buffers. |
| <code>to_empty(*, device[, recurse])</code> | Moves the parameters and buffers to the specified device without copying storage. |
| <code>train([mode])</code> | Sets the module in training mode. |
| <code>type(dst_type)</code> | Casts all parameters and buffers to <code>dst_type</code> . |
| <code>xpu([device])</code> | Moves all model parameters and buffers to the XPU. |
| <code>zero_grad([set_to_none])</code> | Resets gradients of all model parameters. |

Attributes

| | |
|--|--|
| <code>T_destination</code> | alias of <code>TypeVar('T_destination', bound=Dict[str, Any])</code> |
| <code>call_super_init</code> | |
| <code>dump_patches</code> | |
| <code>lazy_buffer_names</code> | |
| <code>lazy_parameter_names</code> | |
| <code>lazy_parmeters_determined</code> | Returns if all lazy parameters are determined. |

Parameters

- **`in_channels`** (*int*) –
- **`args`** (*Any*) –
- **`kwargs`** (*Any*) –

pytorch_pfn_extras.nn.LazyConv3d

class `pytorch_pfn_extras.nn.LazyConv3d`(*in_channels*, **args*, ***kwargs*)

Bases: `_LazyConvNd`, `Conv3d`

Conv3d module with lazy weight initialization.

When `in_channels` is `None`, it is determined at the first time of the forward step.

Initializes internal Module state, shared by both `nn.Module` and `ScriptModule`.

Methods

| | |
|--|---|
| <code>__init__(in_channels, *args, **kwargs)</code> | Initializes internal Module state, shared by both <code>nn.Module</code> and <code>ScriptModule</code> . |
| <code>add_module(name, module)</code> | Adds a child module to the current module. |
| <code>apply(fn)</code> | Applies <code>fn</code> recursively to every submodule (as returned by <code>.children()</code>) as well as self. |
| <code>bfloat16()</code> | Casts all floating point parameters and buffers to <code>bfloat16</code> datatype. |
| <code>buffers([recurse])</code> | Returns an iterator over module buffers. |
| <code>children()</code> | Returns an iterator over immediate children modules. |
| <code>compile(*args, **kwargs)</code> | Compile this Module's forward using <code>torch.compile()</code> . |
| <code>cpu()</code> | Moves all model parameters and buffers to the CPU. |
| <code>cuda([device])</code> | Moves all model parameters and buffers to the GPU. |
| <code>double()</code> | Casts all floating point parameters and buffers to <code>double</code> datatype. |
| <code>eval()</code> | Sets the module in evaluation mode. |
| <code>extra_repr()</code> | Set the extra representation of the module |
| <code>float()</code> | Casts all floating point parameters and buffers to <code>float</code> datatype. |
| <code>forward(input)</code> | Defines the computation performed at every call. |
| <code>get_buffer(target)</code> | Returns the buffer given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_extra_state()</code> | Returns any extra state to include in the module's <code>state_dict</code> . |
| <code>get_parameter(target)</code> | Returns the parameter given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_submodule(target)</code> | Returns the submodule given by <code>target</code> if it exists, otherwise throws an error. |
| <code>half()</code> | Casts all floating point parameters and buffers to <code>half</code> datatype. |
| <code>ipu([device])</code> | Moves all model parameters and buffers to the IPU. |
| <code>load_state_dict(state_dict[, strict, assign])</code> | Copies parameters and buffers from <code>state_dict</code> into this module and its descendants. |
| <code>modules()</code> | Returns an iterator over all modules in the network. |
| <code>named_buffers([prefix, recurse, ...])</code> | Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself. |
| <code>named_children()</code> | Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself. |
| <code>named_modules([memo, prefix, remove_duplicate])</code> | Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself. |
| <code>named_parameters([prefix, recurse, ...])</code> | Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself. |
| <code>parameters([recurse])</code> | Returns an iterator over module parameters. |
| <code>register_backward_hook(hook)</code> | Registers a backward hook on the module. |
| <code>register_buffer(name, tensor[, persistent])</code> | Adds a buffer to the module. |
| <code>register_forward_hook(hook, *[, prepend, ...])</code> | Registers a forward hook on the module. |

continues on next page

Table 8 – continued from previous page

| | |
|---|---|
| <code>register_forward_pre_hook(hook, *[, ...])</code> | Registers a forward pre-hook on the module. |
| <code>register_full_backward_hook(hook[, prepend])</code> | Registers a backward hook on the module. |
| <code>register_full_backward_pre_hook(hook[, prepend])</code> | Registers a backward pre-hook on the module. |
| <code>register_load_state_dict_post_hook(hook)</code> | Registers a post hook to be run after module's <code>load_state_dict</code> is called. |
| <code>register_module(name, module)</code> | Alias for <code>add_module()</code> . |
| <code>register_parameter(name, param)</code> | Adds a parameter to the module. |
| <code>register_state_dict_pre_hook(hook)</code> | These hooks will be called with arguments: <code>self</code> , <code>prefix</code> , and <code>keep_vars</code> before calling <code>state_dict</code> on <code>self</code> . |
| <code>requires_grad_([requires_grad])</code> | Change if autograd should record operations on parameters in this module. |
| <code>reset_parameters()</code> | |
| <code>set_extra_state(state)</code> | This function is called from <code>load_state_dict()</code> to handle any extra state found within the <code>state_dict</code> . |
| <code>share_memory()</code> | See <code>torch.Tensor.share_memory_()</code> |
| <code>state_dict(*args, **kwargs)</code> | Returns a dictionary containing a whole state of the module. |
| <code>to(*args, **kwargs)</code> | Moves and/or casts the parameters and buffers. |
| <code>to_empty(*, device[, recurse])</code> | Moves the parameters and buffers to the specified device without copying storage. |
| <code>train([mode])</code> | Sets the module in training mode. |
| <code>type(dst_type)</code> | Casts all parameters and buffers to <code>dst_type</code> . |
| <code>xpu([device])</code> | Moves all model parameters and buffers to the XPU. |
| <code>zero_grad([set_to_none])</code> | Resets gradients of all model parameters. |

Attributes

| | |
|--|--|
| <code>T_destination</code> | alias of <code>TypeVar('T_destination', bound=Dict[str, Any])</code> |
| <code>call_super_init</code> | |
| <code>dump_patches</code> | |
| <code>lazy_buffer_names</code> | |
| <code>lazy_parameter_names</code> | |
| <code>lazy_parmeters_determined</code> | Returns if all lazy parameters are determined. |

Parameters

- `in_channels` (*int*) –
- `args` (*Any*) –
- `kwargs` (*Any*) –

pytorch_pfn_extras.nn.LazyLinear

class pytorch_pfn_extras.nn.LazyLinear(*in_features*, *args, **kwargs)

Bases: [LazyInitializationMixin](#), [Linear](#)

Linear module with lazy weight initialization.

When *in_features* is None, it is determined at the first time of the forward step.

Initializes internal Module state, shared by both nn.Module and ScriptModule.

Methods

| | |
|--|---|
| __init__ (<i>in_features</i> , *args, **kwargs) | Initializes internal Module state, shared by both nn.Module and ScriptModule. |
| add_module (name, module) | Adds a child module to the current module. |
| apply (fn) | Applies <i>fn</i> recursively to every submodule (as returned by .children()) as well as self. |
| bfloat16 () | Casts all floating point parameters and buffers to bfloat16 datatype. |
| buffers ([recurse]) | Returns an iterator over module buffers. |
| children () | Returns an iterator over immediate children modules. |
| compile (*args, **kwargs) | Compile this Module's forward using torch.compile() . |
| cpu () | Moves all model parameters and buffers to the CPU. |
| cuda ([device]) | Moves all model parameters and buffers to the GPU. |
| double () | Casts all floating point parameters and buffers to double datatype. |
| eval () | Sets the module in evaluation mode. |
| extra_repr () | Set the extra representation of the module |
| float () | Casts all floating point parameters and buffers to float datatype. |
| forward (input) | Defines the computation performed at every call. |
| get_buffer (target) | Returns the buffer given by <i>target</i> if it exists, otherwise throws an error. |
| get_extra_state () | Returns any extra state to include in the module's <i>state_dict</i> . |
| get_parameter (target) | Returns the parameter given by <i>target</i> if it exists, otherwise throws an error. |
| get_submodule (target) | Returns the submodule given by <i>target</i> if it exists, otherwise throws an error. |
| half () | Casts all floating point parameters and buffers to half datatype. |
| ipu ([device]) | Moves all model parameters and buffers to the IPU. |
| load_state_dict (state_dict[, strict, assign]) | Copies parameters and buffers from <i>state_dict</i> into this module and its descendants. |
| modules () | Returns an iterator over all modules in the network. |
| named_buffers ([prefix, recurse, ...]) | Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself. |
| named_children () | Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself. |

continues on next page

Table 9 – continued from previous page

| | |
|---|---|
| <code>named_modules([memo, prefix, remove_duplicate])</code> | Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself. |
| <code>named_parameters([prefix, recurse, ...])</code> | Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself. |
| <code>parameters([recurse])</code> | Returns an iterator over module parameters. |
| <code>register_backward_hook(hook)</code> | Registers a backward hook on the module. |
| <code>register_buffer(name, tensor[, persistent])</code> | Adds a buffer to the module. |
| <code>register_forward_hook(hook, *[, prepend, ...])</code> | Registers a forward hook on the module. |
| <code>register_forward_pre_hook(hook, *[, ...])</code> | Registers a forward pre-hook on the module. |
| <code>register_full_backward_hook(hook[, prepend])</code> | Registers a backward hook on the module. |
| <code>register_full_backward_pre_hook(hook[, prepend])</code> | Registers a backward pre-hook on the module. |
| <code>register_load_state_dict_post_hook(hook)</code> | Registers a post hook to be run after module's <code>load_state_dict</code> is called. |
| <code>register_module(name, module)</code> | Alias for <code>add_module()</code> . |
| <code>register_parameter(name, param)</code> | Adds a parameter to the module. |
| <code>register_state_dict_pre_hook(hook)</code> | These hooks will be called with arguments: <code>self</code> , <code>prefix</code> , and <code>keep_vars</code> before calling <code>state_dict</code> on <code>self</code> . |
| <code>requires_grad_([requires_grad])</code> | Change if autograd should record operations on parameters in this module. |
| <code><i>reset_parameters()</i></code> | |
| <code>set_extra_state(state)</code> | This function is called from <code>load_state_dict()</code> to handle any extra state found within the <code>state_dict</code> . |
| <code>share_memory()</code> | See <code>torch.Tensor.share_memory_()</code> |
| <code>state_dict(*args, **kwargs)</code> | Returns a dictionary containing a whole state of the module. |
| <code>to(*args, **kwargs)</code> | Moves and/or casts the parameters and buffers. |
| <code>to_empty(*, device[, recurse])</code> | Moves the parameters and buffers to the specified device without copying storage. |
| <code>train([mode])</code> | Sets the module in training mode. |
| <code>type(dst_type)</code> | Casts all parameters and buffers to <code>dst_type</code> . |
| <code>xpu([device])</code> | Moves all model parameters and buffers to the XPU. |
| <code>zero_grad([set_to_none])</code> | Resets gradients of all model parameters. |

Attributes

| | |
|---|--|
| <code>T_destination</code> | alias of <code>TypeVar('T_destination', bound=Dict[str, Any])</code> |
| <code>call_super_init</code> | |
| <code>dump_patches</code> | |
| <code>lazy_buffer_names</code> | |
| <code>lazy_parameter_names</code> | |
| <code>lazy_parameters_determined</code> | Returns if all lazy parameters are determined. |

Parameters

- **`in_features`** (*int*) –
- **`args`** (*Any*) –
- **`kwargs`** (*Any*) –

`__init__`(*in_features*, **args*, ***kwargs*)

Initializes internal Module state, shared by both `nn.Module` and `ScriptModule`.

Parameters

- **`in_features`** (*Optional[int]*) –
- **`args`** (*Any*) –
- **`kwargs`** (*Any*) –

Return type

None

`forward`(*input*)

Defines the computation performed at every call.

Should be overridden by all subclasses.

Note: Although the recipe for forward pass needs to be defined within this function, one should call the `Module` instance afterwards instead of this since the former takes care of running the registered hooks while the latter silently ignores them.

Parameters

`input` (*Tensor*) –

Return type

Tensor

`lazy_parameter_names`: `Tuple[str, ...] = ('weight',)`

`reset_parameters`()

Return type

None

Modules

`pytorch_pfn_extras.nn.modules`

`pytorch_pfn_extras.nn.parallel`

pytorch_pfn_extras.nn.modules

Modules

`pytorch_pfn_extras.nn.modules.ensure_shape`

`pytorch_pfn_extras.nn.modules.
extended_sequential`

`pytorch_pfn_extras.nn.modules.lazy`

`pytorch_pfn_extras.nn.modules.
lazy_batchnorm`

`pytorch_pfn_extras.nn.modules.lazy_conv`

`pytorch_pfn_extras.nn.modules.lazy_linear`

pytorch_pfn_extras.nn.modules.ensure_shape

Functions

| | |
|---|--|
| <code>pytorch_pfn_extras.nn.modules. ensure_shape.ensure(tensor)</code> | Checks the shape and type of a tensor. |
|---|--|

pytorch_pfn_extras.nn.modules.ensure_shape.ensure

`pytorch_pfn_extras.nn.modules.ensure_shape.ensure(tensor, shape=None, dtype=None,
broadcastable=False, can_cast=False)`

Checks the shape and type of a tensor.

Parameters

- **shape** (*Optional[Tuple[Optional[int], ...]]*) – Tuple with the desired shape. If the input tensor shape is not compatible, *ValueError* will be raised. If *None* is set as a dimension value, that dimension will be ignored.
- **dtype** (*Optional[dtype]*) – Checks if the *dtype* of the input tensor matches the provided one.
- **broadcastable** (*bool*) – Check if the shapes are compatible using broadcasting rules.
- **can_cast** (*bool*) – Check if the input tensor can be casted to the provided type.

- **tensor** (*Tensor*) –

Return type

None

Classes

| | |
|--|--|
| <code>pytorch_pfn_extras.nn.modules. ensure_shape.Ensure(*)</code> | Module to check the shape of a tensor. |
|--|--|

pytorch_pfn_extras.nn.modules.ensure_shape.Ensure

class `pytorch_pfn_extras.nn.modules.ensure_shape.Ensure`(**, shape=None, dtype=None, broadcastable=False, can_cast=False*)

Bases: Module

Module to check the shape of a tensor.

Parameters

- **shape** (*Optional[Tuple[Optional[int], ...]]*) – Tuple with the desired shape. If the input tensor shape is not compatible, *ValueError* will be raised. If *None* is set as a dimension value, that dimension will be ignored.
- **dtype** (*Optional[dtype]*) – Checks if the *dtype* of the input tensor matches the provided one.
- **broadcastable** (*bool*) – Check if the shapes are compatible using broadcasting rules.
- **can_cast** (*bool*) – Check if the input tensor can be casted to the provided type.

Initializes internal Module state, shared by both nn.Module and ScriptModule.

Methods

| | |
|--|--|
| <code>__init__</code> (<i>*, shape, dtype, broadcastable, ...</i>) | Initializes internal Module state, shared by both nn.Module and ScriptModule. |
| <code>add_module</code> (name, module) | Adds a child module to the current module. |
| <code>apply</code> (fn) | Applies <i>fn</i> recursively to every submodule (as returned by <code>.children()</code>) as well as self. |
| <code>bfloat16</code> () | Casts all floating point parameters and buffers to bfloat16 datatype. |
| <code>buffers</code> ([recurse]) | Returns an iterator over module buffers. |
| <code>children</code> () | Returns an iterator over immediate children modules. |
| <code>compile</code> (<i>*args, **kwargs</i>) | Compile this Module's forward using <code>torch.compile()</code> . |
| <code>cpu</code> () | Moves all model parameters and buffers to the CPU. |
| <code>cuda</code> ([device]) | Moves all model parameters and buffers to the GPU. |
| <code>double</code> () | Casts all floating point parameters and buffers to double datatype. |
| <code>eval</code> () | Sets the module in evaluation mode. |
| <code>extra_repr</code> () | Set the extra representation of the module |

continues on next page

Table 10 – continued from previous page

| | |
|---|---|
| <code>float()</code> | Casts all floating point parameters and buffers to <code>float</code> datatype. |
| <code>forward(input)</code> | Defines the computation performed at every call. |
| <code>get_buffer(target)</code> | Returns the buffer given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_extra_state()</code> | Returns any extra state to include in the module's <code>state_dict</code> . |
| <code>get_parameter(target)</code> | Returns the parameter given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_submodule(target)</code> | Returns the submodule given by <code>target</code> if it exists, otherwise throws an error. |
| <code>half()</code> | Casts all floating point parameters and buffers to <code>half</code> datatype. |
| <code>ipu([device])</code> | Moves all model parameters and buffers to the IPU. |
| <code>load_state_dict(state_dict[, strict, assign])</code> | Copies parameters and buffers from <code>state_dict</code> into this module and its descendants. |
| <code>modules()</code> | Returns an iterator over all modules in the network. |
| <code>named_buffers([prefix, recurse, ...])</code> | Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself. |
| <code>named_children()</code> | Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself. |
| <code>named_modules([memo, prefix, remove_duplicate])</code> | Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself. |
| <code>named_parameters([prefix, recurse, ...])</code> | Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself. |
| <code>parameters([recurse])</code> | Returns an iterator over module parameters. |
| <code>register_backward_hook(hook)</code> | Registers a backward hook on the module. |
| <code>register_buffer(name, tensor[, persistent])</code> | Adds a buffer to the module. |
| <code>register_forward_hook(hook, *[, prepend, ...])</code> | Registers a forward hook on the module. |
| <code>register_forward_pre_hook(hook, *[, ...])</code> | Registers a forward pre-hook on the module. |
| <code>register_full_backward_hook(hook[, prepend])</code> | Registers a backward hook on the module. |
| <code>register_full_backward_pre_hook(hook[, prepend])</code> | Registers a backward pre-hook on the module. |
| <code>register_load_state_dict_post_hook(hook)</code> | Registers a post hook to be run after module's <code>load_state_dict</code> is called. |
| <code>register_module(name, module)</code> | Alias for <code>add_module()</code> . |
| <code>register_parameter(name, param)</code> | Adds a parameter to the module. |
| <code>register_state_dict_pre_hook(hook)</code> | These hooks will be called with arguments: <code>self</code> , <code>prefix</code> , and <code>keep_vars</code> before calling <code>state_dict</code> on <code>self</code> . |
| <code>requires_grad_([requires_grad])</code> | Change if autograd should record operations on parameters in this module. |
| <code>set_extra_state(state)</code> | This function is called from <code>load_state_dict()</code> to handle any extra state found within the <code>state_dict</code> . |
| <code>share_memory()</code> | See <code>torch.Tensor.share_memory_()</code> |
| <code>state_dict(*args[, destination, prefix, ...])</code> | Returns a dictionary containing references to the whole state of the module. |

continues on next page

Table 10 – continued from previous page

| | |
|---|---|
| <code>to(*args, **kwargs)</code> | Moves and/or casts the parameters and buffers. |
| <code>to_empty(*, device[, recurse])</code> | Moves the parameters and buffers to the specified device without copying storage. |
| <code>train([mode])</code> | Sets the module in training mode. |
| <code>type(dst_type)</code> | Casts all parameters and buffers to <code>dst_type</code> . |
| <code>xpu([device])</code> | Moves all model parameters and buffers to the XPU. |
| <code>zero_grad([set_to_none])</code> | Resets gradients of all model parameters. |

Attributes

| | |
|------------------------------|--|
| <code>T_destination</code> | alias of <code>TypeVar('T_destination', bound=Dict[str, Any])</code> |
| <code>call_super_init</code> | |
| <code>dump_patches</code> | |

`__init__`(**, shape=None, dtype=None, broadcastable=False, can_cast=False*)

Initializes internal Module state, shared by both `nn.Module` and `ScriptModule`.

Parameters

- **`shape`** (*Optional[Tuple[Optional[int], ...]]*) –
- **`dtype`** (*Optional[dtype]*) –
- **`broadcastable`** (*bool*) –
- **`can_cast`** (*bool*) –

`forward`(*input*)

Defines the computation performed at every call.

Should be overridden by all subclasses.

Note: Although the recipe for forward pass needs to be defined within this function, one should call the `Module` instance afterwards instead of this since the former takes care of running the registered hooks while the latter silently ignores them.

Parameters

`input` (*Tensor*) –

Return type

Tensor

`training`: `bool`

pytorch_pfn_extras.nn.modules.extended_sequential**Classes**

| | |
|---|--|
| <code>pytorch_pfn_extras.nn.modules.extended_sequential.ExtendedSequential()</code> | Sequential module with extended features from chainer. |
| <code>pytorch_pfn_extras.nn.modules.extended_sequential.TypeVar(...)</code> | Type variable. |

pytorch_pfn_extras.nn.modules.extended_sequential.ExtendedSequential

class pytorch_pfn_extras.nn.modules.extended_sequential.**ExtendedSequential**(*args: Module)

class pytorch_pfn_extras.nn.modules.extended_sequential.**ExtendedSequential**(arg: `OrderedDict[str, Module]`)

Bases: `Sequential`

Sequential module with extended features from chainer.

Initializes internal Module state, shared by both nn.Module and ScriptModule.

Methods

| | |
|---------------------------------------|--|
| <code>__init__(*args)</code> | Initializes internal Module state, shared by both nn.Module and ScriptModule. |
| <code>add_module(name, module)</code> | Adds a child module to the current module. |
| <code>append(module)</code> | Appends a given module to the end. |
| <code>apply(fn)</code> | Applies <code>fn</code> recursively to every submodule (as returned by <code>.children()</code>) as well as self. |
| <code>bfloat16()</code> | Casts all floating point parameters and buffers to <code>bfloat16</code> datatype. |
| <code>buffers([recurse])</code> | Returns an iterator over module buffers. |
| <code>children()</code> | Returns an iterator over immediate children modules. |
| <code>compile(*args, **kwargs)</code> | Compile this Module's forward using <code>torch.compile()</code> . |
| <code>cpu()</code> | Moves all model parameters and buffers to the CPU. |
| <code>cuda([device])</code> | Moves all model parameters and buffers to the GPU. |
| <code>double()</code> | Casts all floating point parameters and buffers to <code>double</code> datatype. |
| <code>eval()</code> | Sets the module in evaluation mode. |
| <code>extend(sequential)</code> | |
| <code>extra_repr()</code> | Set the extra representation of the module |
| <code>float()</code> | Casts all floating point parameters and buffers to <code>float</code> datatype. |
| <code>forward(input)</code> | Defines the computation performed at every call. |
| <code>get_buffer(target)</code> | Returns the buffer given by <code>target</code> if it exists, otherwise throws an error. |

continues on next page

Table 11 – continued from previous page

| | |
|---|---|
| <code>get_extra_state()</code> | Returns any extra state to include in the module's <code>state_dict</code> . |
| <code>get_parameter(target)</code> | Returns the parameter given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_submodule(target)</code> | Returns the submodule given by <code>target</code> if it exists, otherwise throws an error. |
| <code>half()</code> | Casts all floating point parameters and buffers to <code>half</code> datatype. |
| <code>insert(index, module)</code> | |
| <code>ipu([device])</code> | Moves all model parameters and buffers to the IPU. |
| <code>load_state_dict(state_dict[, strict, assign])</code> | Copies parameters and buffers from <code>state_dict</code> into this module and its descendants. |
| <code>modules()</code> | Returns an iterator over all modules in the network. |
| <code>named_buffers([prefix, recurse, ...])</code> | Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself. |
| <code>named_children()</code> | Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself. |
| <code>named_modules([memo, prefix, remove_duplicate])</code> | Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself. |
| <code>named_parameters([prefix, recurse, ...])</code> | Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself. |
| <code>parameters([recurse])</code> | Returns an iterator over module parameters. |
| <code>pop(key)</code> | |
| <code>register_backward_hook(hook)</code> | Registers a backward hook on the module. |
| <code>register_buffer(name, tensor[, persistent])</code> | Adds a buffer to the module. |
| <code>register_forward_hook(hook, *[, prepend, ...])</code> | Registers a forward hook on the module. |
| <code>register_forward_pre_hook(hook, *[, ...])</code> | Registers a forward pre-hook on the module. |
| <code>register_full_backward_hook(hook[, prepend])</code> | Registers a backward hook on the module. |
| <code>register_full_backward_pre_hook(hook[, prepend])</code> | Registers a backward pre-hook on the module. |
| <code>register_load_state_dict_post_hook(hook)</code> | Registers a post hook to be run after module's <code>load_state_dict</code> is called. |
| <code>register_module(name, module)</code> | Alias for <code>add_module()</code> . |
| <code>register_parameter(name, param)</code> | Adds a parameter to the module. |
| <code>register_state_dict_pre_hook(hook)</code> | These hooks will be called with arguments: <code>self</code> , <code>prefix</code> , and <code>keep_vars</code> before calling <code>state_dict</code> on <code>self</code> . |
| <code>repeat(n_repeat[, mode])</code> | Repeats this Sequential multiple times. |
| <code>requires_grad_([requires_grad])</code> | Change if autograd should record operations on parameters in this module. |
| <code>set_extra_state(state)</code> | This function is called from <code>load_state_dict()</code> to handle any extra state found within the <code>state_dict</code> . |
| <code>share_memory()</code> | See <code>torch.Tensor.share_memory_()</code> |
| <code>state_dict(*args[, destination, prefix, ...])</code> | Returns a dictionary containing references to the whole state of the module. |

continues on next page

Table 11 – continued from previous page

| | |
|---|---|
| <code>to(*args, **kwargs)</code> | Moves and/or casts the parameters and buffers. |
| <code>to_empty(*, device[, recurse])</code> | Moves the parameters and buffers to the specified device without copying storage. |
| <code>train([model])</code> | Sets the module in training mode. |
| <code>type(dst_type)</code> | Casts all parameters and buffers to <code>dst_type</code> . |
| <code>xpu([device])</code> | Moves all model parameters and buffers to the XPU. |
| <code>zero_grad([set_to_none])</code> | Resets gradients of all model parameters. |

Attributes

| | |
|------------------------------|--|
| <code>T_destination</code> | alias of <code>TypeVar('T_destination', bound=Dict[str, Any])</code> |
| <code>call_super_init</code> | |
| <code>dump_patches</code> | |

repeat(*n_repeat*, *mode*='init')

Repeats this Sequential multiple times.

This method returns a `Sequential` object which has original *Sequential* multiple times repeatedly. The *mode* argument means how to copy this sequential to repeat.

The functions is supposed to behave the same way as *repeat* in *chainer*.

When the mode is set to *init*, the default value, modules will be copied and reinitialized by calling `reset_parameters` (or `_reset_parameters`) method.

To repeat user-defined modules, which have parameters or buffers, with `mode='init'` in this `Sequential`, you need to implement the `reset_parameters` or `_reset_parameters` method to the module to reinitialize parameters and (if necessary) buffers; otherwise the initialization cannot be performed and a warning message will be shown.

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.

Return type

`ExtendedSequential`

training: `bool`

pytorch_pfn_extras.nn.modules.extended_sequential.TypeVar

```
class pytorch_pfn_extras.nn.modules.extended_sequential.TypeVar(name, *constraints,  
                                                                bound=None, covariant=False,  
                                                                contravariant=False)
```

Bases: `_Final`, `_Immutable`

Type variable.

Usage:

```
T = TypeVar('T') # Can be anything  
A = TypeVar('A', str, bytes) # Must be str or bytes
```

Type variables exist primarily for the benefit of static type checkers. They serve as the parameters for generic types as well as for generic function definitions. See class `Generic` for more information on generic types. Generic functions work as follows:

```
def repeat(x: T, n: int) -> List[T]:  
    """Return a list containing n references to x.""" return [x]*n  
  
def longest(x: A, y: A) -> A:  
    """Return the longest of two strings.""" return x if len(x) >= len(y) else y
```

The latter example's signature is essentially the overloading of `(str, str) -> str` and `(bytes, bytes) -> bytes`. Also note that if the arguments are instances of some subclass of `str`, the return type is still plain `str`.

At runtime, `isinstance(x, T)` and `issubclass(C, T)` will raise `TypeError`.

Type variables defined with `covariant=True` or `contravariant=True` can be used to declare covariant or contravariant generic types. See PEP 484 for more details. By default generic types are invariant in all type variables.

Type variables can be introspected. e.g.:

```
T.__name__ == 'T' T.__constraints__ == () T.__covariant__ == False T.__contravariant__ = False  
A.__constraints__ == (str, bytes)
```

Note that only type variables defined in global scope can be pickled.

Methods

```
__init__(name, *constraints[, bound, ...])
```

```
__init__(name, *constraints, bound=None, covariant=False, contravariant=False)
```

pytorch_pfn_extras.nn.modules.lazy**Classes**

| | |
|--|--|
| <code>pytorch_pfn_extras.nn.modules.lazy.LazyInitializationMixin(...)</code> | A mixin for modules that lazily initialize buffers and parameters. |
| <code>pytorch_pfn_extras.nn.modules.lazy.UninitializedParameter(...)</code> | |

pytorch_pfn_extras.nn.modules.lazy.LazyInitializationMixin

class pytorch_pfn_extras.nn.modules.lazy.LazyInitializationMixin(*args, **kwargs)

Bases: object

A mixin for modules that lazily initialize buffers and parameters.

Unlike regular modules, subclasses of this module can initialize buffers and parameters outside of the constructor (`__init__`). This allows you to, for example, initialize parameters in `forward` method to determine the shape of the weight based on the initial input.

Be sure to run “dummy” forward once to initialize all parameters that should be trained, before passing `module.parameters()` to an optimizer; otherwise weights initialized after `module.parameters()` (e.g., in `forward` function) will never be trained.

Note that lazy modules cannot validate if the shape is correct during deserialization. Also note that the initial weights may become different from the original (non-lazy) module even if the random seed is manually configured, as the order of initialization is different from the original one; especially, `module.cuda()` may cause the initialization to run on a GPU.

The default value of lazy buffers and parameters are `torch.Tensor([])` and `UninitializedParameter()`, respectively.

Methods

`__init__`(*args, **kwargs)

`state_dict`(*args, **kwargs)

Returns a dictionary containing a whole state of the module.

Attributes

`lazy_buffer_names`

`lazy_parameter_names`

`lazy_parameters_determined`

Returns if all lazy parameters are determined.

Parameters

- **args** (*Any*) –
- **kwargs** (*Any*) –

`__init__`(*args, **kwargs)

Parameters

- **self** (*Any*) –
- **args** (*Any*) –
- **kwargs** (*Any*) –

Return type

None

lazy_buffer_names: `Tuple[str, ...] = ()`

lazy_parameter_names: `Tuple[str, ...] = ()`

property lazy_parameters_determined: `bool`

Returns if all lazy parameters are determined.

Subclasses can perform parameters initialization after all lazy parameters are determined. Note that this may be called during `__init__`.

state_dict(*args, **kwargs)

Returns a dictionary containing a whole state of the module.

This function overrides the default behavior to exclude uninitialized parameter from serialization. This is needed because we need to discriminate lazy parameters (`UninitializedParameter()`) and initialized empty parameters (`torch.nn.Parameter(torch.Tensor())`) during deserialization.

See comments of `_lazy_load_hook` for details.

Parameters

- **self** (*Any*) –
- **args** (*Any*) –
- **kwargs** (*Any*) –

Return type

Dict[str, Any]

pytorch_pfn_extras.nn.modules.lazy.UninitializedParameter

class `pytorch_pfn_extras.nn.modules.lazy.UninitializedParameter`(*data=None*,
requires_grad=True)

Bases: `Parameter`

Methods

| | |
|--|---|
| <code>__init__()</code> | |
| <code>abs()</code> | See <code>torch.abs()</code> |
| <code>abs_()</code> | In-place version of <code>abs()</code> |
| <code>absolute()</code> | Alias for <code>abs()</code> |
| <code>absolute_()</code> | In-place version of <code>absolute()</code> Alias for <code>abs_()</code> |
| <code>acos()</code> | See <code>torch.acos()</code> |
| <code>acos_()</code> | In-place version of <code>acos()</code> |
| <code>acosh()</code> | See <code>torch.acosh()</code> |
| <code>acosh_()</code> | In-place version of <code>acosh()</code> |
| <code>add(other, *, alpha)</code> | Add a scalar or tensor to <code>self</code> tensor. |
| <code>add_(other, *, alpha)</code> | In-place version of <code>add()</code> |
| <code>addbmm(batch1, batch2, *, beta, alpha)</code> | See <code>torch.addbmm()</code> |
| <code>addbmm_(batch1, batch2, *, beta, alpha)</code> | In-place version of <code>addbmm()</code> |
| <code>addcddiv(tensor1, tensor2, *, value)</code> | See <code>torch.addcddiv()</code> |

continues on next page

Table 12 – continued from previous page

| | |
|---|--|
| <code>addcdiv_(tensor1, tensor2, *, value)</code> | In-place version of <code>addcdiv()</code> |
| <code>addcmul(tensor1, tensor2, *, value)</code> | See <code>torch.addcmul()</code> |
| <code>addcmul_(tensor1, tensor2, *, value)</code> | In-place version of <code>addcmul()</code> |
| <code>addmm(mat1, mat2, *, beta, alpha)</code> | See <code>torch.addmm()</code> |
| <code>addmm_(mat1, mat2, *, beta, alpha)</code> | In-place version of <code>addmm()</code> |
| <code>addmv(mat, vec, *, beta, alpha)</code> | See <code>torch.addmv()</code> |
| <code>addmv_(mat, vec, *, beta, alpha)</code> | In-place version of <code>addmv()</code> |
| <code>addr(vec1, vec2, *, beta, alpha)</code> | See <code>torch.addr()</code> |
| <code>addr_(vec1, vec2, *, beta, alpha)</code> | In-place version of <code>addr()</code> |
| <code>adjoint()</code> | Alias for <code>adjoint()</code> |
| <code>align_as(other)</code> | Permutates the dimensions of the <code>self</code> tensor to match the dimension order in the <code>other</code> tensor, adding size-one dims for any new names. |
| <code>align_to(*names)</code> | Permutates the dimensions of the <code>self</code> tensor to match the order specified in <code>names</code> , adding size-one dims for any new names. |
| <code>all([dim, keepdim])</code> | See <code>torch.all()</code> |
| <code>allclose(other[, rtol, atol, equal_nan])</code> | See <code>torch.allclose()</code> |
| <code>amax([dim, keepdim])</code> | See <code>torch.amax()</code> |
| <code>amin([dim, keepdim])</code> | See <code>torch.amin()</code> |
| <code>aminmax(*[, dim, keepdim])</code> | See <code>torch.aminmax()</code> |
| <code>angle()</code> | See <code>torch.angle()</code> |
| <code>any([dim, keepdim])</code> | See <code>torch.any()</code> |
| <code>apply_(callable)</code> | Applies the function <code>callable</code> to each element in the tensor, replacing each element with the value returned by <code>callable</code> . |
| <code>arccos()</code> | See <code>torch.arccos()</code> |
| <code>arccos_()</code> | In-place version of <code>arccos()</code> |
| <code>arccosh</code> | <code>acosh()</code> -> Tensor |
| <code>arccosh_</code> | <code>acosh_()</code> -> Tensor |
| <code>arcsin()</code> | See <code>torch.arcsin()</code> |
| <code>arcsin_()</code> | In-place version of <code>arcsin()</code> |
| <code>arcsinh()</code> | See <code>torch.arcsinh()</code> |
| <code>arcsinh_()</code> | In-place version of <code>arcsinh()</code> |
| <code>arctan()</code> | See <code>torch.arctan()</code> |
| <code>arctan2(other)</code> | See <code>torch.arctan2()</code> |
| <code>arctan2_</code> | <code>atan2_(other)</code> -> Tensor |
| <code>arctan_()</code> | In-place version of <code>arctan()</code> |
| <code>arctanh()</code> | See <code>torch.arctanh()</code> |
| <code>arctanh_(other)</code> | In-place version of <code>arctanh()</code> |
| <code>argmax([dim, keepdim])</code> | See <code>torch.argmax()</code> |
| <code>argmin([dim, keepdim])</code> | See <code>torch.argmin()</code> |
| <code>argsort([dim, descending])</code> | See <code>torch.argsort()</code> |
| <code>argwhere()</code> | See <code>torch.argwhere()</code> |
| <code>as_strided(size, stride[, storage_offset])</code> | See <code>torch.as_strided()</code> |
| <code>as_strided_(size, stride[, storage_offset])</code> | In-place version of <code>as_strided()</code> |
| <code>as_strided_scatter(src, size, stride[, ...])</code> | See <code>torch.as_strided_scatter()</code> |
| <code>as_subclass(cls)</code> | Makes a <code>cls</code> instance with the same data pointer as <code>self</code> . |
| <code>asin()</code> | See <code>torch.asin()</code> |

continues on next page

Table 12 – continued from previous page

| | |
|---|---|
| <code>asin_()</code> | In-place version of <code>asin()</code> |
| <code>asinh()</code> | See <code>torch.asinh()</code> |
| <code>asinh_()</code> | In-place version of <code>asinh()</code> |
| <code>atan()</code> | See <code>torch.atan()</code> |
| <code>atan2(other)</code> | See <code>torch.atan2()</code> |
| <code>atan2_(other)</code> | In-place version of <code>atan2()</code> |
| <code>atan_()</code> | In-place version of <code>atan()</code> |
| <code>atanh()</code> | See <code>torch.atanh()</code> |
| <code>atanh_(other)</code> | In-place version of <code>atanh()</code> |
| <code>backward([gradient, retain_graph, ...])</code> | Computes the gradient of current tensor wrt graph leaves. |
| <code>baddbmm(batch1, batch2, *[, beta, alpha])</code> | See <code>torch.baddbmm()</code> |
| <code>baddbmm_(batch1, batch2, *[, beta, alpha])</code> | In-place version of <code>baddbmm()</code> |
| <code>bernoulli(*[, generator])</code> | Returns a result tensor where each <code>result[i]</code> is independently sampled from <code>Bernoulli(self[i])</code> . |
| <code>bernoulli_(p, generator)</code> | Fills each location of <code>self</code> with an independent sample from <code>Bernoulli(p)</code> . |
| <code>bfloat16([memory_format])</code> | <code>self.bfloat16()</code> is equivalent to <code>self.to(torch.bfloat16)</code> . |
| <code>bincount([weights, minlength])</code> | See <code>torch.bincount()</code> |
| <code>bitwise_and()</code> | See <code>torch.bitwise_and()</code> |
| <code>bitwise_and_()</code> | In-place version of <code>bitwise_and()</code> |
| <code>bitwise_left_shift(other)</code> | See <code>torch.bitwise_left_shift()</code> |
| <code>bitwise_left_shift_(other)</code> | In-place version of <code>bitwise_left_shift()</code> |
| <code>bitwise_not()</code> | See <code>torch.bitwise_not()</code> |
| <code>bitwise_not_()</code> | In-place version of <code>bitwise_not()</code> |
| <code>bitwise_or()</code> | See <code>torch.bitwise_or()</code> |
| <code>bitwise_or_()</code> | In-place version of <code>bitwise_or()</code> |
| <code>bitwise_right_shift(other)</code> | See <code>torch.bitwise_right_shift()</code> |
| <code>bitwise_right_shift_(other)</code> | In-place version of <code>bitwise_right_shift()</code> |
| <code>bitwise_xor()</code> | See <code>torch.bitwise_xor()</code> |
| <code>bitwise_xor_()</code> | In-place version of <code>bitwise_xor()</code> |
| <code>bmm(batch2)</code> | See <code>torch.bmm()</code> |
| <code>bool([memory_format])</code> | <code>self.bool()</code> is equivalent to <code>self.to(torch.bool)</code> . |
| <code>broadcast_to(shape)</code> | See <code>torch.broadcast_to()</code> . |
| <code>byte([memory_format])</code> | <code>self.byte()</code> is equivalent to <code>self.to(torch.uint8)</code> . |
| <code>cauchy_([median, sigma, generator])</code> | Fills the tensor with numbers drawn from the Cauchy distribution: |
| <code>ccol_indices</code> | |
| <code>cdouble([memory_format])</code> | <code>self.cdouble()</code> is equivalent to <code>self.to(torch.complex128)</code> . |
| <code>ceil()</code> | See <code>torch.ceil()</code> |
| <code>ceil_()</code> | In-place version of <code>ceil()</code> |
| <code>cfloat([memory_format])</code> | <code>self.cfloat()</code> is equivalent to <code>self.to(torch.complex64)</code> . |
| <code>chalf([memory_format])</code> | <code>self.chalf()</code> is equivalent to <code>self.to(torch.complex32)</code> . |

continues on next page

Table 12 – continued from previous page

| | |
|--|--|
| <code>char([memory_format])</code> | <code>self.char()</code> is equivalent to <code>self.to(torch.int8)</code> . |
| <code>cholesky([upper])</code> | See <code>torch.cholesky()</code> |
| <code>cholesky_inverse([upper])</code> | See <code>torch.cholesky_inverse()</code> |
| <code>cholesky_solve(input2[, upper])</code> | See <code>torch.cholesky_solve()</code> |
| <code>chunk(chunks[, dim])</code> | See <code>torch.chunk()</code> |
| <code>clamp([min, max])</code> | See <code>torch.clamp()</code> |
| <code>clamp_([min, max])</code> | In-place version of <code>clamp()</code> |
| <code>clamp_max</code> | |
| <code>clamp_max_</code> | |
| <code>clamp_min</code> | |
| <code>clamp_min_</code> | |
| <code>clip([min, max])</code> | Alias for <code>clamp()</code> . |
| <code>clip_([min, max])</code> | Alias for <code>clamp_()</code> . |
| <code>clone(*[, memory_format])</code> | See <code>torch.clone()</code> |
| <code>coalesce()</code> | Returns a coalesced copy of <code>self</code> if <code>self</code> is an un-coalesced tensor. |
| <code>col_indices()</code> | Returns the tensor containing the column indices of the <code>self</code> tensor when <code>self</code> is a sparse CSR tensor of layout <code>sparse_csr</code> . |
| <code>conj()</code> | See <code>torch.conj()</code> |
| <code>conj_physical()</code> | See <code>torch.conj_physical()</code> |
| <code>conj_physical_()</code> | In-place version of <code>conj_physical()</code> |
| <code>contiguous([memory_format])</code> | Returns a contiguous in memory tensor containing the same data as <code>self</code> tensor. |
| <code>copy_(src[, non_blocking])</code> | Copies the elements from <code>src</code> into <code>self</code> tensor and returns <code>self</code> . |
| <code>copysign(other)</code> | See <code>torch.copysign()</code> |
| <code>copysign_(other)</code> | In-place version of <code>copysign()</code> |
| <code>corrcoef()</code> | See <code>torch.corrcoef()</code> |
| <code>cos()</code> | See <code>torch.cos()</code> |
| <code>cos_()</code> | In-place version of <code>cos()</code> |
| <code>cosh()</code> | See <code>torch.cosh()</code> |
| <code>cosh_()</code> | In-place version of <code>cosh()</code> |
| <code>count_nonzero([dim])</code> | See <code>torch.count_nonzero()</code> |
| <code>cov(*[, correction, fweights, aweights])</code> | See <code>torch.cov()</code> |
| <code>cpu([memory_format])</code> | Returns a copy of this object in CPU memory. |
| <code>cross(other[, dim])</code> | See <code>torch.cross()</code> |
| <code>crow_indices()</code> | Returns the tensor containing the compressed row indices of the <code>self</code> tensor when <code>self</code> is a sparse CSR tensor of layout <code>sparse_csr</code> . |
| <code>cuda([device, non_blocking, memory_format])</code> | Returns a copy of this object in CUDA memory. |
| <code>cummax(dim)</code> | See <code>torch.cummax()</code> |
| <code>cummin(dim)</code> | See <code>torch.cummin()</code> |
| <code>cumprod(dim[, dtype])</code> | See <code>torch.cumprod()</code> |
| <code>cumprod_(dim[, dtype])</code> | In-place version of <code>cumprod()</code> |
| <code>cumsum(dim[, dtype])</code> | See <code>torch.cumsum()</code> |

continues on next page

Table 12 – continued from previous page

| | |
|--|---|
| <code>cumsum_(dim[, dtype])</code> | In-place version of <code>cumsum()</code> |
| <code>data_ptr()</code> | Returns the address of the first element of <code>self</code> tensor. |
| <code>deg2rad()</code> | See <code>torch.deg2rad()</code> |
| <code>deg2rad_()</code> | In-place version of <code>deg2rad()</code> |
| <code>dense_dim()</code> | Return the number of dense dimensions in a sparse tensor <code>self</code> . |
| <code>dequantize()</code> | Given a quantized Tensor, dequantize it and return the dequantized float Tensor. |
| <code>det()</code> | See <code>torch.det()</code> |
| <code>detach</code> | Returns a new Tensor, detached from the current graph. |
| <code>detach_</code> | Detaches the Tensor from the graph that created it, making it a leaf. |
| <code>diag([diagonal])</code> | See <code>torch.diag()</code> |
| <code>diag_embed([offset, dim1, dim2])</code> | See <code>torch.diag_embed()</code> |
| <code>diagflat([offset])</code> | See <code>torch.diagflat()</code> |
| <code>diagonal([offset, dim1, dim2])</code> | See <code>torch.diagonal()</code> |
| <code>diagonal_scatter(src[, offset, dim1, dim2])</code> | See <code>torch.diagonal_scatter()</code> |
| <code>diff([n, dim, prepend, append])</code> | See <code>torch.diff()</code> |
| <code>digamma()</code> | See <code>torch.digamma()</code> |
| <code>digamma_()</code> | In-place version of <code>digamma()</code> |
| <code>dim()</code> | Returns the number of dimensions of <code>self</code> tensor. |
| <code>dim_order()</code> | Returns a tuple of int describing the dim order or physical layout of <code>self</code> . |
| <code>dist(other[, p])</code> | See <code>torch.dist()</code> |
| <code>div(value, *[, rounding_mode])</code> | See <code>torch.div()</code> |
| <code>div_(value, *[, rounding_mode])</code> | In-place version of <code>div()</code> |
| <code>divide(value, *[, rounding_mode])</code> | See <code>torch.divide()</code> |
| <code>divide_(value, *[, rounding_mode])</code> | In-place version of <code>divide()</code> |
| <code>dot(other)</code> | See <code>torch.dot()</code> |
| <code>double([memory_format])</code> | <code>self.double()</code> is equivalent to <code>self.to(torch.float64)</code> . |
| <code>dsplit(split_size_or_sections)</code> | See <code>torch.dsplit()</code> |
| <code>eig([eigenvectors])</code> | |
| <code>element_size()</code> | Returns the size in bytes of an individual element. |
| <code>eq(other)</code> | See <code>torch.eq()</code> |
| <code>eq_(other)</code> | In-place version of <code>eq()</code> |
| <code>equal(other)</code> | See <code>torch.equal()</code> |
| <code>erf()</code> | See <code>torch.erf()</code> |
| <code>erf_()</code> | In-place version of <code>erf()</code> |
| <code>erfc()</code> | See <code>torch.erfc()</code> |
| <code>erfc_()</code> | In-place version of <code>erfc()</code> |
| <code>erfinv()</code> | See <code>torch.erfinv()</code> |
| <code>erfinv_()</code> | In-place version of <code>erfinv()</code> |
| <code>exp()</code> | See <code>torch.exp()</code> |
| <code>exp2()</code> | See <code>torch.exp2()</code> |
| <code>exp2_()</code> | In-place version of <code>exp2()</code> |
| <code>exp_()</code> | In-place version of <code>exp()</code> |

continues on next page

Table 12 – continued from previous page

| | |
|---|---|
| <code>expand(*sizes)</code> | Returns a new view of the <code>self</code> tensor with singleton dimensions expanded to a larger size. |
| <code>expand_as(other)</code> | Expand this tensor to the same size as <code>other</code> . |
| <code>expm1()</code> | See <code>torch.expm1()</code> |
| <code>expm1_()</code> | In-place version of <code>expm1()</code> |
| <code>exponential_([lambd, generator])</code> | Fills <code>self</code> tensor with elements drawn from the exponential distribution: |
| <code>fill_(value)</code> | Fills <code>self</code> tensor with the specified value. |
| <code>fill_diagonal_(fill_value[, wrap])</code> | Fill the main diagonal of a tensor that has at least 2-dimensions. |
| <code>fix()</code> | See <code>torch.fix()</code> . |
| <code>fix_()</code> | In-place version of <code>fix()</code> |
| <code>flatten([start_dim, end_dim])</code> | See <code>torch.flatten()</code> |
| <code>flip(dims)</code> | See <code>torch.flip()</code> |
| <code>flipr()</code> | See <code>torch.flipr()</code> |
| <code>flipud()</code> | See <code>torch.flipud()</code> |
| <code>float([memory_format])</code> | <code>self.float()</code> is equivalent to <code>self.to(torch.float32)</code> . |
| <code>float_power(exponent)</code> | See <code>torch.float_power()</code> |
| <code>float_power_(exponent)</code> | In-place version of <code>float_power()</code> |
| <code>floor()</code> | See <code>torch.floor()</code> |
| <code>floor_()</code> | In-place version of <code>floor()</code> |
| <code>floor_divide(value)</code> | See <code>torch.floor_divide()</code> |
| <code>floor_divide_(value)</code> | In-place version of <code>floor_divide()</code> |
| <code>fmax(other)</code> | See <code>torch.fmax()</code> |
| <code>fmin(other)</code> | See <code>torch.fmin()</code> |
| <code>fmod(divisor)</code> | See <code>torch.fmod()</code> |
| <code>fmod_(divisor)</code> | In-place version of <code>fmod()</code> |
| <code>frac()</code> | See <code>torch.frac()</code> |
| <code>frac_()</code> | In-place version of <code>frac()</code> |
| <code>frexp(input)</code> | See <code>torch.frexp()</code> |
| <code>gather(dim, index)</code> | See <code>torch.gather()</code> |
| <code>gcd(other)</code> | See <code>torch.gcd()</code> |
| <code>gcd_(other)</code> | In-place version of <code>gcd()</code> |
| <code>ge(other)</code> | See <code>torch.ge()</code> . |
| <code>ge_(other)</code> | In-place version of <code>ge()</code> . |
| <code>geometric_(p, *[, generator])</code> | Fills <code>self</code> tensor with elements drawn from the geometric distribution: |
| <code>geqrf()</code> | See <code>torch.geqrf()</code> |
| <code>ger(vec2)</code> | See <code>torch.ger()</code> |
| <code>get_device()</code> | For CUDA tensors, this function returns the device ordinal of the GPU on which the tensor resides. |
| <code>greater(other)</code> | See <code>torch.greater()</code> . |
| <code>greater_(other)</code> | In-place version of <code>greater()</code> . |
| <code>greater_equal(other)</code> | See <code>torch.greater_equal()</code> . |
| <code>greater_equal_(other)</code> | In-place version of <code>greater_equal()</code> . |
| <code>gt(other)</code> | See <code>torch.gt()</code> . |
| <code>gt_(other)</code> | In-place version of <code>gt()</code> . |
| <code>half([memory_format])</code> | <code>self.half()</code> is equivalent to <code>self.to(torch.float16)</code> . |

continues on next page

Table 12 – continued from previous page

| | |
|--|---|
| <code>hardshrink([lambd])</code> | See <code>torch.nn.functional.hardshrink()</code> |
| <code>has_names</code> | Is True if any of this tensor's dimensions are named. |
| <code>heaviside(values)</code> | See <code>torch.heaviside()</code> |
| <code>heaviside_(values)</code> | In-place version of <code>heaviside()</code> |
| <code>histc([bins, min, max])</code> | See <code>torch.histc()</code> |
| <code>histogram(input, bins, *[, range, weight, ...])</code> | See <code>torch.histogram()</code> |
| <code>hsplit(split_size_or_sections)</code> | See <code>torch.hsplit()</code> |
| <code>hypot(other)</code> | See <code>torch.hypot()</code> |
| <code>hypot_(other)</code> | In-place version of <code>hypot()</code> |
| <code>i0()</code> | See <code>torch.i0()</code> |
| <code>i0_()</code> | In-place version of <code>i0()</code> |
| <code>igamma(other)</code> | See <code>torch.igamma()</code> |
| <code>igamma_(other)</code> | In-place version of <code>igamma()</code> |
| <code>igammac(other)</code> | See <code>torch.igammac()</code> |
| <code>igammac_(other)</code> | In-place version of <code>igammac()</code> |
| <code>index_add(dim, index, source, *[, alpha])</code> | Out-of-place version of <code>torch.Tensor.index_add_()</code> . |
| <code>index_add_(dim, index, source, *[, alpha])</code> | Accumulate the elements of <code>alpha times source</code> into the <code>self</code> tensor by adding to the indices in the order given in <code>index</code> . |
| <code>index_copy(dim, index, tensor2)</code> | Out-of-place version of <code>torch.Tensor.index_copy_()</code> . |
| <code>index_copy_(dim, index, tensor)</code> | Copies the elements of <code>tensor</code> into the <code>self</code> tensor by selecting the indices in the order given in <code>index</code> . |
| <code>index_fill(dim, index, value)</code> | Out-of-place version of <code>torch.Tensor.index_fill_()</code> . |
| <code>index_fill_(dim, index, value)</code> | Fills the elements of the <code>self</code> tensor with <code>value</code> by selecting the indices in the order given in <code>index</code> . |
| <code>index_put(indices, values[, accumulate])</code> | Out-place version of <code>index_put_()</code> . |
| <code>index_put_(indices, values[, accumulate])</code> | Puts values from the tensor values into the tensor <code>self</code> using the indices specified in <code>indices</code> (which is a tuple of Tensors). |
| <code>index_reduce</code> | |
| <code>index_reduce_(dim, index, source, reduce, *)</code> | Accumulate the elements of <code>source</code> into the <code>self</code> tensor by accumulating to the indices in the order given in <code>index</code> using the reduction given by the <code>reduce</code> argument. |
| <code>index_select(dim, index)</code> | See <code>torch.index_select()</code> |
| <code>indices()</code> | Return the indices tensor of a sparse COO tensor. |
| <code>inner(other)</code> | See <code>torch.inner()</code> . |
| <code>int([memory_format])</code> | <code>self.int()</code> is equivalent to <code>self.to(torch.int32)</code> . |
| <code>int_repr()</code> | Given a quantized Tensor, <code>self.int_repr()</code> returns a CPU Tensor with <code>uint8_t</code> as data type that stores the underlying <code>uint8_t</code> values of the given Tensor. |
| <code>inverse()</code> | See <code>torch.inverse()</code> |
| <code>ipu([device, non_blocking, memory_format])</code> | Returns a copy of this object in IPU memory. |

continues on next page

Table 12 – continued from previous page

| | |
|--|---|
| <code>is_coalesced()</code> | Returns True if <code>self</code> is a sparse COO tensor that is coalesced, False otherwise. |
| <code>is_complex()</code> | Returns True if the data type of <code>self</code> is a complex data type. |
| <code>is_conj()</code> | Returns True if the conjugate bit of <code>self</code> is set to true. |
| <code>is_contiguous([memory_format])</code> | Returns True if <code>self</code> tensor is contiguous in memory in the order specified by memory format. |
| <code>is_distributed</code> | |
| <code>is_floating_point()</code> | Returns True if the data type of <code>self</code> is a floating point data type. |
| <code>is_inference()</code> | See <code>torch.is_inference()</code> |
| <code>is_neg()</code> | Returns True if the negative bit of <code>self</code> is set to true. |
| <code>is_nonzero</code> | |
| <code>is_pinned</code> | Returns true if this tensor resides in pinned memory. |
| <code>is_same_size</code> | |
| <code>is_set_to(tensor)</code> | Returns True if both tensors are pointing to the exact same memory (same storage, offset, size and stride). |
| <code>is_shared()</code> | Checks if tensor is in shared memory. |
| <code>is_signed()</code> | Returns True if the data type of <code>self</code> is a signed data type. |
| <code>isclose(other[, rtol, atol, equal_nan])</code> | See <code>torch.isclose()</code> |
| <code>isfinite()</code> | See <code>torch.isfinite()</code> |
| <code>isinf()</code> | See <code>torch.isinf()</code> |
| <code>isnan()</code> | See <code>torch.isnan()</code> |
| <code>isneginf()</code> | See <code>torch.isneginf()</code> |
| <code>isposinf()</code> | See <code>torch.isposinf()</code> |
| <code>isreal()</code> | See <code>torch.isreal()</code> |
| <code>istft(n_fft[, hop_length, win_length, ...])</code> | See <code>torch.istft()</code> |
| <code>item()</code> | Returns the value of this tensor as a standard Python number. |
| <code>kron(other)</code> | See <code>torch.kron()</code> |
| <code>kthvalue(k[, dim, keepdim])</code> | See <code>torch.kthvalue()</code> |
| <code>lcm(other)</code> | See <code>torch.lcm()</code> |
| <code>lcm_(other)</code> | In-place version of <code>lcm()</code> |
| <code>ldexp(other)</code> | See <code>torch.ldexp()</code> |
| <code>ldexp_(other)</code> | In-place version of <code>ldexp()</code> |
| <code>le(other)</code> | See <code>torch.le()</code> . |
| <code>le_(other)</code> | In-place version of <code>le()</code> . |
| <code>lerp(end, weight)</code> | See <code>torch.lerp()</code> |
| <code>lerp_(end, weight)</code> | In-place version of <code>lerp()</code> |
| <code>less</code> | <code>lt(other) -> Tensor</code> |
| <code>less_(other)</code> | In-place version of <code>less()</code> . |
| <code>less_equal(other)</code> | See <code>torch.less_equal()</code> . |
| <code>less_equal_(other)</code> | In-place version of <code>less_equal()</code> . |
| <code>lgamma()</code> | See <code>torch.lgamma()</code> |
| <code>lgamma_()</code> | In-place version of <code>lgamma()</code> |
| <code>log()</code> | See <code>torch.log()</code> |
| <code>log10()</code> | See <code>torch.log10()</code> |

continues on next page

Table 12 – continued from previous page

| | |
|--|--|
| <code>log10_()</code> | In-place version of <code>log10()</code> |
| <code>log1p()</code> | See <code>torch.log1p()</code> |
| <code>log1p_()</code> | In-place version of <code>log1p()</code> |
| <code>log2()</code> | See <code>torch.log2()</code> |
| <code>log2_()</code> | In-place version of <code>log2()</code> |
| <code>log_()</code> | In-place version of <code>log()</code> |
| <code>log_normal_([mean, std, generator])</code> | Fills <code>self</code> tensor with numbers samples from the log-normal distribution parameterized by the given mean μ and standard deviation σ . |
| <code>log_softmax</code> | |
| <code>logaddexp(other)</code> | See <code>torch.logaddexp()</code> |
| <code>logaddexp2(other)</code> | See <code>torch.logaddexp2()</code> |
| <code>logcumsumexp(dim)</code> | See <code>torch.logcumsumexp()</code> |
| <code>logdet()</code> | See <code>torch.logdet()</code> |
| <code>logical_and()</code> | See <code>torch.logical_and()</code> |
| <code>logical_and_()</code> | In-place version of <code>logical_and()</code> |
| <code>logical_not()</code> | See <code>torch.logical_not()</code> |
| <code>logical_not_()</code> | In-place version of <code>logical_not()</code> |
| <code>logical_or()</code> | See <code>torch.logical_or()</code> |
| <code>logical_or_()</code> | In-place version of <code>logical_or()</code> |
| <code>logical_xor()</code> | See <code>torch.logical_xor()</code> |
| <code>logical_xor_()</code> | In-place version of <code>logical_xor()</code> |
| <code>logit()</code> | See <code>torch.logit()</code> |
| <code>logit_()</code> | In-place version of <code>logit()</code> |
| <code>logsumexp(dim[, keepdim])</code> | See <code>torch.logsumexp()</code> |
| <code>long([memory_format])</code> | <code>self.long()</code> is equivalent to <code>self.to(torch.int64)</code> . |
| <code>lstsq(other)</code> | |
| <code>lt(other)</code> | See <code>torch.lt()</code> . |
| <code>lt_(other)</code> | In-place version of <code>lt()</code> . |
| <code>lu([pivot, get_infos])</code> | See <code>torch.lu()</code> |
| <code>lu_solve(LU_data, LU_pivots)</code> | See <code>torch.lu_solve()</code> |
| <code>map2_</code> | |
| <code>map_(tensor, callable)</code> | Applies <code>callable</code> for each element in <code>self</code> tensor and the given <code>tensor</code> and stores the results in <code>self</code> tensor. |
| <code>masked_fill(mask, value)</code> | Out-of-place version of <code>torch.Tensor.masked_fill_()</code> |
| <code>masked_fill_(mask, value)</code> | Fills elements of <code>self</code> tensor with <code>value</code> where <code>mask</code> is <code>True</code> . |
| <code>masked_scatter(mask, tensor)</code> | Out-of-place version of <code>torch.Tensor.masked_scatter_()</code> |
| <code>masked_scatter_(mask, source)</code> | Copies elements from <code>source</code> into <code>self</code> tensor at positions where the <code>mask</code> is <code>True</code> . |
| <code>masked_select(mask)</code> | See <code>torch.masked_select()</code> |
| <code>materialize(shape[, device, dtype])</code> | Create a <code>Parameter</code> with the same properties of the uninitialized one. |
| <code>matmul(tensor2)</code> | See <code>torch.matmul()</code> |

continues on next page

Table 12 – continued from previous page

| | |
|--|---|
| <code>matrix_exp()</code> | See <code>torch.matrix_exp()</code> |
| <code>matrix_power(n)</code> | |
| | Note: <code>matrix_power()</code> is deprecated, use <code>torch.linalg.matrix_power()</code> instead. |
| <code>max([dim, keepdim])</code> | See <code>torch.max()</code> |
| <code>maximum(other)</code> | See <code>torch.maximum()</code> |
| <code>mean([dim, keepdim, dtype])</code> | See <code>torch.mean()</code> |
| <code>median([dim, keepdim])</code> | See <code>torch.median()</code> |
| <code>min([dim, keepdim])</code> | See <code>torch.min()</code> |
| <code>minimum(other)</code> | See <code>torch.minimum()</code> |
| <code>mm(mat2)</code> | See <code>torch.mm()</code> |
| <code>mode([dim, keepdim])</code> | See <code>torch.mode()</code> |
| <code>moveaxis(source, destination)</code> | See <code>torch.moveaxis()</code> |
| <code>movedim(source, destination)</code> | See <code>torch.movedim()</code> |
| <code>msort()</code> | See <code>torch.msort()</code> |
| <code>mul(value)</code> | See <code>torch.mul()</code> . |
| <code>mul_(value)</code> | In-place version of <code>mul()</code> . |
| <code>multinomial(num_samples[, replacement, ...])</code> | See <code>torch.multinomial()</code> |
| <code>multiply(value)</code> | See <code>torch.multiply()</code> . |
| <code>multiply_(value)</code> | In-place version of <code>multiply()</code> . |
| <code>mv(vec)</code> | See <code>torch.mv()</code> |
| <code>mvlgamma(p)</code> | See <code>torch.mvlgamma()</code> |
| <code>mvlgamma_(p)</code> | In-place version of <code>mvlgamma()</code> |
| <code>nan_to_num([nan, posinf, neginf])</code> | See <code>torch.nan_to_num()</code> . |
| <code>nan_to_num_([nan, posinf, neginf])</code> | In-place version of <code>nan_to_num()</code> . |
| <code>nanmean([dim, keepdim, dtype])</code> | See <code>torch.nanmean()</code> |
| <code>nanmedian([dim, keepdim])</code> | See <code>torch.nanmedian()</code> |
| <code>nanquantile(q[, dim, keepdim, interpolation])</code> | See <code>torch.nanquantile()</code> |
| <code>nansum([dim, keepdim, dtype])</code> | See <code>torch.nansum()</code> |
| <code>narrow(dimension, start, length)</code> | See <code>torch.narrow()</code> . |
| <code>narrow_copy(dimension, start, length)</code> | See <code>torch.narrow_copy()</code> . |
| <code>ndimension()</code> | Alias for <code>dim()</code> |
| <code>ne(other)</code> | See <code>torch.ne()</code> . |
| <code>ne_(other)</code> | In-place version of <code>ne()</code> . |
| <code>neg()</code> | See <code>torch.neg()</code> |
| <code>neg_()</code> | In-place version of <code>neg()</code> |
| <code>negative()</code> | See <code>torch.negative()</code> |
| <code>negative_()</code> | In-place version of <code>negative()</code> |
| <code>nelement()</code> | Alias for <code>numel()</code> |
| <code>new</code> | |
| <code>new_empty(size, *[, dtype, device, ...])</code> | Returns a Tensor of size <code>size</code> filled with uninitialized data. |
| <code>new_empty_strided(size, stride[, dtype, ...])</code> | Returns a Tensor of size <code>size</code> and strides <code>stride</code> filled with uninitialized data. |
| <code>new_full(size, fill_value, *[, dtype, ...])</code> | Returns a Tensor of size <code>size</code> filled with <code>fill_value</code> . |
| <code>new_ones(size, *[, dtype, device, ...])</code> | Returns a Tensor of size <code>size</code> filled with 1. |

continues on next page

Table 12 – continued from previous page

| | |
|---|--|
| <code>new_tensor(data, *, dtype, device, ...)</code> | Returns a new Tensor with data as the tensor data. |
| <code>new_zeros(size, *, dtype, device, ...)</code> | Returns a Tensor of size <code>size</code> filled with 0. |
| <code>nextafter(other)</code> | See <code>torch.nextafter()</code> |
| <code>nextafter_(other)</code> | In-place version of <code>nextafter()</code> |
| <code>nonzero()</code> | See <code>torch.nonzero()</code> |
| <code>nonzero_static(input, *, size[, fill_value])</code> | Returns a 2-D tensor where each row is the index for a non-zero value. |
| <code>norm([p, dim, keepdim, dtype])</code> | See <code>torch.norm()</code> |
| <code>normal_([mean, std, generator])</code> | Fills self tensor with elements samples from the normal distribution parameterized by mean and std. |
| <code>not_equal(other)</code> | See <code>torch.not_equal()</code> . |
| <code>not_equal_(other)</code> | In-place version of <code>not_equal()</code> . |
| <code>numel()</code> | See <code>torch.numel()</code> |
| <code>numpy(*[, force])</code> | Returns the tensor as a NumPy ndarray. |
| <code>orgqr(input2)</code> | See <code>torch.orgqr()</code> |
| <code>ormqr(input2, input3[, left, transpose])</code> | See <code>torch.ormqr()</code> |
| <code>outer(vec2)</code> | See <code>torch.outer()</code> . |
| <code>permute(*dims)</code> | See <code>torch.permute()</code> |
| <code>pin_memory()</code> | Copies the tensor to pinned memory, if it's not already pinned. |
| <code>pinverse()</code> | See <code>torch.pinverse()</code> |
| <code>polygamma(n)</code> | See <code>torch.polygamma()</code> |
| <code>polygamma_(n)</code> | In-place version of <code>polygamma()</code> |
| <code>positive()</code> | See <code>torch.positive()</code> |
| <code>pow(exponent)</code> | See <code>torch.pow()</code> |
| <code>pow_(exponent)</code> | In-place version of <code>pow()</code> |
| <code>prelu</code> | |
| <code>prod([dim, keepdim, dtype])</code> | See <code>torch.prod()</code> |
| <code>put(input, index, source[, accumulate])</code> | Out-of-place version of <code>torch.Tensor.put_()</code> . |
| <code>put_(index, source[, accumulate])</code> | Copies the elements from <code>source</code> into the positions specified by <code>index</code> . |
| <code>q_per_channel_axis()</code> | Given a Tensor quantized by linear (affine) per-channel quantization, returns the index of dimension on which per-channel quantization is applied. |
| <code>q_per_channel_scales()</code> | Given a Tensor quantized by linear (affine) per-channel quantization, returns a Tensor of scales of the underlying quantizer. |
| <code>q_per_channel_zero_points()</code> | Given a Tensor quantized by linear (affine) per-channel quantization, returns a tensor of zero_points of the underlying quantizer. |
| <code>q_scale()</code> | Given a Tensor quantized by linear(affine) quantization, returns the scale of the underlying quantizer(). |
| <code>q_zero_point()</code> | Given a Tensor quantized by linear(affine) quantization, returns the zero_point of the underlying quantizer(). |
| <code>qr([some])</code> | See <code>torch.qr()</code> |
| <code>qscheme()</code> | Returns the quantization scheme of a given QTensor. |
| <code>quantile(q[, dim, keepdim, interpolation])</code> | See <code>torch.quantile()</code> |
| <code>rad2deg()</code> | See <code>torch.rad2deg()</code> |
| <code>rad2deg_()</code> | In-place version of <code>rad2deg()</code> |

continues on next page

Table 12 – continued from previous page

| | |
|---|--|
| <code>random_([from, to, generator])</code> | Fills <code>self</code> tensor with numbers sampled from the discrete uniform distribution over <code>[from, to - 1]</code> . |
| <code>ravel()</code> | see <code>torch.ravel()</code> |
| <code>reciprocal()</code> | See <code>torch.reciprocal()</code> |
| <code>reciprocal_()</code> | In-place version of <code>reciprocal()</code> |
| <code>record_stream(stream)</code> | Ensures that the tensor memory is not reused for another tensor until all current work queued on <code>stream</code> are complete. |
| <code>refine_names(*names)</code> | Refines the dimension names of <code>self</code> according to <code>names</code> . |
| <code>register_hook(hook)</code> | Registers a backward hook. |
| <code>register_post_accumulate_grad_hook(hook)</code> | Registers a backward hook that runs after grad accumulation. |
| <code>reinforce(reward)</code> | |
| <code>relu</code> | |
| <code>relu_</code> | |
| <code>remainder(divisor)</code> | See <code>torch.remainder()</code> |
| <code>remainder_(divisor)</code> | In-place version of <code>remainder()</code> |
| <code>rename(*names, **rename_map)</code> | Renames dimension names of <code>self</code> . |
| <code>rename_(*names, **rename_map)</code> | In-place version of <code>rename()</code> . |
| <code>renorm(p, dim, maxnorm)</code> | See <code>torch.renorm()</code> |
| <code>renorm_(p, dim, maxnorm)</code> | In-place version of <code>renorm()</code> |
| <code>repeat(*sizes)</code> | Repeats this tensor along the specified dimensions. |
| <code>repeat_interleave(repeats[, dim, output_size])</code> | See <code>torch.repeat_interleave()</code> . |
| <code>requires_grad_(requires_grad)</code> | Change if autograd should record operations on this tensor: sets this tensor's <code>requires_grad</code> attribute in-place. |
| <code>reshape(*shape)</code> | Returns a tensor with the same data and number of elements as <code>self</code> but with the specified shape. |
| <code>reshape_as(other)</code> | Returns this tensor as the same shape as <code>other</code> . |
| <code>resize(*sizes)</code> | |
| <code>resize_(*sizes[, memory_format])</code> | Resizes <code>self</code> tensor to the specified size. |
| <code>resize_as(tensor)</code> | |
| <code>resize_as_(tensor[, memory_format])</code> | Resizes the <code>self</code> tensor to be the same size as the specified tensor. |
| <code>resize_as_sparse_</code> | |
| <code>resolve_conj()</code> | See <code>torch.resolve_conj()</code> |
| <code>resolve_neg()</code> | See <code>torch.resolve_neg()</code> |
| <code>retain_grad()</code> | Enables this Tensor to have their grad populated during <code>backward()</code> . |
| <code>roll(shifts, dims)</code> | See <code>torch.roll()</code> |
| <code>rot90(k, dims)</code> | See <code>torch.rot90()</code> |
| <code>round([decimals])</code> | See <code>torch.round()</code> |
| <code>round_([decimals])</code> | In-place version of <code>round()</code> |

continues on next page

Table 12 – continued from previous page

| | |
|--|---|
| row_indices | |
| <code>rsqrt()</code> | See <code>torch.rsqrt()</code> |
| <code>rsqrt_()</code> | In-place version of <code>rsqrt()</code> |
| <code>scatter(dim, index, src)</code> | Out-of-place version of <code>torch.Tensor.scatter_()</code> |
| <code>scatter_(dim, index, src[, reduce])</code> | Writes all values from the tensor <code>src</code> into <code>self</code> at the indices specified in the <code>index</code> tensor. |
| <code>scatter_add(dim, index, src)</code> | Out-of-place version of <code>torch.Tensor.scatter_add_()</code> |
| <code>scatter_add_(dim, index, src)</code> | Adds all values from the tensor <code>src</code> into <code>self</code> at the indices specified in the <code>index</code> tensor in a similar fashion as <code>scatter_()</code> . |
| <code>scatter_reduce(dim, index, src, reduce, *[, ...])</code> | Out-of-place version of <code>torch.Tensor.scatter_reduce_()</code> |
| <code>scatter_reduce_(dim, index, src, reduce, *)</code> | Reduces all values from the <code>src</code> tensor to the indices specified in the <code>index</code> tensor in the <code>self</code> tensor using the applied reduction defined via the <code>reduce</code> argument ("sum", "prod", "mean", "amax", "amin"). |
| <code>select(dim, index)</code> | See <code>torch.select()</code> |
| <code>select_scatter(src, dim, index)</code> | See <code>torch.select_scatter()</code> |
| <code>set_([source, storage_offset, size, stride])</code> | Sets the underlying storage, size, and strides. |
| <code>sgn()</code> | See <code>torch.sgn()</code> |
| <code>sgn_()</code> | In-place version of <code>sgn()</code> |
| <code>share_memory_()</code> | Moves the underlying storage to shared memory. |
| <code>short([memory_format])</code> | <code>self.short()</code> is equivalent to <code>self.to(torch.int16)</code> . |
| <code>sigmoid()</code> | See <code>torch.sigmoid()</code> |
| <code>sigmoid_()</code> | In-place version of <code>sigmoid()</code> |
| <code>sign()</code> | See <code>torch.sign()</code> |
| <code>sign_()</code> | In-place version of <code>sign()</code> |
| <code>signbit()</code> | See <code>torch.signbit()</code> |
| <code>sin()</code> | See <code>torch.sin()</code> |
| <code>sin_()</code> | In-place version of <code>sin()</code> |
| <code>sinc()</code> | See <code>torch.sinc()</code> |
| <code>sinc_()</code> | In-place version of <code>sinc()</code> |
| <code>sinh()</code> | See <code>torch.sinh()</code> |
| <code>sinh_()</code> | In-place version of <code>sinh()</code> |
| <code>size([dim])</code> | Returns the size of the <code>self</code> tensor. |
| <code>slice_scatter(src[, dim, start, end, step])</code> | See <code>torch.slice_scatter()</code> |
| <code>slogdet()</code> | See <code>torch.slogdet()</code> |
| <code>smm(mat)</code> | See <code>torch.smm()</code> |
| <code>softmax(dim)</code> | Alias for <code>torch.nn.functional.softmax()</code> . |
| <code>solve(other)</code> | |
| <code>sort([dim, descending])</code> | See <code>torch.sort()</code> |
| <code>sparse_dim()</code> | Return the number of sparse dimensions in a sparse tensor <code>self</code> . |
| <code>sparse_mask(mask)</code> | Returns a new sparse tensor with values from a strided tensor <code>self</code> filtered by the indices of the sparse tensor <code>mask</code> . |

continues on next page

Table 12 – continued from previous page

| | |
|--|--|
| <code>sparse_resize_(size, sparse_dim, dense_dim)</code> | Resizes <code>self</code> sparse tensor to the desired size and the number of sparse and dense dimensions. |
| <code>sparse_resize_and_clear_(size, sparse_dim, ...)</code> | Removes all specified elements from a sparse tensor <code>self</code> and resizes <code>self</code> to the desired size and the number of sparse and dense dimensions. |
| <code>split(split_size[, dim])</code> | See <code>torch.split()</code> |
| <code>split_with_sizes</code> | |
| <code>sqrt()</code> | See <code>torch.sqrt()</code> |
| <code>sqrt_()</code> | In-place version of <code>sqrt()</code> |
| <code>square()</code> | See <code>torch.square()</code> |
| <code>square_()</code> | In-place version of <code>square()</code> |
| <code>squeeze([dim])</code> | See <code>torch.squeeze()</code> |
| <code>squeeze_([dim])</code> | In-place version of <code>squeeze()</code> |
| <code>sspaddmm(mat1, mat2, *[, beta, alpha])</code> | See <code>torch.sspaddmm()</code> |
| <code>std([dim, correction, keepdim])</code> | See <code>torch.std()</code> |
| <code>stft(n_fft[, hop_length, win_length, ...])</code> | See <code>torch.stft()</code> |
| <code>storage()</code> | Returns the underlying <code>TypedStorage</code> . |
| <code>storage_offset()</code> | Returns <code>self</code> tensor's offset in the underlying storage in terms of number of storage elements (not bytes). |
| <code>storage_type()</code> | Returns the type of the underlying storage. |
| <code>stride(dim)</code> | Returns the stride of <code>self</code> tensor. |
| <code>sub(other, *[, alpha])</code> | See <code>torch.sub()</code> . |
| <code>sub_(other, *[, alpha])</code> | In-place version of <code>sub()</code> |
| <code>subtract(other, *[, alpha])</code> | See <code>torch.subtract()</code> . |
| <code>subtract_(other, *[, alpha])</code> | In-place version of <code>subtract()</code> . |
| <code>sum([dim, keepdim, dtype])</code> | See <code>torch.sum()</code> |
| <code>sum_to_size(*size)</code> | Sum this tensor to size. |
| <code>svd([some, compute_uv])</code> | See <code>torch.svd()</code> |
| <code>swapaxes(axis0, axis1)</code> | See <code>torch.swapaxes()</code> |
| <code>swapaxes_(axis0, axis1)</code> | In-place version of <code>swapaxes()</code> |
| <code>swapdims(dim0, dim1)</code> | See <code>torch.swapdims()</code> |
| <code>swapdims_(dim0, dim1)</code> | In-place version of <code>swapdims()</code> |
| <code>syમેig([eigenvectors])</code> | |
| <code>t()</code> | See <code>torch.t()</code> |
| <code>t_()</code> | In-place version of <code>t()</code> |
| <code>take(indices)</code> | See <code>torch.take()</code> |
| <code>take_along_dim(indices, dim)</code> | See <code>torch.take_along_dim()</code> |
| <code>tan()</code> | See <code>torch.tan()</code> |
| <code>tan_()</code> | In-place version of <code>tan()</code> |
| <code>tanh()</code> | See <code>torch.tanh()</code> |
| <code>tanh_()</code> | In-place version of <code>tanh()</code> |
| <code>tensor_split(indices_or_sections[, dim])</code> | See <code>torch.tensor_split()</code> |
| <code>tile(dims)</code> | See <code>torch.tile()</code> |
| <code>to(*args, **kwargs)</code> | Performs Tensor dtype and/or device conversion. |
| <code>to_dense([dtype, masked_grad])</code> | Creates a strided copy of <code>self</code> if <code>self</code> is not a strided tensor, otherwise returns <code>self</code> . |
| <code>to_mkldnn()</code> | Returns a copy of the tensor in <code>torch.mkldnn</code> layout. |
| <code>to_padded_tensor(padding[, output_size])</code> | See <code>to_padded_tensor()</code> |

continues on next page

Table 12 – continued from previous page

| | |
|---|---|
| <code>to_sparse(sparseDims)</code> | Returns a sparse copy of the tensor. |
| <code>to_sparse_bsc(blocksize, dense_dim)</code> | Convert a tensor to a block sparse column (BSC) storage format of given blocksize. |
| <code>to_sparse_bsr(blocksize, dense_dim)</code> | Convert a tensor to a block sparse row (BSR) storage format of given blocksize. |
| <code>to_sparse_coo()</code> | Convert a tensor to coordinate format. |
| <code>to_sparse_csc()</code> | Convert a tensor to compressed column storage (CSC) format. |
| <code>to_sparse_csr([dense_dim])</code> | Convert a tensor to compressed row storage format (CSR). |
| <code>tolist()</code> | Returns the tensor as a (nested) list. |
| <code>topk(k[, dim, largest, sorted])</code> | See <code>torch.topk()</code> |
| <code>trace()</code> | See <code>torch.trace()</code> |
| <code>transpose(dim0, dim1)</code> | See <code>torch.transpose()</code> |
| <code>transpose_(dim0, dim1)</code> | In-place version of <code>transpose()</code> |
| <code>triangular_solve(A[, upper, transpose, ...])</code> | See <code>torch.triangular_solve()</code> |
| <code>tril([diagonal])</code> | See <code>torch.tril()</code> |
| <code>tril_([diagonal])</code> | In-place version of <code>tril()</code> |
| <code>triu([diagonal])</code> | See <code>torch.triu()</code> |
| <code>triu_([diagonal])</code> | In-place version of <code>triu()</code> |
| <code>true_divide(value)</code> | See <code>torch.true_divide()</code> |
| <code>true_divide_(value)</code> | In-place version of <code>true_divide_()</code> |
| <code>trunc()</code> | See <code>torch.trunc()</code> |
| <code>trunc_()</code> | In-place version of <code>trunc()</code> |
| <code>type([dtype, non_blocking])</code> | Returns the type if <i>dtype</i> is not provided, else casts this object to the specified type. |
| <code>type_as(tensor)</code> | Returns this tensor cast to the type of the given tensor. |
| <code>unbind([dim])</code> | See <code>torch.unbind()</code> |
| <code>unflatten(dim, sizes)</code> | See <code>torch.unflatten()</code> . |
| <code>unfold(dimension, size, step)</code> | Returns a view of the original tensor which contains all slices of size <i>size</i> from <i>self</i> tensor in the dimension <i>dimension</i> . |
| <code>uniform_([from, to, generator])</code> | Fills <i>self</i> tensor with numbers sampled from the continuous uniform distribution: |
| <code>unique([sorted, return_inverse, ...])</code> | Returns the unique elements of the input tensor. |
| <code>unique_consecutive([return_inverse, ...])</code> | Eliminates all but the first element from every consecutive group of equivalent elements. |
| <code>unsafe_chunk(chunks[, dim])</code> | See <code>torch.unsafe_chunk()</code> |
| <code>unsafe_split(split_size[, dim])</code> | See <code>torch.unsafe_split()</code> |
| <code>unsafe_split_with_sizes</code> | |
| <code>unsqueeze(dim)</code> | See <code>torch.unsqueeze()</code> |
| <code>unsqueeze_(dim)</code> | In-place version of <code>unsqueeze()</code> |
| <code>untyped_storage()</code> | Returns the underlying <code>UntypedStorage</code> . |
| <code>values()</code> | Return the values tensor of a sparse COO tensor. |
| <code>var([dim, correction, keepdim])</code> | See <code>torch.var()</code> |
| <code>vdot(other)</code> | See <code>torch.vdot()</code> |
| <code>view(*shape)</code> | Returns a new tensor with the same data as the <i>self</i> tensor but of a different <i>shape</i> . |
| <code>view_as(other)</code> | View this tensor as the same size as <i>other</i> . |
| <code>vsplit(split_size_or_sections)</code> | See <code>torch.vsplit()</code> |

continues on next page

Table 12 – continued from previous page

| | |
|---|---|
| <code>where(condition, y)</code> | <code>self.where(condition, y)</code> is equivalent to <code>torch.where(condition, self, y)</code> . |
| <code>xlogy(other)</code> | See <code>torch.xlogy()</code> |
| <code>xlogy_(other)</code> | In-place version of <code>xlogy()</code> |
| <code>xpu([device, non_blocking, memory_format])</code> | Returns a copy of this object in XPU memory. |
| <code>zero_()</code> | Fills <code>self</code> tensor with zeros. |

Attributes

| | |
|----------------------------|---|
| <code>H</code> | Returns a view of a matrix (2-D tensor) conjugated and transposed. |
| <code>T</code> | Returns a view of this tensor with its dimensions reversed. |
| <code>data</code> | |
| <code>device</code> | Is the <code>torch.device</code> where this Tensor is. |
| <code>dtype</code> | |
| <code>grad</code> | This attribute is <code>None</code> by default and becomes a Tensor the first time a call to <code>backward()</code> computes gradients for <code>self</code> . |
| <code>grad_fn</code> | |
| <code>imag</code> | Returns a new tensor containing imaginary values of the <code>self</code> tensor. |
| <code>is_cpu</code> | Is <code>True</code> if the Tensor is stored on the CPU, <code>False</code> otherwise. |
| <code>is_cuda</code> | Is <code>True</code> if the Tensor is stored on the GPU, <code>False</code> otherwise. |
| <code>is_ipu</code> | Is <code>True</code> if the Tensor is stored on the IPU, <code>False</code> otherwise. |
| <code>is_leaf</code> | All Tensors that have <code>requires_grad</code> which is <code>False</code> will be leaf Tensors by convention. |
| <code>is_meta</code> | Is <code>True</code> if the Tensor is a meta tensor, <code>False</code> otherwise. |
| <code>is_mkldnn</code> | |
| <code>is_mps</code> | Is <code>True</code> if the Tensor is stored on the MPS device, <code>False</code> otherwise. |
| <code>is_nested</code> | |
| <code>is_ort</code> | |
| <code>is_quantized</code> | Is <code>True</code> if the Tensor is quantized, <code>False</code> otherwise. |
| <code>is_sparse</code> | Is <code>True</code> if the Tensor uses sparse COO storage layout, <code>False</code> otherwise. |
| <code>is_sparse_csr</code> | Is <code>True</code> if the Tensor uses sparse CSR storage layout, <code>False</code> otherwise. |

continues on next page

Table 13 – continued from previous page

| | |
|----------------------------|---|
| <code>is_vulkan</code> | |
| <code>is_xla</code> | Is True if the Tensor is stored on an XLA device, False otherwise. |
| <code>is_xpu</code> | Is True if the Tensor is stored on the XPU, False otherwise. |
| <code>itemsizes</code> | Alias for <code>element_size()</code> |
| <code>layout</code> | |
| <code>mH</code> | Accessing this property is equivalent to calling <code>adjoint()</code> . |
| <code>mT</code> | Returns a view of this tensor with the last two dimensions transposed. |
| <code>name</code> | |
| <code>names</code> | Stores names for each of this tensor's dimensions. |
| <code>nbytes</code> | Returns the number of bytes consumed by the "view" of elements of the Tensor if the Tensor does not use sparse storage layout. |
| <code>ndim</code> | Alias for <code>dim()</code> |
| <code>output_nr</code> | |
| <code>real</code> | Returns a new tensor containing real values of the <code>self</code> tensor for a complex-valued input tensor. |
| <code>requires_grad</code> | Is True if gradients need to be computed for this Tensor, False otherwise. |
| <code>retains_grad</code> | Is True if this Tensor is non-leaf and its <code>grad</code> is enabled to be populated during <code>backward()</code> , False otherwise. |
| <code>shape</code> | Returns the size of the <code>self</code> tensor. |
| <code>volatile</code> | |

property `is_leaf`: bool

All Tensors that have `requires_grad` which is False will be leaf Tensors by convention.

For Tensors that have `requires_grad` which is True, they will be leaf Tensors if they were created by the user. This means that they are not the result of an operation and so `grad_fn` is None.

Only leaf Tensors will have their `grad` populated during a call to `backward()`. To get `grad` populated for non-leaf Tensors, you can use `retain_grad()`.

Example:

```
>>> a = torch.rand(10, requires_grad=True)
>>> a.is_leaf
True
>>> b = torch.rand(10, requires_grad=True).cuda()
>>> b.is_leaf
False
# b was created by the operation that cast a cpu Tensor into a cuda Tensor
>>> c = torch.rand(10, requires_grad=True) + 2
>>> c.is_leaf
```

(continues on next page)

(continued from previous page)

```

False
# c was created by the addition operation
>>> d = torch.rand(10).cuda()
>>> d.is_leaf
True
# d does not require gradients and so has no operation creating it (that is,
→ tracked by the autograd engine)
>>> e = torch.rand(10).cuda().requires_grad_()
>>> e.is_leaf
True
# e requires gradients and has no operations creating it
>>> f = torch.rand(10, requires_grad=True, device="cuda")
>>> f.is_leaf
True
# f requires grad, has no operation creating it

```

materialize(*shape*, *device=None*, *dtype=None*)

Create a Parameter with the same properties of the uninitialized one. Given a shape, it materializes a parameter in the same device and with the same *dtype* as the current one or the specified ones in the arguments.

Parameters

- **shape** (*Tuple[int, ...]*) – (tuple): the shape for the materialized tensor.
- **device** (*torch.device*) – the desired device of the parameters and buffers in this module. Optional.
- **dtype** (*torch.dtype*) – the desired floating point type of the floating point parameters and buffers in this module. Optional.

Return type

None

share_memory_()

Moves the underlying storage to shared memory.

This is a no-op if the underlying storage is already in shared memory and for CUDA tensors. Tensors in shared memory cannot be resized.

Return type

[UninitializedParameter](#)

pytorch_pfn_extras.nn.modules.lazy_batchnorm

Classes

| | |
|---|--|
| <code>pytorch_pfn_extras.nn.modules.lazy_batchnorm.LazyBatchNorm1d(...)</code> | BatchNorm1d module with lazy weight initialization. |
| <code>pytorch_pfn_extras.nn.modules.lazy_batchnorm.LazyBatchNorm2d(...)</code> | BatchNorm2d module with lazy weight initialization. |
| <code>pytorch_pfn_extras.nn.modules.lazy_batchnorm.LazyBatchNorm3d(...)</code> | BatchNorm3d module with lazy weight initialization. |
| <code>pytorch_pfn_extras.nn.modules.lazy_batchnorm.LazyInitializationMixin(...)</code> | A mixin for modules that lazily initialize buffers and parameters. |
| <code>pytorch_pfn_extras.nn.modules.lazy_batchnorm.UninitializedParameter([...])</code> | |

pytorch_pfn_extras.nn.modules.lazy_batchnorm.LazyBatchNorm1d

class `pytorch_pfn_extras.nn.modules.lazy_batchnorm.LazyBatchNorm1d`(*num_features*, *args, **kwargs)

Bases: `_LazyBatchNorm`, `BatchNorm1d`

BatchNorm1d module with lazy weight initialization.

When *num_features* is `None`, it is determined at the first time of the forward step.

Initializes internal Module state, shared by both `nn.Module` and `ScriptModule`.

Methods

| | |
|--|--|
| <code>__init__(num_features, *args, **kwargs)</code> | Initializes internal Module state, shared by both <code>nn.Module</code> and <code>ScriptModule</code> . |
| <code>add_module(name, module)</code> | Adds a child module to the current module. |
| <code>apply(fn)</code> | Applies <i>fn</i> recursively to every submodule (as returned by <code>.children()</code>) as well as self. |
| <code>bfloat16()</code> | Casts all floating point parameters and buffers to <code>bfloat16</code> datatype. |
| <code>buffers([recurse])</code> | Returns an iterator over module buffers. |
| <code>children()</code> | Returns an iterator over immediate children modules. |
| <code>compile(*args, **kwargs)</code> | Compile this Module's forward using <code>torch.compile()</code> . |
| <code>cpu()</code> | Moves all model parameters and buffers to the CPU. |
| <code>cuda([device])</code> | Moves all model parameters and buffers to the GPU. |
| <code>double()</code> | Casts all floating point parameters and buffers to <code>double</code> datatype. |
| <code>eval()</code> | Sets the module in evaluation mode. |
| <code>extra_repr()</code> | Set the extra representation of the module |
| <code>float()</code> | Casts all floating point parameters and buffers to <code>float</code> datatype. |
| <code>forward(input)</code> | Defines the computation performed at every call. |

continues on next page

Table 14 – continued from previous page

| | |
|---|---|
| <code>get_buffer(target)</code> | Returns the buffer given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_extra_state()</code> | Returns any extra state to include in the module's <code>state_dict</code> . |
| <code>get_parameter(target)</code> | Returns the parameter given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_submodule(target)</code> | Returns the submodule given by <code>target</code> if it exists, otherwise throws an error. |
| <code>half()</code> | Casts all floating point parameters and buffers to <code>half</code> datatype. |
| <code>ipu([device])</code> | Moves all model parameters and buffers to the IPU. |
| <code>load_state_dict(state_dict[, strict, assign])</code> | Copies parameters and buffers from <code>state_dict</code> into this module and its descendants. |
| <code>modules()</code> | Returns an iterator over all modules in the network. |
| <code>named_buffers([prefix, recurse, ...])</code> | Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself. |
| <code>named_children()</code> | Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself. |
| <code>named_modules([memo, prefix, remove_duplicate])</code> | Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself. |
| <code>named_parameters([prefix, recurse, ...])</code> | Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself. |
| <code>parameters([recurse])</code> | Returns an iterator over module parameters. |
| <code>register_backward_hook(hook)</code> | Registers a backward hook on the module. |
| <code>register_buffer(name, tensor[, persistent])</code> | Adds a buffer to the module. |
| <code>register_forward_hook(hook, *[, prepend, ...])</code> | Registers a forward hook on the module. |
| <code>register_forward_pre_hook(hook, *[, ...])</code> | Registers a forward pre-hook on the module. |
| <code>register_full_backward_hook(hook[, prepend])</code> | Registers a backward hook on the module. |
| <code>register_full_backward_pre_hook(hook[, prepend])</code> | Registers a backward pre-hook on the module. |
| <code>register_load_state_dict_post_hook(hook)</code> | Registers a post hook to be run after module's <code>load_state_dict</code> is called. |
| <code>register_module(name, module)</code> | Alias for <code>add_module()</code> . |
| <code>register_parameter(name, param)</code> | Adds a parameter to the module. |
| <code>register_state_dict_pre_hook(hook)</code> | These hooks will be called with arguments: <code>self</code> , <code>prefix</code> , and <code>keep_vars</code> before calling <code>state_dict</code> on <code>self</code> . |
| <code>requires_grad_([requires_grad])</code> | Change if autograd should record operations on parameters in this module. |
| <code>reset_parameters()</code> | |
| <code>reset_running_stats()</code> | |
| <code>set_extra_state(state)</code> | This function is called from <code>load_state_dict()</code> to handle any extra state found within the <code>state_dict</code> . |
| <code>share_memory()</code> | See <code>torch.Tensor.share_memory_()</code> |

continues on next page

Table 14 – continued from previous page

| | |
|---|---|
| <code>state_dict(*args, **kwargs)</code> | Returns a dictionary containing a whole state of the module. |
| <code>to(*args, **kwargs)</code> | Moves and/or casts the parameters and buffers. |
| <code>to_empty(*, device[, recurse])</code> | Moves the parameters and buffers to the specified device without copying storage. |
| <code>train([mode])</code> | Sets the module in training mode. |
| <code>type(dst_type)</code> | Casts all parameters and buffers to <code>dst_type</code> . |
| <code>xpu([device])</code> | Moves all model parameters and buffers to the XPU. |
| <code>zero_grad([set_to_none])</code> | Resets gradients of all model parameters. |

Attributes

| | |
|---|--|
| <code>T_destination</code> | alias of <code>TypeVar('T_destination', bound=Dict[str, Any])</code> |
| <code>call_super_init</code> | |
| <code>dump_patches</code> | |
| <code>lazy_buffer_names</code> | |
| <code>lazy_parameter_names</code> | |
| <code>lazy_parameters_determined</code> | Returns if all lazy parameters are determined. |

Parameters

- `num_features` (*Optional[int]*) –
- `args` (*Any*) –
- `kwargs` (*Any*) –

affine: `bool`**eps:** `float`**momentum:** `float`**num_batches_tracked:** `Optional[Tensor]`**num_features:** `Optional[int]`**running_mean:** `Any`**running_var:** `Any`**track_running_stats:** `bool`**training:** `bool`

pytorch_pfn_extras.nn.modules.lazy_batchnorm.LazyBatchNorm2d

```
class pytorch_pfn_extras.nn.modules.lazy_batchnorm.LazyBatchNorm2d(num_features, *args,
                                                                    **kwargs)
```

Bases: `_LazyBatchNorm`, `BatchNorm2d`

BatchNorm2d module with lazy weight initialization.

When `num_features` is `None`, it is determined at the first time of the forward step.

Initializes internal Module state, shared by both `nn.Module` and `ScriptModule`.

Methods

| | |
|--|--|
| <code>__init__(num_features, *args, **kwargs)</code> | Initializes internal Module state, shared by both <code>nn.Module</code> and <code>ScriptModule</code> . |
| <code>add_module(name, module)</code> | Adds a child module to the current module. |
| <code>apply(fn)</code> | Applies <code>fn</code> recursively to every submodule (as returned by <code>.children()</code>) as well as self. |
| <code>bfloat16()</code> | Casts all floating point parameters and buffers to <code>bfloat16</code> datatype. |
| <code>buffers([recurse])</code> | Returns an iterator over module buffers. |
| <code>children()</code> | Returns an iterator over immediate children modules. |
| <code>compile(*args, **kwargs)</code> | Compile this Module's forward using <code>torch.compile()</code> . |
| <code>cpu()</code> | Moves all model parameters and buffers to the CPU. |
| <code>cuda([device])</code> | Moves all model parameters and buffers to the GPU. |
| <code>double()</code> | Casts all floating point parameters and buffers to <code>double</code> datatype. |
| <code>eval()</code> | Sets the module in evaluation mode. |
| <code>extra_repr()</code> | Set the extra representation of the module |
| <code>float()</code> | Casts all floating point parameters and buffers to <code>float</code> datatype. |
| <code>forward(input)</code> | Defines the computation performed at every call. |
| <code>get_buffer(target)</code> | Returns the buffer given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_extra_state()</code> | Returns any extra state to include in the module's <code>state_dict</code> . |
| <code>get_parameter(target)</code> | Returns the parameter given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_submodule(target)</code> | Returns the submodule given by <code>target</code> if it exists, otherwise throws an error. |
| <code>half()</code> | Casts all floating point parameters and buffers to <code>half</code> datatype. |
| <code>ipu([device])</code> | Moves all model parameters and buffers to the IPU. |
| <code>load_state_dict(state_dict[, strict, assign])</code> | Copies parameters and buffers from <code>state_dict</code> into this module and its descendants. |
| <code>modules()</code> | Returns an iterator over all modules in the network. |
| <code>named_buffers([prefix, recurse, ...])</code> | Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself. |

continues on next page

Table 15 – continued from previous page

| | |
|---|---|
| <code>named_children()</code> | Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself. |
| <code>named_modules([memo, prefix, remove_duplicate])</code> | Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself. |
| <code>named_parameters([prefix, recurse, ...])</code> | Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself. |
| <code>parameters([recurse])</code> | Returns an iterator over module parameters. |
| <code>register_backward_hook(hook)</code> | Registers a backward hook on the module. |
| <code>register_buffer(name, tensor[, persistent])</code> | Adds a buffer to the module. |
| <code>register_forward_hook(hook, *[, prepend, ...])</code> | Registers a forward hook on the module. |
| <code>register_forward_pre_hook(hook, *[, ...])</code> | Registers a forward pre-hook on the module. |
| <code>register_full_backward_hook(hook[, prepend])</code> | Registers a backward hook on the module. |
| <code>register_full_backward_pre_hook(hook[, prepend])</code> | Registers a backward pre-hook on the module. |
| <code>register_load_state_dict_post_hook(hook)</code> | Registers a post hook to be run after module's <code>load_state_dict</code> is called. |
| <code>register_module(name, module)</code> | Alias for <code>add_module()</code> . |
| <code>register_parameter(name, param)</code> | Adds a parameter to the module. |
| <code>register_state_dict_pre_hook(hook)</code> | These hooks will be called with arguments: <code>self</code> , <code>prefix</code> , and <code>keep_vars</code> before calling <code>state_dict</code> on <code>self</code> . |
| <code>requires_grad_([requires_grad])</code> | Change if autograd should record operations on parameters in this module. |
| <code>reset_parameters()</code> | |
| <code>reset_running_stats()</code> | |
| <code>set_extra_state(state)</code> | This function is called from <code>load_state_dict()</code> to handle any extra state found within the <code>state_dict</code> . |
| <code>share_memory()</code> | See <code>torch.Tensor.share_memory_()</code> |
| <code>state_dict(*args, **kwargs)</code> | Returns a dictionary containing a whole state of the module. |
| <code>to(*args, **kwargs)</code> | Moves and/or casts the parameters and buffers. |
| <code>to_empty(*, device[, recurse])</code> | Moves the parameters and buffers to the specified device without copying storage. |
| <code>train([mode])</code> | Sets the module in training mode. |
| <code>type(dst_type)</code> | Casts all parameters and buffers to <code>dst_type</code> . |
| <code>xpu([device])</code> | Moves all model parameters and buffers to the XPU. |
| <code>zero_grad([set_to_none])</code> | Resets gradients of all model parameters. |

Attributes

| | |
|---|--|
| <code>T_destination</code> | alias of <code>TypeVar('T_destination', bound=Dict[str, Any])</code> |
| <code>call_super_init</code> | |
| <code>dump_patches</code> | |
| <code>lazy_buffer_names</code> | |
| <code>lazy_parameter_names</code> | |
| <code>lazy_parameters_determined</code> | Returns if all lazy parameters are determined. |

Parameters

- **num_features** (*Optional[int]*) –
- **args** (*Any*) –
- **kwargs** (*Any*) –

affine: bool

eps: float

momentum: float

num_batches_tracked: *Optional[Tensor]*

num_features: *Optional[int]*

running_mean: *Any*

running_var: *Any*

track_running_stats: bool

training: bool

pytorch_pfn_extras.nn.modules.lazy_batchnorm.LazyBatchNorm3d

```
class pytorch_pfn_extras.nn.modules.lazy_batchnorm.LazyBatchNorm3d(num_features, *args,
                                                                    **kwargs)
```

Bases: `_LazyBatchNorm`, `BatchNorm3d`

BatchNorm3d module with lazy weight initialization.

When `num_features` is `None`, it is determined at the first time of the forward step.

Initializes internal Module state, shared by both `nn.Module` and `ScriptModule`.

Methods

| | |
|--|---|
| <code>__init__(num_features, *args, **kwargs)</code> | Initializes internal Module state, shared by both <code>nn.Module</code> and <code>ScriptModule</code> . |
| <code>add_module(name, module)</code> | Adds a child module to the current module. |
| <code>apply(fn)</code> | Applies <code>fn</code> recursively to every submodule (as returned by <code>.children()</code>) as well as self. |
| <code>bfloat16()</code> | Casts all floating point parameters and buffers to <code>bfloat16</code> datatype. |
| <code>buffers([recurse])</code> | Returns an iterator over module buffers. |
| <code>children()</code> | Returns an iterator over immediate children modules. |
| <code>compile(*args, **kwargs)</code> | Compile this Module's forward using <code>torch.compile()</code> . |
| <code>cpu()</code> | Moves all model parameters and buffers to the CPU. |
| <code>cuda([device])</code> | Moves all model parameters and buffers to the GPU. |
| <code>double()</code> | Casts all floating point parameters and buffers to <code>double</code> datatype. |
| <code>eval()</code> | Sets the module in evaluation mode. |
| <code>extra_repr()</code> | Set the extra representation of the module |
| <code>float()</code> | Casts all floating point parameters and buffers to <code>float</code> datatype. |
| <code>forward(input)</code> | Defines the computation performed at every call. |
| <code>get_buffer(target)</code> | Returns the buffer given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_extra_state()</code> | Returns any extra state to include in the module's <code>state_dict</code> . |
| <code>get_parameter(target)</code> | Returns the parameter given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_submodule(target)</code> | Returns the submodule given by <code>target</code> if it exists, otherwise throws an error. |
| <code>half()</code> | Casts all floating point parameters and buffers to <code>half</code> datatype. |
| <code>ipu([device])</code> | Moves all model parameters and buffers to the IPU. |
| <code>load_state_dict(state_dict[, strict, assign])</code> | Copies parameters and buffers from <code>state_dict</code> into this module and its descendants. |
| <code>modules()</code> | Returns an iterator over all modules in the network. |
| <code>named_buffers([prefix, recurse, ...])</code> | Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself. |
| <code>named_children()</code> | Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself. |
| <code>named_modules([memo, prefix, remove_duplicate])</code> | Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself. |
| <code>named_parameters([prefix, recurse, ...])</code> | Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself. |
| <code>parameters([recurse])</code> | Returns an iterator over module parameters. |
| <code>register_backward_hook(hook)</code> | Registers a backward hook on the module. |
| <code>register_buffer(name, tensor[, persistent])</code> | Adds a buffer to the module. |
| <code>register_forward_hook(hook, *[, prepend, ...])</code> | Registers a forward hook on the module. |

continues on next page

Table 16 – continued from previous page

| | |
|---|---|
| <code>register_forward_pre_hook(hook, *[, ...])</code> | Registers a forward pre-hook on the module. |
| <code>register_full_backward_hook(hook[, prepend])</code> | Registers a backward hook on the module. |
| <code>register_full_backward_pre_hook(hook[, prepend])</code> | Registers a backward pre-hook on the module. |
| <code>register_load_state_dict_post_hook(hook)</code> | Registers a post hook to be run after module's <code>load_state_dict</code> is called. |
| <code>register_module(name, module)</code> | Alias for <code>add_module()</code> . |
| <code>register_parameter(name, param)</code> | Adds a parameter to the module. |
| <code>register_state_dict_pre_hook(hook)</code> | These hooks will be called with arguments: <code>self</code> , <code>prefix</code> , and <code>keep_vars</code> before calling <code>state_dict</code> on <code>self</code> . |
| <code>requires_grad_([requires_grad])</code> | Change if autograd should record operations on parameters in this module. |
| <code>reset_parameters()</code> | |
| <code>reset_running_stats()</code> | |
| <code>set_extra_state(state)</code> | This function is called from <code>load_state_dict()</code> to handle any extra state found within the <code>state_dict</code> . |
| <code>share_memory()</code> | See <code>torch.Tensor.share_memory_()</code> |
| <code>state_dict(*args, **kwargs)</code> | Returns a dictionary containing a whole state of the module. |
| <code>to(*args, **kwargs)</code> | Moves and/or casts the parameters and buffers. |
| <code>to_empty(*, device[, recurse])</code> | Moves the parameters and buffers to the specified device without copying storage. |
| <code>train([mode])</code> | Sets the module in training mode. |
| <code>type(dst_type)</code> | Casts all parameters and buffers to <code>dst_type</code> . |
| <code>xpu([device])</code> | Moves all model parameters and buffers to the XPU. |
| <code>zero_grad([set_to_none])</code> | Resets gradients of all model parameters. |

Attributes

| | |
|---|--|
| <code>T_destination</code> | alias of <code>TypeVar('T_destination', bound=Dict[str, Any])</code> |
| <code>call_super_init</code> | |
| <code>dump_patches</code> | |
| <code>lazy_buffer_names</code> | |
| <code>lazy_parameter_names</code> | |
| <code>lazy_parameters_determined</code> | Returns if all lazy parameters are determined. |

Parameters

- `num_features` (*Optional[int]*) –
- `args` (*Any*) –
- `kwargs` (*Any*) –

```
affine: bool
eps: float
momentum: float
num_batches_tracked: Optional[Tensor]
num_features: Optional[int]
running_mean: Any
running_var: Any
track_running_stats: bool
training: bool
```

pytorch_pfn_extras.nn.modules.lazy_batchnorm.LazyInitializationMixin

```
class pytorch_pfn_extras.nn.modules.lazy_batchnorm.LazyInitializationMixin(*args, **kwargs)
```

Bases: object

A mixin for modules that lazily initialize buffers and parameters.

Unlike regular modules, subclasses of this module can initialize buffers and parameters outside of the constructor (`__init__`). This allows you to, for example, initialize parameters in `forward` method to determine the shape of the weight based on the initial input.

Be sure to run “dummy” forward once to initialize all parameters that should be trained, before passing `module.parameters()` to an optimizer; otherwise weights initialized after `module.parameters()` (e.g., in `forward` function) will never be trained.

Note that lazy modules cannot validate if the shape is correct during deserialization. Also note that the initial weights may become different from the original (non-lazy) module even if the random seed is manually configured, as the order of initialization is different from the original one; especially, `module.cuda()` may cause the initialization to run on a GPU.

The default value of lazy buffers and parameters are `torch.Tensor([])` and `UninitializedParameter()`, respectively.

Methods

`__init__`(*args, **kwargs)

`state_dict`(*args, **kwargs)

Returns a dictionary containing a whole state of the module.

Attributes

lazy_buffer_names

lazy_parameter_names

lazy_parameters_determined Returns if all lazy parameters are determined.

Parameters

- **args** (*Any*) –
- **kwargs** (*Any*) –

__init__(*args, **kwargs)

Parameters

- **self** (*Any*) –
- **args** (*Any*) –
- **kwargs** (*Any*) –

Return type

None

lazy_buffer_names: `Tuple[str, ...] = ()`

lazy_parameter_names: `Tuple[str, ...] = ()`

property lazy_parameters_determined: `bool`

Returns if all lazy parameters are determined.

Subclasses can perform parameters initialization after all lazy parameters are determined. Note that this may be called during `__init__`.

state_dict(*args, **kwargs)

Returns a dictionary containing a whole state of the module.

This function overrides the default behavior to exclude uninitialized parameter from serialization. This is needed because we need to discriminate lazy parameters (`UninitializedParameter()`) and initialized empty parameters (`torch.nn.Parameter(torch.Tensor())`) during deserialization.

See comments of `_lazy_load_hook` for details.

Parameters

- **self** (*Any*) –
- **args** (*Any*) –
- **kwargs** (*Any*) –

Return type

`Dict[str, Any]`

pytorch_pfn_extras.nn.modules.lazy_batchnorm.UninitializedParameter

class pytorch_pfn_extras.nn.modules.lazy_batchnorm.**UninitializedParameter**(*data=None, requires_grad=True*)

Bases: Parameter

Methods

| | |
|---|--|
| <code>__init__()</code> | |
| <code>abs()</code> | See <code>torch.abs()</code> |
| <code>abs_()</code> | In-place version of <code>abs()</code> |
| <code>absolute()</code> | Alias for <code>abs()</code> |
| <code>absolute_()</code> | In-place version of <code>absolute()</code> Alias for <code>abs_()</code> |
| <code>acos()</code> | See <code>torch.acos()</code> |
| <code>acos_()</code> | In-place version of <code>acos()</code> |
| <code>acosh()</code> | See <code>torch.acosh()</code> |
| <code>acosh_()</code> | In-place version of <code>acosh()</code> |
| <code>add(other, *, alpha)</code> | Add a scalar or tensor to <code>self</code> tensor. |
| <code>add_(other, *, alpha)</code> | In-place version of <code>add()</code> |
| <code>addbmm(batch1, batch2, *, beta, alpha)</code> | See <code>torch.addbmm()</code> |
| <code>addbmm_(batch1, batch2, *, beta, alpha)</code> | In-place version of <code>addbmm()</code> |
| <code>addcdiv(tensor1, tensor2, *, value)</code> | See <code>torch.addcdiv()</code> |
| <code>addcdiv_(tensor1, tensor2, *, value)</code> | In-place version of <code>addcdiv()</code> |
| <code>addcmul(tensor1, tensor2, *, value)</code> | See <code>torch.addcmul()</code> |
| <code>addcmul_(tensor1, tensor2, *, value)</code> | In-place version of <code>addcmul()</code> |
| <code>addmm(mat1, mat2, *, beta, alpha)</code> | See <code>torch.addmm()</code> |
| <code>addmm_(mat1, mat2, *, beta, alpha)</code> | In-place version of <code>addmm()</code> |
| <code>addmv(mat, vec, *, beta, alpha)</code> | See <code>torch.addmv()</code> |
| <code>addmv_(mat, vec, *, beta, alpha)</code> | In-place version of <code>addmv()</code> |
| <code>addr(vec1, vec2, *, beta, alpha)</code> | See <code>torch.addr()</code> |
| <code>addr_(vec1, vec2, *, beta, alpha)</code> | In-place version of <code>addr()</code> |
| <code>adjoint()</code> | Alias for <code>adjoint()</code> |
| <code>align_as(other)</code> | Permutates the dimensions of the <code>self</code> tensor to match the dimension order in the <code>other</code> tensor, adding size-one dims for any new names. |
| <code>align_to(*names)</code> | Permutates the dimensions of the <code>self</code> tensor to match the order specified in <code>names</code> , adding size-one dims for any new names. |
| <code>all([dim, keepdim])</code> | See <code>torch.all()</code> |
| <code>allclose(other[, rtol, atol, equal_nan])</code> | See <code>torch.allclose()</code> |
| <code>amax([dim, keepdim])</code> | See <code>torch.amax()</code> |
| <code>amin([dim, keepdim])</code> | See <code>torch.amin()</code> |
| <code>aminmax(*[, dim, keepdim])</code> | See <code>torch.aminmax()</code> |
| <code>angle()</code> | See <code>torch.angle()</code> |
| <code>any([dim, keepdim])</code> | See <code>torch.any()</code> |
| <code>apply_(callable)</code> | Applies the function <code>callable</code> to each element in the tensor, replacing each element with the value returned by <code>callable</code> . |
| <code>arccos()</code> | See <code>torch.arccos()</code> |

continues on next page

Table 17 – continued from previous page

| | |
|---|---|
| <code>arccos_()</code> | In-place version of <code>arccos()</code> |
| <code>arccosh</code> | <code>acosh()</code> -> Tensor |
| <code>arccosh_</code> | <code>acosh_()</code> -> Tensor |
| <code>arcsin()</code> | See <code>torch.arcsin()</code> |
| <code>arcsin_()</code> | In-place version of <code>arcsin()</code> |
| <code>arcsinh()</code> | See <code>torch.arcsinh()</code> |
| <code>arcsinh_()</code> | In-place version of <code>arcsinh()</code> |
| <code>arctan()</code> | See <code>torch.arctan()</code> |
| <code>arctan2(other)</code> | See <code>torch.arctan2()</code> |
| <code>arctan2_</code> | <code>atan2_(other)</code> -> Tensor |
| <code>arctan_()</code> | In-place version of <code>arctan()</code> |
| <code>arctanh()</code> | See <code>torch.arctanh()</code> |
| <code>arctanh_(other)</code> | In-place version of <code>arctanh()</code> |
| <code>argmax([dim, keepdim])</code> | See <code>torch.argmax()</code> |
| <code>argmin([dim, keepdim])</code> | See <code>torch.argmin()</code> |
| <code>argsort([dim, descending])</code> | See <code>torch.argsort()</code> |
| <code>argwhere()</code> | See <code>torch.argwhere()</code> |
| <code>as_strided(size, stride[, storage_offset])</code> | See <code>torch.as_strided()</code> |
| <code>as_strided_(size, stride[, storage_offset])</code> | In-place version of <code>as_strided()</code> |
| <code>as_strided_scatter(src, size, stride[, ...])</code> | See <code>torch.as_strided_scatter()</code> |
| <code>as_subclass(cls)</code> | Makes a <code>cls</code> instance with the same data pointer as <code>self</code> . |
| <code>asin()</code> | See <code>torch.asin()</code> |
| <code>asin_()</code> | In-place version of <code>asin()</code> |
| <code>asinh()</code> | See <code>torch.asinh()</code> |
| <code>asinh_()</code> | In-place version of <code>asinh()</code> |
| <code>atan()</code> | See <code>torch.atan()</code> |
| <code>atan2(other)</code> | See <code>torch.atan2()</code> |
| <code>atan2_(other)</code> | In-place version of <code>atan2()</code> |
| <code>atan_()</code> | In-place version of <code>atan()</code> |
| <code>atanh()</code> | See <code>torch.atanh()</code> |
| <code>atanh_(other)</code> | In-place version of <code>atanh()</code> |
| <code>backward([gradient, retain_graph, ...])</code> | Computes the gradient of current tensor wrt graph leaves. |
| <code>baddbmm(batch1, batch2, *[, beta, alpha])</code> | See <code>torch.baddbmm()</code> |
| <code>baddbmm_(batch1, batch2, *[, beta, alpha])</code> | In-place version of <code>baddbmm()</code> |
| <code>bernoulli(*[, generator])</code> | Returns a result tensor where each <code>result[i]</code> is independently sampled from <code>Bernoulli(self[i])</code> . |
| <code>bernoulli_(p, generator)</code> | Fills each location of <code>self</code> with an independent sample from <code>Bernoulli(p)</code> . |
| <code>bfloat16([memory_format])</code> | <code>self.bfloat16()</code> is equivalent to <code>self.to(torch.bfloat16)</code> . |
| <code>bincount([weights, minlength])</code> | See <code>torch.bincount()</code> |
| <code>bitwise_and()</code> | See <code>torch.bitwise_and()</code> |
| <code>bitwise_and_()</code> | In-place version of <code>bitwise_and()</code> |
| <code>bitwise_left_shift(other)</code> | See <code>torch.bitwise_left_shift()</code> |
| <code>bitwise_left_shift_(other)</code> | In-place version of <code>bitwise_left_shift()</code> |
| <code>bitwise_not()</code> | See <code>torch.bitwise_not()</code> |
| <code>bitwise_not_()</code> | In-place version of <code>bitwise_not()</code> |
| <code>bitwise_or()</code> | See <code>torch.bitwise_or()</code> |

continues on next page

Table 17 – continued from previous page

| | |
|--|--|
| <code>bitwise_or_()</code> | In-place version of <code>bitwise_or()</code> |
| <code>bitwise_right_shift(other)</code> | See <code>torch.bitwise_right_shift()</code> |
| <code>bitwise_right_shift_(other)</code> | In-place version of <code>bitwise_right_shift()</code> |
| <code>bitwise_xor()</code> | See <code>torch.bitwise_xor()</code> |
| <code>bitwise_xor_()</code> | In-place version of <code>bitwise_xor()</code> |
| <code>bmm(batch2)</code> | See <code>torch.bmm()</code> |
| <code>bool([memory_format])</code> | <code>self.bool()</code> is equivalent to <code>self.to(torch.bool)</code> . |
| <code>broadcast_to(shape)</code> | See <code>torch.broadcast_to()</code> . |
| <code>byte([memory_format])</code> | <code>self.byte()</code> is equivalent to <code>self.to(torch.uint8)</code> . |
| <code>cauchy_([median, sigma, generator])</code> | Fills the tensor with numbers drawn from the Cauchy distribution: |
| <code>ccol_indices</code> | |
| <code>cdouble([memory_format])</code> | <code>self.cdouble()</code> is equivalent to <code>self.to(torch.complex128)</code> . |
| <code>ceil()</code> | See <code>torch.ceil()</code> |
| <code>ceil_()</code> | In-place version of <code>ceil()</code> |
| <code>cfloat([memory_format])</code> | <code>self.cfloat()</code> is equivalent to <code>self.to(torch.complex64)</code> . |
| <code>chalf([memory_format])</code> | <code>self.chalf()</code> is equivalent to <code>self.to(torch.complex32)</code> . |
| <code>char([memory_format])</code> | <code>self.char()</code> is equivalent to <code>self.to(torch.int8)</code> . |
| <code>cholesky([upper])</code> | See <code>torch.cholesky()</code> |
| <code>cholesky_inverse([upper])</code> | See <code>torch.cholesky_inverse()</code> |
| <code>cholesky_solve(input2[, upper])</code> | See <code>torch.cholesky_solve()</code> |
| <code>chunk(chunks[, dim])</code> | See <code>torch.chunk()</code> |
| <code>clamp([min, max])</code> | See <code>torch.clamp()</code> |
| <code>clamp_([min, max])</code> | In-place version of <code>clamp()</code> |
| <code>clamp_max</code> | |
| <code>clamp_max_</code> | |
| <code>clamp_min</code> | |
| <code>clamp_min_</code> | |
| <code>clip([min, max])</code> | Alias for <code>clamp()</code> . |
| <code>clip_([min, max])</code> | Alias for <code>clamp_()</code> . |
| <code>clone(*[, memory_format])</code> | See <code>torch.clone()</code> |
| <code>coalesce()</code> | Returns a coalesced copy of <code>self</code> if <code>self</code> is an uncoalesced tensor. |
| <code>col_indices()</code> | Returns the tensor containing the column indices of the <code>self</code> tensor when <code>self</code> is a sparse CSR tensor of layout <code>sparse_csr</code> . |
| <code>conj()</code> | See <code>torch.conj()</code> |
| <code>conj_physical()</code> | See <code>torch.conj_physical()</code> |
| <code>conj_physical_()</code> | In-place version of <code>conj_physical()</code> |

continues on next page

Table 17 – continued from previous page

| | |
|--|--|
| <code>contiguous([memory_format])</code> | Returns a contiguous in memory tensor containing the same data as <code>self</code> tensor. |
| <code>copy_(src[, non_blocking])</code> | Copies the elements from <code>src</code> into <code>self</code> tensor and returns <code>self</code> . |
| <code>copysign(other)</code> | See <code>torch.copysign()</code> |
| <code>copysign_(other)</code> | In-place version of <code>copysign()</code> |
| <code>corrcoef()</code> | See <code>torch.corrcoef()</code> |
| <code>cos()</code> | See <code>torch.cos()</code> |
| <code>cos_()</code> | In-place version of <code>cos()</code> |
| <code>cosh()</code> | See <code>torch.cosh()</code> |
| <code>cosh_()</code> | In-place version of <code>cosh()</code> |
| <code>count_nonzero([dim])</code> | See <code>torch.count_nonzero()</code> |
| <code>cov(*[, correction, fweights, aweights])</code> | See <code>torch.cov()</code> |
| <code>cpu([memory_format])</code> | Returns a copy of this object in CPU memory. |
| <code>cross(other[, dim])</code> | See <code>torch.cross()</code> |
| <code>crow_indices()</code> | Returns the tensor containing the compressed row indices of the <code>self</code> tensor when <code>self</code> is a sparse CSR tensor of layout <code>sparse_csr</code> . |
| <code>cuda([device, non_blocking, memory_format])</code> | Returns a copy of this object in CUDA memory. |
| <code>cummax(dim)</code> | See <code>torch.cummax()</code> |
| <code>cummin(dim)</code> | See <code>torch.cummin()</code> |
| <code>cumprod(dim[, dtype])</code> | See <code>torch.cumprod()</code> |
| <code>cumprod_(dim[, dtype])</code> | In-place version of <code>cumprod()</code> |
| <code>cumsum(dim[, dtype])</code> | See <code>torch.cumsum()</code> |
| <code>cumsum_(dim[, dtype])</code> | In-place version of <code>cumsum()</code> |
| <code>data_ptr()</code> | Returns the address of the first element of <code>self</code> tensor. |
| <code>deg2rad()</code> | See <code>torch.deg2rad()</code> |
| <code>deg2rad_()</code> | In-place version of <code>deg2rad()</code> |
| <code>dense_dim()</code> | Return the number of dense dimensions in a sparse tensor <code>self</code> . |
| <code>dequantize()</code> | Given a quantized Tensor, dequantize it and return the dequantized float Tensor. |
| <code>det()</code> | See <code>torch.det()</code> |
| <code>detach</code> | Returns a new Tensor, detached from the current graph. |
| <code>detach_</code> | Detaches the Tensor from the graph that created it, making it a leaf. |
| <code>diag([diagonal])</code> | See <code>torch.diag()</code> |
| <code>diag_embed([offset, dim1, dim2])</code> | See <code>torch.diag_embed()</code> |
| <code>diagflat([offset])</code> | See <code>torch.diagflat()</code> |
| <code>diagonal([offset, dim1, dim2])</code> | See <code>torch.diagonal()</code> |
| <code>diagonal_scatter(src[, offset, dim1, dim2])</code> | See <code>torch.diagonal_scatter()</code> |
| <code>diff([n, dim, prepend, append])</code> | See <code>torch.diff()</code> |
| <code>digamma()</code> | See <code>torch.digamma()</code> |
| <code>digamma_()</code> | In-place version of <code>digamma()</code> |
| <code>dim()</code> | Returns the number of dimensions of <code>self</code> tensor. |
| <code>dim_order()</code> | Returns a tuple of int describing the dim order or physical layout of <code>self</code> . |
| <code>dist(other[, p])</code> | See <code>torch.dist()</code> |
| <code>div(value, *[, rounding_mode])</code> | See <code>torch.div()</code> |

continues on next page

Table 17 – continued from previous page

| | |
|---|---|
| <code>div_(value, *, rounding_mode)</code> | In-place version of <code>div()</code> |
| <code>divide(value, *, rounding_mode)</code> | See <code>torch.divide()</code> |
| <code>divide_(value, *, rounding_mode)</code> | In-place version of <code>divide()</code> |
| <code>dot(other)</code> | See <code>torch.dot()</code> |
| <code>double([memory_format])</code> | <code>self.double()</code> is equivalent to <code>self.to(torch.float64)</code> . |
| <code>dsplit(split_size_or_sections)</code> | See <code>torch.dsplit()</code> |
| <code>eig([eigenvectors])</code> | |
| <code>element_size()</code> | Returns the size in bytes of an individual element. |
| <code>eq(other)</code> | See <code>torch.eq()</code> |
| <code>eq_(other)</code> | In-place version of <code>eq()</code> |
| <code>equal(other)</code> | See <code>torch.equal()</code> |
| <code>erf()</code> | See <code>torch.erf()</code> |
| <code>erf_()</code> | In-place version of <code>erf()</code> |
| <code>erfc()</code> | See <code>torch.erfc()</code> |
| <code>erfc_()</code> | In-place version of <code>erfc()</code> |
| <code>erfinv()</code> | See <code>torch.erfinv()</code> |
| <code>erfinv_()</code> | In-place version of <code>erfinv()</code> |
| <code>exp()</code> | See <code>torch.exp()</code> |
| <code>exp2()</code> | See <code>torch.exp2()</code> |
| <code>exp2_()</code> | In-place version of <code>exp2()</code> |
| <code>exp_()</code> | In-place version of <code>exp()</code> |
| <code>expand(*sizes)</code> | Returns a new view of the <code>self</code> tensor with singleton dimensions expanded to a larger size. |
| <code>expand_as(other)</code> | Expand this tensor to the same size as <code>other</code> . |
| <code>expm1()</code> | See <code>torch.expm1()</code> |
| <code>expm1_()</code> | In-place version of <code>expm1()</code> |
| <code>exponential_([lambd, generator])</code> | Fills <code>self</code> tensor with elements drawn from the exponential distribution: |
| <code>fill_(value)</code> | Fills <code>self</code> tensor with the specified value. |
| <code>fill_diagonal_(fill_value[, wrap])</code> | Fill the main diagonal of a tensor that has at least 2-dimensions. |
| <code>fix()</code> | See <code>torch.fix()</code> . |
| <code>fix_()</code> | In-place version of <code>fix()</code> |
| <code>flatten([start_dim, end_dim])</code> | See <code>torch.flatten()</code> |
| <code>flip(dims)</code> | See <code>torch.flip()</code> |
| <code>flipud()</code> | See <code>torch.flipud()</code> |
| <code>float([memory_format])</code> | <code>self.float()</code> is equivalent to <code>self.to(torch.float32)</code> . |
| <code>float_power(exponent)</code> | See <code>torch.float_power()</code> |
| <code>float_power_(exponent)</code> | In-place version of <code>float_power()</code> |
| <code>floor()</code> | See <code>torch.floor()</code> |
| <code>floor_()</code> | In-place version of <code>floor()</code> |
| <code>floor_divide(value)</code> | See <code>torch.floor_divide()</code> |
| <code>floor_divide_(value)</code> | In-place version of <code>floor_divide()</code> |
| <code>fmax(other)</code> | See <code>torch.fmax()</code> |
| <code>fmin(other)</code> | See <code>torch.fmin()</code> |
| <code>fmod(divisor)</code> | See <code>torch.fmod()</code> |

continues on next page

Table 17 – continued from previous page

| | |
|--|---|
| <code>fmod_(divisor)</code> | In-place version of <code>fmod()</code> |
| <code>frac()</code> | See <code>torch.frac()</code> |
| <code>frac_()</code> | In-place version of <code>frac()</code> |
| <code>frexp(input)</code> | See <code>torch.frexp()</code> |
| <code>gather(dim, index)</code> | See <code>torch.gather()</code> |
| <code>gcd(other)</code> | See <code>torch.gcd()</code> |
| <code>gcd_(other)</code> | In-place version of <code>gcd()</code> |
| <code>ge(other)</code> | See <code>torch.ge()</code> . |
| <code>ge_(other)</code> | In-place version of <code>ge()</code> . |
| <code>geometric_(p, *[, generator])</code> | Fills <code>self</code> tensor with elements drawn from the geometric distribution: |
| <code>geqrf()</code> | See <code>torch.geqrf()</code> |
| <code>ger(vec2)</code> | See <code>torch.ger()</code> |
| <code>get_device()</code> | For CUDA tensors, this function returns the device ordinal of the GPU on which the tensor resides. |
| <code>greater(other)</code> | See <code>torch.greater()</code> . |
| <code>greater_(other)</code> | In-place version of <code>greater()</code> . |
| <code>greater_equal(other)</code> | See <code>torch.greater_equal()</code> . |
| <code>greater_equal_(other)</code> | In-place version of <code>greater_equal()</code> . |
| <code>gt(other)</code> | See <code>torch.gt()</code> . |
| <code>gt_(other)</code> | In-place version of <code>gt()</code> . |
| <code>half([memory_format])</code> | <code>self.half()</code> is equivalent to <code>self.to(torch.float16)</code> . |
| <code>hardshrink([lambd])</code> | See <code>torch.nn.functional.hardshrink()</code> |
| <code>has_names</code> | Is True if any of this tensor's dimensions are named. |
| <code>heaviside(values)</code> | See <code>torch.heaviside()</code> |
| <code>heaviside_(values)</code> | In-place version of <code>heaviside()</code> |
| <code>histc([bins, min, max])</code> | See <code>torch.histc()</code> |
| <code>histogram(input, bins, *[, range, weight, ...])</code> | See <code>torch.histogram()</code> |
| <code>hsplit(split_size_or_sections)</code> | See <code>torch.hsplit()</code> |
| <code>hypot(other)</code> | See <code>torch.hypot()</code> |
| <code>hypot_(other)</code> | In-place version of <code>hypot()</code> |
| <code>i0()</code> | See <code>torch.i0()</code> |
| <code>i0_()</code> | In-place version of <code>i0()</code> |
| <code>igamma(other)</code> | See <code>torch.igamma()</code> |
| <code>igamma_(other)</code> | In-place version of <code>igamma()</code> |
| <code>igammac(other)</code> | See <code>torch.igammac()</code> |
| <code>igammac_(other)</code> | In-place version of <code>igammac()</code> |
| <code>index_add(dim, index, source, *[, alpha])</code> | Out-of-place version of <code>torch.Tensor.index_add_()</code> . |
| <code>index_add_(dim, index, source, *[, alpha])</code> | Accumulate the elements of <code>alpha</code> times <code>source</code> into the <code>self</code> tensor by adding to the indices in the order given in <code>index</code> . |
| <code>index_copy(dim, index, tensor2)</code> | Out-of-place version of <code>torch.Tensor.index_copy_()</code> . |
| <code>index_copy_(dim, index, tensor)</code> | Copies the elements of <code>tensor</code> into the <code>self</code> tensor by selecting the indices in the order given in <code>index</code> . |
| <code>index_fill(dim, index, value)</code> | Out-of-place version of <code>torch.Tensor.index_fill_()</code> . |

continues on next page

Table 17 – continued from previous page

| | |
|---|---|
| <code>index_fill_(dim, index, value)</code> | Fills the elements of the <code>self</code> tensor with value <code>value</code> by selecting the indices in the order given in <code>index</code> . |
| <code>index_put(indices, values[, accumulate])</code> | Out-place version of <code>index_put_()</code> . |
| <code>index_put_(indices, values[, accumulate])</code> | Puts values from the tensor <code>values</code> into the tensor <code>self</code> using the indices specified in <code>indices</code> (which is a tuple of Tensors). |
| <code>index_reduce</code> | |
| <code>index_reduce_(dim, index, source, reduce, *)</code> | Accumulate the elements of <code>source</code> into the <code>self</code> tensor by accumulating to the indices in the order given in <code>index</code> using the reduction given by the <code>reduce</code> argument. |
| <code>index_select(dim, index)</code> | See <code>torch.index_select()</code> |
| <code>indices()</code> | Return the indices tensor of a sparse COO tensor. |
| <code>inner(other)</code> | See <code>torch.inner()</code> . |
| <code>int([memory_format])</code> | <code>self.int()</code> is equivalent to <code>self.to(torch.int32)</code> . |
| <code>int_repr()</code> | Given a quantized Tensor, <code>self.int_repr()</code> returns a CPU Tensor with <code>uint8_t</code> as data type that stores the underlying <code>uint8_t</code> values of the given Tensor. |
| <code>inverse()</code> | See <code>torch.inverse()</code> |
| <code>ipu([device, non_blocking, memory_format])</code> | Returns a copy of this object in IPU memory. |
| <code>is_coalesced()</code> | Returns True if <code>self</code> is a sparse COO tensor that is coalesced, False otherwise. |
| <code>is_complex()</code> | Returns True if the data type of <code>self</code> is a complex data type. |
| <code>is_conj()</code> | Returns True if the conjugate bit of <code>self</code> is set to true. |
| <code>is_contiguous([memory_format])</code> | Returns True if <code>self</code> tensor is contiguous in memory in the order specified by memory format. |
| <code>is_distributed</code> | |
| <code>is_floating_point()</code> | Returns True if the data type of <code>self</code> is a floating point data type. |
| <code>is_inference()</code> | See <code>torch.is_inference()</code> |
| <code>is_neg()</code> | Returns True if the negative bit of <code>self</code> is set to true. |
| <code>is_nonzero</code> | |
| <code>is_pinned</code> | Returns true if this tensor resides in pinned memory. |
| <code>is_same_size</code> | |
| <code>is_set_to(tensor)</code> | Returns True if both tensors are pointing to the exact same memory (same storage, offset, size and stride). |
| <code>is_shared()</code> | Checks if tensor is in shared memory. |
| <code>is_signed()</code> | Returns True if the data type of <code>self</code> is a signed data type. |
| <code>isclose(other[, rtol, atol, equal_nan])</code> | See <code>torch.isclose()</code> |
| <code>isfinite()</code> | See <code>torch.isfinite()</code> |
| <code>isinf()</code> | See <code>torch.isinf()</code> |
| <code>isnan()</code> | See <code>torch.isnan()</code> |

continues on next page

Table 17 – continued from previous page

| | |
|--|--|
| <code>isneginf()</code> | See <code>torch.isneginf()</code> |
| <code>isposinf()</code> | See <code>torch.isposinf()</code> |
| <code>isreal()</code> | See <code>torch.isreal()</code> |
| <code>istft(n_fft[, hop_length, win_length, ...])</code> | See <code>torch.istft()</code> |
| <code>item()</code> | Returns the value of this tensor as a standard Python number. |
| <code>kron(other)</code> | See <code>torch.kron()</code> |
| <code>kthvalue(k[, dim, keepdim])</code> | See <code>torch.kthvalue()</code> |
| <code>lcm(other)</code> | See <code>torch.lcm()</code> |
| <code>lcm_(other)</code> | In-place version of <code>lcm()</code> |
| <code>ldexp(other)</code> | See <code>torch.ldexp()</code> |
| <code>ldexp_(other)</code> | In-place version of <code>ldexp()</code> |
| <code>le(other)</code> | See <code>torch.le()</code> . |
| <code>le_(other)</code> | In-place version of <code>le()</code> . |
| <code>lerp(end, weight)</code> | See <code>torch.lerp()</code> |
| <code>lerp_(end, weight)</code> | In-place version of <code>lerp()</code> |
| <code>less</code> | <code>lt(other) -> Tensor</code> |
| <code>less_(other)</code> | In-place version of <code>less()</code> . |
| <code>less_equal(other)</code> | See <code>torch.less_equal()</code> . |
| <code>less_equal_(other)</code> | In-place version of <code>less_equal()</code> . |
| <code>lgamma()</code> | See <code>torch.lgamma()</code> |
| <code>lgamma_()</code> | In-place version of <code>lgamma()</code> |
| <code>log()</code> | See <code>torch.log()</code> |
| <code>log10()</code> | See <code>torch.log10()</code> |
| <code>log10_()</code> | In-place version of <code>log10()</code> |
| <code>log1p()</code> | See <code>torch.log1p()</code> |
| <code>log1p_()</code> | In-place version of <code>log1p()</code> |
| <code>log2()</code> | See <code>torch.log2()</code> |
| <code>log2_()</code> | In-place version of <code>log2()</code> |
| <code>log_()</code> | In-place version of <code>log()</code> |
| <code>log_normal_([mean, std, generator])</code> | Fills <code>self</code> tensor with numbers samples from the log-normal distribution parameterized by the given mean μ and standard deviation σ . |
| log_softmax | |
| <code>logaddexp(other)</code> | See <code>torch.logaddexp()</code> |
| <code>logaddexp2(other)</code> | See <code>torch.logaddexp2()</code> |
| <code>logcumsumexp(dim)</code> | See <code>torch.logcumsumexp()</code> |
| <code>logdet()</code> | See <code>torch.logdet()</code> |
| <code>logical_and()</code> | See <code>torch.logical_and()</code> |
| <code>logical_and_()</code> | In-place version of <code>logical_and()</code> |
| <code>logical_not()</code> | See <code>torch.logical_not()</code> |
| <code>logical_not_()</code> | In-place version of <code>logical_not()</code> |
| <code>logical_or()</code> | See <code>torch.logical_or()</code> |
| <code>logical_or_()</code> | In-place version of <code>logical_or()</code> |
| <code>logical_xor()</code> | See <code>torch.logical_xor()</code> |
| <code>logical_xor_()</code> | In-place version of <code>logical_xor()</code> |
| <code>logit()</code> | See <code>torch.logit()</code> |
| <code>logit_()</code> | In-place version of <code>logit()</code> |
| <code>logsumexp(dim[, keepdim])</code> | See <code>torch.logsumexp()</code> |

continues on next page

Table 17 – continued from previous page

| | |
|---|---|
| <code>long([memory_format])</code> | <code>self.long()</code> is equivalent to <code>self.to(torch.int64)</code> . |
| <code>lstsq(other)</code> | |
| <code>lt(other)</code> | See <code>torch.lt()</code> . |
| <code>lt_(other)</code> | In-place version of <code>lt()</code> . |
| <code>lu([pivot, get_infos])</code> | See <code>torch.lu()</code> |
| <code>lu_solve(LU_data, LU_pivots)</code> | See <code>torch.lu_solve()</code> |
| <code>map2_</code> | |
| <code>map_(tensor, callable)</code> | Applies callable for each element in <code>self</code> tensor and the given <code>tensor</code> and stores the results in <code>self</code> tensor. |
| <code>masked_fill(mask, value)</code> | Out-of-place version of <code>torch.Tensor.masked_fill_()</code> |
| <code>masked_fill_(mask, value)</code> | Fills elements of <code>self</code> tensor with <code>value</code> where <code>mask</code> is <code>True</code> . |
| <code>masked_scatter(mask, tensor)</code> | Out-of-place version of <code>torch.Tensor.masked_scatter_()</code> |
| <code>masked_scatter_(mask, source)</code> | Copies elements from <code>source</code> into <code>self</code> tensor at positions where the <code>mask</code> is <code>True</code> . |
| <code>masked_select(mask)</code> | See <code>torch.masked_select()</code> |
| <code>materialize(shape[, device, dtype])</code> | Create a <code>Parameter</code> with the same properties of the uninitialized one. |
| <code>matmul(tensor2)</code> | See <code>torch.matmul()</code> |
| <code>matrix_exp()</code> | See <code>torch.matrix_exp()</code> |
| <code>matrix_power(n)</code> | |
| Note: <code>matrix_power()</code> is deprecated, use <code>torch.linalg.matrix_power()</code> instead. | |
| <code>max([dim, keepdim])</code> | See <code>torch.max()</code> |
| <code>maximum(other)</code> | See <code>torch.maximum()</code> |
| <code>mean([dim, keepdim, dtype])</code> | See <code>torch.mean()</code> |
| <code>median([dim, keepdim])</code> | See <code>torch.median()</code> |
| <code>min([dim, keepdim])</code> | See <code>torch.min()</code> |
| <code>minimum(other)</code> | See <code>torch.minimum()</code> |
| <code>mm(mat2)</code> | See <code>torch.mm()</code> |
| <code>mode([dim, keepdim])</code> | See <code>torch.mode()</code> |
| <code>moveaxis(source, destination)</code> | See <code>torch.moveaxis()</code> |
| <code>movedim(source, destination)</code> | See <code>torch.movedim()</code> |
| <code>msort()</code> | See <code>torch.msort()</code> |
| <code>mul(value)</code> | See <code>torch.mul()</code> . |
| <code>mul_(value)</code> | In-place version of <code>mul()</code> . |
| <code>multinomial(num_samples[, replacement, ...])</code> | See <code>torch.multinomial()</code> |
| <code>multiply(value)</code> | See <code>torch.multiply()</code> . |
| <code>multiply_(value)</code> | In-place version of <code>multiply()</code> . |
| <code>mv(vec)</code> | See <code>torch.mv()</code> |
| <code>mvlgamma(p)</code> | See <code>torch.mvlgamma()</code> |
| <code>mvlgamma_(p)</code> | In-place version of <code>mvlgamma()</code> |

continues on next page

Table 17 – continued from previous page

| | |
|--|---|
| <code>nan_to_num([nan, posinf, neginf])</code> | See <code>torch.nan_to_num()</code> . |
| <code>nan_to_num_([nan, posinf, neginf])</code> | In-place version of <code>nan_to_num()</code> . |
| <code>nanmean([dim, keepdim, dtype])</code> | See <code>torch.nanmean()</code> |
| <code>nanmedian([dim, keepdim])</code> | See <code>torch.nanmedian()</code> |
| <code>nanquantile(q[, dim, keepdim, interpolation])</code> | See <code>torch.nanquantile()</code> |
| <code>nansum([dim, keepdim, dtype])</code> | See <code>torch.nansum()</code> |
| <code>narrow(dimension, start, length)</code> | See <code>torch.narrow()</code> . |
| <code>narrow_copy(dimension, start, length)</code> | See <code>torch.narrow_copy()</code> . |
| <code>ndimension()</code> | Alias for <code>dim()</code> |
| <code>ne(other)</code> | See <code>torch.ne()</code> . |
| <code>ne_(other)</code> | In-place version of <code>ne()</code> . |
| <code>neg()</code> | See <code>torch.neg()</code> |
| <code>neg_()</code> | In-place version of <code>neg()</code> |
| <code>negative()</code> | See <code>torch.negative()</code> |
| <code>negative_()</code> | In-place version of <code>negative()</code> |
| <code>nelement()</code> | Alias for <code>numel()</code> |
| <code>new</code> | |
| <code>new_empty(size, *[, dtype, device, ...])</code> | Returns a Tensor of size <code>size</code> filled with uninitialized data. |
| <code>new_empty_strided(size, stride[, dtype, ...])</code> | Returns a Tensor of size <code>size</code> and strides <code>stride</code> filled with uninitialized data. |
| <code>new_full(size, fill_value, *[, dtype, ...])</code> | Returns a Tensor of size <code>size</code> filled with <code>fill_value</code> . |
| <code>new_ones(size, *[, dtype, device, ...])</code> | Returns a Tensor of size <code>size</code> filled with 1. |
| <code>new_tensor(data, *[, dtype, device, ...])</code> | Returns a new Tensor with data as the tensor data. |
| <code>new_zeros(size, *[, dtype, device, ...])</code> | Returns a Tensor of size <code>size</code> filled with 0. |
| <code>nextafter(other)</code> | See <code>torch.nextafter()</code> |
| <code>nextafter_(other)</code> | In-place version of <code>nextafter()</code> |
| <code>nonzero()</code> | See <code>torch.nonzero()</code> |
| <code>nonzero_static(input, *, size[, fill_value])</code> | Returns a 2-D tensor where each row is the index for a non-zero value. |
| <code>norm([p, dim, keepdim, dtype])</code> | See <code>torch.norm()</code> |
| <code>normal_([mean, std, generator])</code> | Fills <code>self</code> tensor with elements samples from the normal distribution parameterized by <code>mean</code> and <code>std</code> . |
| <code>not_equal(other)</code> | See <code>torch.not_equal()</code> . |
| <code>not_equal_(other)</code> | In-place version of <code>not_equal()</code> . |
| <code>numel()</code> | See <code>torch.numel()</code> |
| <code>numpy(*[, force])</code> | Returns the tensor as a NumPy ndarray. |
| <code>orgqr(input2)</code> | See <code>torch.orgqr()</code> |
| <code>ormqr(input2, input3[, left, transpose])</code> | See <code>torch.ormqr()</code> |
| <code>outer(vec2)</code> | See <code>torch.outer()</code> . |
| <code>permute(*dims)</code> | See <code>torch.permute()</code> |
| <code>pin_memory()</code> | Copies the tensor to pinned memory, if it's not already pinned. |
| <code>pinverse()</code> | See <code>torch.pinverse()</code> |
| <code>polygamma(n)</code> | See <code>torch.polygamma()</code> |
| <code>polygamma_(n)</code> | In-place version of <code>polygamma()</code> |
| <code>positive()</code> | See <code>torch.positive()</code> |
| <code>pow(exponent)</code> | See <code>torch.pow()</code> |

continues on next page

Table 17 – continued from previous page

| | |
|---|--|
| <code>pow_(exponent)</code> | In-place version of <code>pow()</code> |
| <code>prelu</code> | |
| <code>prod([dim, keepdim, dtype])</code> | See <code>torch.prod()</code> |
| <code>put(input, index, source[, accumulate])</code> | Out-of-place version of <code>torch.Tensor.put_()</code> . |
| <code>put_(index, source[, accumulate])</code> | Copies the elements from <code>source</code> into the positions specified by <code>index</code> . |
| <code>q_per_channel_axis()</code> | Given a Tensor quantized by linear (affine) per-channel quantization, returns the index of dimension on which per-channel quantization is applied. |
| <code>q_per_channel_scales()</code> | Given a Tensor quantized by linear (affine) per-channel quantization, returns a Tensor of scales of the underlying quantizer. |
| <code>q_per_channel_zero_points()</code> | Given a Tensor quantized by linear (affine) per-channel quantization, returns a tensor of zero_points of the underlying quantizer. |
| <code>q_scale()</code> | Given a Tensor quantized by linear(affine) quantization, returns the scale of the underlying quantizer(). |
| <code>q_zero_point()</code> | Given a Tensor quantized by linear(affine) quantization, returns the zero_point of the underlying quantizer(). |
| <code>qr([some])</code> | See <code>torch.qr()</code> |
| <code>qscheme()</code> | Returns the quantization scheme of a given QTensor. |
| <code>quantile(q[, dim, keepdim, interpolation])</code> | See <code>torch.quantile()</code> |
| <code>rad2deg()</code> | See <code>torch.rad2deg()</code> |
| <code>rad2deg_()</code> | In-place version of <code>rad2deg()</code> |
| <code>random_([from, to, generator])</code> | Fills <code>self</code> tensor with numbers sampled from the discrete uniform distribution over <code>[from, to - 1]</code> . |
| <code>ravel()</code> | see <code>torch.ravel()</code> |
| <code>reciprocal()</code> | See <code>torch.reciprocal()</code> |
| <code>reciprocal_()</code> | In-place version of <code>reciprocal()</code> |
| <code>record_stream(stream)</code> | Ensures that the tensor memory is not reused for another tensor until all current work queued on <code>stream</code> are complete. |
| <code>refine_names(*names)</code> | Refines the dimension names of <code>self</code> according to <code>names</code> . |
| <code>register_hook(hook)</code> | Registers a backward hook. |
| <code>register_post_accumulate_grad_hook(hook)</code> | Registers a backward hook that runs after grad accumulation. |
| <code>reinforce(reward)</code> | |
| <code>relu</code> | |
| <code>relu_</code> | |
| <code>remainder(divisor)</code> | See <code>torch.remainder()</code> |
| <code>remainder_(divisor)</code> | In-place version of <code>remainder()</code> |
| <code>rename(*names, **rename_map)</code> | Renames dimension names of <code>self</code> . |
| <code>rename_(*names, **rename_map)</code> | In-place version of <code>rename()</code> . |
| <code>renorm(p, dim, maxnorm)</code> | See <code>torch.renorm()</code> |
| <code>renorm_(p, dim, maxnorm)</code> | In-place version of <code>renorm()</code> |

continues on next page

Table 17 – continued from previous page

| | |
|--|---|
| <code>repeat(*sizes)</code> | Repeats this tensor along the specified dimensions. |
| <code>repeat_interleave(repeats[, dim, output_size])</code> | See <code>torch.repeat_interleave()</code> . |
| <code>requires_grad_([requires_grad])</code> | Change if autograd should record operations on this tensor: sets this tensor's <code>requires_grad</code> attribute in-place. |
| <code>reshape(*shape)</code> | Returns a tensor with the same data and number of elements as <code>self</code> but with the specified shape. |
| <code>reshape_as(other)</code> | Returns this tensor as the same shape as <code>other</code> . |
| <code>resize(*sizes)</code> | |
| <code>resize_(*sizes[, memory_format])</code> | Resizes <code>self</code> tensor to the specified size. |
| <code>resize_as(tensor)</code> | |
| <code>resize_as_(tensor[, memory_format])</code> | Resizes the <code>self</code> tensor to be the same size as the specified tensor. |
| <code>resize_as_sparse_</code> | |
| <code>resolve_conj()</code> | See <code>torch.resolve_conj()</code> |
| <code>resolve_neg()</code> | See <code>torch.resolve_neg()</code> |
| <code>retain_grad()</code> | Enables this Tensor to have their grad populated during <code>backward()</code> . |
| <code>roll(shifts, dims)</code> | See <code>torch.roll()</code> |
| <code>rot90(k, dims)</code> | See <code>torch.rot90()</code> |
| <code>round([decimals])</code> | See <code>torch.round()</code> |
| <code>round_([decimals])</code> | In-place version of <code>round()</code> |
| <code>row_indices</code> | |
| <code>rsqrt()</code> | See <code>torch.rsqrt()</code> |
| <code>rsqrt_()</code> | In-place version of <code>rsqrt()</code> |
| <code>scatter(dim, index, src)</code> | Out-of-place version of <code>torch.Tensor.scatter_()</code> |
| <code>scatter_(dim, index, src[, reduce])</code> | Writes all values from the tensor <code>src</code> into <code>self</code> at the indices specified in the <code>index</code> tensor. |
| <code>scatter_add(dim, index, src)</code> | Out-of-place version of <code>torch.Tensor.scatter_add_()</code> |
| <code>scatter_add_(dim, index, src)</code> | Adds all values from the tensor <code>src</code> into <code>self</code> at the indices specified in the <code>index</code> tensor in a similar fashion as <code>scatter_()</code> . |
| <code>scatter_reduce(dim, index, src, reduce, *[, ...])</code> | Out-of-place version of <code>torch.Tensor.scatter_reduce_()</code> |
| <code>scatter_reduce_(dim, index, src, reduce, *)</code> | Reduces all values from the <code>src</code> tensor to the indices specified in the <code>index</code> tensor in the <code>self</code> tensor using the applied reduction defined via the <code>reduce</code> argument ("sum", "prod", "mean", "amax", "amin"). |
| <code>select(dim, index)</code> | See <code>torch.select()</code> |
| <code>select_scatter(src, dim, index)</code> | See <code>torch.select_scatter()</code> |
| <code>set_([source, storage_offset, size, stride])</code> | Sets the underlying storage, size, and strides. |
| <code>sgn()</code> | See <code>torch.sgn()</code> |
| <code>sgn_()</code> | In-place version of <code>sgn()</code> |
| <code>share_memory_()</code> | Moves the underlying storage to shared memory. |

continues on next page

Table 17 – continued from previous page

| | |
|--|--|
| <code>short([memory_format])</code> | <code>self.short()</code> is equivalent to <code>self.to(torch.int16)</code> . |
| <code>sigmoid()</code> | See <code>torch.sigmoid()</code> |
| <code>sigmoid_()</code> | In-place version of <code>sigmoid()</code> |
| <code>sign()</code> | See <code>torch.sign()</code> |
| <code>sign_()</code> | In-place version of <code>sign()</code> |
| <code>signbit()</code> | See <code>torch.signbit()</code> |
| <code>sin()</code> | See <code>torch.sin()</code> |
| <code>sin_()</code> | In-place version of <code>sin()</code> |
| <code>sinc()</code> | See <code>torch.sinc()</code> |
| <code>sinc_()</code> | In-place version of <code>sinc()</code> |
| <code>sinh()</code> | See <code>torch.sinh()</code> |
| <code>sinh_()</code> | In-place version of <code>sinh()</code> |
| <code>size([dim])</code> | Returns the size of the <code>self</code> tensor. |
| <code>slice_scatter(src[, dim, start, end, step])</code> | See <code>torch.slice_scatter()</code> |
| <code>slogdet()</code> | See <code>torch.slogdet()</code> |
| <code>smm(mat)</code> | See <code>torch.smm()</code> |
| <code>softmax(dim)</code> | Alias for <code>torch.nn.functional.softmax()</code> . |
| <code>solve(other)</code> | |
| <code>sort([dim, descending])</code> | See <code>torch.sort()</code> |
| <code>sparse_dim()</code> | Return the number of sparse dimensions in a sparse tensor <code>self</code> . |
| <code>sparse_mask(mask)</code> | Returns a new sparse tensor with values from a strided tensor <code>self</code> filtered by the indices of the sparse tensor <code>mask</code> . |
| <code>sparse_resize_(size, sparse_dim, dense_dim)</code> | Resizes <code>self</code> sparse tensor to the desired size and the number of sparse and dense dimensions. |
| <code>sparse_resize_and_clear_(size, sparse_dim, ...)</code> | Removes all specified elements from a sparse tensor <code>self</code> and resizes <code>self</code> to the desired size and the number of sparse and dense dimensions. |
| <code>split(split_size[, dim])</code> | See <code>torch.split()</code> |
| <code>split_with_sizes</code> | |
| <code>sqrt()</code> | See <code>torch.sqrt()</code> |
| <code>sqrt_()</code> | In-place version of <code>sqrt()</code> |
| <code>square()</code> | See <code>torch.square()</code> |
| <code>square_()</code> | In-place version of <code>square()</code> |
| <code>squeeze([dim])</code> | See <code>torch.squeeze()</code> |
| <code>squeeze_([dim])</code> | In-place version of <code>squeeze()</code> |
| <code>sspaddmm(mat1, mat2, *[, beta, alpha])</code> | See <code>torch.sspaddmm()</code> |
| <code>std([dim, correction, keepdim])</code> | See <code>torch.std()</code> |
| <code>stft(n_fft[, hop_length, win_length, ...])</code> | See <code>torch.stft()</code> |
| <code>storage()</code> | Returns the underlying <code>TypedStorage</code> . |
| <code>storage_offset()</code> | Returns <code>self</code> tensor's offset in the underlying storage in terms of number of storage elements (not bytes). |
| <code>storage_type()</code> | Returns the type of the underlying storage. |
| <code>stride(dim)</code> | Returns the stride of <code>self</code> tensor. |
| <code>sub(other, *[, alpha])</code> | See <code>torch.sub()</code> . |
| <code>sub_(other, *[, alpha])</code> | In-place version of <code>sub()</code> |
| <code>subtract(other, *[, alpha])</code> | See <code>torch.subtract()</code> . |

continues on next page

Table 17 – continued from previous page

| | |
|---|---|
| <code>subtract_(other, *, alpha)</code> | In-place version of <code>subtract()</code> . |
| <code>sum([dim, keepdim, dtype])</code> | See <code>torch.sum()</code> |
| <code>sum_to_size(*size)</code> | Sum this tensor to size. |
| <code>svd([some, compute_uv])</code> | See <code>torch.svd()</code> |
| <code>swapaxes(axis0, axis1)</code> | See <code>torch.swapaxes()</code> |
| <code>swapaxes_(axis0, axis1)</code> | In-place version of <code>swapaxes()</code> |
| <code>swapdims(dim0, dim1)</code> | See <code>torch.swapdims()</code> |
| <code>swapdims_(dim0, dim1)</code> | In-place version of <code>swapdims()</code> |
| <code>symeig([eigenvectors])</code> | |
| <code>t()</code> | See <code>torch.t()</code> |
| <code>t_()</code> | In-place version of <code>t()</code> |
| <code>take(indices)</code> | See <code>torch.take()</code> |
| <code>take_along_dim(indices, dim)</code> | See <code>torch.take_along_dim()</code> |
| <code>tan()</code> | See <code>torch.tan()</code> |
| <code>tan_()</code> | In-place version of <code>tan()</code> |
| <code>tanh()</code> | See <code>torch.tanh()</code> |
| <code>tanh_()</code> | In-place version of <code>tanh()</code> |
| <code>tensor_split(indices_or_sections[, dim])</code> | See <code>torch.tensor_split()</code> |
| <code>tile(dims)</code> | See <code>torch.tile()</code> |
| <code>to(*args, **kwargs)</code> | Performs Tensor dtype and/or device conversion. |
| <code>to_dense([dtype, masked_grad])</code> | Creates a strided copy of <code>self</code> if <code>self</code> is not a strided tensor, otherwise returns <code>self</code> . |
| <code>to_mkldnn()</code> | Returns a copy of the tensor in <code>torch.mkldnn</code> layout. |
| <code>to_padded_tensor(padding[, output_size])</code> | See <code>to_padded_tensor()</code> |
| <code>to_sparse(sparseDims)</code> | Returns a sparse copy of the tensor. |
| <code>to_sparse_bsc(blocksize, dense_dim)</code> | Convert a tensor to a block sparse column (BSC) storage format of given blocksize. |
| <code>to_sparse_bsr(blocksize, dense_dim)</code> | Convert a tensor to a block sparse row (BSR) storage format of given blocksize. |
| <code>to_sparse_coo()</code> | Convert a tensor to coordinate format. |
| <code>to_sparse_csc()</code> | Convert a tensor to compressed column storage (CSC) format. |
| <code>to_sparse_csr([dense_dim])</code> | Convert a tensor to compressed row storage format (CSR). |
| <code>tolist()</code> | Returns the tensor as a (nested) list. |
| <code>topk(k[, dim, largest, sorted])</code> | See <code>torch.topk()</code> |
| <code>trace()</code> | See <code>torch.trace()</code> |
| <code>transpose(dim0, dim1)</code> | See <code>torch.transpose()</code> |
| <code>transpose_(dim0, dim1)</code> | In-place version of <code>transpose()</code> |
| <code>triangular_solve(A[, upper, transpose, ...])</code> | See <code>torch.triangular_solve()</code> |
| <code>tril([diagonal])</code> | See <code>torch.tril()</code> |
| <code>tril_([diagonal])</code> | In-place version of <code>tril()</code> |
| <code>triu([diagonal])</code> | See <code>torch.triu()</code> |
| <code>triu_([diagonal])</code> | In-place version of <code>triu()</code> |
| <code>true_divide(value)</code> | See <code>torch.true_divide()</code> |
| <code>true_divide_(value)</code> | In-place version of <code>true_divide_()</code> |
| <code>trunc()</code> | See <code>torch.trunc()</code> |
| <code>trunc_()</code> | In-place version of <code>trunc()</code> |

continues on next page

Table 17 – continued from previous page

| | |
|---|---|
| <code>type([dtype, non_blocking])</code> | Returns the type if <i>dtype</i> is not provided, else casts this object to the specified type. |
| <code>type_as(tensor)</code> | Returns this tensor cast to the type of the given tensor. |
| <code>unbind([dim])</code> | See <code>torch.unbind()</code> |
| <code>unflatten(dim, sizes)</code> | See <code>torch.unflatten()</code> . |
| <code>unfold(dimension, size, step)</code> | Returns a view of the original tensor which contains all slices of size <i>size</i> from <i>self</i> tensor in the dimension <i>dimension</i> . |
| <code>uniform_([from, to, generator])</code> | Fills <i>self</i> tensor with numbers sampled from the continuous uniform distribution: |
| <code>unique([sorted, return_inverse, ...])</code> | Returns the unique elements of the input tensor. |
| <code>unique_consecutive([return_inverse, ...])</code> | Eliminates all but the first element from every consecutive group of equivalent elements. |
| <code>unsafe_chunk(chunks[, dim])</code> | See <code>torch.unsafe_chunk()</code> |
| <code>unsafe_split(split_size[, dim])</code> | See <code>torch.unsafe_split()</code> |
| <code>unsafe_split_with_sizes</code> | |
| <code>unsqueeze(dim)</code> | See <code>torch.unsqueeze()</code> |
| <code>unsqueeze_(dim)</code> | In-place version of <code>unsqueeze()</code> |
| <code>untyped_storage()</code> | Returns the underlying <code>UntypedStorage</code> . |
| <code>values()</code> | Return the values tensor of a sparse COO tensor. |
| <code>var([dim, correction, keepdim])</code> | See <code>torch.var()</code> |
| <code>vdot(other)</code> | See <code>torch.vdot()</code> |
| <code>view(*shape)</code> | Returns a new tensor with the same data as the <i>self</i> tensor but of a different <i>shape</i> . |
| <code>view_as(other)</code> | View this tensor as the same size as <i>other</i> . |
| <code>vsplit(split_size_or_sections)</code> | See <code>torch.vsplit()</code> |
| <code>where(condition, y)</code> | <code>self.where(condition, y)</code> is equivalent to <code>torch.where(condition, self, y)</code> . |
| <code>xlogy(other)</code> | See <code>torch.xlogy()</code> |
| <code>xlogy_(other)</code> | In-place version of <code>xlogy()</code> |
| <code>xpu([device, non_blocking, memory_format])</code> | Returns a copy of this object in XPU memory. |
| <code>zero_()</code> | Fills <i>self</i> tensor with zeros. |

Attributes

| | |
|---------------------|---|
| <code>H</code> | Returns a view of a matrix (2-D tensor) conjugated and transposed. |
| <code>T</code> | Returns a view of this tensor with its dimensions reversed. |
| <code>data</code> | |
| <code>device</code> | Is the <code>torch.device</code> where this Tensor is. |
| <code>dtype</code> | |
| <code>grad</code> | This attribute is <code>None</code> by default and becomes a Tensor the first time a call to <code>backward()</code> computes gradients for <i>self</i> . |

continues on next page

Table 18 – continued from previous page

| | |
|----------------------------|--|
| <code>grad_fn</code> | |
| <code>imag</code> | Returns a new tensor containing imaginary values of the <code>self</code> tensor. |
| <code>is_cpu</code> | Is True if the Tensor is stored on the CPU, False otherwise. |
| <code>is_cuda</code> | Is True if the Tensor is stored on the GPU, False otherwise. |
| <code>is_ipu</code> | Is True if the Tensor is stored on the IPU, False otherwise. |
| <code>is_leaf</code> | All Tensors that have <code>requires_grad</code> which is False will be leaf Tensors by convention. |
| <code>is_meta</code> | Is True if the Tensor is a meta tensor, False otherwise. |
| <code>is_mkldnn</code> | |
| <code>is_mps</code> | Is True if the Tensor is stored on the MPS device, False otherwise. |
| <code>is_nested</code> | |
| <code>is_ort</code> | |
| <code>is_quantized</code> | Is True if the Tensor is quantized, False otherwise. |
| <code>is_sparse</code> | Is True if the Tensor uses sparse COO storage layout, False otherwise. |
| <code>is_sparse_csr</code> | Is True if the Tensor uses sparse CSR storage layout, False otherwise. |
| <code>is_vulkan</code> | |
| <code>is_xla</code> | Is True if the Tensor is stored on an XLA device, False otherwise. |
| <code>is_xpu</code> | Is True if the Tensor is stored on the XPU, False otherwise. |
| <code>itemsize</code> | Alias for <code>element_size()</code> |
| <code>layout</code> | |
| <code>mH</code> | Accessing this property is equivalent to calling <code>adjoint()</code> . |
| <code>mT</code> | Returns a view of this tensor with the last two dimensions transposed. |
| <code>name</code> | |
| <code>names</code> | Stores names for each of this tensor's dimensions. |
| <code>nbytes</code> | Returns the number of bytes consumed by the "view" of elements of the Tensor if the Tensor does not use sparse storage layout. |
| <code>ndim</code> | Alias for <code>dim()</code> |
| <code>output_nr</code> | |
| <code>real</code> | Returns a new tensor containing real values of the <code>self</code> tensor for a complex-valued input tensor. |

continues on next page

Table 18 – continued from previous page

| | |
|----------------------------|---|
| <code>requires_grad</code> | Is <code>True</code> if gradients need to be computed for this Tensor, <code>False</code> otherwise. |
| <code>retains_grad</code> | Is <code>True</code> if this Tensor is non-leaf and its <code>grad</code> is enabled to be populated during <code>backward()</code> , <code>False</code> otherwise. |
| <code>shape</code> | Returns the size of the <code>self</code> tensor. |
| <code>volatile</code> | |

property `is_leaf`: bool

All Tensors that have `requires_grad` which is `False` will be leaf Tensors by convention.

For Tensors that have `requires_grad` which is `True`, they will be leaf Tensors if they were created by the user. This means that they are not the result of an operation and so `grad_fn` is `None`.

Only leaf Tensors will have their `grad` populated during a call to `backward()`. To get `grad` populated for non-leaf Tensors, you can use `retain_grad()`.

Example:

```
>>> a = torch.rand(10, requires_grad=True)
>>> a.is_leaf
True
>>> b = torch.rand(10, requires_grad=True).cuda()
>>> b.is_leaf
False
# b was created by the operation that cast a cpu Tensor into a cuda Tensor
>>> c = torch.rand(10, requires_grad=True) + 2
>>> c.is_leaf
False
# c was created by the addition operation
>>> d = torch.rand(10).cuda()
>>> d.is_leaf
True
# d does not require gradients and so has no operation creating it (that is,
↳ tracked by the autograd engine)
>>> e = torch.rand(10).cuda().requires_grad_()
>>> e.is_leaf
True
# e requires gradients and has no operations creating it
>>> f = torch.rand(10, requires_grad=True, device="cuda")
>>> f.is_leaf
True
# f requires grad, has no operation creating it
```

`materialize(shape, device=None, dtype=None)`

Create a Parameter with the same properties of the uninitialized one. Given a shape, it materializes a parameter in the same device and with the same `dtype` as the current one or the specified ones in the arguments.

Parameters

- **shape** (`Tuple[int, ...]`) – (tuple): the shape for the materialized tensor.
- **device** (`torch.device`) – the desired device of the parameters and buffers in this module. Optional.

- **dtype** (`torch.dtype`) – the desired floating point type of the floating point parameters and buffers in this module. Optional.

Return type

None

share_memory_()

Moves the underlying storage to shared memory.

This is a no-op if the underlying storage is already in shared memory and for CUDA tensors. Tensors in shared memory cannot be resized.

Return type`UninitializedParameter`**pytorch_pfn_extras.nn.modules.lazy_conv****Classes**

| | |
|--|--|
| <code>pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv1d(...)</code> | Conv1d module with lazy weight initialization. |
| <code>pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv2d(...)</code> | Conv2d module with lazy weight initialization. |
| <code>pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv3d(...)</code> | Conv3d module with lazy weight initialization. |
| <code>pytorch_pfn_extras.nn.modules.lazy_conv.LazyInitializationMixin(...)</code> | A mixin for modules that lazily initialize buffers and parameters. |
| <code>pytorch_pfn_extras.nn.modules.lazy_conv.UninitializedParameter([...])</code> | |

pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv1d

class `pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv1d(in_channels, *args, **kwargs)`

Bases: `_LazyConvNd`, `Conv1d`

Conv1d module with lazy weight initialization.

When `in_channels` is `None`, it is determined at the first time of the forward step.

Initializes internal Module state, shared by both `nn.Module` and `ScriptModule`.

Methods

| | |
|---|--|
| <code>__init__(in_channels, *args, **kwargs)</code> | Initializes internal Module state, shared by both <code>nn.Module</code> and <code>ScriptModule</code> . |
| <code>add_module(name, module)</code> | Adds a child module to the current module. |
| <code>apply(fn)</code> | Applies <code>fn</code> recursively to every submodule (as returned by <code>.children()</code>) as well as self. |
| <code>bfloat16()</code> | Casts all floating point parameters and buffers to <code>bfloat16</code> datatype. |
| <code>buffers([recurse])</code> | Returns an iterator over module buffers. |

continues on next page

Table 19 – continued from previous page

| | |
|---|---|
| <code>children()</code> | Returns an iterator over immediate children modules. |
| <code>compile(*args, **kwargs)</code> | Compile this Module's forward using <code>torch.compile()</code> . |
| <code>cpu()</code> | Moves all model parameters and buffers to the CPU. |
| <code>cuda([device])</code> | Moves all model parameters and buffers to the GPU. |
| <code>double()</code> | Casts all floating point parameters and buffers to <code>double</code> datatype. |
| <code>eval()</code> | Sets the module in evaluation mode. |
| <code>extra_repr()</code> | Set the extra representation of the module |
| <code>float()</code> | Casts all floating point parameters and buffers to <code>float</code> datatype. |
| <code>forward(input)</code> | Defines the computation performed at every call. |
| <code>get_buffer(target)</code> | Returns the buffer given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_extra_state()</code> | Returns any extra state to include in the module's <code>state_dict</code> . |
| <code>get_parameter(target)</code> | Returns the parameter given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_submodule(target)</code> | Returns the submodule given by <code>target</code> if it exists, otherwise throws an error. |
| <code>half()</code> | Casts all floating point parameters and buffers to <code>half</code> datatype. |
| <code>ipu([device])</code> | Moves all model parameters and buffers to the IPU. |
| <code>load_state_dict(state_dict[, strict, assign])</code> | Copies parameters and buffers from <code>state_dict</code> into this module and its descendants. |
| <code>modules()</code> | Returns an iterator over all modules in the network. |
| <code>named_buffers([prefix, recurse, ...])</code> | Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself. |
| <code>named_children()</code> | Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself. |
| <code>named_modules([memo, prefix, remove_duplicate])</code> | Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself. |
| <code>named_parameters([prefix, recurse, ...])</code> | Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself. |
| <code>parameters([recurse])</code> | Returns an iterator over module parameters. |
| <code>register_backward_hook(hook)</code> | Registers a backward hook on the module. |
| <code>register_buffer(name, tensor[, persistent])</code> | Adds a buffer to the module. |
| <code>register_forward_hook(hook, *[, prepend, ...])</code> | Registers a forward hook on the module. |
| <code>register_forward_pre_hook(hook, *[, ...])</code> | Registers a forward pre-hook on the module. |
| <code>register_full_backward_hook(hook[, prepend])</code> | Registers a backward hook on the module. |
| <code>register_full_backward_pre_hook(hook[, prepend])</code> | Registers a backward pre-hook on the module. |
| <code>register_load_state_dict_post_hook(hook)</code> | Registers a post hook to be run after module's <code>load_state_dict</code> is called. |
| <code>register_module(name, module)</code> | Alias for <code>add_module()</code> . |
| <code>register_parameter(name, param)</code> | Adds a parameter to the module. |

continues on next page

Table 19 – continued from previous page

| | |
|---|---|
| <code>register_state_dict_pre_hook(hook)</code> | These hooks will be called with arguments: <code>self</code> , <code>prefix</code> , and <code>keep_vars</code> before calling <code>state_dict</code> on <code>self</code> . |
| <code>requires_grad_([requires_grad])</code> | Change if autograd should record operations on parameters in this module. |
| <code>reset_parameters()</code> | |
| <code>set_extra_state(state)</code> | This function is called from <code>load_state_dict()</code> to handle any extra state found within the <code>state_dict</code> . |
| <code>share_memory()</code> | See <code>torch.Tensor.share_memory_()</code> |
| <code>state_dict(*args, **kwargs)</code> | Returns a dictionary containing a whole state of the module. |
| <code>to(*args, **kwargs)</code> | Moves and/or casts the parameters and buffers. |
| <code>to_empty(*, device[, recurse])</code> | Moves the parameters and buffers to the specified device without copying storage. |
| <code>train([mode])</code> | Sets the module in training mode. |
| <code>type(dst_type)</code> | Casts all parameters and buffers to <code>dst_type</code> . |
| <code>xpu([device])</code> | Moves all model parameters and buffers to the XPU. |
| <code>zero_grad([set_to_none])</code> | Resets gradients of all model parameters. |

Attributes

| | |
|---|--|
| <code>T_destination</code> | alias of <code>TypeVar('T_destination', bound=Dict[str, Any])</code> |
| <code>call_super_init</code> | |
| <code>dump_patches</code> | |
| <code>lazy_buffer_names</code> | |
| <code>lazy_parameter_names</code> | |
| <code>lazy_parameters_determined</code> | Returns if all lazy parameters are determined. |

Parameters

- **`in_channels`** (*Optional[int]*) –
- **`args`** (*Any*) –
- **`kwargs`** (*Any*) –

`bias`: `Optional[Tensor]`

`dilation`: `Tuple[int, ...]`

`groups`: `int`

`in_channels`: `Optional[int]`

`kernel_size`: `Tuple[int, ...]`

```
out_channels: int
output_padding: Tuple[int, ...]
padding: Union[str, Tuple[int, ...]]
padding_mode: str
stride: Tuple[int, ...]
training: bool
transposed: bool
weight: Tensor
```

pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv2d

class pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv2d(*in_channels*, **args*, ***kwargs*)

Bases: `_LazyConvNd`, `Conv2d`

Conv2d module with lazy weight initialization.

When *in_channels* is `None`, it is determined at the first time of the forward step.

Initializes internal Module state, shared by both `nn.Module` and `ScriptModule`.

Methods

| | |
|---|--|
| <code>__init__(in_channels, *args, **kwargs)</code> | Initializes internal Module state, shared by both <code>nn.Module</code> and <code>ScriptModule</code> . |
| <code>add_module(name, module)</code> | Adds a child module to the current module. |
| <code>apply(fn)</code> | Applies <code>fn</code> recursively to every submodule (as returned by <code>.children()</code>) as well as self. |
| <code>bfloat16()</code> | Casts all floating point parameters and buffers to <code>bfloat16</code> datatype. |
| <code>buffers([recurse])</code> | Returns an iterator over module buffers. |
| <code>children()</code> | Returns an iterator over immediate children modules. |
| <code>compile(*args, **kwargs)</code> | Compile this Module's forward using <code>torch.compile()</code> . |
| <code>cpu()</code> | Moves all model parameters and buffers to the CPU. |
| <code>cuda([device])</code> | Moves all model parameters and buffers to the GPU. |
| <code>double()</code> | Casts all floating point parameters and buffers to <code>double</code> datatype. |
| <code>eval()</code> | Sets the module in evaluation mode. |
| <code>extra_repr()</code> | Set the extra representation of the module |
| <code>float()</code> | Casts all floating point parameters and buffers to <code>float</code> datatype. |
| <code>forward(input)</code> | Defines the computation performed at every call. |
| <code>get_buffer(target)</code> | Returns the buffer given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_extra_state()</code> | Returns any extra state to include in the module's <code>state_dict</code> . |

continues on next page

Table 20 – continued from previous page

| | |
|---|---|
| <code>get_parameter(target)</code> | Returns the parameter given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_submodule(target)</code> | Returns the submodule given by <code>target</code> if it exists, otherwise throws an error. |
| <code>half()</code> | Casts all floating point parameters and buffers to <code>half</code> datatype. |
| <code>ipu([device])</code> | Moves all model parameters and buffers to the IPU. |
| <code>load_state_dict(state_dict[, strict, assign])</code> | Copies parameters and buffers from <code>state_dict</code> into this module and its descendants. |
| <code>modules()</code> | Returns an iterator over all modules in the network. |
| <code>named_buffers([prefix, recurse, ...])</code> | Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself. |
| <code>named_children()</code> | Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself. |
| <code>named_modules([memo, prefix, remove_duplicate])</code> | Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself. |
| <code>named_parameters([prefix, recurse, ...])</code> | Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself. |
| <code>parameters([recurse])</code> | Returns an iterator over module parameters. |
| <code>register_backward_hook(hook)</code> | Registers a backward hook on the module. |
| <code>register_buffer(name, tensor[, persistent])</code> | Adds a buffer to the module. |
| <code>register_forward_hook(hook, *[, prepend, ...])</code> | Registers a forward hook on the module. |
| <code>register_forward_pre_hook(hook, *[, ...])</code> | Registers a forward pre-hook on the module. |
| <code>register_full_backward_hook(hook[, prepend])</code> | Registers a backward hook on the module. |
| <code>register_full_backward_pre_hook(hook[, prepend])</code> | Registers a backward pre-hook on the module. |
| <code>register_load_state_dict_post_hook(hook)</code> | Registers a post hook to be run after module's <code>load_state_dict</code> is called. |
| <code>register_module(name, module)</code> | Alias for <code>add_module()</code> . |
| <code>register_parameter(name, param)</code> | Adds a parameter to the module. |
| <code>register_state_dict_pre_hook(hook)</code> | These hooks will be called with arguments: <code>self</code> , <code>prefix</code> , and <code>keep_vars</code> before calling <code>state_dict</code> on <code>self</code> . |
| <code>requires_grad_([requires_grad])</code> | Change if autograd should record operations on parameters in this module. |
| <code>reset_parameters()</code> | |
| <code>set_extra_state(state)</code> | This function is called from <code>load_state_dict()</code> to handle any extra state found within the <code>state_dict</code> . |
| <code>share_memory()</code> | See <code>torch.Tensor.share_memory_()</code> |
| <code>state_dict(*args, **kwargs)</code> | Returns a dictionary containing a whole state of the module. |
| <code>to(*args, **kwargs)</code> | Moves and/or casts the parameters and buffers. |
| <code>to_empty(*, device[, recurse])</code> | Moves the parameters and buffers to the specified device without copying storage. |
| <code>train([mode])</code> | Sets the module in training mode. |
| <code>type(dst_type)</code> | Casts all parameters and buffers to <code>dst_type</code> . |

continues on next page

Table 20 – continued from previous page

| | |
|---------------------------------------|--|
| <code>xpu([device])</code> | Moves all model parameters and buffers to the XPU. |
| <code>zero_grad([set_to_none])</code> | Resets gradients of all model parameters. |

Attributes

| | |
|--|--|
| <code>T_destination</code> | alias of <code>TypeVar('T_destination', bound=Dict[str, Any])</code> |
| <code>call_super_init</code> | |
| <code>dump_patches</code> | |
| <code>lazy_buffer_names</code> | |
| <code>lazy_parameter_names</code> | |
| <code>lazy_parmeters_determined</code> | Returns if all lazy parameters are determined. |

Parameters

- `in_channels` (*Optional[int]*) –
- `args` (*Any*) –
- `kwargs` (*Any*) –

`bias: Optional[Tensor]``dilation: Tuple[int, ...]``groups: int``in_channels: Optional[int]``kernel_size: Tuple[int, ...]``out_channels: int``output_padding: Tuple[int, ...]``padding: Union[str, Tuple[int, ...]]``padding_mode: str``stride: Tuple[int, ...]``training: bool``transposed: bool``weight: Tensor`

pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv3d

class pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv3d(*in_channels*, *args, **kwargs)

Bases: _LazyConvNd, Conv3d

Conv3d module with lazy weight initialization.

When *in_channels* is None, it is determined at the first time of the forward step.

Initializes internal Module state, shared by both nn.Module and ScriptModule.

Methods

| | |
|--|---|
| <code>__init__(in_channels, *args, **kwargs)</code> | Initializes internal Module state, shared by both nn.Module and ScriptModule. |
| <code>add_module(name, module)</code> | Adds a child module to the current module. |
| <code>apply(fn)</code> | Applies <code>fn</code> recursively to every submodule (as returned by <code>.children()</code>) as well as self. |
| <code>bfloat16()</code> | Casts all floating point parameters and buffers to <code>bfloat16</code> datatype. |
| <code>buffers([recurse])</code> | Returns an iterator over module buffers. |
| <code>children()</code> | Returns an iterator over immediate children modules. |
| <code>compile(*args, **kwargs)</code> | Compile this Module's forward using <code>torch.compile()</code> . |
| <code>cpu()</code> | Moves all model parameters and buffers to the CPU. |
| <code>cuda([device])</code> | Moves all model parameters and buffers to the GPU. |
| <code>double()</code> | Casts all floating point parameters and buffers to <code>double</code> datatype. |
| <code>eval()</code> | Sets the module in evaluation mode. |
| <code>extra_repr()</code> | Set the extra representation of the module |
| <code>float()</code> | Casts all floating point parameters and buffers to <code>float</code> datatype. |
| <code>forward(input)</code> | Defines the computation performed at every call. |
| <code>get_buffer(target)</code> | Returns the buffer given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_extra_state()</code> | Returns any extra state to include in the module's <code>state_dict</code> . |
| <code>get_parameter(target)</code> | Returns the parameter given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_submodule(target)</code> | Returns the submodule given by <code>target</code> if it exists, otherwise throws an error. |
| <code>half()</code> | Casts all floating point parameters and buffers to <code>half</code> datatype. |
| <code>ipu([device])</code> | Moves all model parameters and buffers to the IPU. |
| <code>load_state_dict(state_dict[, strict, assign])</code> | Copies parameters and buffers from <code>state_dict</code> into this module and its descendants. |
| <code>modules()</code> | Returns an iterator over all modules in the network. |
| <code>named_buffers([prefix, recurse, ...])</code> | Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself. |
| <code>named_children()</code> | Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself. |

continues on next page

Table 21 – continued from previous page

| | |
|---|---|
| <code>named_modules([memo, prefix, remove_duplicate])</code> | Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself. |
| <code>named_parameters([prefix, recurse, ...])</code> | Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself. |
| <code>parameters([recurse])</code> | Returns an iterator over module parameters. |
| <code>register_backward_hook(hook)</code> | Registers a backward hook on the module. |
| <code>register_buffer(name, tensor[, persistent])</code> | Adds a buffer to the module. |
| <code>register_forward_hook(hook, *[, prepend, ...])</code> | Registers a forward hook on the module. |
| <code>register_forward_pre_hook(hook, *[, ...])</code> | Registers a forward pre-hook on the module. |
| <code>register_full_backward_hook(hook[, prepend])</code> | Registers a backward hook on the module. |
| <code>register_full_backward_pre_hook(hook[, prepend])</code> | Registers a backward pre-hook on the module. |
| <code>register_load_state_dict_post_hook(hook)</code> | Registers a post hook to be run after module's <code>load_state_dict</code> is called. |
| <code>register_module(name, module)</code> | Alias for <code>add_module()</code> . |
| <code>register_parameter(name, param)</code> | Adds a parameter to the module. |
| <code>register_state_dict_pre_hook(hook)</code> | These hooks will be called with arguments: <code>self</code> , <code>prefix</code> , and <code>keep_vars</code> before calling <code>state_dict</code> on <code>self</code> . |
| <code>requires_grad_([requires_grad])</code> | Change if autograd should record operations on parameters in this module. |
| <code>reset_parameters()</code> | |
| <code>set_extra_state(state)</code> | This function is called from <code>load_state_dict()</code> to handle any extra state found within the <code>state_dict</code> . |
| <code>share_memory()</code> | See <code>torch.Tensor.share_memory_()</code> |
| <code>state_dict(*args, **kwargs)</code> | Returns a dictionary containing a whole state of the module. |
| <code>to(*args, **kwargs)</code> | Moves and/or casts the parameters and buffers. |
| <code>to_empty(*, device[, recurse])</code> | Moves the parameters and buffers to the specified device without copying storage. |
| <code>train([mode])</code> | Sets the module in training mode. |
| <code>type(dst_type)</code> | Casts all parameters and buffers to <code>dst_type</code> . |
| <code>xpu([device])</code> | Moves all model parameters and buffers to the XPU. |
| <code>zero_grad([set_to_none])</code> | Resets gradients of all model parameters. |

Attributes

| | |
|--|--|
| <code>T_destination</code> | alias of <code>TypeVar('T_destination', bound=Dict[str, Any])</code> |
| <code>call_super_init</code> | |
| <code>dump_patches</code> | |
| <code>lazy_buffer_names</code> | |
| <code>lazy_parameter_names</code> | |
| <code>lazy_parmeters_determined</code> | Returns if all lazy parameters are determined. |

Parameters

- `in_channels` (*Optional[int]*) –
- `args` (*Any*) –
- `kwargs` (*Any*) –

`bias`: `Optional[Tensor]`

`dilation`: `Tuple[int, ...]`

`groups`: `int`

`in_channels`: `Optional[int]`

`kernel_size`: `Tuple[int, ...]`

`out_channels`: `int`

`output_padding`: `Tuple[int, ...]`

`padding`: `Union[str, Tuple[int, ...]]`

`padding_mode`: `str`

`stride`: `Tuple[int, ...]`

`training`: `bool`

`transposed`: `bool`

`weight`: `Tensor`

pytorch_pfn_extras.nn.modules.lazy_conv.LazyInitializationMixin**class** pytorch_pfn_extras.nn.modules.lazy_conv.LazyInitializationMixin(*args, **kwargs)

Bases: object

A mixin for modules that lazily initialize buffers and parameters.

Unlike regular modules, subclasses of this module can initialize buffers and parameters outside of the constructor (`__init__`). This allows you to, for example, initialize parameters in `forward` method to determine the shape of the weight based on the initial input.

Be sure to run “dummy” forward once to initialize all parameters that should be trained, before passing `module.parameters()` to an optimizer; otherwise weights initialized after `module.parameters()` (e.g., in `forward` function) will never be trained.

Note that lazy modules cannot validate if the shape is correct during deserialization. Also note that the initial weights may become different from the original (non-lazy) module even if the random seed is manually configured, as the order of initialization is different from the original one; especially, `module.cuda()` may cause the initialization to run on a GPU.

The default value of lazy buffers and parameters are `torch.Tensor([])` and `UninitializedParameter()`, respectively.

Methods

`__init__`(*args, **kwargs)

`state_dict`(*args, **kwargs)Returns a dictionary containing a whole state of the module.

Attributes

`lazy_buffer_names`

`lazy_parameter_names`

`lazy_parameters_determined`Returns if all lazy parameters are determined.

Parameters

- **args** (*Any*) –
- **kwargs** (*Any*) –

`__init__`(*args, **kwargs)**Parameters**

- **self** (*Any*) –
- **args** (*Any*) –
- **kwargs** (*Any*) –

Return type

None

lazy_buffer_names: `Tuple[str, ...] = ()`

lazy_parameter_names: `Tuple[str, ...] = ()`

property lazy_parameters_determined: `bool`

Returns if all lazy parameters are determined.

Subclasses can perform parameters initialization after all lazy parameters are determined. Note that this may be called during `__init__`.

state_dict(*args, **kwargs)

Returns a dictionary containing a whole state of the module.

This function overrides the default behavior to exclude uninitialized parameter from serialization. This is needed because we need to discriminate lazy parameters (`UninitializedParameter()`) and initialized empty parameters (`torch.nn.Parameter(torch.Tensor())`) during deserialization.

See comments of `_lazy_load_hook` for details.

Parameters

- **self** (*Any*) –
- **args** (*Any*) –
- **kwargs** (*Any*) –

Return type

Dict[str, Any]

pytorch_pfn_extras.nn.modules.lazy_conv.UninitializedParameter

class `pytorch_pfn_extras.nn.modules.lazy_conv.UninitializedParameter`(*data=None, requires_grad=True*)

Bases: `Parameter`

Methods

| | |
|--|---|
| <code>__init__()</code> | |
| <code>abs()</code> | See <code>torch.abs()</code> |
| <code>abs_()</code> | In-place version of <code>abs()</code> |
| <code>absolute()</code> | Alias for <code>abs()</code> |
| <code>absolute_()</code> | In-place version of <code>absolute()</code> Alias for <code>abs_()</code> |
| <code>acos()</code> | See <code>torch.acos()</code> |
| <code>acos_()</code> | In-place version of <code>acos()</code> |
| <code>acosh()</code> | See <code>torch.acosh()</code> |
| <code>acosh_()</code> | In-place version of <code>acosh()</code> |
| <code>add(other, *, alpha)</code> | Add a scalar or tensor to <code>self</code> tensor. |
| <code>add_(other, *, alpha)</code> | In-place version of <code>add()</code> |
| <code>addbmm(batch1, batch2, *, beta, alpha)</code> | See <code>torch.addbmm()</code> |
| <code>addbmm_(batch1, batch2, *, beta, alpha)</code> | In-place version of <code>addbmm()</code> |
| <code>addcddiv(tensor1, tensor2, *, value)</code> | See <code>torch.addcddiv()</code> |

continues on next page

Table 22 – continued from previous page

| | |
|---|--|
| <code>addcdiv_(tensor1, tensor2, *, value)</code> | In-place version of <code>addcdiv()</code> |
| <code>addcmul(tensor1, tensor2, *, value)</code> | See <code>torch.addcmul()</code> |
| <code>addcmul_(tensor1, tensor2, *, value)</code> | In-place version of <code>addcmul()</code> |
| <code>addmm(mat1, mat2, *, beta, alpha)</code> | See <code>torch.addmm()</code> |
| <code>addmm_(mat1, mat2, *, beta, alpha)</code> | In-place version of <code>addmm()</code> |
| <code>addmv(mat, vec, *, beta, alpha)</code> | See <code>torch.addmv()</code> |
| <code>addmv_(mat, vec, *, beta, alpha)</code> | In-place version of <code>addmv()</code> |
| <code>addr(vec1, vec2, *, beta, alpha)</code> | See <code>torch.addr()</code> |
| <code>addr_(vec1, vec2, *, beta, alpha)</code> | In-place version of <code>addr()</code> |
| <code>adjoint()</code> | Alias for <code>adjoint()</code> |
| <code>align_as(other)</code> | Permutates the dimensions of the <code>self</code> tensor to match the dimension order in the <code>other</code> tensor, adding size-one dims for any new names. |
| <code>align_to(*names)</code> | Permutates the dimensions of the <code>self</code> tensor to match the order specified in <code>names</code> , adding size-one dims for any new names. |
| <code>all([dim, keepdim])</code> | See <code>torch.all()</code> |
| <code>allclose(other[, rtol, atol, equal_nan])</code> | See <code>torch.allclose()</code> |
| <code>amax([dim, keepdim])</code> | See <code>torch.amax()</code> |
| <code>amin([dim, keepdim])</code> | See <code>torch.amin()</code> |
| <code>aminmax(*[, dim, keepdim])</code> | See <code>torch.aminmax()</code> |
| <code>angle()</code> | See <code>torch.angle()</code> |
| <code>any([dim, keepdim])</code> | See <code>torch.any()</code> |
| <code>apply_(callable)</code> | Applies the function <code>callable</code> to each element in the tensor, replacing each element with the value returned by <code>callable</code> . |
| <code>arccos()</code> | See <code>torch.arccos()</code> |
| <code>arccos_()</code> | In-place version of <code>arccos()</code> |
| <code>arccosh</code> | <code>acosh()</code> -> Tensor |
| <code>arccosh_</code> | <code>acosh_()</code> -> Tensor |
| <code>arcsin()</code> | See <code>torch.arcsin()</code> |
| <code>arcsin_()</code> | In-place version of <code>arcsin()</code> |
| <code>arcsinh()</code> | See <code>torch.arcsinh()</code> |
| <code>arcsinh_()</code> | In-place version of <code>arcsinh()</code> |
| <code>arctan()</code> | See <code>torch.arctan()</code> |
| <code>arctan2(other)</code> | See <code>torch.arctan2()</code> |
| <code>arctan2_</code> | <code>atan2_(other)</code> -> Tensor |
| <code>arctan_()</code> | In-place version of <code>arctan()</code> |
| <code>arctanh()</code> | See <code>torch.arctanh()</code> |
| <code>arctanh_(other)</code> | In-place version of <code>arctanh()</code> |
| <code>argmax([dim, keepdim])</code> | See <code>torch.argmax()</code> |
| <code>argmin([dim, keepdim])</code> | See <code>torch.argmin()</code> |
| <code>argsort([dim, descending])</code> | See <code>torch.argsort()</code> |
| <code>argwhere()</code> | See <code>torch.argwhere()</code> |
| <code>as_strided(size, stride[, storage_offset])</code> | See <code>torch.as_strided()</code> |
| <code>as_strided_(size, stride[, storage_offset])</code> | In-place version of <code>as_strided()</code> |
| <code>as_strided_scatter(src, size, stride[, ...])</code> | See <code>torch.as_strided_scatter()</code> |
| <code>as_subclass(cls)</code> | Makes a <code>cls</code> instance with the same data pointer as <code>self</code> . |
| <code>asin()</code> | See <code>torch.asin()</code> |

continues on next page

Table 22 – continued from previous page

| | |
|---|---|
| <code>asin_()</code> | In-place version of <code>asin()</code> |
| <code>asinh()</code> | See <code>torch.asinh()</code> |
| <code>asinh_()</code> | In-place version of <code>asinh()</code> |
| <code>atan()</code> | See <code>torch.atan()</code> |
| <code>atan2(other)</code> | See <code>torch.atan2()</code> |
| <code>atan2_(other)</code> | In-place version of <code>atan2()</code> |
| <code>atan_()</code> | In-place version of <code>atan()</code> |
| <code>atanh()</code> | See <code>torch.atanh()</code> |
| <code>atanh_(other)</code> | In-place version of <code>atanh()</code> |
| <code>backward([gradient, retain_graph, ...])</code> | Computes the gradient of current tensor wrt graph leaves. |
| <code>baddbmm(batch1, batch2, *[, beta, alpha])</code> | See <code>torch.baddbmm()</code> |
| <code>baddbmm_(batch1, batch2, *[, beta, alpha])</code> | In-place version of <code>baddbmm()</code> |
| <code>bernoulli(*[, generator])</code> | Returns a result tensor where each <code>result[i]</code> is independently sampled from <code>Bernoulli(self[i])</code> . |
| <code>bernoulli_(p, generator)</code> | Fills each location of <code>self</code> with an independent sample from <code>Bernoulli(p)</code> . |
| <code>bfloat16([memory_format])</code> | <code>self.bfloat16()</code> is equivalent to <code>self.to(torch.bfloat16)</code> . |
| <code>bincount([weights, minlength])</code> | See <code>torch.bincount()</code> |
| <code>bitwise_and()</code> | See <code>torch.bitwise_and()</code> |
| <code>bitwise_and_()</code> | In-place version of <code>bitwise_and()</code> |
| <code>bitwise_left_shift(other)</code> | See <code>torch.bitwise_left_shift()</code> |
| <code>bitwise_left_shift_(other)</code> | In-place version of <code>bitwise_left_shift()</code> |
| <code>bitwise_not()</code> | See <code>torch.bitwise_not()</code> |
| <code>bitwise_not_()</code> | In-place version of <code>bitwise_not()</code> |
| <code>bitwise_or()</code> | See <code>torch.bitwise_or()</code> |
| <code>bitwise_or_()</code> | In-place version of <code>bitwise_or()</code> |
| <code>bitwise_right_shift(other)</code> | See <code>torch.bitwise_right_shift()</code> |
| <code>bitwise_right_shift_(other)</code> | In-place version of <code>bitwise_right_shift()</code> |
| <code>bitwise_xor()</code> | See <code>torch.bitwise_xor()</code> |
| <code>bitwise_xor_()</code> | In-place version of <code>bitwise_xor()</code> |
| <code>bmm(batch2)</code> | See <code>torch.bmm()</code> |
| <code>bool([memory_format])</code> | <code>self.bool()</code> is equivalent to <code>self.to(torch.bool)</code> . |
| <code>broadcast_to(shape)</code> | See <code>torch.broadcast_to()</code> . |
| <code>byte([memory_format])</code> | <code>self.byte()</code> is equivalent to <code>self.to(torch.uint8)</code> . |
| <code>cauchy_([median, sigma, generator])</code> | Fills the tensor with numbers drawn from the Cauchy distribution: |
| <code>ccol_indices</code> | |
| <code>cdouble([memory_format])</code> | <code>self.cdouble()</code> is equivalent to <code>self.to(torch.complex128)</code> . |
| <code>ceil()</code> | See <code>torch.ceil()</code> |
| <code>ceil_()</code> | In-place version of <code>ceil()</code> |
| <code>cfloat([memory_format])</code> | <code>self.cfloat()</code> is equivalent to <code>self.to(torch.complex64)</code> . |
| <code>chalf([memory_format])</code> | <code>self.chalf()</code> is equivalent to <code>self.to(torch.complex32)</code> . |

continues on next page

Table 22 – continued from previous page

| | |
|--|--|
| <code>char([memory_format])</code> | <code>self.char()</code> is equivalent to <code>self.to(torch.int8)</code> . |
| <code>cholesky([upper])</code> | See <code>torch.cholesky()</code> |
| <code>cholesky_inverse([upper])</code> | See <code>torch.cholesky_inverse()</code> |
| <code>cholesky_solve(input2[, upper])</code> | See <code>torch.cholesky_solve()</code> |
| <code>chunk(chunks[, dim])</code> | See <code>torch.chunk()</code> |
| <code>clamp([min, max])</code> | See <code>torch.clamp()</code> |
| <code>clamp_([min, max])</code> | In-place version of <code>clamp()</code> |
| <code>clamp_max</code> | |
| <code>clamp_max_</code> | |
| <code>clamp_min</code> | |
| <code>clamp_min_</code> | |
| <code>clip([min, max])</code> | Alias for <code>clamp()</code> . |
| <code>clip_([min, max])</code> | Alias for <code>clamp_()</code> . |
| <code>clone(*[, memory_format])</code> | See <code>torch.clone()</code> |
| <code>coalesce()</code> | Returns a coalesced copy of <code>self</code> if <code>self</code> is an un-coalesced tensor. |
| <code>col_indices()</code> | Returns the tensor containing the column indices of the <code>self</code> tensor when <code>self</code> is a sparse CSR tensor of layout <code>sparse_csr</code> . |
| <code>conj()</code> | See <code>torch.conj()</code> |
| <code>conj_physical()</code> | See <code>torch.conj_physical()</code> |
| <code>conj_physical_()</code> | In-place version of <code>conj_physical()</code> |
| <code>contiguous([memory_format])</code> | Returns a contiguous in memory tensor containing the same data as <code>self</code> tensor. |
| <code>copy_(src[, non_blocking])</code> | Copies the elements from <code>src</code> into <code>self</code> tensor and returns <code>self</code> . |
| <code>copysign(other)</code> | See <code>torch.copysign()</code> |
| <code>copysign_(other)</code> | In-place version of <code>copysign()</code> |
| <code>corrcoef()</code> | See <code>torch.corrcoef()</code> |
| <code>cos()</code> | See <code>torch.cos()</code> |
| <code>cos_()</code> | In-place version of <code>cos()</code> |
| <code>cosh()</code> | See <code>torch.cosh()</code> |
| <code>cosh_()</code> | In-place version of <code>cosh()</code> |
| <code>count_nonzero([dim])</code> | See <code>torch.count_nonzero()</code> |
| <code>cov(*[, correction, fweights, aweights])</code> | See <code>torch.cov()</code> |
| <code>cpu([memory_format])</code> | Returns a copy of this object in CPU memory. |
| <code>cross(other[, dim])</code> | See <code>torch.cross()</code> |
| <code>crow_indices()</code> | Returns the tensor containing the compressed row indices of the <code>self</code> tensor when <code>self</code> is a sparse CSR tensor of layout <code>sparse_csr</code> . |
| <code>cuda([device, non_blocking, memory_format])</code> | Returns a copy of this object in CUDA memory. |
| <code>cummax(dim)</code> | See <code>torch.cummax()</code> |
| <code>cummin(dim)</code> | See <code>torch.cummin()</code> |
| <code>cumprod(dim[, dtype])</code> | See <code>torch.cumprod()</code> |
| <code>cumprod_(dim[, dtype])</code> | In-place version of <code>cumprod()</code> |
| <code>cumsum(dim[, dtype])</code> | See <code>torch.cumsum()</code> |

continues on next page

Table 22 – continued from previous page

| | |
|--|---|
| <code>cumsum_(dim[, dtype])</code> | In-place version of <code>cumsum()</code> |
| <code>data_ptr()</code> | Returns the address of the first element of <code>self</code> tensor. |
| <code>deg2rad()</code> | See <code>torch.deg2rad()</code> |
| <code>deg2rad_()</code> | In-place version of <code>deg2rad()</code> |
| <code>dense_dim()</code> | Return the number of dense dimensions in a sparse tensor <code>self</code> . |
| <code>dequantize()</code> | Given a quantized Tensor, dequantize it and return the dequantized float Tensor. |
| <code>det()</code> | See <code>torch.det()</code> |
| <code>detach</code> | Returns a new Tensor, detached from the current graph. |
| <code>detach_</code> | Detaches the Tensor from the graph that created it, making it a leaf. |
| <code>diag([diagonal])</code> | See <code>torch.diag()</code> |
| <code>diag_embed([offset, dim1, dim2])</code> | See <code>torch.diag_embed()</code> |
| <code>diagflat([offset])</code> | See <code>torch.diagflat()</code> |
| <code>diagonal([offset, dim1, dim2])</code> | See <code>torch.diagonal()</code> |
| <code>diagonal_scatter(src[, offset, dim1, dim2])</code> | See <code>torch.diagonal_scatter()</code> |
| <code>diff([n, dim, prepend, append])</code> | See <code>torch.diff()</code> |
| <code>digamma()</code> | See <code>torch.digamma()</code> |
| <code>digamma_()</code> | In-place version of <code>digamma()</code> |
| <code>dim()</code> | Returns the number of dimensions of <code>self</code> tensor. |
| <code>dim_order()</code> | Returns a tuple of int describing the dim order or physical layout of <code>self</code> . |
| <code>dist(other[, p])</code> | See <code>torch.dist()</code> |
| <code>div(value, *[, rounding_mode])</code> | See <code>torch.div()</code> |
| <code>div_(value, *[, rounding_mode])</code> | In-place version of <code>div()</code> |
| <code>divide(value, *[, rounding_mode])</code> | See <code>torch.divide()</code> |
| <code>divide_(value, *[, rounding_mode])</code> | In-place version of <code>divide()</code> |
| <code>dot(other)</code> | See <code>torch.dot()</code> |
| <code>double([memory_format])</code> | <code>self.double()</code> is equivalent to <code>self.to(torch.float64)</code> . |
| <code>dsplit(split_size_or_sections)</code> | See <code>torch.dsplit()</code> |
| <code>eig([eigenvectors])</code> | |
| <code>element_size()</code> | Returns the size in bytes of an individual element. |
| <code>eq(other)</code> | See <code>torch.eq()</code> |
| <code>eq_(other)</code> | In-place version of <code>eq()</code> |
| <code>equal(other)</code> | See <code>torch.equal()</code> |
| <code>erf()</code> | See <code>torch.erf()</code> |
| <code>erf_()</code> | In-place version of <code>erf()</code> |
| <code>erfc()</code> | See <code>torch.erfc()</code> |
| <code>erfc_()</code> | In-place version of <code>erfc()</code> |
| <code>erfinv()</code> | See <code>torch.erfinv()</code> |
| <code>erfinv_()</code> | In-place version of <code>erfinv()</code> |
| <code>exp()</code> | See <code>torch.exp()</code> |
| <code>exp2()</code> | See <code>torch.exp2()</code> |
| <code>exp2_()</code> | In-place version of <code>exp2()</code> |
| <code>exp_()</code> | In-place version of <code>exp()</code> |

continues on next page

Table 22 – continued from previous page

| | |
|---|---|
| <code>expand(*sizes)</code> | Returns a new view of the <code>self</code> tensor with singleton dimensions expanded to a larger size. |
| <code>expand_as(other)</code> | Expand this tensor to the same size as <code>other</code> . |
| <code>expm1()</code> | See <code>torch.expm1()</code> |
| <code>expm1_()</code> | In-place version of <code>expm1()</code> |
| <code>exponential_([lambd, generator])</code> | Fills <code>self</code> tensor with elements drawn from the exponential distribution: |
| <code>fill_(value)</code> | Fills <code>self</code> tensor with the specified value. |
| <code>fill_diagonal_(fill_value[, wrap])</code> | Fill the main diagonal of a tensor that has at least 2-dimensions. |
| <code>fix()</code> | See <code>torch.fix()</code> . |
| <code>fix_()</code> | In-place version of <code>fix()</code> |
| <code>flatten([start_dim, end_dim])</code> | See <code>torch.flatten()</code> |
| <code>flip(dims)</code> | See <code>torch.flip()</code> |
| <code>flipr()</code> | See <code>torch.flipr()</code> |
| <code>flipud()</code> | See <code>torch.flipud()</code> |
| <code>float([memory_format])</code> | <code>self.float()</code> is equivalent to <code>self.to(torch.float32)</code> . |
| <code>float_power(exponent)</code> | See <code>torch.float_power()</code> |
| <code>float_power_(exponent)</code> | In-place version of <code>float_power()</code> |
| <code>floor()</code> | See <code>torch.floor()</code> |
| <code>floor_()</code> | In-place version of <code>floor()</code> |
| <code>floor_divide(value)</code> | See <code>torch.floor_divide()</code> |
| <code>floor_divide_(value)</code> | In-place version of <code>floor_divide()</code> |
| <code>fmax(other)</code> | See <code>torch.fmax()</code> |
| <code>fmin(other)</code> | See <code>torch.fmin()</code> |
| <code>fmod(divisor)</code> | See <code>torch.fmod()</code> |
| <code>fmod_(divisor)</code> | In-place version of <code>fmod()</code> |
| <code>frac()</code> | See <code>torch.frac()</code> |
| <code>frac_()</code> | In-place version of <code>frac()</code> |
| <code>frexp(input)</code> | See <code>torch.frexp()</code> |
| <code>gather(dim, index)</code> | See <code>torch.gather()</code> |
| <code>gcd(other)</code> | See <code>torch.gcd()</code> |
| <code>gcd_(other)</code> | In-place version of <code>gcd()</code> |
| <code>ge(other)</code> | See <code>torch.ge()</code> . |
| <code>ge_(other)</code> | In-place version of <code>ge()</code> . |
| <code>geometric_(p, *[, generator])</code> | Fills <code>self</code> tensor with elements drawn from the geometric distribution: |
| <code>geqrf()</code> | See <code>torch.geqrf()</code> |
| <code>ger(vec2)</code> | See <code>torch.ger()</code> |
| <code>get_device()</code> | For CUDA tensors, this function returns the device ordinal of the GPU on which the tensor resides. |
| <code>greater(other)</code> | See <code>torch.greater()</code> . |
| <code>greater_(other)</code> | In-place version of <code>greater()</code> . |
| <code>greater_equal(other)</code> | See <code>torch.greater_equal()</code> . |
| <code>greater_equal_(other)</code> | In-place version of <code>greater_equal()</code> . |
| <code>gt(other)</code> | See <code>torch.gt()</code> . |
| <code>gt_(other)</code> | In-place version of <code>gt()</code> . |
| <code>half([memory_format])</code> | <code>self.half()</code> is equivalent to <code>self.to(torch.float16)</code> . |

continues on next page

Table 22 – continued from previous page

| | |
|--|---|
| <code>hardshrink([lambd])</code> | See <code>torch.nn.functional.hardshrink()</code> |
| <code>has_names</code> | Is True if any of this tensor's dimensions are named. |
| <code>heaviside(values)</code> | See <code>torch.heaviside()</code> |
| <code>heaviside_(values)</code> | In-place version of <code>heaviside()</code> |
| <code>histc([bins, min, max])</code> | See <code>torch.histc()</code> |
| <code>histogram(input, bins, *[, range, weight, ...])</code> | See <code>torch.histogram()</code> |
| <code>hsplit(split_size_or_sections)</code> | See <code>torch.hsplit()</code> |
| <code>hypot(other)</code> | See <code>torch.hypot()</code> |
| <code>hypot_(other)</code> | In-place version of <code>hypot()</code> |
| <code>i0()</code> | See <code>torch.i0()</code> |
| <code>i0_()</code> | In-place version of <code>i0()</code> |
| <code>igamma(other)</code> | See <code>torch.igamma()</code> |
| <code>igamma_(other)</code> | In-place version of <code>igamma()</code> |
| <code>igammac(other)</code> | See <code>torch.igammac()</code> |
| <code>igammac_(other)</code> | In-place version of <code>igammac()</code> |
| <code>index_add(dim, index, source, *[, alpha])</code> | Out-of-place version of <code>torch.Tensor.index_add_()</code> . |
| <code>index_add_(dim, index, source, *[, alpha])</code> | Accumulate the elements of <code>alpha times source</code> into the <code>self</code> tensor by adding to the indices in the order given in <code>index</code> . |
| <code>index_copy(dim, index, tensor2)</code> | Out-of-place version of <code>torch.Tensor.index_copy_()</code> . |
| <code>index_copy_(dim, index, tensor)</code> | Copies the elements of <code>tensor</code> into the <code>self</code> tensor by selecting the indices in the order given in <code>index</code> . |
| <code>index_fill(dim, index, value)</code> | Out-of-place version of <code>torch.Tensor.index_fill_()</code> . |
| <code>index_fill_(dim, index, value)</code> | Fills the elements of the <code>self</code> tensor with <code>value</code> by selecting the indices in the order given in <code>index</code> . |
| <code>index_put(indices, values[, accumulate])</code> | Out-place version of <code>index_put_()</code> . |
| <code>index_put_(indices, values[, accumulate])</code> | Puts values from the tensor values into the tensor <code>self</code> using the indices specified in <code>indices</code> (which is a tuple of Tensors). |
| <code>index_reduce</code> | |
| <code>index_reduce_(dim, index, source, reduce, *)</code> | Accumulate the elements of <code>source</code> into the <code>self</code> tensor by accumulating to the indices in the order given in <code>index</code> using the reduction given by the <code>reduce</code> argument. |
| <code>index_select(dim, index)</code> | See <code>torch.index_select()</code> |
| <code>indices()</code> | Return the indices tensor of a sparse COO tensor. |
| <code>inner(other)</code> | See <code>torch.inner()</code> . |
| <code>int([memory_format])</code> | <code>self.int()</code> is equivalent to <code>self.to(torch.int32)</code> . |
| <code>int_repr()</code> | Given a quantized Tensor, <code>self.int_repr()</code> returns a CPU Tensor with <code>uint8_t</code> as data type that stores the underlying <code>uint8_t</code> values of the given Tensor. |
| <code>inverse()</code> | See <code>torch.inverse()</code> |
| <code>ipu([device, non_blocking, memory_format])</code> | Returns a copy of this object in IPU memory. |

continues on next page

Table 22 – continued from previous page

| | |
|--|---|
| <code>is_coalesced()</code> | Returns True if <code>self</code> is a sparse COO tensor that is coalesced, False otherwise. |
| <code>is_complex()</code> | Returns True if the data type of <code>self</code> is a complex data type. |
| <code>is_conj()</code> | Returns True if the conjugate bit of <code>self</code> is set to true. |
| <code>is_contiguous([memory_format])</code> | Returns True if <code>self</code> tensor is contiguous in memory in the order specified by memory format. |
| <code>is_distributed</code> | |
| <code>is_floating_point()</code> | Returns True if the data type of <code>self</code> is a floating point data type. |
| <code>is_inference()</code> | See <code>torch.is_inference()</code> |
| <code>is_neg()</code> | Returns True if the negative bit of <code>self</code> is set to true. |
| <code>is_nonzero</code> | |
| <code>is_pinned</code> | Returns true if this tensor resides in pinned memory. |
| <code>is_same_size</code> | |
| <code>is_set_to(tensor)</code> | Returns True if both tensors are pointing to the exact same memory (same storage, offset, size and stride). |
| <code>is_shared()</code> | Checks if tensor is in shared memory. |
| <code>is_signed()</code> | Returns True if the data type of <code>self</code> is a signed data type. |
| <code>isclose(other[, rtol, atol, equal_nan])</code> | See <code>torch.isclose()</code> |
| <code>isfinite()</code> | See <code>torch.isfinite()</code> |
| <code>isinf()</code> | See <code>torch.isinf()</code> |
| <code>isnan()</code> | See <code>torch.isnan()</code> |
| <code>isneginf()</code> | See <code>torch.isneginf()</code> |
| <code>isposinf()</code> | See <code>torch.isposinf()</code> |
| <code>isreal()</code> | See <code>torch.isreal()</code> |
| <code>istft(n_fft[, hop_length, win_length, ...])</code> | See <code>torch.istft()</code> |
| <code>item()</code> | Returns the value of this tensor as a standard Python number. |
| <code>kron(other)</code> | See <code>torch.kron()</code> |
| <code>kthvalue(k[, dim, keepdim])</code> | See <code>torch.kthvalue()</code> |
| <code>lcm(other)</code> | See <code>torch.lcm()</code> |
| <code>lcm_(other)</code> | In-place version of <code>lcm()</code> |
| <code>ldexp(other)</code> | See <code>torch.ldexp()</code> |
| <code>ldexp_(other)</code> | In-place version of <code>ldexp()</code> |
| <code>le(other)</code> | See <code>torch.le()</code> . |
| <code>le_(other)</code> | In-place version of <code>le()</code> . |
| <code>lerp(end, weight)</code> | See <code>torch.lerp()</code> |
| <code>lerp_(end, weight)</code> | In-place version of <code>lerp()</code> |
| <code>less</code> | <code>lt(other) -> Tensor</code> |
| <code>less_(other)</code> | In-place version of <code>less()</code> . |
| <code>less_equal(other)</code> | See <code>torch.less_equal()</code> . |
| <code>less_equal_(other)</code> | In-place version of <code>less_equal()</code> . |
| <code>lgamma()</code> | See <code>torch.lgamma()</code> |
| <code>lgamma_()</code> | In-place version of <code>lgamma()</code> |
| <code>log()</code> | See <code>torch.log()</code> |
| <code>log10()</code> | See <code>torch.log10()</code> |

continues on next page

Table 22 – continued from previous page

| | |
|--|--|
| <code>log10_()</code> | In-place version of <code>log10()</code> |
| <code>log1p()</code> | See <code>torch.log1p()</code> |
| <code>log1p_()</code> | In-place version of <code>log1p()</code> |
| <code>log2()</code> | See <code>torch.log2()</code> |
| <code>log2_()</code> | In-place version of <code>log2()</code> |
| <code>log_()</code> | In-place version of <code>log()</code> |
| <code>log_normal_([mean, std, generator])</code> | Fills <code>self</code> tensor with numbers samples from the log-normal distribution parameterized by the given mean μ and standard deviation σ . |
| <code>log_softmax</code> | |
| <code>logaddexp(other)</code> | See <code>torch.logaddexp()</code> |
| <code>logaddexp2(other)</code> | See <code>torch.logaddexp2()</code> |
| <code>logcumsumexp(dim)</code> | See <code>torch.logcumsumexp()</code> |
| <code>logdet()</code> | See <code>torch.logdet()</code> |
| <code>logical_and()</code> | See <code>torch.logical_and()</code> |
| <code>logical_and_()</code> | In-place version of <code>logical_and()</code> |
| <code>logical_not()</code> | See <code>torch.logical_not()</code> |
| <code>logical_not_()</code> | In-place version of <code>logical_not()</code> |
| <code>logical_or()</code> | See <code>torch.logical_or()</code> |
| <code>logical_or_()</code> | In-place version of <code>logical_or()</code> |
| <code>logical_xor()</code> | See <code>torch.logical_xor()</code> |
| <code>logical_xor_()</code> | In-place version of <code>logical_xor()</code> |
| <code>logit()</code> | See <code>torch.logit()</code> |
| <code>logit_()</code> | In-place version of <code>logit()</code> |
| <code>logsumexp(dim[, keepdim])</code> | See <code>torch.logsumexp()</code> |
| <code>long([memory_format])</code> | <code>self.long()</code> is equivalent to <code>self.to(torch.int64)</code> . |
| <code>lstsq(other)</code> | |
| <code>lt(other)</code> | See <code>torch.lt()</code> . |
| <code>lt_(other)</code> | In-place version of <code>lt()</code> . |
| <code>lu([pivot, get_infos])</code> | See <code>torch.lu()</code> |
| <code>lu_solve(LU_data, LU_pivots)</code> | See <code>torch.lu_solve()</code> |
| <code>map2_</code> | |
| <code>map_(tensor, callable)</code> | Applies <code>callable</code> for each element in <code>self</code> tensor and the given <code>tensor</code> and stores the results in <code>self</code> tensor. |
| <code>masked_fill(mask, value)</code> | Out-of-place version of <code>torch.Tensor.masked_fill_()</code> |
| <code>masked_fill_(mask, value)</code> | Fills elements of <code>self</code> tensor with <code>value</code> where <code>mask</code> is <code>True</code> . |
| <code>masked_scatter(mask, tensor)</code> | Out-of-place version of <code>torch.Tensor.masked_scatter_()</code> |
| <code>masked_scatter_(mask, source)</code> | Copies elements from <code>source</code> into <code>self</code> tensor at positions where the <code>mask</code> is <code>True</code> . |
| <code>masked_select(mask)</code> | See <code>torch.masked_select()</code> |
| <code>materialize(shape[, device, dtype])</code> | Create a <code>Parameter</code> with the same properties of the uninitialized one. |
| <code>matmul(tensor2)</code> | See <code>torch.matmul()</code> |

continues on next page

Table 22 – continued from previous page

| | |
|---|--|
| <code>matrix_exp()</code> | See <code>torch.matrix_exp()</code> |
| <code>matrix_power(n)</code> | |
| Note: <code>matrix_power()</code> is deprecated, use <code>torch.linalg.matrix_power()</code> instead. | |
| <code>max([dim, keepdim])</code> | See <code>torch.max()</code> |
| <code>maximum(other)</code> | See <code>torch.maximum()</code> |
| <code>mean([dim, keepdim, dtype])</code> | See <code>torch.mean()</code> |
| <code>median([dim, keepdim])</code> | See <code>torch.median()</code> |
| <code>min([dim, keepdim])</code> | See <code>torch.min()</code> |
| <code>minimum(other)</code> | See <code>torch.minimum()</code> |
| <code>mm(mat2)</code> | See <code>torch.mm()</code> |
| <code>mode([dim, keepdim])</code> | See <code>torch.mode()</code> |
| <code>moveaxis(source, destination)</code> | See <code>torch.moveaxis()</code> |
| <code>movedim(source, destination)</code> | See <code>torch.movedim()</code> |
| <code>msort()</code> | See <code>torch.msort()</code> |
| <code>mul(value)</code> | See <code>torch.mul()</code> . |
| <code>mul_(value)</code> | In-place version of <code>mul()</code> . |
| <code>multinomial(num_samples[, replacement, ...])</code> | See <code>torch.multinomial()</code> |
| <code>multiply(value)</code> | See <code>torch.multiply()</code> . |
| <code>multiply_(value)</code> | In-place version of <code>multiply()</code> . |
| <code>mv(vec)</code> | See <code>torch.mv()</code> |
| <code>mvlgamma(p)</code> | See <code>torch.mvlgamma()</code> |
| <code>mvlgamma_(p)</code> | In-place version of <code>mvlgamma()</code> |
| <code>nan_to_num([nan, posinf, neginf])</code> | See <code>torch.nan_to_num()</code> . |
| <code>nan_to_num_([nan, posinf, neginf])</code> | In-place version of <code>nan_to_num()</code> . |
| <code>nanmean([dim, keepdim, dtype])</code> | See <code>torch.nanmean()</code> |
| <code>nanmedian([dim, keepdim])</code> | See <code>torch.nanmedian()</code> |
| <code>nanquantile(q[, dim, keepdim, interpolation])</code> | See <code>torch.nanquantile()</code> |
| <code>nansum([dim, keepdim, dtype])</code> | See <code>torch.nansum()</code> |
| <code>narrow(dimension, start, length)</code> | See <code>torch.narrow()</code> . |
| <code>narrow_copy(dimension, start, length)</code> | See <code>torch.narrow_copy()</code> . |
| <code>ndimension()</code> | Alias for <code>dim()</code> |
| <code>ne(other)</code> | See <code>torch.ne()</code> . |
| <code>ne_(other)</code> | In-place version of <code>ne()</code> . |
| <code>neg()</code> | See <code>torch.neg()</code> |
| <code>neg_()</code> | In-place version of <code>neg()</code> |
| <code>negative()</code> | See <code>torch.negative()</code> |
| <code>negative_()</code> | In-place version of <code>negative()</code> |
| <code>nelement()</code> | Alias for <code>numel()</code> |
| <code>new</code> | |
| <code>new_empty(size, *[, dtype, device, ...])</code> | Returns a Tensor of size <code>size</code> filled with uninitialized data. |
| <code>new_empty_strided(size, stride[, dtype, ...])</code> | Returns a Tensor of size <code>size</code> and strides <code>stride</code> filled with uninitialized data. |
| <code>new_full(size, fill_value, *[, dtype, ...])</code> | Returns a Tensor of size <code>size</code> filled with <code>fill_value</code> . |
| <code>new_ones(size, *[, dtype, device, ...])</code> | Returns a Tensor of size <code>size</code> filled with 1. |

continues on next page

Table 22 – continued from previous page

| | |
|---|--|
| <code>new_tensor(data, *, dtype, device, ...)</code> | Returns a new Tensor with data as the tensor data. |
| <code>new_zeros(size, *, dtype, device, ...)</code> | Returns a Tensor of size <code>size</code> filled with 0. |
| <code>nextafter(other)</code> | See <code>torch.nextafter()</code> |
| <code>nextafter_(other)</code> | In-place version of <code>nextafter()</code> |
| <code>nonzero()</code> | See <code>torch.nonzero()</code> |
| <code>nonzero_static(input, *, size[, fill_value])</code> | Returns a 2-D tensor where each row is the index for a non-zero value. |
| <code>norm([p, dim, keepdim, dtype])</code> | See <code>torch.norm()</code> |
| <code>normal_([mean, std, generator])</code> | Fills self tensor with elements samples from the normal distribution parameterized by mean and std. |
| <code>not_equal(other)</code> | See <code>torch.not_equal()</code> . |
| <code>not_equal_(other)</code> | In-place version of <code>not_equal()</code> . |
| <code>numel()</code> | See <code>torch.numel()</code> |
| <code>numpy(*[, force])</code> | Returns the tensor as a NumPy ndarray. |
| <code>orgqr(input2)</code> | See <code>torch.orgqr()</code> |
| <code>ormqr(input2, input3[, left, transpose])</code> | See <code>torch.ormqr()</code> |
| <code>outer(vec2)</code> | See <code>torch.outer()</code> . |
| <code>permute(*dims)</code> | See <code>torch.permute()</code> |
| <code>pin_memory()</code> | Copies the tensor to pinned memory, if it's not already pinned. |
| <code>pinverse()</code> | See <code>torch.pinverse()</code> |
| <code>polygamma(n)</code> | See <code>torch.polygamma()</code> |
| <code>polygamma_(n)</code> | In-place version of <code>polygamma()</code> |
| <code>positive()</code> | See <code>torch.positive()</code> |
| <code>pow(exponent)</code> | See <code>torch.pow()</code> |
| <code>pow_(exponent)</code> | In-place version of <code>pow()</code> |
| <code>prelu</code> | |
| <code>prod([dim, keepdim, dtype])</code> | See <code>torch.prod()</code> |
| <code>put(input, index, source[, accumulate])</code> | Out-of-place version of <code>torch.Tensor.put_()</code> . |
| <code>put_(index, source[, accumulate])</code> | Copies the elements from <code>source</code> into the positions specified by <code>index</code> . |
| <code>q_per_channel_axis()</code> | Given a Tensor quantized by linear (affine) per-channel quantization, returns the index of dimension on which per-channel quantization is applied. |
| <code>q_per_channel_scales()</code> | Given a Tensor quantized by linear (affine) per-channel quantization, returns a Tensor of scales of the underlying quantizer. |
| <code>q_per_channel_zero_points()</code> | Given a Tensor quantized by linear (affine) per-channel quantization, returns a tensor of zero_points of the underlying quantizer. |
| <code>q_scale()</code> | Given a Tensor quantized by linear(affine) quantization, returns the scale of the underlying quantizer(). |
| <code>q_zero_point()</code> | Given a Tensor quantized by linear(affine) quantization, returns the zero_point of the underlying quantizer(). |
| <code>qr([some])</code> | See <code>torch.qr()</code> |
| <code>qscheme()</code> | Returns the quantization scheme of a given QTensor. |
| <code>quantile(q[, dim, keepdim, interpolation])</code> | See <code>torch.quantile()</code> |
| <code>rad2deg()</code> | See <code>torch.rad2deg()</code> |
| <code>rad2deg_()</code> | In-place version of <code>rad2deg()</code> |

continues on next page

Table 22 – continued from previous page

| | |
|---|--|
| <code>random_([from, to, generator])</code> | Fills <code>self</code> tensor with numbers sampled from the discrete uniform distribution over <code>[from, to - 1]</code> . |
| <code>ravel()</code> | see <code>torch.ravel()</code> |
| <code>reciprocal()</code> | See <code>torch.reciprocal()</code> |
| <code>reciprocal_()</code> | In-place version of <code>reciprocal()</code> |
| <code>record_stream(stream)</code> | Ensures that the tensor memory is not reused for another tensor until all current work queued on <code>stream</code> are complete. |
| <code>refine_names(*names)</code> | Refines the dimension names of <code>self</code> according to <code>names</code> . |
| <code>register_hook(hook)</code> | Registers a backward hook. |
| <code>register_post_accumulate_grad_hook(hook)</code> | Registers a backward hook that runs after grad accumulation. |
| <code>reinforce(reward)</code> | |
| <code>relu</code> | |
| <code>relu_</code> | |
| <code>remainder(divisor)</code> | See <code>torch.remainder()</code> |
| <code>remainder_(divisor)</code> | In-place version of <code>remainder()</code> |
| <code>rename(*names, **rename_map)</code> | Renames dimension names of <code>self</code> . |
| <code>rename_(*names, **rename_map)</code> | In-place version of <code>rename()</code> . |
| <code>renorm(p, dim, maxnorm)</code> | See <code>torch.renorm()</code> |
| <code>renorm_(p, dim, maxnorm)</code> | In-place version of <code>renorm()</code> |
| <code>repeat(*sizes)</code> | Repeats this tensor along the specified dimensions. |
| <code>repeat_interleave(repeats[, dim, output_size])</code> | See <code>torch.repeat_interleave()</code> . |
| <code>requires_grad_([requires_grad])</code> | Change if autograd should record operations on this tensor: sets this tensor's <code>requires_grad</code> attribute in-place. |
| <code>reshape(*shape)</code> | Returns a tensor with the same data and number of elements as <code>self</code> but with the specified shape. |
| <code>reshape_as(other)</code> | Returns this tensor as the same shape as <code>other</code> . |
| <code>resize(*sizes)</code> | |
| <code>resize_(*sizes[, memory_format])</code> | Resizes <code>self</code> tensor to the specified size. |
| <code>resize_as(tensor)</code> | |
| <code>resize_as_(tensor[, memory_format])</code> | Resizes the <code>self</code> tensor to be the same size as the specified tensor. |
| <code>resize_as_sparse_</code> | |
| <code>resolve_conj()</code> | See <code>torch.resolve_conj()</code> |
| <code>resolve_neg()</code> | See <code>torch.resolve_neg()</code> |
| <code>retain_grad()</code> | Enables this Tensor to have their grad populated during <code>backward()</code> . |
| <code>roll(shifts, dims)</code> | See <code>torch.roll()</code> |
| <code>rot90(k, dims)</code> | See <code>torch.rot90()</code> |
| <code>round([decimals])</code> | See <code>torch.round()</code> |
| <code>round_([decimals])</code> | In-place version of <code>round()</code> |

continues on next page

Table 22 – continued from previous page

| | |
|--|---|
| row_indices | |
| <code>rsqrt()</code> | See <code>torch.rsqrt()</code> |
| <code>rsqrt_()</code> | In-place version of <code>rsqrt()</code> |
| <code>scatter(dim, index, src)</code> | Out-of-place version of <code>torch.Tensor.scatter_()</code> |
| <code>scatter_(dim, index, src[, reduce])</code> | Writes all values from the tensor <code>src</code> into <code>self</code> at the indices specified in the <code>index</code> tensor. |
| <code>scatter_add(dim, index, src)</code> | Out-of-place version of <code>torch.Tensor.scatter_add_()</code> |
| <code>scatter_add_(dim, index, src)</code> | Adds all values from the tensor <code>src</code> into <code>self</code> at the indices specified in the <code>index</code> tensor in a similar fashion as <code>scatter_()</code> . |
| <code>scatter_reduce(dim, index, src, reduce, *[, ...])</code> | Out-of-place version of <code>torch.Tensor.scatter_reduce_()</code> |
| <code>scatter_reduce_(dim, index, src, reduce, *)</code> | Reduces all values from the <code>src</code> tensor to the indices specified in the <code>index</code> tensor in the <code>self</code> tensor using the applied reduction defined via the <code>reduce</code> argument ("sum", "prod", "mean", "amax", "amin"). |
| <code>select(dim, index)</code> | See <code>torch.select()</code> |
| <code>select_scatter(src, dim, index)</code> | See <code>torch.select_scatter()</code> |
| <code>set_([source, storage_offset, size, stride])</code> | Sets the underlying storage, size, and strides. |
| <code>sgn()</code> | See <code>torch.sgn()</code> |
| <code>sgn_()</code> | In-place version of <code>sgn()</code> |
| <code>share_memory_()</code> | Moves the underlying storage to shared memory. |
| <code>short([memory_format])</code> | <code>self.short()</code> is equivalent to <code>self.to(torch.int16)</code> . |
| <code>sigmoid()</code> | See <code>torch.sigmoid()</code> |
| <code>sigmoid_()</code> | In-place version of <code>sigmoid()</code> |
| <code>sign()</code> | See <code>torch.sign()</code> |
| <code>sign_()</code> | In-place version of <code>sign()</code> |
| <code>signbit()</code> | See <code>torch.signbit()</code> |
| <code>sin()</code> | See <code>torch.sin()</code> |
| <code>sin_()</code> | In-place version of <code>sin()</code> |
| <code>sinc()</code> | See <code>torch.sinc()</code> |
| <code>sinc_()</code> | In-place version of <code>sinc()</code> |
| <code>sinh()</code> | See <code>torch.sinh()</code> |
| <code>sinh_()</code> | In-place version of <code>sinh()</code> |
| <code>size([dim])</code> | Returns the size of the <code>self</code> tensor. |
| <code>slice_scatter(src[, dim, start, end, step])</code> | See <code>torch.slice_scatter()</code> |
| <code>slogdet()</code> | See <code>torch.slogdet()</code> |
| <code>smm(mat)</code> | See <code>torch.smm()</code> |
| <code>softmax(dim)</code> | Alias for <code>torch.nn.functional.softmax()</code> . |
| <code>solve(other)</code> | |
| <code>sort([dim, descending])</code> | See <code>torch.sort()</code> |
| <code>sparse_dim()</code> | Return the number of sparse dimensions in a sparse tensor <code>self</code> . |
| <code>sparse_mask(mask)</code> | Returns a new sparse tensor with values from a strided tensor <code>self</code> filtered by the indices of the sparse tensor <code>mask</code> . |

continues on next page

Table 22 – continued from previous page

| | |
|--|--|
| <code>sparse_resize_(size, sparse_dim, dense_dim)</code> | Resizes <code>self</code> sparse tensor to the desired size and the number of sparse and dense dimensions. |
| <code>sparse_resize_and_clear_(size, sparse_dim, ...)</code> | Removes all specified elements from a sparse tensor <code>self</code> and resizes <code>self</code> to the desired size and the number of sparse and dense dimensions. |
| <code>split(split_size[, dim])</code> | See <code>torch.split()</code> |
| <code>split_with_sizes</code> | |
| <code>sqrt()</code> | See <code>torch.sqrt()</code> |
| <code>sqrt_()</code> | In-place version of <code>sqrt()</code> |
| <code>square()</code> | See <code>torch.square()</code> |
| <code>square_()</code> | In-place version of <code>square()</code> |
| <code>squeeze([dim])</code> | See <code>torch.squeeze()</code> |
| <code>squeeze_([dim])</code> | In-place version of <code>squeeze()</code> |
| <code>sspaddmm(mat1, mat2, *[, beta, alpha])</code> | See <code>torch.sspaddmm()</code> |
| <code>std([dim, correction, keepdim])</code> | See <code>torch.std()</code> |
| <code>stft(n_fft[, hop_length, win_length, ...])</code> | See <code>torch.stft()</code> |
| <code>storage()</code> | Returns the underlying <code>TypedStorage</code> . |
| <code>storage_offset()</code> | Returns <code>self</code> tensor's offset in the underlying storage in terms of number of storage elements (not bytes). |
| <code>storage_type()</code> | Returns the type of the underlying storage. |
| <code>stride(dim)</code> | Returns the stride of <code>self</code> tensor. |
| <code>sub(other, *[, alpha])</code> | See <code>torch.sub()</code> . |
| <code>sub_(other, *[, alpha])</code> | In-place version of <code>sub()</code> |
| <code>subtract(other, *[, alpha])</code> | See <code>torch.subtract()</code> . |
| <code>subtract_(other, *[, alpha])</code> | In-place version of <code>subtract()</code> . |
| <code>sum([dim, keepdim, dtype])</code> | See <code>torch.sum()</code> |
| <code>sum_to_size(*size)</code> | Sum this tensor to size. |
| <code>svd([some, compute_uv])</code> | See <code>torch.svd()</code> |
| <code>swapaxes(axis0, axis1)</code> | See <code>torch.swapaxes()</code> |
| <code>swapaxes_(axis0, axis1)</code> | In-place version of <code>swapaxes()</code> |
| <code>swapdims(dim0, dim1)</code> | See <code>torch.swapdims()</code> |
| <code>swapdims_(dim0, dim1)</code> | In-place version of <code>swapdims()</code> |
| <code>symeig([eigenvectors])</code> | |
| <code>t()</code> | See <code>torch.t()</code> |
| <code>t_()</code> | In-place version of <code>t()</code> |
| <code>take(indices)</code> | See <code>torch.take()</code> |
| <code>take_along_dim(indices, dim)</code> | See <code>torch.take_along_dim()</code> |
| <code>tan()</code> | See <code>torch.tan()</code> |
| <code>tan_()</code> | In-place version of <code>tan()</code> |
| <code>tanh()</code> | See <code>torch.tanh()</code> |
| <code>tanh_()</code> | In-place version of <code>tanh()</code> |
| <code>tensor_split(indices_or_sections[, dim])</code> | See <code>torch.tensor_split()</code> |
| <code>tile(dims)</code> | See <code>torch.tile()</code> |
| <code>to(*args, **kwargs)</code> | Performs Tensor dtype and/or device conversion. |
| <code>to_dense([dtype, masked_grad])</code> | Creates a strided copy of <code>self</code> if <code>self</code> is not a strided tensor, otherwise returns <code>self</code> . |
| <code>to_mkldnn()</code> | Returns a copy of the tensor in <code>torch.mkldnn</code> layout. |
| <code>to_padded_tensor(padding[, output_size])</code> | See <code>to_padded_tensor()</code> |

continues on next page

Table 22 – continued from previous page

| | |
|---|---|
| <code>to_sparse(sparseDims)</code> | Returns a sparse copy of the tensor. |
| <code>to_sparse_bsc(blocksize, dense_dim)</code> | Convert a tensor to a block sparse column (BSC) storage format of given blocksize. |
| <code>to_sparse_bsr(blocksize, dense_dim)</code> | Convert a tensor to a block sparse row (BSR) storage format of given blocksize. |
| <code>to_sparse_coo()</code> | Convert a tensor to coordinate format. |
| <code>to_sparse_csc()</code> | Convert a tensor to compressed column storage (CSC) format. |
| <code>to_sparse_csr([dense_dim])</code> | Convert a tensor to compressed row storage format (CSR). |
| <code>tolist()</code> | Returns the tensor as a (nested) list. |
| <code>topk(k[, dim, largest, sorted])</code> | See <code>torch.topk()</code> |
| <code>trace()</code> | See <code>torch.trace()</code> |
| <code>transpose(dim0, dim1)</code> | See <code>torch.transpose()</code> |
| <code>transpose_(dim0, dim1)</code> | In-place version of <code>transpose()</code> |
| <code>triangular_solve(A[, upper, transpose, ...])</code> | See <code>torch.triangular_solve()</code> |
| <code>tril([diagonal])</code> | See <code>torch.tril()</code> |
| <code>tril_([diagonal])</code> | In-place version of <code>tril()</code> |
| <code>triu([diagonal])</code> | See <code>torch.triu()</code> |
| <code>triu_([diagonal])</code> | In-place version of <code>triu()</code> |
| <code>true_divide(value)</code> | See <code>torch.true_divide()</code> |
| <code>true_divide_(value)</code> | In-place version of <code>true_divide_()</code> |
| <code>trunc()</code> | See <code>torch.trunc()</code> |
| <code>trunc_()</code> | In-place version of <code>trunc()</code> |
| <code>type([dtype, non_blocking])</code> | Returns the type if <i>dtype</i> is not provided, else casts this object to the specified type. |
| <code>type_as(tensor)</code> | Returns this tensor cast to the type of the given tensor. |
| <code>unbind([dim])</code> | See <code>torch.unbind()</code> |
| <code>unflatten(dim, sizes)</code> | See <code>torch.unflatten()</code> . |
| <code>unfold(dimension, size, step)</code> | Returns a view of the original tensor which contains all slices of size <i>size</i> from <i>self</i> tensor in the dimension <i>dimension</i> . |
| <code>uniform_([from, to, generator])</code> | Fills <i>self</i> tensor with numbers sampled from the continuous uniform distribution: |
| <code>unique([sorted, return_inverse, ...])</code> | Returns the unique elements of the input tensor. |
| <code>unique_consecutive([return_inverse, ...])</code> | Eliminates all but the first element from every consecutive group of equivalent elements. |
| <code>unsafe_chunk(chunks[, dim])</code> | See <code>torch.unsafe_chunk()</code> |
| <code>unsafe_split(split_size[, dim])</code> | See <code>torch.unsafe_split()</code> |
| <code>unsafe_split_with_sizes</code> | |
| <code>unsqueeze(dim)</code> | See <code>torch.unsqueeze()</code> |
| <code>unsqueeze_(dim)</code> | In-place version of <code>unsqueeze()</code> |
| <code>untyped_storage()</code> | Returns the underlying <code>UntypedStorage</code> . |
| <code>values()</code> | Return the values tensor of a sparse COO tensor. |
| <code>var([dim, correction, keepdim])</code> | See <code>torch.var()</code> |
| <code>vdot(other)</code> | See <code>torch.vdot()</code> |
| <code>view(*shape)</code> | Returns a new tensor with the same data as the <i>self</i> tensor but of a different <i>shape</i> . |
| <code>view_as(other)</code> | View this tensor as the same size as <i>other</i> . |
| <code>vsplit(split_size_or_sections)</code> | See <code>torch.vsplit()</code> |

continues on next page

Table 22 – continued from previous page

| | |
|---|---|
| <code>where(condition, y)</code> | <code>self.where(condition, y)</code> is equivalent to <code>torch.where(condition, self, y)</code> . |
| <code>xlogy(other)</code> | See <code>torch.xlogy()</code> |
| <code>xlogy_(other)</code> | In-place version of <code>xlogy()</code> |
| <code>xpu([device, non_blocking, memory_format])</code> | Returns a copy of this object in XPU memory. |
| <code>zero_()</code> | Fills <code>self</code> tensor with zeros. |

Attributes

| | |
|----------------------------|---|
| <code>H</code> | Returns a view of a matrix (2-D tensor) conjugated and transposed. |
| <code>T</code> | Returns a view of this tensor with its dimensions reversed. |
| <code>data</code> | |
| <code>device</code> | Is the <code>torch.device</code> where this Tensor is. |
| <code>dtype</code> | |
| <code>grad</code> | This attribute is <code>None</code> by default and becomes a Tensor the first time a call to <code>backward()</code> computes gradients for <code>self</code> . |
| <code>grad_fn</code> | |
| <code>imag</code> | Returns a new tensor containing imaginary values of the <code>self</code> tensor. |
| <code>is_cpu</code> | Is <code>True</code> if the Tensor is stored on the CPU, <code>False</code> otherwise. |
| <code>is_cuda</code> | Is <code>True</code> if the Tensor is stored on the GPU, <code>False</code> otherwise. |
| <code>is_ipu</code> | Is <code>True</code> if the Tensor is stored on the IPU, <code>False</code> otherwise. |
| <code>is_leaf</code> | All Tensors that have <code>requires_grad</code> which is <code>False</code> will be leaf Tensors by convention. |
| <code>is_meta</code> | Is <code>True</code> if the Tensor is a meta tensor, <code>False</code> otherwise. |
| <code>is_mkldnn</code> | |
| <code>is_mps</code> | Is <code>True</code> if the Tensor is stored on the MPS device, <code>False</code> otherwise. |
| <code>is_nested</code> | |
| <code>is_ort</code> | |
| <code>is_quantized</code> | Is <code>True</code> if the Tensor is quantized, <code>False</code> otherwise. |
| <code>is_sparse</code> | Is <code>True</code> if the Tensor uses sparse COO storage layout, <code>False</code> otherwise. |
| <code>is_sparse_csr</code> | Is <code>True</code> if the Tensor uses sparse CSR storage layout, <code>False</code> otherwise. |

continues on next page

Table 23 – continued from previous page

| | |
|----------------------------|---|
| <code>is_vulkan</code> | |
| <code>is_xla</code> | Is True if the Tensor is stored on an XLA device, False otherwise. |
| <code>is_xpu</code> | Is True if the Tensor is stored on the XPU, False otherwise. |
| <code>itemsize</code> | Alias for <code>element_size()</code> |
| <code>layout</code> | |
| <code>mH</code> | Accessing this property is equivalent to calling <code>adjoint()</code> . |
| <code>mT</code> | Returns a view of this tensor with the last two dimensions transposed. |
| <code>name</code> | |
| <code>names</code> | Stores names for each of this tensor's dimensions. |
| <code>nbytes</code> | Returns the number of bytes consumed by the "view" of elements of the Tensor if the Tensor does not use sparse storage layout. |
| <code>ndim</code> | Alias for <code>dim()</code> |
| <code>output_nr</code> | |
| <code>real</code> | Returns a new tensor containing real values of the <code>self</code> tensor for a complex-valued input tensor. |
| <code>requires_grad</code> | Is True if gradients need to be computed for this Tensor, False otherwise. |
| <code>retains_grad</code> | Is True if this Tensor is non-leaf and its <code>grad</code> is enabled to be populated during <code>backward()</code> , False otherwise. |
| <code>shape</code> | Returns the size of the <code>self</code> tensor. |
| <code>volatile</code> | |

property `is_leaf`: bool

All Tensors that have `requires_grad` which is False will be leaf Tensors by convention.

For Tensors that have `requires_grad` which is True, they will be leaf Tensors if they were created by the user. This means that they are not the result of an operation and so `grad_fn` is None.

Only leaf Tensors will have their `grad` populated during a call to `backward()`. To get `grad` populated for non-leaf Tensors, you can use `retain_grad()`.

Example:

```
>>> a = torch.rand(10, requires_grad=True)
>>> a.is_leaf
True
>>> b = torch.rand(10, requires_grad=True).cuda()
>>> b.is_leaf
False
# b was created by the operation that cast a cpu Tensor into a cuda Tensor
>>> c = torch.rand(10, requires_grad=True) + 2
>>> c.is_leaf
```

(continues on next page)

(continued from previous page)

```

False
# c was created by the addition operation
>>> d = torch.rand(10).cuda()
>>> d.is_leaf
True
# d does not require gradients and so has no operation creating it (that is,
→ tracked by the autograd engine)
>>> e = torch.rand(10).cuda().requires_grad_()
>>> e.is_leaf
True
# e requires gradients and has no operations creating it
>>> f = torch.rand(10, requires_grad=True, device="cuda")
>>> f.is_leaf
True
# f requires grad, has no operation creating it

```

materialize(*shape*, *device=None*, *dtype=None*)

Create a Parameter with the same properties of the uninitialized one. Given a shape, it materializes a parameter in the same device and with the same *dtype* as the current one or the specified ones in the arguments.

Parameters

- **shape** (*Tuple[int, ...]*) – (tuple): the shape for the materialized tensor.
- **device** (*torch.device*) – the desired device of the parameters and buffers in this module. Optional.
- **dtype** (*torch.dtype*) – the desired floating point type of the floating point parameters and buffers in this module. Optional.

Return type

None

share_memory_()

Moves the underlying storage to shared memory.

This is a no-op if the underlying storage is already in shared memory and for CUDA tensors. Tensors in shared memory cannot be resized.

Return type

[UninitializedParameter](#)

pytorch_pfn_extras.nn.modules.lazy_linear

Classes

| | |
|--|--|
| pytorch_pfn_extras.nn.modules.lazy_linear.LazyInitializationMixin(...) | A mixin for modules that lazily initialize buffers and parameters. |
| pytorch_pfn_extras.nn.modules.lazy_linear.LazyLinear(...) | Linear module with lazy weight initialization. |
| pytorch_pfn_extras.nn.modules.lazy_linear.UninitializedParameter(...) | |

pytorch_pfn_extras.nn.modules.lazy_linear.LazyInitializationMixin

class pytorch_pfn_extras.nn.modules.lazy_linear.LazyInitializationMixin(*args, **kwargs)

Bases: object

A mixin for modules that lazily initialize buffers and parameters.

Unlike regular modules, subclasses of this module can initialize buffers and parameters outside of the constructor (`__init__`). This allows you to, for example, initialize parameters in `forward` method to determine the shape of the weight based on the initial input.

Be sure to run “dummy” forward once to initialize all parameters that should be trained, before passing `module.parameters()` to an optimizer; otherwise weights initialized after `module.parameters()` (e.g., in `forward` function) will never be trained.

Note that lazy modules cannot validate if the shape is correct during deserialization. Also note that the initial weights may become different from the original (non-lazy) module even if the random seed is manually configured, as the order of initialization is different from the original one; especially, `module.cuda()` may cause the initialization to run on a GPU.

The default value of lazy buffers and parameters are `torch.Tensor([])` and `UninitializedParameter()`, respectively.

Methods

`__init__`(*args, **kwargs)

`state_dict`(*args, **kwargs)

Returns a dictionary containing a whole state of the module.

Attributes

`lazy_buffer_names`

`lazy_parameter_names`

`lazy_parameters_determined`

Returns if all lazy parameters are determined.

Parameters

- **args** (*Any*) –
- **kwargs** (*Any*) –

`__init__`(*args, **kwargs)

Parameters

- **self** (*Any*) –
- **args** (*Any*) –
- **kwargs** (*Any*) –

Return type

None

lazy_buffer_names: `Tuple[str, ...] = ()`

lazy_parameter_names: `Tuple[str, ...] = ()`

property lazy_parameters_determined: `bool`

Returns if all lazy parameters are determined.

Subclasses can perform parameters initialization after all lazy parameters are determined. Note that this may be called during `__init__`.

state_dict(*args, **kwargs)

Returns a dictionary containing a whole state of the module.

This function overrides the default behavior to exclude uninitialized parameter from serialization. This is needed because we need to discriminate lazy parameters (`UninitializedParameter()`) and initialized empty parameters (`torch.nn.Parameter(torch.Tensor())`) during deserialization.

See comments of `_lazy_load_hook` for details.

Parameters

- **self** (*Any*) –
- **args** (*Any*) –
- **kwargs** (*Any*) –

Return type

Dict[str, Any]

pytorch_pfn_extras.nn.modules.lazy_linear.LazyLinear

class `pytorch_pfn_extras.nn.modules.lazy_linear.LazyLinear`(*in_features*, *args, **kwargs)

Bases: [*LazyInitializationMixin*](#), [*Linear*](#)

Linear module with lazy weight initialization.

When *in_features* is `None`, it is determined at the first time of the forward step.

Initializes internal Module state, shared by both `nn.Module` and `ScriptModule`.

Methods

| | |
|---|--|
| <code>__init__(in_features, *args, **kwargs)</code> | Initializes internal Module state, shared by both <code>nn.Module</code> and <code>ScriptModule</code> . |
| <code>add_module(name, module)</code> | Adds a child module to the current module. |
| <code>apply(fn)</code> | Applies <code>fn</code> recursively to every submodule (as returned by <code>.children()</code>) as well as self. |
| <code>bfloat16()</code> | Casts all floating point parameters and buffers to <code>bfloat16</code> datatype. |
| <code>buffers([recurse])</code> | Returns an iterator over module buffers. |
| <code>children()</code> | Returns an iterator over immediate children modules. |
| <code>compile(*args, **kwargs)</code> | Compile this Module's forward using <code>torch.compile()</code> . |
| <code>cpu()</code> | Moves all model parameters and buffers to the CPU. |

continues on next page

Table 24 – continued from previous page

| | |
|---|---|
| <code>cuda([device])</code> | Moves all model parameters and buffers to the GPU. |
| <code>double()</code> | Casts all floating point parameters and buffers to double datatype. |
| <code>eval()</code> | Sets the module in evaluation mode. |
| <code>extra_repr()</code> | Set the extra representation of the module |
| <code>float()</code> | Casts all floating point parameters and buffers to float datatype. |
| <code>forward(input)</code> | Defines the computation performed at every call. |
| <code>get_buffer(target)</code> | Returns the buffer given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_extra_state()</code> | Returns any extra state to include in the module's <code>state_dict</code> . |
| <code>get_parameter(target)</code> | Returns the parameter given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_submodule(target)</code> | Returns the submodule given by <code>target</code> if it exists, otherwise throws an error. |
| <code>half()</code> | Casts all floating point parameters and buffers to half datatype. |
| <code>ipu([device])</code> | Moves all model parameters and buffers to the IPU. |
| <code>load_state_dict(state_dict[, strict, assign])</code> | Copies parameters and buffers from <code>state_dict</code> into this module and its descendants. |
| <code>modules()</code> | Returns an iterator over all modules in the network. |
| <code>named_buffers([prefix, recurse, ...])</code> | Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself. |
| <code>named_children()</code> | Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself. |
| <code>named_modules([memo, prefix, remove_duplicate])</code> | Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself. |
| <code>named_parameters([prefix, recurse, ...])</code> | Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself. |
| <code>parameters([recurse])</code> | Returns an iterator over module parameters. |
| <code>register_backward_hook(hook)</code> | Registers a backward hook on the module. |
| <code>register_buffer(name, tensor[, persistent])</code> | Adds a buffer to the module. |
| <code>register_forward_hook(hook, *[, prepend, ...])</code> | Registers a forward hook on the module. |
| <code>register_forward_pre_hook(hook, *[, ...])</code> | Registers a forward pre-hook on the module. |
| <code>register_full_backward_hook(hook[, prepend])</code> | Registers a backward hook on the module. |
| <code>register_full_backward_pre_hook(hook[, prepend])</code> | Registers a backward pre-hook on the module. |
| <code>register_load_state_dict_post_hook(hook)</code> | Registers a post hook to be run after module's <code>load_state_dict</code> is called. |
| <code>register_module(name, module)</code> | Alias for <code>add_module()</code> . |
| <code>register_parameter(name, param)</code> | Adds a parameter to the module. |
| <code>register_state_dict_pre_hook(hook)</code> | These hooks will be called with arguments: <code>self</code> , <code>prefix</code> , and <code>keep_vars</code> before calling <code>state_dict</code> on <code>self</code> . |
| <code>requires_grad_([requires_grad])</code> | Change if autograd should record operations on parameters in this module. |

continues on next page

Table 24 – continued from previous page

| | |
|---|--|
| <code>reset_parameters()</code> | |
| <code>set_extra_state(state)</code> | This function is called from <code>load_state_dict()</code> to handle any extra state found within the <code>state_dict</code> . |
| <code>share_memory()</code> | See <code>torch.Tensor.share_memory_()</code> |
| <code>state_dict(*args, **kwargs)</code> | Returns a dictionary containing a whole state of the module. |
| <code>to(*args, **kwargs)</code> | Moves and/or casts the parameters and buffers. |
| <code>to_empty(*, device[, recurse])</code> | Moves the parameters and buffers to the specified device without copying storage. |
| <code>train([mode])</code> | Sets the module in training mode. |
| <code>type(dst_type)</code> | Casts all parameters and buffers to <code>dst_type</code> . |
| <code>xpu([device])</code> | Moves all model parameters and buffers to the XPU. |
| <code>zero_grad([set_to_none])</code> | Resets gradients of all model parameters. |

Attributes

| | |
|---|--|
| <code>T_destination</code> | alias of <code>TypeVar('T_destination', bound=Dict[str, Any])</code> |
| <code>call_super_init</code> | |
| <code>dump_patches</code> | |
| <code>lazy_buffer_names</code> | |
| <code>lazy_parameter_names</code> | |
| <code>lazy_parameters_determined</code> | Returns if all lazy parameters are determined. |

Parameters

- **in_features** (*Optional[int]*) –
- **args** (*Any*) –
- **kwargs** (*Any*) –

`__init__(in_features, *args, **kwargs)`

Initializes internal Module state, shared by both `nn.Module` and `ScriptModule`.

Parameters

- **in_features** (*Optional[int]*) –
- **args** (*Any*) –
- **kwargs** (*Any*) –

Return type

None

forward(*input*)

Defines the computation performed at every call.

Should be overridden by all subclasses.

Note: Although the recipe for forward pass needs to be defined within this function, one should call the Module instance afterwards instead of this since the former takes care of running the registered hooks while the latter silently ignores them.

Parameters**input** (*Tensor*) –**Return type***Tensor***in_features:** `int`**lazy_parameter_names:** `Tuple[str, ...] = ('weight',)`**out_features:** `int`**reset_parameters()****Return type**`None`**training:** `bool`**weight:** `Tensor`**pytorch_pfn_extras.nn.modules.lazy_linear.UninitializedParameter**

```
class pytorch_pfn_extras.nn.modules.lazy_linear.UninitializedParameter(data=None,
                                                                       requires_grad=True)
```

Bases: `Parameter`**Methods**

| | |
|--|---|
| <code>__init__()</code> | |
| <code>abs()</code> | See <code>torch.abs()</code> |
| <code>abs_()</code> | In-place version of <code>abs()</code> |
| <code>absolute()</code> | Alias for <code>abs()</code> |
| <code>absolute_()</code> | In-place version of <code>absolute()</code> Alias for <code>abs_()</code> |
| <code>acos()</code> | See <code>torch.acos()</code> |
| <code>acos_()</code> | In-place version of <code>acos()</code> |
| <code>acosh()</code> | See <code>torch.acosh()</code> |
| <code>acosh_()</code> | In-place version of <code>acosh()</code> |
| <code>add(other, *, alpha)</code> | Add a scalar or tensor to <code>self</code> tensor. |
| <code>add_(other, *, alpha)</code> | In-place version of <code>add()</code> |
| <code>addbmm(batch1, batch2, *, beta, alpha)</code> | See <code>torch.addbmm()</code> |
| <code>addbmm_(batch1, batch2, *, beta, alpha)</code> | In-place version of <code>addbmm()</code> |
| <code>addcddiv(tensor1, tensor2, *, value)</code> | See <code>torch.addcddiv()</code> |
| <code>addcddiv_(tensor1, tensor2, *, value)</code> | In-place version of <code>addcddiv()</code> |
| <code>addcmul(tensor1, tensor2, *, value)</code> | See <code>torch.addcmul()</code> |

continues on next page

Table 25 – continued from previous page

| | |
|---|--|
| <code>addcmul_(tensor1, tensor2, *, value)</code> | In-place version of <code>addcmul()</code> |
| <code>addmm(mat1, mat2, *, beta, alpha)</code> | See <code>torch.addmm()</code> |
| <code>addmm_(mat1, mat2, *, beta, alpha)</code> | In-place version of <code>addmm()</code> |
| <code>addmv(mat, vec, *, beta, alpha)</code> | See <code>torch.addmv()</code> |
| <code>addmv_(mat, vec, *, beta, alpha)</code> | In-place version of <code>addmv()</code> |
| <code>addr(vec1, vec2, *, beta, alpha)</code> | See <code>torch.addr()</code> |
| <code>addr_(vec1, vec2, *, beta, alpha)</code> | In-place version of <code>addr()</code> |
| <code>adjoint()</code> | Alias for <code>adjoint()</code> |
| <code>align_as(other)</code> | Permutates the dimensions of the <code>self</code> tensor to match the dimension order in the <code>other</code> tensor, adding size-one dims for any new names. |
| <code>align_to(*names)</code> | Permutates the dimensions of the <code>self</code> tensor to match the order specified in <code>names</code> , adding size-one dims for any new names. |
| <code>all([dim, keepdim])</code> | See <code>torch.all()</code> |
| <code>allclose(other[, rtol, atol, equal_nan])</code> | See <code>torch.allclose()</code> |
| <code>amax([dim, keepdim])</code> | See <code>torch.amax()</code> |
| <code>amin([dim, keepdim])</code> | See <code>torch.amin()</code> |
| <code>aminmax(*[, dim, keepdim])</code> | See <code>torch.aminmax()</code> |
| <code>angle()</code> | See <code>torch.angle()</code> |
| <code>any([dim, keepdim])</code> | See <code>torch.any()</code> |
| <code>apply_(callable)</code> | Applies the function <code>callable</code> to each element in the tensor, replacing each element with the value returned by <code>callable</code> . |
| <code>arccos()</code> | See <code>torch.arccos()</code> |
| <code>arccos_()</code> | In-place version of <code>arccos()</code> |
| <code>arccosh</code> | <code>acosh()</code> -> Tensor |
| <code>arccosh_</code> | <code>acosh_()</code> -> Tensor |
| <code>arcsin()</code> | See <code>torch.arcsin()</code> |
| <code>arcsin_()</code> | In-place version of <code>arcsin()</code> |
| <code>arcsinh()</code> | See <code>torch.arcsinh()</code> |
| <code>arcsinh_()</code> | In-place version of <code>arcsinh()</code> |
| <code>arctan()</code> | See <code>torch.arctan()</code> |
| <code>arctan2(other)</code> | See <code>torch.arctan2()</code> |
| <code>arctan2_</code> | <code>atan2_(other)</code> -> Tensor |
| <code>arctan_()</code> | In-place version of <code>arctan()</code> |
| <code>arctanh()</code> | See <code>torch.arctanh()</code> |
| <code>arctanh_(other)</code> | In-place version of <code>arctanh()</code> |
| <code>argmax([dim, keepdim])</code> | See <code>torch.argmax()</code> |
| <code>argmin([dim, keepdim])</code> | See <code>torch.argmin()</code> |
| <code>argsort([dim, descending])</code> | See <code>torch.argsort()</code> |
| <code>argwhere()</code> | See <code>torch.argwhere()</code> |
| <code>as_strided(size, stride[, storage_offset])</code> | See <code>torch.as_strided()</code> |
| <code>as_strided_(size, stride[, storage_offset])</code> | In-place version of <code>as_strided()</code> |
| <code>as_strided_scatter(src, size, stride[, ...])</code> | See <code>torch.as_strided_scatter()</code> |
| <code>as_subclass(cls)</code> | Makes a <code>cls</code> instance with the same data pointer as <code>self</code> . |
| <code>asin()</code> | See <code>torch.asin()</code> |
| <code>asin_()</code> | In-place version of <code>asin()</code> |
| <code>asinh()</code> | See <code>torch.asinh()</code> |

continues on next page

Table 25 – continued from previous page

| | |
|---|---|
| <code>asinh_()</code> | In-place version of <code>asinh()</code> |
| <code>atan()</code> | See <code>torch.atan()</code> |
| <code>atan2(other)</code> | See <code>torch.atan2()</code> |
| <code>atan2_(other)</code> | In-place version of <code>atan2()</code> |
| <code>atan_()</code> | In-place version of <code>atan()</code> |
| <code>atanh()</code> | See <code>torch.atanh()</code> |
| <code>atanh_(other)</code> | In-place version of <code>atanh()</code> |
| <code>backward([gradient, retain_graph, ...])</code> | Computes the gradient of current tensor wrt graph leaves. |
| <code>baddbmm(batch1, batch2, *[, beta, alpha])</code> | See <code>torch.baddbmm()</code> |
| <code>baddbmm_(batch1, batch2, *[, beta, alpha])</code> | In-place version of <code>baddbmm()</code> |
| <code>bernoulli(*[, generator])</code> | Returns a result tensor where each <code>result[i]</code> is independently sampled from <code>Bernoulli(self[i])</code> . |
| <code>bernoulli_(p, generator)</code> | Fills each location of <code>self</code> with an independent sample from <code>Bernoulli(p)</code> . |
| <code>bfloat16([memory_format])</code> | <code>self.bfloat16()</code> is equivalent to <code>self.to(torch.bfloat16)</code> . |
| <code>bincount([weights, minlength])</code> | See <code>torch.bincount()</code> |
| <code>bitwise_and()</code> | See <code>torch.bitwise_and()</code> |
| <code>bitwise_and_()</code> | In-place version of <code>bitwise_and()</code> |
| <code>bitwise_left_shift(other)</code> | See <code>torch.bitwise_left_shift()</code> |
| <code>bitwise_left_shift_(other)</code> | In-place version of <code>bitwise_left_shift()</code> |
| <code>bitwise_not()</code> | See <code>torch.bitwise_not()</code> |
| <code>bitwise_not_()</code> | In-place version of <code>bitwise_not()</code> |
| <code>bitwise_or()</code> | See <code>torch.bitwise_or()</code> |
| <code>bitwise_or_()</code> | In-place version of <code>bitwise_or()</code> |
| <code>bitwise_right_shift(other)</code> | See <code>torch.bitwise_right_shift()</code> |
| <code>bitwise_right_shift_(other)</code> | In-place version of <code>bitwise_right_shift()</code> |
| <code>bitwise_xor()</code> | See <code>torch.bitwise_xor()</code> |
| <code>bitwise_xor_()</code> | In-place version of <code>bitwise_xor()</code> |
| <code>bmm(batch2)</code> | See <code>torch.bmm()</code> |
| <code>bool([memory_format])</code> | <code>self.bool()</code> is equivalent to <code>self.to(torch.bool)</code> . |
| <code>broadcast_to(shape)</code> | See <code>torch.broadcast_to()</code> . |
| <code>byte([memory_format])</code> | <code>self.byte()</code> is equivalent to <code>self.to(torch.uint8)</code> . |
| <code>cauchy_([median, sigma, generator])</code> | Fills the tensor with numbers drawn from the Cauchy distribution: |
| <code>ccol_indices</code> | |
| <code>cdouble([memory_format])</code> | <code>self.cdouble()</code> is equivalent to <code>self.to(torch.complex128)</code> . |
| <code>ceil()</code> | See <code>torch.ceil()</code> |
| <code>ceil_()</code> | In-place version of <code>ceil()</code> |
| <code>cfloat([memory_format])</code> | <code>self.cfloat()</code> is equivalent to <code>self.to(torch.complex64)</code> . |
| <code>chalf([memory_format])</code> | <code>self.chalf()</code> is equivalent to <code>self.to(torch.complex32)</code> . |
| <code>char([memory_format])</code> | <code>self.char()</code> is equivalent to <code>self.to(torch.int8)</code> . |
| <code>cholesky([upper])</code> | See <code>torch.cholesky()</code> |

continues on next page

Table 25 – continued from previous page

| | |
|--|--|
| <code>cholesky_inverse([upper])</code> | See <code>torch.cholesky_inverse()</code> |
| <code>cholesky_solve(input2[, upper])</code> | See <code>torch.cholesky_solve()</code> |
| <code>chunk(chunks[, dim])</code> | See <code>torch.chunk()</code> |
| <code>clamp([min, max])</code> | See <code>torch.clamp()</code> |
| <code>clamp_([min, max])</code> | In-place version of <code>clamp()</code> |
| <code>clamp_max</code> | |
| <code>clamp_max_</code> | |
| <code>clamp_min</code> | |
| <code>clamp_min_</code> | |
| <code>clip([min, max])</code> | Alias for <code>clamp()</code> . |
| <code>clip_([min, max])</code> | Alias for <code>clamp_()</code> . |
| <code>clone(*[, memory_format])</code> | See <code>torch.clone()</code> |
| <code>coalesce()</code> | Returns a coalesced copy of <code>self</code> if <code>self</code> is an un-coalesced tensor. |
| <code>col_indices()</code> | Returns the tensor containing the column indices of the <code>self</code> tensor when <code>self</code> is a sparse CSR tensor of layout <code>sparse_csr</code> . |
| <code>conj()</code> | See <code>torch.conj()</code> |
| <code>conj_physical()</code> | See <code>torch.conj_physical()</code> |
| <code>conj_physical_()</code> | In-place version of <code>conj_physical()</code> |
| <code>contiguous([memory_format])</code> | Returns a contiguous in memory tensor containing the same data as <code>self</code> tensor. |
| <code>copy_(src[, non_blocking])</code> | Copies the elements from <code>src</code> into <code>self</code> tensor and returns <code>self</code> . |
| <code>copysign(other)</code> | See <code>torch.copysign()</code> |
| <code>copysign_(other)</code> | In-place version of <code>copysign()</code> |
| <code>corrcoef()</code> | See <code>torch.corrcoef()</code> |
| <code>cos()</code> | See <code>torch.cos()</code> |
| <code>cos_()</code> | In-place version of <code>cos()</code> |
| <code>cosh()</code> | See <code>torch.cosh()</code> |
| <code>cosh_()</code> | In-place version of <code>cosh()</code> |
| <code>count_nonzero([dim])</code> | See <code>torch.count_nonzero()</code> |
| <code>cov(*[, correction, fweights, aweights])</code> | See <code>torch.cov()</code> |
| <code>cpu([memory_format])</code> | Returns a copy of this object in CPU memory. |
| <code>cross(other[, dim])</code> | See <code>torch.cross()</code> |
| <code>crow_indices()</code> | Returns the tensor containing the compressed row indices of the <code>self</code> tensor when <code>self</code> is a sparse CSR tensor of layout <code>sparse_csr</code> . |
| <code>cuda([device, non_blocking, memory_format])</code> | Returns a copy of this object in CUDA memory. |
| <code>cummax(dim)</code> | See <code>torch.cummax()</code> |
| <code>cummin(dim)</code> | See <code>torch.cummin()</code> |
| <code>cumprod(dim[, dtype])</code> | See <code>torch.cumprod()</code> |
| <code>cumprod_(dim[, dtype])</code> | In-place version of <code>cumprod()</code> |
| <code>cumsum(dim[, dtype])</code> | See <code>torch.cumsum()</code> |
| <code>cumsum_(dim[, dtype])</code> | In-place version of <code>cumsum()</code> |
| <code>data_ptr()</code> | Returns the address of the first element of <code>self</code> tensor. |

continues on next page

Table 25 – continued from previous page

| | |
|--|---|
| <code>deg2rad()</code> | See <code>torch.deg2rad()</code> |
| <code>deg2rad_()</code> | In-place version of <code>deg2rad()</code> |
| <code>dense_dim()</code> | Return the number of dense dimensions in a sparse tensor <code>self</code> . |
| <code>dequantize()</code> | Given a quantized Tensor, dequantize it and return the dequantized float Tensor. |
| <code>det()</code> | See <code>torch.det()</code> |
| <code>detach</code> | Returns a new Tensor, detached from the current graph. |
| <code>detach_</code> | Detaches the Tensor from the graph that created it, making it a leaf. |
| <code>diag([diagonal])</code> | See <code>torch.diag()</code> |
| <code>diag_embed([offset, dim1, dim2])</code> | See <code>torch.diag_embed()</code> |
| <code>diagflat([offset])</code> | See <code>torch.diagflat()</code> |
| <code>diagonal([offset, dim1, dim2])</code> | See <code>torch.diagonal()</code> |
| <code>diagonal_scatter(src[, offset, dim1, dim2])</code> | See <code>torch.diagonal_scatter()</code> |
| <code>diff([n, dim, prepend, append])</code> | See <code>torch.diff()</code> |
| <code>digamma()</code> | See <code>torch.digamma()</code> |
| <code>digamma_()</code> | In-place version of <code>digamma()</code> |
| <code>dim()</code> | Returns the number of dimensions of <code>self</code> tensor. |
| <code>dim_order()</code> | Returns a tuple of int describing the dim order or physical layout of <code>self</code> . |
| <code>dist(other[, p])</code> | See <code>torch.dist()</code> |
| <code>div(value, *[, rounding_mode])</code> | See <code>torch.div()</code> |
| <code>div_(value, *[, rounding_mode])</code> | In-place version of <code>div()</code> |
| <code>divide(value, *[, rounding_mode])</code> | See <code>torch.divide()</code> |
| <code>divide_(value, *[, rounding_mode])</code> | In-place version of <code>divide()</code> |
| <code>dot(other)</code> | See <code>torch.dot()</code> |
| <code>double([memory_format])</code> | <code>self.double()</code> is equivalent to <code>self.to(torch.float64)</code> . |
| <code>dsplit(split_size_or_sections)</code> | See <code>torch.dsplit()</code> |
| <code>eig([eigenvectors])</code> | |
| <code>element_size()</code> | Returns the size in bytes of an individual element. |
| <code>eq(other)</code> | See <code>torch.eq()</code> |
| <code>eq_(other)</code> | In-place version of <code>eq()</code> |
| <code>equal(other)</code> | See <code>torch.equal()</code> |
| <code>erf()</code> | See <code>torch.erf()</code> |
| <code>erf_()</code> | In-place version of <code>erf()</code> |
| <code>erfc()</code> | See <code>torch.erfc()</code> |
| <code>erfc_()</code> | In-place version of <code>erfc()</code> |
| <code>erfinv()</code> | See <code>torch.erfinv()</code> |
| <code>erfinv_()</code> | In-place version of <code>erfinv()</code> |
| <code>exp()</code> | See <code>torch.exp()</code> |
| <code>exp2()</code> | See <code>torch.exp2()</code> |
| <code>exp2_()</code> | In-place version of <code>exp2()</code> |
| <code>exp_()</code> | In-place version of <code>exp()</code> |
| <code>expand(*sizes)</code> | Returns a new view of the <code>self</code> tensor with singleton dimensions expanded to a larger size. |
| <code>expand_as(other)</code> | Expand this tensor to the same size as <code>other</code> . |
| <code>expm1()</code> | See <code>torch.expm1()</code> |

continues on next page

Table 25 – continued from previous page

| | |
|---|--|
| <code>expm1_()</code> | In-place version of <code>expm1()</code> |
| <code>exponential_([lambd, generator])</code> | Fills <code>self</code> tensor with elements drawn from the exponential distribution: |
| <code>fill_(value)</code> | Fills <code>self</code> tensor with the specified value. |
| <code>fill_diagonal_(fill_value[, wrap])</code> | Fill the main diagonal of a tensor that has at least 2-dimensions. |
| <code>fix()</code> | See <code>torch.fix()</code> . |
| <code>fix_()</code> | In-place version of <code>fix()</code> |
| <code>flatten([start_dim, end_dim])</code> | See <code>torch.flatten()</code> |
| <code>flip(dims)</code> | See <code>torch.flip()</code> |
| <code>fliplr()</code> | See <code>torch.fliplr()</code> |
| <code>flipud()</code> | See <code>torch.flipud()</code> |
| <code>float([memory_format])</code> | <code>self.float()</code> is equivalent to <code>self.to(torch.float32)</code> . |
| <code>float_power(exponent)</code> | See <code>torch.float_power()</code> |
| <code>float_power_(exponent)</code> | In-place version of <code>float_power()</code> |
| <code>floor()</code> | See <code>torch.floor()</code> |
| <code>floor_()</code> | In-place version of <code>floor()</code> |
| <code>floor_divide(value)</code> | See <code>torch.floor_divide()</code> |
| <code>floor_divide_(value)</code> | In-place version of <code>floor_divide()</code> |
| <code>fmax(other)</code> | See <code>torch.fmax()</code> |
| <code>fmin(other)</code> | See <code>torch.fmin()</code> |
| <code>fmod(divisor)</code> | See <code>torch.fmod()</code> |
| <code>fmod_(divisor)</code> | In-place version of <code>fmod()</code> |
| <code>frac()</code> | See <code>torch.frac()</code> |
| <code>frac_()</code> | In-place version of <code>frac()</code> |
| <code>frexp(input)</code> | See <code>torch.frexp()</code> |
| <code>gather(dim, index)</code> | See <code>torch.gather()</code> |
| <code>gcd(other)</code> | See <code>torch.gcd()</code> |
| <code>gcd_(other)</code> | In-place version of <code>gcd()</code> |
| <code>ge(other)</code> | See <code>torch.ge()</code> . |
| <code>ge_(other)</code> | In-place version of <code>ge()</code> . |
| <code>geometric_(p, *[, generator])</code> | Fills <code>self</code> tensor with elements drawn from the geometric distribution: |
| <code>geqrf()</code> | See <code>torch.geqrf()</code> |
| <code>ger(vec2)</code> | See <code>torch.ger()</code> |
| <code>get_device()</code> | For CUDA tensors, this function returns the device ordinal of the GPU on which the tensor resides. |
| <code>greater(other)</code> | See <code>torch.greater()</code> . |
| <code>greater_(other)</code> | In-place version of <code>greater()</code> . |
| <code>greater_equal(other)</code> | See <code>torch.greater_equal()</code> . |
| <code>greater_equal_(other)</code> | In-place version of <code>greater_equal()</code> . |
| <code>gt(other)</code> | See <code>torch.gt()</code> . |
| <code>gt_(other)</code> | In-place version of <code>gt()</code> . |
| <code>half([memory_format])</code> | <code>self.half()</code> is equivalent to <code>self.to(torch.float16)</code> . |
| <code>hardshrink([lambd])</code> | See <code>torch.nn.functional.hardshrink()</code> |
| <code>has_names</code> | Is True if any of this tensor's dimensions are named. |
| <code>heaviside(values)</code> | See <code>torch.heaviside()</code> |
| <code>heaviside_(values)</code> | In-place version of <code>heaviside()</code> |

continues on next page

Table 25 – continued from previous page

| | |
|--|---|
| <code>histc([bins, min, max])</code> | See <code>torch.histc()</code> |
| <code>histogram(input, bins, *[, range, weight, ...])</code> | See <code>torch.histogram()</code> |
| <code>hsplit(split_size_or_sections)</code> | See <code>torch.hsplit()</code> |
| <code>hypot(other)</code> | See <code>torch.hypot()</code> |
| <code>hypot_(other)</code> | In-place version of <code>hypot()</code> |
| <code>i0()</code> | See <code>torch.i0()</code> |
| <code>i0_()</code> | In-place version of <code>i0()</code> |
| <code>igamma(other)</code> | See <code>torch.igamma()</code> |
| <code>igamma_(other)</code> | In-place version of <code>igamma()</code> |
| <code>igammac(other)</code> | See <code>torch.igammac()</code> |
| <code>igammac_(other)</code> | In-place version of <code>igammac()</code> |
| <code>index_add(dim, index, source, *[, alpha])</code> | Out-of-place version of <code>torch.Tensor.index_add_()</code> . |
| <code>index_add_(dim, index, source, *[, alpha])</code> | Accumulate the elements of <code>alpha times source</code> into the <code>self</code> tensor by adding to the indices in the order given in <code>index</code> . |
| <code>index_copy(dim, index, tensor2)</code> | Out-of-place version of <code>torch.Tensor.index_copy_()</code> . |
| <code>index_copy_(dim, index, tensor)</code> | Copies the elements of <code>tensor</code> into the <code>self</code> tensor by selecting the indices in the order given in <code>index</code> . |
| <code>index_fill(dim, index, value)</code> | Out-of-place version of <code>torch.Tensor.index_fill_()</code> . |
| <code>index_fill_(dim, index, value)</code> | Fills the elements of the <code>self</code> tensor with <code>value</code> by selecting the indices in the order given in <code>index</code> . |
| <code>index_put(indices, values[, accumulate])</code> | Out-place version of <code>index_put_()</code> . |
| <code>index_put_(indices, values[, accumulate])</code> | Puts values from the tensor <code>values</code> into the tensor <code>self</code> using the indices specified in <code>indices</code> (which is a tuple of Tensors). |
| <code>index_reduce</code> | |
| <code>index_reduce_(dim, index, source, reduce, *)</code> | Accumulate the elements of <code>source</code> into the <code>self</code> tensor by accumulating to the indices in the order given in <code>index</code> using the reduction given by the <code>reduce</code> argument. |
| <code>index_select(dim, index)</code> | See <code>torch.index_select()</code> |
| <code>indices()</code> | Return the indices tensor of a sparse COO tensor. |
| <code>inner(other)</code> | See <code>torch.inner()</code> . |
| <code>int([memory_format])</code> | <code>self.int()</code> is equivalent to <code>self.to(torch.int32)</code> . |
| <code>int_repr()</code> | Given a quantized Tensor, <code>self.int_repr()</code> returns a CPU Tensor with <code>uint8_t</code> as data type that stores the underlying <code>uint8_t</code> values of the given Tensor. |
| <code>inverse()</code> | See <code>torch.inverse()</code> |
| <code>ipu([device, non_blocking, memory_format])</code> | Returns a copy of this object in IPU memory. |
| <code>is_coalesced()</code> | Returns True if <code>self</code> is a sparse COO tensor that is coalesced, False otherwise. |
| <code>is_complex()</code> | Returns True if the data type of <code>self</code> is a complex data type. |
| <code>is_conj()</code> | Returns True if the conjugate bit of <code>self</code> is set to true. |

continues on next page

Table 25 – continued from previous page

| | |
|--|---|
| <code>is_contiguous([memory_format])</code> | Returns True if <code>self</code> tensor is contiguous in memory in the order specified by memory format. |
| <code>is_distributed</code> | |
| <code>is_floating_point()</code> | Returns True if the data type of <code>self</code> is a floating point data type. |
| <code>is_inference()</code> | See <code>torch.is_inference()</code> |
| <code>is_neg()</code> | Returns True if the negative bit of <code>self</code> is set to true. |
| <code>is_nonzero</code> | |
| <code>is_pinned</code> | Returns true if this tensor resides in pinned memory. |
| <code>is_same_size</code> | |
| <code>is_set_to(tensor)</code> | Returns True if both tensors are pointing to the exact same memory (same storage, offset, size and stride). |
| <code>is_shared()</code> | Checks if tensor is in shared memory. |
| <code>is_signed()</code> | Returns True if the data type of <code>self</code> is a signed data type. |
| <code>isclose(other[, rtol, atol, equal_nan])</code> | See <code>torch.isclose()</code> |
| <code>isfinite()</code> | See <code>torch.isfinite()</code> |
| <code>isinf()</code> | See <code>torch.isinf()</code> |
| <code>isnan()</code> | See <code>torch.isnan()</code> |
| <code>isneginf()</code> | See <code>torch.isneginf()</code> |
| <code>isposinf()</code> | See <code>torch.isposinf()</code> |
| <code>isreal()</code> | See <code>torch.isreal()</code> |
| <code>istft(n_fft[, hop_length, win_length, ...])</code> | See <code>torch.istft()</code> |
| <code>item()</code> | Returns the value of this tensor as a standard Python number. |
| <code>kron(other)</code> | See <code>torch.kron()</code> |
| <code>kthvalue(k[, dim, keepdim])</code> | See <code>torch.kthvalue()</code> |
| <code>lcm(other)</code> | See <code>torch.lcm()</code> |
| <code>lcm_(other)</code> | In-place version of <code>lcm()</code> |
| <code>ldexp(other)</code> | See <code>torch.ldexp()</code> |
| <code>ldexp_(other)</code> | In-place version of <code>ldexp()</code> |
| <code>le(other)</code> | See <code>torch.le()</code> . |
| <code>le_(other)</code> | In-place version of <code>le()</code> . |
| <code>lerp(end, weight)</code> | See <code>torch.lerp()</code> |
| <code>lerp_(end, weight)</code> | In-place version of <code>lerp()</code> |
| <code>less</code> | <code>lt(other) -> Tensor</code> |
| <code>less_(other)</code> | In-place version of <code>less()</code> . |
| <code>less_equal(other)</code> | See <code>torch.less_equal()</code> . |
| <code>less_equal_(other)</code> | In-place version of <code>less_equal()</code> . |
| <code>lgamma()</code> | See <code>torch.lgamma()</code> |
| <code>lgamma_()</code> | In-place version of <code>lgamma()</code> |
| <code>log()</code> | See <code>torch.log()</code> |
| <code>log10()</code> | See <code>torch.log10()</code> |
| <code>log10_()</code> | In-place version of <code>log10()</code> |
| <code>log1p()</code> | See <code>torch.log1p()</code> |
| <code>log1p_()</code> | In-place version of <code>log1p()</code> |
| <code>log2()</code> | See <code>torch.log2()</code> |
| <code>log2_()</code> | In-place version of <code>log2()</code> |

continues on next page

Table 25 – continued from previous page

| | |
|--|--|
| <code>log_()</code> | In-place version of <code>log()</code> |
| <code>log_normal_([mean, std, generator])</code> | Fills <code>self</code> tensor with numbers samples from the log-normal distribution parameterized by the given mean μ and standard deviation σ . |
| <code>log_softmax</code> | |
| <code>logaddexp(other)</code> | See <code>torch.logaddexp()</code> |
| <code>logaddexp2(other)</code> | See <code>torch.logaddexp2()</code> |
| <code>logcumsumexp(dim)</code> | See <code>torch.logcumsumexp()</code> |
| <code>logdet()</code> | See <code>torch.logdet()</code> |
| <code>logical_and()</code> | See <code>torch.logical_and()</code> |
| <code>logical_and_()</code> | In-place version of <code>logical_and()</code> |
| <code>logical_not()</code> | See <code>torch.logical_not()</code> |
| <code>logical_not_()</code> | In-place version of <code>logical_not()</code> |
| <code>logical_or()</code> | See <code>torch.logical_or()</code> |
| <code>logical_or_()</code> | In-place version of <code>logical_or()</code> |
| <code>logical_xor()</code> | See <code>torch.logical_xor()</code> |
| <code>logical_xor_()</code> | In-place version of <code>logical_xor()</code> |
| <code>logit()</code> | See <code>torch.logit()</code> |
| <code>logit_()</code> | In-place version of <code>logit()</code> |
| <code>logsumexp(dim[, keepdim])</code> | See <code>torch.logsumexp()</code> |
| <code>long([memory_format])</code> | <code>self.long()</code> is equivalent to <code>self.to(torch.int64)</code> . |
| <code>lstsq(other)</code> | |
| <code>lt(other)</code> | See <code>torch.lt()</code> . |
| <code>lt_(other)</code> | In-place version of <code>lt()</code> . |
| <code>lu([pivot, get_infos])</code> | See <code>torch.lu()</code> |
| <code>lu_solve(LU_data, LU_pivots)</code> | See <code>torch.lu_solve()</code> |
| <code>map2_</code> | |
| <code>map_(tensor, callable)</code> | Applies callable for each element in <code>self</code> tensor and the given tensor and stores the results in <code>self</code> tensor. |
| <code>masked_fill(mask, value)</code> | Out-of-place version of <code>torch.Tensor.masked_fill_()</code> |
| <code>masked_fill_(mask, value)</code> | Fills elements of <code>self</code> tensor with <code>value</code> where <code>mask</code> is <code>True</code> . |
| <code>masked_scatter(mask, tensor)</code> | Out-of-place version of <code>torch.Tensor.masked_scatter_()</code> |
| <code>masked_scatter_(mask, source)</code> | Copies elements from <code>source</code> into <code>self</code> tensor at positions where the <code>mask</code> is <code>True</code> . |
| <code>masked_select(mask)</code> | See <code>torch.masked_select()</code> |
| <code>materialize(shape[, device, dtype])</code> | Create a Parameter with the same properties of the uninitialized one. |
| <code>matmul(tensor2)</code> | See <code>torch.matmul()</code> |
| <code>matrix_exp()</code> | See <code>torch.matrix_exp()</code> |

continues on next page

Table 25 – continued from previous page

| | |
|--|---|
| <code>matrix_power(n)</code> | Note: <code>matrix_power()</code> is deprecated, use <code>torch.linalg.matrix_power()</code> instead. |
| <code>max([dim, keepdim])</code> | See <code>torch.max()</code> |
| <code>maximum(other)</code> | See <code>torch.maximum()</code> |
| <code>mean([dim, keepdim, dtype])</code> | See <code>torch.mean()</code> |
| <code>median([dim, keepdim])</code> | See <code>torch.median()</code> |
| <code>min([dim, keepdim])</code> | See <code>torch.min()</code> |
| <code>minimum(other)</code> | See <code>torch.minimum()</code> |
| <code>mm(mat2)</code> | See <code>torch.mm()</code> |
| <code>mode([dim, keepdim])</code> | See <code>torch.mode()</code> |
| <code>moveaxis(source, destination)</code> | See <code>torch.moveaxis()</code> |
| <code>movedim(source, destination)</code> | See <code>torch.movedim()</code> |
| <code>msort()</code> | See <code>torch.msort()</code> |
| <code>mul(value)</code> | See <code>torch.mul()</code> . |
| <code>mul_(value)</code> | In-place version of <code>mul()</code> . |
| <code>multinomial(num_samples[, replacement, ...])</code> | See <code>torch.multinomial()</code> |
| <code>multiply(value)</code> | See <code>torch.multiply()</code> . |
| <code>multiply_(value)</code> | In-place version of <code>multiply()</code> . |
| <code>mv(vec)</code> | See <code>torch.mv()</code> |
| <code>mvlgamma(p)</code> | See <code>torch.mvlgamma()</code> |
| <code>mvlgamma_(p)</code> | In-place version of <code>mvlgamma()</code> |
| <code>nan_to_num([nan, posinf, neginf])</code> | See <code>torch.nan_to_num()</code> . |
| <code>nan_to_num_([nan, posinf, neginf])</code> | In-place version of <code>nan_to_num()</code> . |
| <code>nanmean([dim, keepdim, dtype])</code> | See <code>torch.nanmean()</code> |
| <code>nanmedian([dim, keepdim])</code> | See <code>torch.nanmedian()</code> |
| <code>nanquantile(q[, dim, keepdim, interpolation])</code> | See <code>torch.nanquantile()</code> |
| <code>nansum([dim, keepdim, dtype])</code> | See <code>torch.nansum()</code> |
| <code>narrow(dimension, start, length)</code> | See <code>torch.narrow()</code> . |
| <code>narrow_copy(dimension, start, length)</code> | See <code>torch.narrow_copy()</code> . |
| <code>ndimension()</code> | Alias for <code>dim()</code> |
| <code>ne(other)</code> | See <code>torch.ne()</code> . |
| <code>ne_(other)</code> | In-place version of <code>ne()</code> . |
| <code>neg()</code> | See <code>torch.neg()</code> |
| <code>neg_()</code> | In-place version of <code>neg()</code> |
| <code>negative()</code> | See <code>torch.negative()</code> |
| <code>negative_()</code> | In-place version of <code>negative()</code> |
| <code>nelement()</code> | Alias for <code>numel()</code> |
| <code>new</code> | |
| <code>new_empty(size, *[, dtype, device, ...])</code> | Returns a Tensor of size <code>size</code> filled with uninitialized data. |
| <code>new_empty_strided(size, stride[, dtype, ...])</code> | Returns a Tensor of size <code>size</code> and strides <code>stride</code> filled with uninitialized data. |
| <code>new_full(size, fill_value, *[, dtype, ...])</code> | Returns a Tensor of size <code>size</code> filled with <code>fill_value</code> . |
| <code>new_ones(size, *[, dtype, device, ...])</code> | Returns a Tensor of size <code>size</code> filled with 1. |
| <code>new_tensor(data, *[, dtype, device, ...])</code> | Returns a new Tensor with data as the tensor data. |

continues on next page

Table 25 – continued from previous page

| | |
|---|--|
| <code>new_zeros(size, *, dtype, device, ...)</code> | Returns a Tensor of size <code>size</code> filled with 0. |
| <code>nextafter(other)</code> | See <code>torch.nextafter()</code> |
| <code>nextafter_(other)</code> | In-place version of <code>nextafter()</code> |
| <code>nonzero()</code> | See <code>torch.nonzero()</code> |
| <code>nonzero_static(input, *, size[, fill_value])</code> | Returns a 2-D tensor where each row is the index for a non-zero value. |
| <code>norm([p, dim, keepdim, dtype])</code> | See <code>torch.norm()</code> |
| <code>normal_([mean, std, generator])</code> | Fills self tensor with elements samples from the normal distribution parameterized by mean and std. |
| <code>not_equal(other)</code> | See <code>torch.not_equal()</code> |
| <code>not_equal_(other)</code> | In-place version of <code>not_equal()</code> |
| <code>numel()</code> | See <code>torch.numel()</code> |
| <code>numpy(*[, force])</code> | Returns the tensor as a NumPy ndarray. |
| <code>orgqr(input2)</code> | See <code>torch.orgqr()</code> |
| <code>ormqr(input2, input3[, left, transpose])</code> | See <code>torch.ormqr()</code> |
| <code>outer(vec2)</code> | See <code>torch.outer()</code> |
| <code>permute(*dims)</code> | See <code>torch.permute()</code> |
| <code>pin_memory()</code> | Copies the tensor to pinned memory, if it's not already pinned. |
| <code>pinverse()</code> | See <code>torch.pinverse()</code> |
| <code>polygamma(n)</code> | See <code>torch.polygamma()</code> |
| <code>polygamma_(n)</code> | In-place version of <code>polygamma()</code> |
| <code>positive()</code> | See <code>torch.positive()</code> |
| <code>pow(exponent)</code> | See <code>torch.pow()</code> |
| <code>pow_(exponent)</code> | In-place version of <code>pow()</code> |
| <code>prelu</code> | |
| <code>prod([dim, keepdim, dtype])</code> | See <code>torch.prod()</code> |
| <code>put(input, index, source[, accumulate])</code> | Out-of-place version of <code>torch.Tensor.put_()</code> . |
| <code>put_(index, source[, accumulate])</code> | Copies the elements from <code>source</code> into the positions specified by <code>index</code> . |
| <code>q_per_channel_axis()</code> | Given a Tensor quantized by linear (affine) per-channel quantization, returns the index of dimension on which per-channel quantization is applied. |
| <code>q_per_channel_scales()</code> | Given a Tensor quantized by linear (affine) per-channel quantization, returns a Tensor of scales of the underlying quantizer. |
| <code>q_per_channel_zero_points()</code> | Given a Tensor quantized by linear (affine) per-channel quantization, returns a tensor of zero_points of the underlying quantizer. |
| <code>q_scale()</code> | Given a Tensor quantized by linear(affine) quantization, returns the scale of the underlying quantizer(). |
| <code>q_zero_point()</code> | Given a Tensor quantized by linear(affine) quantization, returns the zero_point of the underlying quantizer(). |
| <code>qr([some])</code> | See <code>torch.qr()</code> |
| <code>qscheme()</code> | Returns the quantization scheme of a given QTensor. |
| <code>quantile(q[, dim, keepdim, interpolation])</code> | See <code>torch.quantile()</code> |
| <code>rad2deg()</code> | See <code>torch.rad2deg()</code> |
| <code>rad2deg_()</code> | In-place version of <code>rad2deg()</code> |

continues on next page

Table 25 – continued from previous page

| | |
|---|--|
| <code>random_([from, to, generator])</code> | Fills <code>self</code> tensor with numbers sampled from the discrete uniform distribution over <code>[from, to - 1]</code> . |
| <code>ravel()</code> | see <code>torch.ravel()</code> |
| <code>reciprocal()</code> | See <code>torch.reciprocal()</code> |
| <code>reciprocal_()</code> | In-place version of <code>reciprocal()</code> |
| <code>record_stream(stream)</code> | Ensures that the tensor memory is not reused for another tensor until all current work queued on <code>stream</code> are complete. |
| <code>refine_names(*names)</code> | Refines the dimension names of <code>self</code> according to <code>names</code> . |
| <code>register_hook(hook)</code> | Registers a backward hook. |
| <code>register_post_accumulate_grad_hook(hook)</code> | Registers a backward hook that runs after grad accumulation. |
| <code>reinforce(reward)</code> | |
| <code>relu</code> | |
| <code>relu_</code> | |
| <code>remainder(divisor)</code> | See <code>torch.remainder()</code> |
| <code>remainder_(divisor)</code> | In-place version of <code>remainder()</code> |
| <code>rename(*names, **rename_map)</code> | Renames dimension names of <code>self</code> . |
| <code>rename_(*names, **rename_map)</code> | In-place version of <code>rename()</code> . |
| <code>renorm(p, dim, maxnorm)</code> | See <code>torch.renorm()</code> |
| <code>renorm_(p, dim, maxnorm)</code> | In-place version of <code>renorm()</code> |
| <code>repeat(*sizes)</code> | Repeats this tensor along the specified dimensions. |
| <code>repeat_interleave(repeats[, dim, output_size])</code> | See <code>torch.repeat_interleave()</code> . |
| <code>requires_grad_(requires_grad)</code> | Change if autograd should record operations on this tensor: sets this tensor's <code>requires_grad</code> attribute in-place. |
| <code>reshape(*shape)</code> | Returns a tensor with the same data and number of elements as <code>self</code> but with the specified shape. |
| <code>reshape_as(other)</code> | Returns this tensor as the same shape as <code>other</code> . |
| <code>resize(*sizes)</code> | |
| <code>resize_(*sizes[, memory_format])</code> | Resizes <code>self</code> tensor to the specified size. |
| <code>resize_as(tensor)</code> | |
| <code>resize_as_(tensor[, memory_format])</code> | Resizes the <code>self</code> tensor to be the same size as the specified tensor. |
| <code>resize_as_sparse_</code> | |
| <code>resolve_conj()</code> | See <code>torch.resolve_conj()</code> |
| <code>resolve_neg()</code> | See <code>torch.resolve_neg()</code> |
| <code>retain_grad()</code> | Enables this Tensor to have their grad populated during <code>backward()</code> . |
| <code>roll(shifts, dims)</code> | See <code>torch.roll()</code> |
| <code>rot90(k, dims)</code> | See <code>torch.rot90()</code> |
| <code>round([decimals])</code> | See <code>torch.round()</code> |
| <code>round_([decimals])</code> | In-place version of <code>round()</code> |

continues on next page

Table 25 – continued from previous page

| | |
|--|---|
| row_indices | |
| <code>rsqrt()</code> | See <code>torch.rsqrt()</code> |
| <code>rsqrt_()</code> | In-place version of <code>rsqrt()</code> |
| <code>scatter(dim, index, src)</code> | Out-of-place version of <code>torch.Tensor.scatter_()</code> |
| <code>scatter_(dim, index, src[, reduce])</code> | Writes all values from the tensor <code>src</code> into <code>self</code> at the indices specified in the <code>index</code> tensor. |
| <code>scatter_add(dim, index, src)</code> | Out-of-place version of <code>torch.Tensor.scatter_add_()</code> |
| <code>scatter_add_(dim, index, src)</code> | Adds all values from the tensor <code>src</code> into <code>self</code> at the indices specified in the <code>index</code> tensor in a similar fashion as <code>scatter_()</code> . |
| <code>scatter_reduce(dim, index, src, reduce, *[, ...])</code> | Out-of-place version of <code>torch.Tensor.scatter_reduce_()</code> |
| <code>scatter_reduce_(dim, index, src, reduce, *)</code> | Reduces all values from the <code>src</code> tensor to the indices specified in the <code>index</code> tensor in the <code>self</code> tensor using the applied reduction defined via the <code>reduce</code> argument ("sum", "prod", "mean", "amax", "amin"). |
| <code>select(dim, index)</code> | See <code>torch.select()</code> |
| <code>select_scatter(src, dim, index)</code> | See <code>torch.select_scatter()</code> |
| <code>set_([source, storage_offset, size, stride])</code> | Sets the underlying storage, size, and strides. |
| <code>sgn()</code> | See <code>torch.sgn()</code> |
| <code>sgn_()</code> | In-place version of <code>sgn()</code> |
| <code>share_memory_()</code> | Moves the underlying storage to shared memory. |
| <code>short([memory_format])</code> | <code>self.short()</code> is equivalent to <code>self.to(torch.int16)</code> . |
| <code>sigmoid()</code> | See <code>torch.sigmoid()</code> |
| <code>sigmoid_()</code> | In-place version of <code>sigmoid()</code> |
| <code>sign()</code> | See <code>torch.sign()</code> |
| <code>sign_()</code> | In-place version of <code>sign()</code> |
| <code>signbit()</code> | See <code>torch.signbit()</code> |
| <code>sin()</code> | See <code>torch.sin()</code> |
| <code>sin_()</code> | In-place version of <code>sin()</code> |
| <code>sinc()</code> | See <code>torch.sinc()</code> |
| <code>sinc_()</code> | In-place version of <code>sinc()</code> |
| <code>sinh()</code> | See <code>torch.sinh()</code> |
| <code>sinh_()</code> | In-place version of <code>sinh()</code> |
| <code>size([dim])</code> | Returns the size of the <code>self</code> tensor. |
| <code>slice_scatter(src[, dim, start, end, step])</code> | See <code>torch.slice_scatter()</code> |
| <code>slogdet()</code> | See <code>torch.slogdet()</code> |
| <code>smm(mat)</code> | See <code>torch.smm()</code> |
| <code>softmax(dim)</code> | Alias for <code>torch.nn.functional.softmax()</code> . |
| <code>solve(other)</code> | |
| <code>sort([dim, descending])</code> | See <code>torch.sort()</code> |
| <code>sparse_dim()</code> | Return the number of sparse dimensions in a sparse tensor <code>self</code> . |
| <code>sparse_mask(mask)</code> | Returns a new sparse tensor with values from a strided tensor <code>self</code> filtered by the indices of the sparse tensor <code>mask</code> . |

continues on next page

Table 25 – continued from previous page

| | |
|--|--|
| <code>sparse_resize_(size, sparse_dim, dense_dim)</code> | Resizes <code>self</code> sparse tensor to the desired size and the number of sparse and dense dimensions. |
| <code>sparse_resize_and_clear_(size, sparse_dim, ...)</code> | Removes all specified elements from a sparse tensor <code>self</code> and resizes <code>self</code> to the desired size and the number of sparse and dense dimensions. |
| <code>split(split_size[, dim])</code> | See <code>torch.split()</code> |
| <code>split_with_sizes</code> | |
| <code>sqrt()</code> | See <code>torch.sqrt()</code> |
| <code>sqrt_()</code> | In-place version of <code>sqrt()</code> |
| <code>square()</code> | See <code>torch.square()</code> |
| <code>square_()</code> | In-place version of <code>square()</code> |
| <code>squeeze([dim])</code> | See <code>torch.squeeze()</code> |
| <code>squeeze_([dim])</code> | In-place version of <code>squeeze()</code> |
| <code>sspaddmm(mat1, mat2, *[, beta, alpha])</code> | See <code>torch.sspaddmm()</code> |
| <code>std([dim, correction, keepdim])</code> | See <code>torch.std()</code> |
| <code>stft(n_fft[, hop_length, win_length, ...])</code> | See <code>torch.stft()</code> |
| <code>storage()</code> | Returns the underlying <code>TypedStorage</code> . |
| <code>storage_offset()</code> | Returns <code>self</code> tensor's offset in the underlying storage in terms of number of storage elements (not bytes). |
| <code>storage_type()</code> | Returns the type of the underlying storage. |
| <code>stride(dim)</code> | Returns the stride of <code>self</code> tensor. |
| <code>sub(other, *[, alpha])</code> | See <code>torch.sub()</code> . |
| <code>sub_(other, *[, alpha])</code> | In-place version of <code>sub()</code> |
| <code>subtract(other, *[, alpha])</code> | See <code>torch.subtract()</code> . |
| <code>subtract_(other, *[, alpha])</code> | In-place version of <code>subtract()</code> . |
| <code>sum([dim, keepdim, dtype])</code> | See <code>torch.sum()</code> |
| <code>sum_to_size(*size)</code> | Sum this tensor to size. |
| <code>svd([some, compute_uv])</code> | See <code>torch.svd()</code> |
| <code>swapaxes(axis0, axis1)</code> | See <code>torch.swapaxes()</code> |
| <code>swapaxes_(axis0, axis1)</code> | In-place version of <code>swapaxes()</code> |
| <code>swapdims(dim0, dim1)</code> | See <code>torch.swapdims()</code> |
| <code>swapdims_(dim0, dim1)</code> | In-place version of <code>swapdims()</code> |
| <code>syમેig([eigenvectors])</code> | |
| <code>t()</code> | See <code>torch.t()</code> |
| <code>t_()</code> | In-place version of <code>t()</code> |
| <code>take(indices)</code> | See <code>torch.take()</code> |
| <code>take_along_dim(indices, dim)</code> | See <code>torch.take_along_dim()</code> |
| <code>tan()</code> | See <code>torch.tan()</code> |
| <code>tan_()</code> | In-place version of <code>tan()</code> |
| <code>tanh()</code> | See <code>torch.tanh()</code> |
| <code>tanh_()</code> | In-place version of <code>tanh()</code> |
| <code>tensor_split(indices_or_sections[, dim])</code> | See <code>torch.tensor_split()</code> |
| <code>tile(dims)</code> | See <code>torch.tile()</code> |
| <code>to(*args, **kwargs)</code> | Performs Tensor dtype and/or device conversion. |
| <code>to_dense([dtype, masked_grad])</code> | Creates a strided copy of <code>self</code> if <code>self</code> is not a strided tensor, otherwise returns <code>self</code> . |
| <code>to_mkldnn()</code> | Returns a copy of the tensor in <code>torch.mkldnn</code> layout. |
| <code>to_padded_tensor(padding[, output_size])</code> | See <code>to_padded_tensor()</code> |

continues on next page

Table 25 – continued from previous page

| | |
|---|---|
| <code>to_sparse(sparseDims)</code> | Returns a sparse copy of the tensor. |
| <code>to_sparse_bsc(blocksize, dense_dim)</code> | Convert a tensor to a block sparse column (BSC) storage format of given blocksize. |
| <code>to_sparse_bsr(blocksize, dense_dim)</code> | Convert a tensor to a block sparse row (BSR) storage format of given blocksize. |
| <code>to_sparse_coo()</code> | Convert a tensor to coordinate format. |
| <code>to_sparse_csc()</code> | Convert a tensor to compressed column storage (CSC) format. |
| <code>to_sparse_csr([dense_dim])</code> | Convert a tensor to compressed row storage format (CSR). |
| <code>tolist()</code> | Returns the tensor as a (nested) list. |
| <code>topk(k[, dim, largest, sorted])</code> | See <code>torch.topk()</code> |
| <code>trace()</code> | See <code>torch.trace()</code> |
| <code>transpose(dim0, dim1)</code> | See <code>torch.transpose()</code> |
| <code>transpose_(dim0, dim1)</code> | In-place version of <code>transpose()</code> |
| <code>triangular_solve(A[, upper, transpose, ...])</code> | See <code>torch.triangular_solve()</code> |
| <code>tril([diagonal])</code> | See <code>torch.tril()</code> |
| <code>tril_([diagonal])</code> | In-place version of <code>tril()</code> |
| <code>triu([diagonal])</code> | See <code>torch.triu()</code> |
| <code>triu_([diagonal])</code> | In-place version of <code>triu()</code> |
| <code>true_divide(value)</code> | See <code>torch.true_divide()</code> |
| <code>true_divide_(value)</code> | In-place version of <code>true_divide_()</code> |
| <code>trunc()</code> | See <code>torch.trunc()</code> |
| <code>trunc_()</code> | In-place version of <code>trunc()</code> |
| <code>type([dtype, non_blocking])</code> | Returns the type if <i>dtype</i> is not provided, else casts this object to the specified type. |
| <code>type_as(tensor)</code> | Returns this tensor cast to the type of the given tensor. |
| <code>unbind([dim])</code> | See <code>torch.unbind()</code> |
| <code>unflatten(dim, sizes)</code> | See <code>torch.unflatten()</code> . |
| <code>unfold(dimension, size, step)</code> | Returns a view of the original tensor which contains all slices of size <i>size</i> from <i>self</i> tensor in the dimension <i>dimension</i> . |
| <code>uniform_([from, to, generator])</code> | Fills <i>self</i> tensor with numbers sampled from the continuous uniform distribution: |
| <code>unique([sorted, return_inverse, ...])</code> | Returns the unique elements of the input tensor. |
| <code>unique_consecutive([return_inverse, ...])</code> | Eliminates all but the first element from every consecutive group of equivalent elements. |
| <code>unsafe_chunk(chunks[, dim])</code> | See <code>torch.unsafe_chunk()</code> |
| <code>unsafe_split(split_size[, dim])</code> | See <code>torch.unsafe_split()</code> |
| <code>unsafe_split_with_sizes</code> | |
| <code>unsqueeze(dim)</code> | See <code>torch.unsqueeze()</code> |
| <code>unsqueeze_(dim)</code> | In-place version of <code>unsqueeze()</code> |
| <code>untyped_storage()</code> | Returns the underlying <code>UntypedStorage</code> . |
| <code>values()</code> | Return the values tensor of a sparse COO tensor. |
| <code>var([dim, correction, keepdim])</code> | See <code>torch.var()</code> |
| <code>vdot(other)</code> | See <code>torch.vdot()</code> |
| <code>view(*shape)</code> | Returns a new tensor with the same data as the <i>self</i> tensor but of a different <i>shape</i> . |
| <code>view_as(other)</code> | View this tensor as the same size as <i>other</i> . |
| <code>vsplit(split_size_or_sections)</code> | See <code>torch.vsplit()</code> |

continues on next page

Table 25 – continued from previous page

| | |
|---|---|
| <code>where(condition, y)</code> | <code>self.where(condition, y)</code> is equivalent to <code>torch.where(condition, self, y)</code> . |
| <code>xlogy(other)</code> | See <code>torch.xlogy()</code> |
| <code>xlogy_(other)</code> | In-place version of <code>xlogy()</code> |
| <code>xpu([device, non_blocking, memory_format])</code> | Returns a copy of this object in XPU memory. |
| <code>zero_()</code> | Fills <code>self</code> tensor with zeros. |

Attributes

| | |
|----------------------------|---|
| <code>H</code> | Returns a view of a matrix (2-D tensor) conjugated and transposed. |
| <code>T</code> | Returns a view of this tensor with its dimensions reversed. |
| <code>data</code> | |
| <code>device</code> | Is the <code>torch.device</code> where this Tensor is. |
| <code>dtype</code> | |
| <code>grad</code> | This attribute is <code>None</code> by default and becomes a Tensor the first time a call to <code>backward()</code> computes gradients for <code>self</code> . |
| <code>grad_fn</code> | |
| <code>imag</code> | Returns a new tensor containing imaginary values of the <code>self</code> tensor. |
| <code>is_cpu</code> | Is <code>True</code> if the Tensor is stored on the CPU, <code>False</code> otherwise. |
| <code>is_cuda</code> | Is <code>True</code> if the Tensor is stored on the GPU, <code>False</code> otherwise. |
| <code>is_ipu</code> | Is <code>True</code> if the Tensor is stored on the IPU, <code>False</code> otherwise. |
| <code>is_leaf</code> | All Tensors that have <code>requires_grad</code> which is <code>False</code> will be leaf Tensors by convention. |
| <code>is_meta</code> | Is <code>True</code> if the Tensor is a meta tensor, <code>False</code> otherwise. |
| <code>is_mkldnn</code> | |
| <code>is_mps</code> | Is <code>True</code> if the Tensor is stored on the MPS device, <code>False</code> otherwise. |
| <code>is_nested</code> | |
| <code>is_ort</code> | |
| <code>is_quantized</code> | Is <code>True</code> if the Tensor is quantized, <code>False</code> otherwise. |
| <code>is_sparse</code> | Is <code>True</code> if the Tensor uses sparse COO storage layout, <code>False</code> otherwise. |
| <code>is_sparse_csr</code> | Is <code>True</code> if the Tensor uses sparse CSR storage layout, <code>False</code> otherwise. |

continues on next page

Table 26 – continued from previous page

| | |
|----------------------------|---|
| <code>is_vulkan</code> | |
| <code>is_xla</code> | Is True if the Tensor is stored on an XLA device, False otherwise. |
| <code>is_xpu</code> | Is True if the Tensor is stored on the XPU, False otherwise. |
| <code>itemsize</code> | Alias for <code>element_size()</code> |
| <code>layout</code> | |
| <code>mH</code> | Accessing this property is equivalent to calling <code>adjoint()</code> . |
| <code>mT</code> | Returns a view of this tensor with the last two dimensions transposed. |
| <code>name</code> | |
| <code>names</code> | Stores names for each of this tensor's dimensions. |
| <code>nbytes</code> | Returns the number of bytes consumed by the "view" of elements of the Tensor if the Tensor does not use sparse storage layout. |
| <code>ndim</code> | Alias for <code>dim()</code> |
| <code>output_nr</code> | |
| <code>real</code> | Returns a new tensor containing real values of the <code>self</code> tensor for a complex-valued input tensor. |
| <code>requires_grad</code> | Is True if gradients need to be computed for this Tensor, False otherwise. |
| <code>retains_grad</code> | Is True if this Tensor is non-leaf and its <code>grad</code> is enabled to be populated during <code>backward()</code> , False otherwise. |
| <code>shape</code> | Returns the size of the <code>self</code> tensor. |
| <code>volatile</code> | |

property `is_leaf`: bool

All Tensors that have `requires_grad` which is False will be leaf Tensors by convention.

For Tensors that have `requires_grad` which is True, they will be leaf Tensors if they were created by the user. This means that they are not the result of an operation and so `grad_fn` is None.

Only leaf Tensors will have their `grad` populated during a call to `backward()`. To get `grad` populated for non-leaf Tensors, you can use `retain_grad()`.

Example:

```
>>> a = torch.rand(10, requires_grad=True)
>>> a.is_leaf
True
>>> b = torch.rand(10, requires_grad=True).cuda()
>>> b.is_leaf
False
# b was created by the operation that cast a cpu Tensor into a cuda Tensor
>>> c = torch.rand(10, requires_grad=True) + 2
>>> c.is_leaf
```

(continues on next page)

(continued from previous page)

```

False
# c was created by the addition operation
>>> d = torch.rand(10).cuda()
>>> d.is_leaf
True
# d does not require gradients and so has no operation creating it (that is,
→ tracked by the autograd engine)
>>> e = torch.rand(10).cuda().requires_grad_()
>>> e.is_leaf
True
# e requires gradients and has no operations creating it
>>> f = torch.rand(10, requires_grad=True, device="cuda")
>>> f.is_leaf
True
# f requires grad, has no operation creating it

```

materialize(*shape*, *device=None*, *dtype=None*)

Create a Parameter with the same properties of the uninitialized one. Given a shape, it materializes a parameter in the same device and with the same *dtype* as the current one or the specified ones in the arguments.

Parameters

- **shape** (*Tuple[int, ...]*) – (tuple): the shape for the materialized tensor.
- **device** (*torch.device*) – the desired device of the parameters and buffers in this module. Optional.
- **dtype** (*torch.dtype*) – the desired floating point type of the floating point parameters and buffers in this module. Optional.

Return type

None

share_memory_()

Moves the underlying storage to shared memory.

This is a no-op if the underlying storage is already in shared memory and for CUDA tensors. Tensors in shared memory cannot be resized.

Return type

[UninitializedParameter](#)

pytorch_pfn_extras.nn.parallel

Classes

[pytorch_pfn_extras.nn.parallel.
DistributedDataParallel](#)(module)

Module for distributed data parallelism

pytorch_pfn_extras.nn.parallel.DistributedDataParallel

```
class pytorch_pfn_extras.nn.parallel.DistributedDataParallel(module, broadcast_buffers=True,
                                                         negotiate_grads=True,
                                                         process_group=None,
                                                         reduce_function=None,
                                                         broadcast_function=None,
                                                         **kwargs)
```

Bases: Module

Module for distributed data parallelism

This class synchronizes the gradients and the buffers after backward computations.

Parameters

- **module** (*Module*) – torch.nn.Module object to be trained
- **broadcast_buffers** (*bool*) – Boolean flag to broadcast buffers after backward computations. Broadcasting buffers may be helpful when the module includes BatchNormalization. However, it will degrade training throughput. (default: *True*)
- **negotiate_grads** (*bool*) – Boolean flag to choose gradients to be sent before all-reduce. This flag is necessary when the computation graph of the module is dynamic. (default: *True*)
- **process_group** (*Optional[ProcessGroup]*) – Process group used for broadcasting and reducing. (default: *torch.distributed.group.WORLD*)
- **reduce_function** (*Optional[Callable[[Sequence[Tensor], Optional[ProcessGroup]], None]]*) – All-reduce function
- **broadcast_function** (*Optional[Callable[[Sequence[Tensor], Optional[ProcessGroup]], None]]*) – Broadcast function
- **kwargs** (*Any*) –

This module receives keyword arguments for the compatibility with *torch.nn.parallel.DistributedDataParallel*. It shows a warning when setting the ignored arguments.

Methods

| | |
|---|--|
| <code>__init__(module[, broadcast_buffers, ...])</code> | This module receives keyword arguments for the compatibility with <i>torch.nn.parallel.DistributedDataParallel</i> . |
| <code>add_module(name, module)</code> | Adds a child module to the current module. |
| <code>apply(fn)</code> | Applies <i>fn</i> recursively to every submodule (as returned by <code>.children()</code>) as well as self. |
| <code>bfloat16()</code> | Casts all floating point parameters and buffers to <code>bfloat16</code> datatype. |
| <code>buffers([recurse])</code> | Returns an iterator over module buffers. |
| <code>children()</code> | Returns an iterator over immediate children modules. |
| <code>compile(*args, **kwargs)</code> | Compile this Module's forward using <code>torch.compile()</code> . |
| <code>cpu()</code> | Moves all model parameters and buffers to the CPU. |
| <code>cuda([device])</code> | Moves all model parameters and buffers to the GPU. |

continues on next page

Table 27 – continued from previous page

| | |
|---|---|
| <code>double()</code> | Casts all floating point parameters and buffers to <code>double</code> datatype. |
| <code>eval()</code> | Sets the module in evaluation mode. |
| <code>extra_repr()</code> | Set the extra representation of the module |
| <code>float()</code> | Casts all floating point parameters and buffers to <code>float</code> datatype. |
| <code>forward(*args, **kwargs)</code> | Defines the computation performed at every call. |
| <code>get_buffer(target)</code> | Returns the buffer given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_extra_state()</code> | Returns any extra state to include in the module's <code>state_dict</code> . |
| <code>get_parameter(target)</code> | Returns the parameter given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_submodule(target)</code> | Returns the submodule given by <code>target</code> if it exists, otherwise throws an error. |
| <code>half()</code> | Casts all floating point parameters and buffers to <code>half</code> datatype. |
| <code>ipu([device])</code> | Moves all model parameters and buffers to the IPU. |
| <code>load_state_dict(state_dict[, strict])</code> | Copies parameters and buffers from <code>state_dict</code> into this module and its descendants. |
| <code>modules()</code> | Returns an iterator over all modules in the network. |
| <code>named_buffers([prefix, recurse, ...])</code> | Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself. |
| <code>named_children()</code> | Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself. |
| <code>named_modules([memo, prefix, remove_duplicate])</code> | Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself. |
| <code>named_parameters([prefix, recurse, ...])</code> | Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself. |
| <code>no_sync()</code> | A context manager to disable synchronization after backward |
| <code>parameters([recurse])</code> | Returns an iterator over module parameters. |
| <code>register_backward_hook(hook)</code> | Registers a backward hook on the module. |
| <code>register_buffer(name, tensor[, persistent])</code> | Adds a buffer to the module. |
| <code>register_comm_hook(hook)</code> | Registers a hook function. |
| <code>register_forward_hook(hook, *[, prepend, ...])</code> | Registers a forward hook on the module. |
| <code>register_forward_pre_hook(hook, *[, ...])</code> | Registers a forward pre-hook on the module. |
| <code>register_full_backward_hook(hook[, prepend])</code> | Registers a backward hook on the module. |
| <code>register_full_backward_pre_hook(hook[, prepend])</code> | Registers a backward pre-hook on the module. |
| <code>register_load_state_dict_post_hook(hook)</code> | Registers a post hook to be run after module's <code>load_state_dict</code> is called. |
| <code>register_module(name, module)</code> | Alias for <code>add_module()</code> . |
| <code>register_parameter(name, param)</code> | Adds a parameter to the module. |
| <code>register_state_dict_pre_hook(hook)</code> | These hooks will be called with arguments: <code>self</code> , <code>prefix</code> , and <code>keep_vars</code> before calling <code>state_dict</code> on <code>self</code> . |

continues on next page

Table 27 – continued from previous page

| | |
|--|--|
| <code>requires_grad_([requires_grad])</code> | Change if autograd should record operations on parameters in this module. |
| <code>set_extra_state(state)</code> | This function is called from <code>load_state_dict()</code> to handle any extra state found within the <code>state_dict</code> . |
| <code>share_memory()</code> | See <code>torch.Tensor.share_memory_()</code> |
| <code>state_dict()</code> | Returns a dictionary containing references to the whole state of the module. |
| <code>to(*args, **kwargs)</code> | Moves and/or casts the parameters and buffers. |
| <code>to_empty(*, device[, recurse])</code> | Moves the parameters and buffers to the specified device without copying storage. |
| <code>train([mode])</code> | Sets the module in training mode. |
| <code>type(dst_type)</code> | Casts all parameters and buffers to <code>dst_type</code> . |
| <code>xpu([device])</code> | Moves all model parameters and buffers to the XPU. |
| <code>zero_grad([set_to_none])</code> | Resets gradients of all model parameters. |

Attributes

| | |
|------------------------------|--|
| <code>T_destination</code> | alias of <code>TypeVar('T_destination', bound=Mapping[str, Tensor])</code> |
| <code>call_super_init</code> | |
| <code>dump_patches</code> | |

T_destination

alias of `TypeVar('T_destination', bound=Mapping[str, Tensor])`

__init__ (*module*, *broadcast_buffers=True*, *negotiate_grads=True*, *process_group=None*, *reduce_function=None*, *broadcast_function=None*, ***kwargs*)

This module receives keyword arguments for the compatibility with `torch.nn.parallel.DistributedDataParallel`. It shows a warning when setting the ignored arguments.

Parameters

- **module** (*Module*) –
- **broadcast_buffers** (*bool*) –
- **negotiate_grads** (*bool*) –
- **process_group** (*Optional[ProcessGroup]*) –
- **reduce_function** (*Optional[Callable[[Sequence[Tensor], Optional[ProcessGroup]], None]]*) –
- **broadcast_function** (*Optional[Callable[[Sequence[Tensor], Optional[ProcessGroup]], None]]*) –
- **kwargs** (*Any*) –

Return type

None

forward(*args, **kwargs)

Defines the computation performed at every call.

Should be overridden by all subclasses.

Note: Although the recipe for forward pass needs to be defined within this function, one should call the Module instance afterwards instead of this since the former takes care of running the registered hooks while the latter silently ignores them.

Parameters

- **args** (*Any*) –
- **kwargs** (*Any*) –

Return type

Any

load_state_dict(state_dict, strict=True, *args)

Copies parameters and buffers from `state_dict` into this module and its descendants. If `strict` is `True`, then the keys of `state_dict` must exactly match the keys returned by this module's `state_dict()` function.

Warning: If `assign` is `True` the optimizer must be created after the call to `load_state_dict`.

Parameters

- **state_dict** (*dict*) – a dict containing parameters and persistent buffers.
- **strict** (*bool*, *optional*) – whether to strictly enforce that the keys in `state_dict` match the keys returned by this module's `state_dict()` function. Default: `True`
- **assign** (*bool*, *optional*) – whether to assign items in the state dictionary to their corresponding keys in the module instead of copying them inplace into the module's current parameters and buffers. When `False`, the properties of the tensors in the current module are preserved while when `True`, the properties of the Tensors in the state dict are preserved. Default: `False`
- **args** (*Any*) –

Returns

- **missing_keys** is a list of str containing the missing keys
- **unexpected_keys** is a list of str containing the unexpected keys

Return type

NamedTuple with `missing_keys` and `unexpected_keys` fields

Note: If a parameter or buffer is registered as `None` and its corresponding key exists in `state_dict`, `load_state_dict()` will raise a `RuntimeError`.

no_sync()

A context manager to disable synchronization after backward

Return type*Generator[None, None, None]***register_comm_hook(hook)**

Registers a hook function. This module will invoke the hook before starting the synchronization.

Args: hook: Callable object that will be invoked before synchronization

Parameters

hook (*Callable*[[[DistributedDataParallel](#)], None]) –

Return type*RemovableHandle***state_dict()**

Returns a dictionary containing references to the whole state of the module.

Both parameters and persistent buffers (e.g. running averages) are included. Keys are corresponding parameter and buffer names. Parameters and buffers set to None are not included.

Note: The returned object is a shallow copy. It contains references to the module's parameters and buffers.

Warning: Currently `state_dict()` also accepts positional arguments for `destination`, `prefix` and `keep_vars` in order. However, this is being deprecated and keyword arguments will be enforced in future releases.

Warning: Please avoid the use of argument `destination` as it is not designed for end-users.

Parameters

- **destination** (*dict*, *optional*) – If provided, the state of module will be updated into the dict and the same object is returned. Otherwise, an `OrderedDict` will be created and returned. Default: `None`.
- **prefix** (*str*, *optional*) – a prefix added to parameter and buffer names to compose the keys in `state_dict`. Default: `''`.
- **keep_vars** (*bool*, *optional*) – by default the `Tensor`s returned in the state dict are detached from autograd. If it's set to `True`, detaching will not be performed. Default: `False`.

Returns

a dictionary containing a whole state of the module

Return type*dict*

Example:

```
>>> # xdoctest: +SKIP("undefined vars")
>>> module.state_dict().keys()
['bias', 'weight']
```

training: `bool`

Modules

pytorch_pfn_extras.nn.parallel.distributed

pytorch_pfn_extras.nn.parallel.distributed

Functions

| | |
|---|----------------------------|
| <i>pytorch_pfn_extras.nn.parallel.distributed.contextmanager(func)</i> | @contextmanager decorator. |
| <i>pytorch_pfn_extras.nn.parallel.distributed.get_foreach_wrapper()</i> | |
| <i>pytorch_pfn_extras.nn.parallel.distributed.record(tag)</i> | |

pytorch_pfn_extras.nn.parallel.distributed.contextmanager

pytorch_pfn_extras.nn.parallel.distributed.**contextmanager**(*func*)

@contextmanager decorator.

Typical usage:

```
@contextmanager
def some_generator(<arguments>):
    <setup> try:
        yield <value>

    finally:
        <cleanup>
```

This makes this:

```
with some_generator(<arguments>) as <variable>:
    <body>
```

equivalent to this:

```
<setup> try:
    <variable> = <value> <body>

finally:
    <cleanup>
```

pytorch_pfn_extras.nn.parallel.distributed.get_foreach_wrapper

pytorch_pfn_extras.nn.parallel.distributed.get_foreach_wrapper()

Return type

_ForEachWrapper

pytorch_pfn_extras.nn.parallel.distributed.record

pytorch_pfn_extras.nn.parallel.distributed.**record**(tag, metric=None, use_cuda=False, enable=True, device='cpu', trace=False)

Parameters

- **tag** (*Optional[str]*) –
- **metric** (*Optional[str]*) –
- **use_cuda** (*bool*) –
- **enable** (*bool*) –
- **device** (*DeviceLike*) –
- **trace** (*Union[Tracer, bool]*) –

Return type

Generator[_ReportNotification, None, None]

Classes

| | |
|---|--|
| <code>pytorch_pfn_extras.nn.parallel.distributed.DistributedDataParallel(module)</code> | Module for distributed data parallelism |
| <code>pytorch_pfn_extras.nn.parallel.distributed.OrderedDict</code> | Dictionary that remembers insertion order |
| <code>pytorch_pfn_extras.nn.parallel.distributed.TypeVar(...)</code> | Type variable. |
| <code>pytorch_pfn_extras.nn.parallel.distributed.Variable</code> | |
| <code>pytorch_pfn_extras.nn.parallel.distributed.record_function(name)</code> | Context manager/function decorator that adds a label to a block of Python code (or function) when running autograd profiler. |

pytorch_pfn_extras.nn.parallel.distributed.DistributedDataParallel

```
class pytorch_pfn_extras.nn.parallel.distributed.DistributedDataParallel(module, broadcast_buffers=True, negotiate_grads=True, process_group=None, reduce_function=None, broadcast_function=None, **kwargs)
```

Bases: Module

Module for distributed data parallelism

This class synchronizes the gradients and the buffers after backward computations.

Parameters

- **module** (*Module*) – torch.nn.Module object to be trained
- **broadcast_buffers** (*bool*) – Boolean flag to broadcast buffers after backward computations. Broadcasting buffers may be helpful when the module includes BatchNormalization. However, it will degrade training throughput. (default: *True*)
- **negotiate_grads** (*bool*) – Boolean flag to choose gradients to be sent before all-reduce. This flag is necessary when the computation graph of the module is dynamic. (default: *True*)
- **process_group** (*Optional[ProcessGroup]*) – Process group used for broadcasting and reducing. (default: *torch.distributed.group.WORLD*)
- **reduce_function** (*Optional[Callable[[Sequence[Tensor], Optional[ProcessGroup]], None]]*) – All-reduce function
- **broadcast_function** (*Optional[Callable[[Sequence[Tensor], Optional[ProcessGroup]], None]]*) – Broadcast function
- **kwargs** (*Any*) –

This module receives keyword arguments for the compatibility with *torch.nn.parallel.DistributedDataParallel*. It shows a warning when setting the ignored arguments.

Methods

| | |
|---|--|
| <code>__init__(module[, broadcast_buffers, ...])</code> | This module receives keyword arguments for the compatibility with <i>torch.nn.parallel.DistributedDataParallel</i> . |
| <code>add_module(name, module)</code> | Adds a child module to the current module. |
| <code>apply(fn)</code> | Applies <i>fn</i> recursively to every submodule (as returned by <code>.children()</code>) as well as self. |
| <code>bfloat16()</code> | Casts all floating point parameters and buffers to <code>bfloat16</code> datatype. |
| <code>buffers([recurse])</code> | Returns an iterator over module buffers. |
| <code>children()</code> | Returns an iterator over immediate children modules. |

continues on next page

Table 28 – continued from previous page

| | |
|---|---|
| <code>compile(*args, **kwargs)</code> | Compile this Module's forward using <code>torch.compile()</code> . |
| <code>cpu()</code> | Moves all model parameters and buffers to the CPU. |
| <code>cuda([device])</code> | Moves all model parameters and buffers to the GPU. |
| <code>double()</code> | Casts all floating point parameters and buffers to <code>double</code> datatype. |
| <code>eval()</code> | Sets the module in evaluation mode. |
| <code>extra_repr()</code> | Set the extra representation of the module |
| <code>float()</code> | Casts all floating point parameters and buffers to <code>float</code> datatype. |
| <code>forward(*args, **kwargs)</code> | Defines the computation performed at every call. |
| <code>get_buffer(target)</code> | Returns the buffer given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_extra_state()</code> | Returns any extra state to include in the module's <code>state_dict</code> . |
| <code>get_parameter(target)</code> | Returns the parameter given by <code>target</code> if it exists, otherwise throws an error. |
| <code>get_submodule(target)</code> | Returns the submodule given by <code>target</code> if it exists, otherwise throws an error. |
| <code>half()</code> | Casts all floating point parameters and buffers to <code>half</code> datatype. |
| <code>ipu([device])</code> | Moves all model parameters and buffers to the IPU. |
| <code>load_state_dict(state_dict[, strict])</code> | Copies parameters and buffers from <code>state_dict</code> into this module and its descendants. |
| <code>modules()</code> | Returns an iterator over all modules in the network. |
| <code>named_buffers([prefix, recurse, ...])</code> | Returns an iterator over module buffers, yielding both the name of the buffer as well as the buffer itself. |
| <code>named_children()</code> | Returns an iterator over immediate children modules, yielding both the name of the module as well as the module itself. |
| <code>named_modules([memo, prefix, remove_duplicate])</code> | Returns an iterator over all modules in the network, yielding both the name of the module as well as the module itself. |
| <code>named_parameters([prefix, recurse, ...])</code> | Returns an iterator over module parameters, yielding both the name of the parameter as well as the parameter itself. |
| <code>no_sync()</code> | A context manager to disable synchronization after backward |
| <code>parameters([recurse])</code> | Returns an iterator over module parameters. |
| <code>register_backward_hook(hook)</code> | Registers a backward hook on the module. |
| <code>register_buffer(name, tensor[, persistent])</code> | Adds a buffer to the module. |
| <code>register_comm_hook(hook)</code> | Registers a hook function. |
| <code>register_forward_hook(hook, *[, prepend, ...])</code> | Registers a forward hook on the module. |
| <code>register_forward_pre_hook(hook, *[, ...])</code> | Registers a forward pre-hook on the module. |
| <code>register_full_backward_hook(hook[, prepend])</code> | Registers a backward hook on the module. |
| <code>register_full_backward_pre_hook(hook[, prepend])</code> | Registers a backward pre-hook on the module. |
| <code>register_load_state_dict_post_hook(hook)</code> | Registers a post hook to be run after module's <code>load_state_dict</code> is called. |
| <code>register_module(name, module)</code> | Alias for <code>add_module()</code> . |

continues on next page

Table 28 – continued from previous page

| | |
|---|---|
| <code>register_parameter(name, param)</code> | Adds a parameter to the module. |
| <code>register_state_dict_pre_hook(hook)</code> | These hooks will be called with arguments: <code>self</code> , <code>prefix</code> , and <code>keep_vars</code> before calling <code>state_dict</code> on <code>self</code> . |
| <code>requires_grad_([requires_grad])</code> | Change if autograd should record operations on parameters in this module. |
| <code>set_extra_state(state)</code> | This function is called from <code>load_state_dict()</code> to handle any extra state found within the <code>state_dict</code> . |
| <code>share_memory()</code> | See <code>torch.Tensor.share_memory_()</code> |
| <code>state_dict()</code> | Returns a dictionary containing references to the whole state of the module. |
| <code>to(*args, **kwargs)</code> | Moves and/or casts the parameters and buffers. |
| <code>to_empty(*, device[, recurse])</code> | Moves the parameters and buffers to the specified device without copying storage. |
| <code>train([mode])</code> | Sets the module in training mode. |
| <code>type(dst_type)</code> | Casts all parameters and buffers to <code>dst_type</code> . |
| <code>xpu([device])</code> | Moves all model parameters and buffers to the XPU. |
| <code>zero_grad([set_to_none])</code> | Resets gradients of all model parameters. |

Attributes

| | |
|------------------------------|--|
| <code>T_destination</code> | alias of <code>TypeVar('T_destination', bound=Mapping[str, Tensor])</code> |
| <code>call_super_init</code> | |
| <code>dump_patches</code> | |

T_destination

alias of `TypeVar('T_destination', bound=Mapping[str, Tensor])`

__init__ (*module*, *broadcast_buffers=True*, *negotiate_grads=True*, *process_group=None*, *reduce_function=None*, *broadcast_function=None*, ***kwargs*)

This module receives keyword arguments for the compatibility with `torch.nn.parallel.DistributedDataParallel`. It shows a warning when setting the ignored arguments.

Parameters

- **module** (*Module*) –
- **broadcast_buffers** (*bool*) –
- **negotiate_grads** (*bool*) –
- **process_group** (*Optional[ProcessGroup]*) –
- **reduce_function** (*Optional[Callable[[Sequence[Tensor], Optional[ProcessGroup]], None]]*) –
- **broadcast_function** (*Optional[Callable[[Sequence[Tensor], Optional[ProcessGroup]], None]]*) –
- **kwargs** (*Any*) –

Return type

None

forward(*args, **kwargs)

Defines the computation performed at every call.

Should be overridden by all subclasses.

Note: Although the recipe for forward pass needs to be defined within this function, one should call the Module instance afterwards instead of this since the former takes care of running the registered hooks while the latter silently ignores them.

Parameters

- **args** (*Any*) –
- **kwargs** (*Any*) –

Return type*Any***load_state_dict**(state_dict, strict=True, *args)

Copies parameters and buffers from *state_dict* into this module and its descendants. If *strict* is True, then the keys of *state_dict* must exactly match the keys returned by this module's *state_dict()* function.

Warning: If *assign* is True the optimizer must be created after the call to *load_state_dict*.

Parameters

- **state_dict** (*dict*) – a dict containing parameters and persistent buffers.
- **strict** (*bool*, *optional*) – whether to strictly enforce that the keys in *state_dict* match the keys returned by this module's *state_dict()* function. Default: True
- **assign** (*bool*, *optional*) – whether to assign items in the state dictionary to their corresponding keys in the module instead of copying them inplace into the module's current parameters and buffers. When False, the properties of the tensors in the current module are preserved while when True, the properties of the Tensors in the state dict are preserved. Default: False
- **args** (*Any*) –

Returns

- **missing_keys** is a list of str containing the missing keys
- **unexpected_keys** is a list of str containing the unexpected keys

Return typeNamedTuple with *missing_keys* and *unexpected_keys* fields

Note: If a parameter or buffer is registered as None and its corresponding key exists in *state_dict*, *load_state_dict()* will raise a *RuntimeError*.

no_sync()

A context manager to disable synchronization after backward

Return type

Generator[None, None, None]

register_comm_hook(hook)

Registers a hook function. This module will invoke the hook before starting the synchronization.

Args: hook: Callable object that will be invoked before synchronization

Parameters

hook (*Callable*[[*DistributedDataParallel*], None]) –

Return type

RemovableHandle

state_dict()

Returns a dictionary containing references to the whole state of the module.

Both parameters and persistent buffers (e.g. running averages) are included. Keys are corresponding parameter and buffer names. Parameters and buffers set to `None` are not included.

Note: The returned object is a shallow copy. It contains references to the module's parameters and buffers.

Warning: Currently `state_dict()` also accepts positional arguments for `destination`, `prefix` and `keep_vars` in order. However, this is being deprecated and keyword arguments will be enforced in future releases.

Warning: Please avoid the use of argument `destination` as it is not designed for end-users.

Parameters

- **destination** (*dict*, *optional*) – If provided, the state of module will be updated into the dict and the same object is returned. Otherwise, an `OrderedDict` will be created and returned. Default: `None`.
- **prefix** (*str*, *optional*) – a prefix added to parameter and buffer names to compose the keys in `state_dict`. Default: `''`.
- **keep_vars** (*bool*, *optional*) – by default the `Tensor`s returned in the state dict are detached from autograd. If it's set to `True`, detaching will not be performed. Default: `False`.

Returns

a dictionary containing a whole state of the module

Return type

`dict`

Example:


```
>>> # xdoctest: +SKIP("undefined vars")
>>> module.state_dict().keys()
['bias', 'weight']
```

training: bool

pytorch_pfn_extras.nn.parallel.distributed.OrderedDict

class pytorch_pfn_extras.nn.parallel.distributed.**OrderedDict**

Bases: dict

Dictionary that remembers insertion order

Methods

| | |
|---|--|
| <code>__init__(*args, **kwargs)</code> | |
| <code>clear()</code> | |
| <code>copy()</code> | |
| <code>fromkeys([value])</code> | Create a new ordered dictionary with keys from iterable and values set to value. |
| <code>get(key[, default])</code> | Return the value for key if key is in the dictionary, else default. |
| <code>items()</code> | |
| <code>keys()</code> | |
| <code>move_to_end(key[, last])</code> | Move an existing element to the end (or beginning if last is false). |
| <code>pop(k[,d])</code> | value. |
| <code>popitem([last])</code> | Remove and return a (key, value) pair from the dictionary. |
| <code>setdefault(key[, default])</code> | Insert key with a value of default if key is not in the dictionary. |
| <code>update([E,]**F)</code> | If E is present and has a .keys() method, then does: for k in E: D[k] = E[k] If E is present and lacks a .keys() method, then does: for k, v in E: D[k] = v In either case, this is followed by: for k in F: D[k] = F[k] |
| <code>values()</code> | |

`__init__(*args, **kwargs)`

`clear()` → None. Remove all items from od.

`copy()` → a shallow copy of od

`fromkeys(value=None)`

Create a new ordered dictionary with keys from iterable and values set to value.

items() → a set-like object providing a view on D's items

keys() → a set-like object providing a view on D's keys

move_to_end(key, last=True)

Move an existing element to the end (or beginning if last is false).

Raise KeyError if the element does not exist.

pop(k[, d]) → v, remove specified key and return the corresponding value. If key is not found, d is returned if given, otherwise KeyError is raised.

popitem(last=True)

Remove and return a (key, value) pair from the dictionary.

Pairs are returned in LIFO order if last is true or FIFO order if false.

setdefault(key, default=None)

Insert key with a value of default if key is not in the dictionary.

Return the value for key if key is in the dictionary, else default.

update([E], **F) → None. Update D from dict/iterable E and F.

If E is present and has a .keys() method, then does: for k in E: D[k] = E[k] If E is present and lacks a .keys() method, then does: for k, v in E: D[k] = v In either case, this is followed by: for k in F: D[k] = F[k]

values() → an object providing a view on D's values

pytorch_pfn_extras.nn.parallel.distributed.TypeVar

```
class pytorch_pfn_extras.nn.parallel.distributed.TypeVar(name, *constraints, bound=None,
                                                         covariant=False, contravariant=False)
```

Bases: `_Final`, `_Immutable`

Type variable.

Usage:

```
T = TypeVar('T') # Can be anything
A = TypeVar('A', str, bytes) # Must be str or bytes
```

Type variables exist primarily for the benefit of static type checkers. They serve as the parameters for generic types as well as for generic function definitions. See class `Generic` for more information on generic types. Generic functions work as follows:

```
def repeat(x: T, n: int) -> List[T]:
    """Return a list containing n references to x.""" return [x]*n
```

```
def longest(x: A, y: A) -> A:
    """Return the longest of two strings.""" return x if len(x) >= len(y) else y
```

The latter example's signature is essentially the overloading of (str, str) -> str and (bytes, bytes) -> bytes. Also note that if the arguments are instances of some subclass of str, the return type is still plain str.

At runtime, `isinstance(x, T)` and `issubclass(C, T)` will raise `TypeError`.

Type variables defined with `covariant=True` or `contravariant=True` can be used to declare covariant or contravariant generic types. See PEP 484 for more details. By default generic types are invariant in all type variables.

Type variables can be introspected. e.g.:

```
T.__name__ == 'T' T.__constraints__ == () T.__covariant__ == False T.__contravariant__ = False
A.__constraints__ == (str, bytes)
```

Note that only type variables defined in global scope can be pickled.

Methods

```
__init__(name, *constraints[, bound, ...])
```

```
__init__(name, *constraints, bound=None, covariant=False, contravariant=False)
```

pytorch_pfn_extras.nn.parallel.distributed.Variable

```
class pytorch_pfn_extras.nn.parallel.distributed.Variable
```

Bases: `_LegacyVariableBase`

Methods

```
__init__()
```

pytorch_pfn_extras.nn.parallel.distributed.record_function

```
class pytorch_pfn_extras.nn.parallel.distributed.record_function(name, args=None)
```

Bases: `ContextDecorator`

Context manager/function decorator that adds a label to a block of Python code (or function) when running autograd profiler. It is useful when tracing the code profile.

Parameters

- **name** (*str*) – Label assigned to the block of code.
- **node_id** (*int*) – ID of node, for distributed profiling. Unset in
- **cases.** (*non-distributed*) –
- **args** (*Optional[str]*) –

Example

```
>>> # xdoctest: +REQUIRES(env:TORCH_DOCTEST_AUTOGRAD_PROFILER)
>>> x = torch.randn((1, 1), requires_grad=True)
>>> with torch.autograd.profiler.profile() as prof:
...     y = x ** 2
...     with torch.autograd.profiler.record_function("label-z"): # label the block
...         z = y ** 3
...     y.backward()
```

(continues on next page)

(continued from previous page)

```

...
>>> # xdoctest: +IGNORE_WANT
>>> # NOTE: some columns were removed for brevity
>>> print(prof.key_averages().table(sort_by="self_cpu_time_total"))
-----
↪--
Name                Self CPU total %   CPU time avg      Number of
↪Calls
-----
↪--
pow                  60.77%           47.470us          3
mul                  21.73%           25.465us          2
PowBackward0        12.03%           121.891us         1
torch::autograd::AccumulateGrad  2.70%           6.324us           1
label-z             2.13%           12.421us           1
torch::autograd::GraphRoot      0.64%           1.503us           1
-----
↪--
Self CPU time total: 234.344us
CUDA time total: 0.000us

```

Methods

`__init__(name[, args])`

`__init__(name, args=None)`

Parameters

- **name** (*str*) –
- **args** (*Optional[str]*) –

pytorch_pfn_extras.onnx**Functions**

| | |
|--|---|
| <code>pytorch_pfn_extras.onnx.annotate(**attrs)</code> | Annotation parameters to the target function. |
| <code>pytorch_pfn_extras.onnx.apply_annotation(fn, ...)</code> | Annotation applicator to the target function |
| <code>pytorch_pfn_extras.onnx.as_output(name, value)</code> | |
| <code>pytorch_pfn_extras.onnx.export(model, args, f)</code> | Export model into ONNX Graph. |
| <code>pytorch_pfn_extras.onnx.export_testcase(...)</code> | Export model and I/O tensors of the model in protobuf format. |
| <code>pytorch_pfn_extras.onnx.grad(output, inputs)</code> | |
| <code>pytorch_pfn_extras.onnx.is_large_tensor(...)</code> | |
| <code>pytorch_pfn_extras.onnx.load_model(f[, ...])</code> | Load model from ONNX file. |
| <code>pytorch_pfn_extras.onnx.no_grad(fn, *args, ...)</code> | |
| <code>pytorch_pfn_extras.onnx.scoped_anchor(**attrs)</code> | Add anchor node to the scoped modules |

pytorch_pfn_extras.onnx.annotate

`pytorch_pfn_extras.onnx.annotate(**attrs)`

Annotation parameters to the target function.

Usage:

```
>>> class Net(nn.Module):
...     def __init__(self):
...         super(Net, self).__init__()
...         self.conv = nn.Conv2d(1, 6, 3)
...         self.conv2 = nn.Conv2d(6, 12, 3)
...     def forward(self, x):
...         with pytorch_pfn_extras.onnx.annotate(key='value'):
...             h = self.conv(x)
...             h = self.conv2(h)
...         return h
```

Use this annotate function under with statement, then the first Conv operator will be emit with customized attributes. Customized attributes are invalid for ONNX format, so pay attention that some ONNX runtimes cannot run the output ONNX graph.

This annotation is enabled with either `pytorch_pfn_extras.onnx.export_testcase` or `pytorch_pfn_extras.onnx.export`.

Parameters

attrs (*dict*) – annotation parameters

Return type

AbstractContextManager[None]

pytorch_pfn_extras.onnx.apply_annotation

`pytorch_pfn_extras.onnx.apply_annotation(fn, *args, **attrs)`

Annotation applier to the target function

Usage:

```
>>> class Net(nn.Module):
...     def __init__(self):
...         super(Net, self).__init__()
...         self.conv = nn.Conv2d(1, 6, 3)
...         self.conv2 = nn.Conv2d(6, 12, 3)
...     def forward(self, x):
...         def _conv(x):
...             h = self.conv(x)
...             return torch.relu(h)
...         h = pytorch_pfn_extras.onnx.apply_annotation(
...             _conv, key='value')
...         h = self.conv2(h)
...         return h
```

Annotate into all operators emitted from the target function even if included not `nn.Module` function. On the above code, the first Conv and ReLu operator will be emit with customized attributes. Customized attributes are invalid for ONNX format, so pay attention that some ONNX runtimes cannot run the output ONNX graph.

This applier is enabled with either `pytorch_pfn_extras.onnx.export_testcase` or `pytorch_pfn_extras.onnx.export`.

Parameters

- **fn** (*func*) – the target function to be annotated, `args` is used for this function. Cannot pass `kwargs` for the function.
- **args** (*tuple*) – arguments for the target function
- **attrs** (*dict*) – annotation parameters

Return type

Any

pytorch_pfn_extras.onnx.as_output

`pytorch_pfn_extras.onnx.as_output(name, value, add_identity=True)`

Parameters

- **name** (*str*) –
- **value** (*Tensor*) –
- **add_identity** (*bool*) –

Return type

Tensor

pytorch_pfn_extras.onnx.export

`pytorch_pfn_extras.onnx.export(model, args, f, return_output=False, strip_large_tensor_data=False, large_tensor_threshold=100, chrome_tracing="", **kwargs)`

Export model into ONNX Graph.

Parameters

- **f** (*IO*) – A file-like object or a string file path to be written to this file.
- **return_output** (*bool*) – If True, return output values come from the model.
- **strip_large_tensor_data** (*bool*) – If True, this function will strip data of large tensors to reduce ONNX file size for benchmarking
- **large_tensor_threshold** (*int*) – If number of elements of tensor is larger than this value, the tensor is stripped when *strip_large_tensor_data* is True
- **model** (*Module*) –
- **args** (*Sequence[Any]*) –
- **chrome_tracing** (*str*) –
- **kwargs** (*Any*) –

Return type

Any

Warning: This function is not thread safe.

pytorch_pfn_extras.onnx.export_testcase

`pytorch_pfn_extras.onnx.export_testcase(model, args, out_dir, *, output_grad=False, metadata=True, model_overwrite=True, strip_large_tensor_data=False, large_tensor_threshold=100, return_output=False, user_meta=None, export_torch_script=False, export_torch_trace=False, export_chrome_tracing=True, **kwargs)`

Export model and I/O tensors of the model in protobuf format.

Parameters

- **output_grad** (*bool or Tensor*) – If True, this function will output model's gradient with names 'gradient_%.d.pb'. If set Tensor, use it as gradient *input*. The gradient inputs are output as 'gradient_input_%.d.pb' along with gradient.
- **metadata** (*bool*) – If True, output meta information taken from git log.
- **model_overwrite** (*bool*) – If False and model.onnx has already existed, only export input/output data as another test dataset.
- **strip_large_tensor_data** (*bool*) – If True, this function will strip data of large tensors to reduce ONNX file size for benchmarking
- **large_tensor_threshold** (*int*) – If number of elements of tensor is larger than this value, the tensor is stripped when *strip_large_tensor_data* is True
- **return_output** (*bool*) – If True, return output values come from the model.

- **export_torch_script** (*bool*) – Output model_script.pt using torch.jit.script
- **export_torch_trace** (*bool*) – Output model_trace.pt using torch.jit.trace
- **model** (*Union[Module, ScriptModule]*) –
- **args** (*Any*) –
- **out_dir** (*str*) –
- **user_meta** (*Optional[Mapping[str, Any]]*) –
- **export_chrome_tracing** (*bool*) –
- **kwargs** (*Any*) –

Return type*Any*

Warning: This function is not thread safe.

Note: When exporting a model whose forward takes keyword arguments of `torch.Tensor` type, you can pass them by putting a dict as the last element of `args`. When the keyword arguments have default values, you need to explicitly include them into the dict. Also, you must explicitly specify `input_names` that are the names of both positional and keyword arguments.

pytorch_pfn_extras.onnx.grad

`pytorch_pfn_extras.onnx.grad(output, inputs, retain_graph=None, create_graph=False, only_inputs=True, allow_unused=False)`

Parameters

- **output** (*Tensor*) –
- **inputs** (*Tuple[Tensor, ...]*) –
- **retain_graph** (*Optional[bool]*) –
- **create_graph** (*bool*) –
- **only_inputs** (*bool*) –
- **allow_unused** (*bool*) –

Return type*Tuple[Optional[Tensor], ...]*

pytorch_pfn_extras.onnx.is_large_tensor

`pytorch_pfn_extras.onnx.is_large_tensor(tensor, threshold)`

Parameters

- **tensor** (*TensorProto*) –
- **threshold** (*int*) –

Return type

bool

pytorch_pfn_extras.onnx.load_model

`pytorch_pfn_extras.onnx.load_model(f, format=None, load_external_data=True)`

Load model from ONNX file.

This is a wrapper to *onnx.load_model* that automatically falls back to *load_external_data=False* when tensors are stripped.

Parameters

- **f** (*Union[IO, str]*) – A file-like object or a string file path to be written to this file.
- **format** (*Optional[Any]*) – A reserved arg
- **load_external_data** (*bool*) – If True and the external data under the same directory of the model, load the external data

Return type

ModelProto

pytorch_pfn_extras.onnx.no_grad

`pytorch_pfn_extras.onnx.no_grad(fn, *args, **kwargs)`

Parameters

- **fn** (*Callable[[...], Any]*) –
- **args** (*Any*) –
- **kwargs** (*Any*) –

Return type

Any

pytorch_pfn_extras.onnx.scoped_anchor

`pytorch_pfn_extras.onnx.scoped_anchor(**attrs)`

Add anchor node to the scoped modules

Usage:

```

>>> class Net(nn.Module):
...     def __init__(self):
...         super(Net, self).__init__()
...         self.conv = nn.Conv2d(1, 6, 3)
...         self.conv2 = nn.Conv2d(6, 12, 3)
...     def forward(self, x):
...         with pytorch_pfn_extras.onnx.scoped_anchor(key='value'):
...             h = self.conv(x)
...             h = self.conv2(h)
...         return h

```

Use this scoped anchoring under with statement, then dummy Identity nodes are added before/after the first Conv operator with customized attributes.

This anchoring is triggered by `nn.Module` applying function, cannot use this with `torch.*` functions.

This annotation is enabled with either `pytorch_pfn_extras.onnx.export_testcase` or `pytorch_pfn_extras.onnx.export`.

Parameters

attrs (*dict*) – annotation parameters

Return type

AbstractContextManager[None]

Modules

| | |
|---|---|
| <code>pytorch_pfn_extras.onnx.annotate(**attrs)</code> | Annotation parameters to the target function. |
| <code>pytorch_pfn_extras.onnx.export_testcase(...)</code> | Export model and I/O tensors of the model in protobuf format. |
| <code>pytorch_pfn_extras.onnx.load</code> | |
| <code>pytorch_pfn_extras.onnx.pfto_exporter</code> | |
| <code>pytorch_pfn_extras.onnx.strip_large_tensor</code> | |
| <code>pytorch_pfn_extras.onnx.symbolic_registry</code> | |
| <code>pytorch_pfn_extras.onnx.unstrip_tensor</code> | |

pytorch_pfn_extras.onnx.load

Functions

| | |
|--|----------------------------|
| <code>pytorch_pfn_extras.onnx.load.load_model(f[, ...])</code> | Load model from ONNX file. |
|--|----------------------------|

pytorch_pfn_extras.onnx.load.load_model

`pytorch_pfn_extras.onnx.load.load_model(f, format=None, load_external_data=True)`

Load model from ONNX file.

This is a wrapper to `onnx.load_model` that automatically falls back to `load_external_data=False` when tensors are stripped.

Parameters

- **f** (*Union*[*IO*, *str*]) – A file-like object or a string file path to be written to this file.
- **format** (*Optional*[*Any*]) – A reserved arg
- **load_external_data** (*bool*) – If True and the external data under the same directory of the model, load the external data

Return type

ModelProto

Classes

| | |
|--|---|
| <code>pytorch_pfn_extras.onnx.load.IO(*args, **kws)</code> | Generic base class for TextIO and BinaryIO. |
| <code>pytorch_pfn_extras.onnx.load.Path(*args, ...)</code> | PurePath subclass that can make system calls. |
| <code>pytorch_pfn_extras.onnx.load.Text</code> | alias of <code>str</code> |

pytorch_pfn_extras.onnx.load.IO

class `pytorch_pfn_extras.onnx.load.IO(*args, **kws)`

Bases: `Generic`

Generic base class for TextIO and BinaryIO.

This is an abstract, generic version of the return of `open()`.

NOTE: This does not distinguish between the different possible classes (text vs. binary, read vs. write vs. read/write, append-only, unbuffered). The TextIO and BinaryIO subclasses below capture the distinctions between text vs. binary, which is pervasive in the interface; however we currently do not offer a way to track the other distinctions in the type system.

Methods

__init__()

close()

fileno()

flush()

isatty()

read([n])

readable()

readline([limit])

readlines([hint])

seek(offset[, whence])

seekable()

tell()

truncate([size])

writable()

write(s)

writelines(lines)

Attributes

closed

mode

name

abstract close()**Return type**

None

abstract property closed: bool

```

abstract fileno()

    Return type
    int

abstract flush()

    Return type
    None

abstract isatty()

    Return type
    bool

abstract property mode: str

abstract property name: str

abstract read(n=-1)

    Parameters
    n (int) –

    Return type
    AnyStr

abstract readable()

    Return type
    bool

abstract readline(limit=-1)

    Parameters
    limit (int) –

    Return type
    AnyStr

abstract readlines(hint=-1)

    Parameters
    hint (int) –

    Return type
    List

abstract seek(offset, whence=0)

    Parameters
    • offset (int) –
    • whence (int) –

    Return type
    int

abstract seekable()

    Return type
    bool

```

abstract tell()

Return type
int

abstract truncate(*size=None*)

Parameters
size (*Optional[int]*) –

Return type
int

abstract writable()

Return type
bool

abstract write(*s*)

Parameters
s (*AnyStr*) –

Return type
int

abstract writelines(*lines*)

Parameters
lines (*List*) –

Return type
None

pytorch_pfn_extras.onnx.load.Path

class pytorch_pfn_extras.onnx.load.**Path**(*args, **kwargs)

Bases: PurePath

PurePath subclass that can make system calls.

Path represents a filesystem path but unlike PurePath, also offers methods to do system calls on path objects. Depending on your system, instantiating a Path will return either a PosixPath or a WindowsPath object. You can also instantiate a PosixPath or WindowsPath directly, but cannot instantiate a WindowsPath on a POSIX system or vice versa.

Construct a PurePath from one or several strings and or existing PurePath objects. The strings and path objects are combined so as to yield a canonicalized path, which is incorporated into the new PurePath object.

Methods

| | |
|---|---|
| <code>__init__()</code> | |
| <code>absolute()</code> | Return an absolute version of this path. |
| <code>as_posix()</code> | Return the string representation of the path with forward (/) slashes. |
| <code>as_uri()</code> | Return the path as a 'file' URI. |
| <code>chmod(mode)</code> | Change the permissions of the path, like <code>os.chmod()</code> . |
| <code>cwd()</code> | Return a new path pointing to the current working directory (as returned by <code>os.getcwd()</code>). |
| <code>exists()</code> | Whether this path exists. |
| <code>expanduser()</code> | Return a new path with expanded ~ and ~user constructs (as returned by <code>os.path.expanduser</code>) |
| <code>glob(pattern)</code> | Iterate over this subtree and yield all existing files (of any kind, including directories) matching the given relative pattern. |
| <code>group()</code> | Return the group name of the file gid. |
| <code>home()</code> | Return a new path pointing to the user's home directory (as returned by <code>os.path.expanduser('~')</code>). |
| <code>is_absolute()</code> | True if the path is absolute (has both a root and, if applicable, a drive). |
| <code>is_block_device()</code> | Whether this path is a block device. |
| <code>is_char_device()</code> | Whether this path is a character device. |
| <code>is_dir()</code> | Whether this path is a directory. |
| <code>is_fifo()</code> | Whether this path is a FIFO. |
| <code>is_file()</code> | Whether this path is a regular file (also True for symlinks pointing to regular files). |
| <code>is_mount()</code> | Check if this path is a POSIX mount point |
| <code>is_reserved()</code> | Return True if the path contains one of the special names reserved by the system, if any. |
| <code>is_socket()</code> | Whether this path is a socket. |
| <code>is_symlink()</code> | Whether this path is a symbolic link. |
| <code>iterdir()</code> | Iterate over the files in this directory. |
| <code>joinpath(*args)</code> | Combine this path with one or several arguments, and return a new path representing either a subpath (if all arguments are relative paths) or a totally different path (if one of the arguments is anchored). |
| <code>lchmod(mode)</code> | Like <code>chmod()</code> , except if the path points to a symlink, the symlink's permissions are changed, rather than its target's. |
| <code>link_to(target)</code> | Make the target path a hard link pointing to this path. |
| <code>lstat()</code> | Like <code>stat()</code> , except if the path points to a symlink, the symlink's status information is returned, rather than its target's. |
| <code>match(path_pattern)</code> | Return True if this path matches the given pattern. |
| <code>mkdir([mode, parents, exist_ok])</code> | Create a new directory at this given path. |
| <code>open([mode, buffering, encoding, errors, ...])</code> | Open the file pointed by this path and return a file object, as the built-in <code>open()</code> function does. |
| <code>owner()</code> | Return the login name of the file owner. |
| <code>read_bytes()</code> | Open the file in bytes mode, read it, and close the file. |

continues on next page

Table 29 – continued from previous page

| | |
|--|---|
| <code>read_text(encoding, errors)</code> | Open the file in text mode, read it, and close the file. |
| <code>relative_to(*other)</code> | Return the relative path to another path identified by the passed arguments. |
| <code>rename(target)</code> | Rename this path to the target path. |
| <code>replace(target)</code> | Rename this path to the target path, overwriting if that path exists. |
| <code>resolve([strict])</code> | Make the path absolute, resolving all symlinks on the way and also normalizing it (for example turning slashes into backslashes under Windows). |
| <code>rglob(pattern)</code> | Recursively yield all existing files (of any kind, including directories) matching the given relative pattern, anywhere in this subtree. |
| <code>rmdir()</code> | Remove this directory. |
| <code>samefile(other_path)</code> | Return whether other_path is the same or not as this file (as returned by <code>os.path.samefile()</code>). |
| <code>stat()</code> | Return the result of the <code>stat()</code> system call on this path, like <code>os.stat()</code> does. |
| <code>symlink_to(target[, target_is_directory])</code> | Make this path a symlink pointing to the target path. |
| <code>touch([mode, exist_ok])</code> | Create this file with the given access mode, if it doesn't exist. |
| <code>unlink([missing_ok])</code> | Remove this file or link. |
| <code>with_name(name)</code> | Return a new path with the file name changed. |
| <code>with_suffix(suffix)</code> | Return a new path with the file suffix changed. |
| <code>write_bytes(data)</code> | Open the file in bytes mode, write to it, and close the file. |
| <code>write_text(data[, encoding, errors])</code> | Open the file in text mode, write to it, and close the file. |

Attributes

| | |
|-----------------------|--|
| <code>anchor</code> | The concatenation of the drive and root, or "". |
| <code>drive</code> | The drive prefix (letter or UNC path), if any. |
| <code>name</code> | The final path component, if any. |
| <code>parent</code> | The logical parent of the path. |
| <code>parents</code> | A sequence of this path's logical parents. |
| <code>parts</code> | An object providing sequence-like access to the components in the filesystem path. |
| <code>root</code> | The root of the path, if any. |
| <code>stem</code> | The final path component, minus its last suffix. |
| <code>suffix</code> | The final component's last suffix, if any. |
| <code>suffixes</code> | A list of the final component's suffixes, if any. |

`absolute()`

Return an absolute version of this path. This function works even if the path doesn't point to anything.

No normalization is done, i.e. all `'.'` and `'..'` will be kept along. Use `resolve()` to get the canonical path to a file.

`chmod(mode)`

Change the permissions of the path, like `os.chmod()`.

classmethod `cwd()`

Return a new path pointing to the current working directory (as returned by `os.getcwd()`).

exists()

Whether this path exists.

expanduser()

Return a new path with expanded `~` and `~user` constructs (as returned by `os.path.expanduser`)

glob(*pattern*)

Iterate over this subtree and yield all existing files (of any kind, including directories) matching the given relative pattern.

group()

Return the group name of the file gid.

classmethod `home()`

Return a new path pointing to the user's home directory (as returned by `os.path.expanduser('~')`).

is_block_device()

Whether this path is a block device.

is_char_device()

Whether this path is a character device.

is_dir()

Whether this path is a directory.

is_fifo()

Whether this path is a FIFO.

is_file()

Whether this path is a regular file (also True for symlinks pointing to regular files).

is_mount()

Check if this path is a POSIX mount point

is_socket()

Whether this path is a socket.

is_symlink()

Whether this path is a symbolic link.

iterdir()

Iterate over the files in this directory. Does not yield any result for the special paths `'.'` and `'..'`.

lchmod(*mode*)

Like `chmod()`, except if the path points to a symlink, the symlink's permissions are changed, rather than its target's.

link_to(*target*)

Make the target path a hard link pointing to this path.

Note this function does not make this path a hard link to *target*, despite the implication of the function and argument names. The order of arguments (*target*, *link*) is the reverse of `Path.symlink_to`, but matches that of `os.link`.

lstat()

Like `stat()`, except if the path points to a symlink, the symlink's status information is returned, rather than its target's.

makedirs(mode=511, parents=False, exist_ok=False)

Create a new directory at this given path.

open(mode='r', buffering=-1, encoding=None, errors=None, newline=None)

Open the file pointed by this path and return a file object, as the built-in `open()` function does.

owner()

Return the login name of the file owner.

read_bytes()

Open the file in bytes mode, read it, and close the file.

read_text(encoding=None, errors=None)

Open the file in text mode, read it, and close the file.

rename(target)

Rename this path to the target path.

The target path may be absolute or relative. Relative paths are interpreted relative to the current working directory, *not* the directory of the Path object.

Returns the new Path instance pointing to the target path.

replace(target)

Rename this path to the target path, overwriting if that path exists.

The target path may be absolute or relative. Relative paths are interpreted relative to the current working directory, *not* the directory of the Path object.

Returns the new Path instance pointing to the target path.

resolve(strict=False)

Make the path absolute, resolving all symlinks on the way and also normalizing it (for example turning slashes into backslashes under Windows).

rglob(pattern)

Recursively yield all existing files (of any kind, including directories) matching the given relative pattern, anywhere in this subtree.

rmdir()

Remove this directory. The directory must be empty.

samefile(other_path)

Return whether `other_path` is the same or not as this file (as returned by `os.path.samefile()`).

stat()

Return the result of the `stat()` system call on this path, like `os.stat()` does.

symlink_to(target, target_is_directory=False)

Make this path a symlink pointing to the target path. Note the order of arguments (link, target) is the reverse of `os.symlink`.

touch(mode=438, exist_ok=True)

Create this file with the given access mode, if it doesn't exist.

unlink(*missing_ok=False*)

Remove this file or link. If the path is a directory, use `rmdir()` instead.

write_bytes(*data*)

Open the file in bytes mode, write to it, and close the file.

write_text(*data, encoding=None, errors=None*)

Open the file in text mode, write to it, and close the file.

pytorch_pfn_extras.onnx.load.Text

`pytorch_pfn_extras.onnx.load.Text`

alias of `str`

pytorch_pfn_extras.onnx.pfto_exporter

Modules

`pytorch_pfn_extras.onnx.pfto_exporter.`
`export`

pytorch_pfn_extras.onnx.strip_large_tensor

Functions

| | |
|---|--|
| <code>pytorch_pfn_extras.onnx.</code> <code>strip_large_tensor.is_large_tensor(...)</code> | |
| <code>pytorch_pfn_extras.onnx.</code> <code>strip_large_tensor.reduce(...)</code> | Apply a function of two arguments cumulatively to the items of a sequence, from left to right, so as to reduce the sequence to a single value. |

pytorch_pfn_extras.onnx.strip_large_tensor.is_large_tensor

`pytorch_pfn_extras.onnx.strip_large_tensor.is_large_tensor(tensor, threshold)`

Parameters

- **tensor** (*TensorProto*) –
- **threshold** (*int*) –

Return type

`bool`

pytorch_pfn_extras.onnx.strip_large_tensor.reduce

`pytorch_pfn_extras.onnx.strip_large_tensor.reduce(function, sequence[, initial])` → value

Apply a function of two arguments cumulatively to the items of a sequence, from left to right, so as to reduce the sequence to a single value. For example, `reduce(lambda x, y: x+y, [1, 2, 3, 4, 5])` calculates `((((1+2)+3)+4)+5)`. If `initial` is present, it is placed before the items of the sequence in the calculation, and serves as a default when the sequence is empty.

pytorch_pfn_extras.onnx.symbolic_registry

Functions

| | |
|--|-------------------------|
| <code>pytorch_pfn_extras.onnx.symbolic_registry.cast(...)</code> | Cast a value to a type. |
|--|-------------------------|

| |
|---|
| <code>pytorch_pfn_extras.onnx.symbolic_registry.get_registered_op(...)</code> |
|---|

| |
|--|
| <code>pytorch_pfn_extras.onnx.symbolic_registry.is_registered_op(...)</code> |
|--|

| |
|---|
| <code>pytorch_pfn_extras.onnx.symbolic_registry.register_op(...)</code> |
|---|

Classes

| |
|--|
| <code>pytorch_pfn_extras.onnx.symbolic_registry.Value</code> |
|--|

pytorch_pfn_extras.onnx.unstrip_tensor

Functions

| | |
|---|---|
| <code>pytorch_pfn_extras.onnx.unstrip_tensor.unstrip(path)</code> | Unstrip ONNX models and test data(.pb). |
|---|---|

pytorch_pfn_extras.onnx.unstrip_tensor.unstrip

`pytorch_pfn_extras.onnx.unstrip_tensor.unstrip(path, out_path=)`

Unstrip ONNX models and test data(.pb).

Add tensor(raw data) to the target ONNXs (and test data). Values are random following mean and variance written in meta information.

Parameters

- **path** (*str*) – The target directory path, ONNX file, or Tensor (Protobuf) file path.
- **out_path** (*str*) – Output path to be written.

Return type

None

Classes

| | |
|---|---|
| <code>pytorch_pfn_extras.onnx.unstrip_tensor.Path(...)</code> | PurePath subclass that can make system calls. |
|---|---|

`pytorch_pfn_extras.onnx.unstrip_tensor.Path`

class `pytorch_pfn_extras.onnx.unstrip_tensor.Path(*args, **kwargs)`

Bases: `PurePath`

PurePath subclass that can make system calls.

Path represents a filesystem path but unlike `PurePath`, also offers methods to do system calls on path objects. Depending on your system, instantiating a `Path` will return either a `PosixPath` or a `WindowsPath` object. You can also instantiate a `PosixPath` or `WindowsPath` directly, but cannot instantiate a `WindowsPath` on a POSIX system or vice versa.

Construct a `PurePath` from one or several strings and or existing `PurePath` objects. The strings and path objects are combined so as to yield a canonicalized path, which is incorporated into the new `PurePath` object.

Methods

| | |
|--------------------------------|--|
| <code>__init__()</code> | |
| <code>absolute()</code> | Return an absolute version of this path. |
| <code>as_posix()</code> | Return the string representation of the path with forward (/) slashes. |
| <code>as_uri()</code> | Return the path as a 'file' URI. |
| <code>chmod(mode)</code> | Change the permissions of the path, like <code>os.chmod()</code> . |
| <code>cwd()</code> | Return a new path pointing to the current working directory (as returned by <code>os.getcwd()</code>). |
| <code>exists()</code> | Whether this path exists. |
| <code>expanduser()</code> | Return a new path with expanded ~ and ~user constructs (as returned by <code>os.path.expanduser</code>) |
| <code>glob(pattern)</code> | Iterate over this subtree and yield all existing files (of any kind, including directories) matching the given relative pattern. |
| <code>group()</code> | Return the group name of the file gid. |
| <code>home()</code> | Return a new path pointing to the user's home directory (as returned by <code>os.path.expanduser('~')</code>). |
| <code>is_absolute()</code> | True if the path is absolute (has both a root and, if applicable, a drive). |
| <code>is_block_device()</code> | Whether this path is a block device. |
| <code>is_char_device()</code> | Whether this path is a character device. |
| <code>is_dir()</code> | Whether this path is a directory. |
| <code>is_fifo()</code> | Whether this path is a FIFO. |
| <code>is_file()</code> | Whether this path is a regular file (also True for symlinks pointing to regular files). |
| <code>is_mount()</code> | Check if this path is a POSIX mount point |

continues on next page

Table 30 – continued from previous page

| | |
|---|---|
| <code>is_reserved()</code> | Return True if the path contains one of the special names reserved by the system, if any. |
| <code>is_socket()</code> | Whether this path is a socket. |
| <code>is_symlink()</code> | Whether this path is a symbolic link. |
| <code>iterdir()</code> | Iterate over the files in this directory. |
| <code>joinpath(*args)</code> | Combine this path with one or several arguments, and return a new path representing either a subpath (if all arguments are relative paths) or a totally different path (if one of the arguments is anchored). |
| <code>lchmod(mode)</code> | Like <code>chmod()</code> , except if the path points to a symlink, the symlink's permissions are changed, rather than its target's. |
| <code>link_to(target)</code> | Make the target path a hard link pointing to this path. |
| <code>lstat()</code> | Like <code>stat()</code> , except if the path points to a symlink, the symlink's status information is returned, rather than its target's. |
| <code>match(path_pattern)</code> | Return True if this path matches the given pattern. |
| <code>mkdir([mode, parents, exist_ok])</code> | Create a new directory at this given path. |
| <code>open([mode, buffering, encoding, errors, ...])</code> | Open the file pointed by this path and return a file object, as the built-in <code>open()</code> function does. |
| <code>owner()</code> | Return the login name of the file owner. |
| <code>read_bytes()</code> | Open the file in bytes mode, read it, and close the file. |
| <code>read_text([encoding, errors])</code> | Open the file in text mode, read it, and close the file. |
| <code>relative_to(*other)</code> | Return the relative path to another path identified by the passed arguments. |
| <code>rename(target)</code> | Rename this path to the target path. |
| <code>replace(target)</code> | Rename this path to the target path, overwriting if that path exists. |
| <code>resolve([strict])</code> | Make the path absolute, resolving all symlinks on the way and also normalizing it (for example turning slashes into backslashes under Windows). |
| <code>rglob(pattern)</code> | Recursively yield all existing files (of any kind, including directories) matching the given relative pattern, anywhere in this subtree. |
| <code>rmdir()</code> | Remove this directory. |
| <code>samefile(other_path)</code> | Return whether <code>other_path</code> is the same or not as this file (as returned by <code>os.path.samefile()</code>). |
| <code>stat()</code> | Return the result of the <code>stat()</code> system call on this path, like <code>os.stat()</code> does. |
| <code>symlink_to(target[, target_is_directory])</code> | Make this path a symlink pointing to the target path. |
| <code>touch([mode, exist_ok])</code> | Create this file with the given access mode, if it doesn't exist. |
| <code>unlink([missing_ok])</code> | Remove this file or link. |
| <code>with_name(name)</code> | Return a new path with the file name changed. |
| <code>with_suffix(suffix)</code> | Return a new path with the file suffix changed. |
| <code>write_bytes(data)</code> | Open the file in bytes mode, write to it, and close the file. |
| <code>write_text(data[, encoding, errors])</code> | Open the file in text mode, write to it, and close the file. |

Attributes

| | |
|-----------------|--|
| anchor | The concatenation of the drive and root, or "". |
| drive | The drive prefix (letter or UNC path), if any. |
| name | The final path component, if any. |
| parent | The logical parent of the path. |
| parents | A sequence of this path's logical parents. |
| parts | An object providing sequence-like access to the components in the filesystem path. |
| root | The root of the path, if any. |
| stem | The final path component, minus its last suffix. |
| suffix | The final component's last suffix, if any. |
| suffixes | A list of the final component's suffixes, if any. |

absolute()

Return an absolute version of this path. This function works even if the path doesn't point to anything.

No normalization is done, i.e. all '.' and '..' will be kept along. Use resolve() to get the canonical path to a file.

chmod(mode)

Change the permissions of the path, like os.chmod().

classmethod cwd()

Return a new path pointing to the current working directory (as returned by os.getcwd()).

exists()

Whether this path exists.

expanduser()

Return a new path with expanded ~ and ~user constructs (as returned by os.path.expanduser)

glob(pattern)

Iterate over this subtree and yield all existing files (of any kind, including directories) matching the given relative pattern.

group()

Return the group name of the file gid.

classmethod home()

Return a new path pointing to the user's home directory (as returned by os.path.expanduser('~')).

is_block_device()

Whether this path is a block device.

is_char_device()

Whether this path is a character device.

is_dir()

Whether this path is a directory.

is_fifo()

Whether this path is a FIFO.

is_file()

Whether this path is a regular file (also True for symlinks pointing to regular files).

is_mount()

Check if this path is a POSIX mount point

is_socket()

Whether this path is a socket.

is_symlink()

Whether this path is a symbolic link.

iterdir()

Iterate over the files in this directory. Does not yield any result for the special paths '.' and '..'.

lchmod(mode)

Like chmod(), except if the path points to a symlink, the symlink's permissions are changed, rather than its target's.

link_to(target)

Make the target path a hard link pointing to this path.

Note this function does not make this path a hard link to *target*, despite the implication of the function and argument names. The order of arguments (target, link) is the reverse of Path.symlink_to, but matches that of os.link.

lstat()

Like stat(), except if the path points to a symlink, the symlink's status information is returned, rather than its target's.

makedirs(mode=511, parents=False, exist_ok=False)

Create a new directory at this given path.

open(mode='r', buffering=-1, encoding=None, errors=None, newline=None)

Open the file pointed by this path and return a file object, as the built-in open() function does.

owner()

Return the login name of the file owner.

read_bytes()

Open the file in bytes mode, read it, and close the file.

read_text(encoding=None, errors=None)

Open the file in text mode, read it, and close the file.

rename(target)

Rename this path to the target path.

The target path may be absolute or relative. Relative paths are interpreted relative to the current working directory, *not* the directory of the Path object.

Returns the new Path instance pointing to the target path.

replace(target)

Rename this path to the target path, overwriting if that path exists.

The target path may be absolute or relative. Relative paths are interpreted relative to the current working directory, *not* the directory of the Path object.

Returns the new Path instance pointing to the target path.

resolve(*strict=False*)

Make the path absolute, resolving all symlinks on the way and also normalizing it (for example turning slashes into backslashes under Windows).

rglob(*pattern*)

Recursively yield all existing files (of any kind, including directories) matching the given relative pattern, anywhere in this subtree.

rmdir()

Remove this directory. The directory must be empty.

samefile(*other_path*)

Return whether *other_path* is the same or not as this file (as returned by `os.path.samefile()`).

stat()

Return the result of the `stat()` system call on this path, like `os.stat()` does.

symlink_to(*target, target_is_directory=False*)

Make this path a symlink pointing to the target path. Note the order of arguments (link, target) is the reverse of `os.symlink`.

touch(*mode=438, exist_ok=True*)

Create this file with the given access mode, if it doesn't exist.

unlink(*missing_ok=False*)

Remove this file or link. If the path is a directory, use `rmdir()` instead.

write_bytes(*data*)

Open the file in bytes mode, write to it, and close the file.

write_text(*data, encoding=None, errors=None*)

Open the file in text mode, write to it, and close the file.

pytorch_pfn_extras.ops

Functions

| | |
|---|--|
| <code>pytorch_pfn_extras.ops.register(name, ...)</code> | Register a custom op under <code>torch.ops.ppe.name</code> |
|---|--|

pytorch_pfn_extras.ops.register

`pytorch_pfn_extras.ops.register(name, fwd_op, bwd_op)`

Register a custom op under `torch.ops.ppe.name`

The function appears as a primitive op in the forward and backward `torch.fx.Graph`'s after compiling torch code with ``aot_autograd`` backend. Note that for backward functions, all the arguments of the backward pass together with the forward arguments are passed to it. This means if forward had ``fwd_op(x, y)`` *x, y* arguments, the custom `bwd_op` needs to have a signature like ``bwd_op(grad_output, x, y)``

Parameters

- **name** (*str*) – name of the op, shows how it is registered in `torch.ops.ppe`.
- **fwd_op** (*ppe.ops.OpDesc*) – code that is executed in the forward pass

- **bwd_op** (*ppe.ops.OpDesc*) – code that is executed in the backward pass

Return type

None

Classes

| | |
|--|--|
| <code>pytorch_pfn_extras.ops.OpDesc</code> (op, meta, ...) | Metadata to register an op to torch.library. |
|--|--|

pytorch_pfn_extras.ops.OpDesc**class** `pytorch_pfn_extras.ops.OpDesc`(op, meta, signature)

Bases: object

Metadata to register an op to torch.library.

Parameters

- **op** (*Callable*[[...], Any]) –
- **meta** (*Callable*[[...], Any]) –
- **signature** (*str*) –

op

code to be executed in the forward/backward of the op.

Type

callable

meta

function to perform shape inference for forward/backward passes.

Type

callable

signature

Arguments and return type of the function "(Tensor a, Tensor b) -> Tensor[]".

Type

str

Methods

`__init__`(op, meta, signature)

`__init__`(op, meta, signature)**Parameters**

- **op** (*Callable*[[...], Any]) –
- **meta** (*Callable*[[...], Any]) –
- **signature** (*str*) –

Return type

None

Modules

| | |
|---|--|
| <code>pytorch_pfn_extras.ops.register(name, ...)</code> | Register a custom op under <code>torch.ops.ppe.name</code> |
|---|--|

pytorch_pfn_extras.profiler**Functions**

| | |
|---|---|
| <code>pytorch_pfn_extras.profiler.clear_tracer()</code> | Resets the status of the global tracer. |
| <code>pytorch_pfn_extras.profiler.enable_global_trace(enable)</code> | Enable or disable tracing for all the threads. |
| <code>pytorch_pfn_extras.profiler.enable_thread_trace(enable)</code> | Enable or disable tracing for the current thread. |
| <code>pytorch_pfn_extras.profiler.get_time_summary()</code> | |
| <code>pytorch_pfn_extras.profiler.get_tracer([...])</code> | Gets the current global tracer. |
| <code>pytorch_pfn_extras.profiler.load_chrome_trace_as_json(...)</code> | |
| <code>pytorch_pfn_extras.profiler.record(tag[, ...])</code> | |
| <code>pytorch_pfn_extras.profiler.record_function(tag)</code> | |
| <code>pytorch_pfn_extras.profiler.record_iterable(...)</code> | |

pytorch_pfn_extras.profiler.clear_tracer`pytorch_pfn_extras.profiler.clear_tracer()`

Resets the status of the global tracer.

Return type

None

pytorch_pfn_extras.profiler.enable_global_trace`pytorch_pfn_extras.profiler.enable_global_trace(enable)`

Enable or disable tracing for all the threads.

Parameters**enable** (*bool*) – Enable or disable flag.**Return type**

None

pytorch_pfn_extras.profiler.enable_thread_trace

```
pytorch_pfn_extras.profiler.enable_thread_trace(enable)
```

Enable or disable tracing for the current thread.

Parameters

enable (*bool*) – Enable or disable flag.

Return type

None

pytorch_pfn_extras.profiler.get_time_summary

```
pytorch_pfn_extras.profiler.get_time_summary()
```

Return type

[TimeSummary](#)

pytorch_pfn_extras.profiler.get_tracer

```
pytorch_pfn_extras.profiler.get_tracer(tracer_cls=<class  
                                         'pytorch_pfn_extras.profiler._tracing.ChromeTracer'>, *params)
```

Gets the current global tracer.

Parameters

- **tracer_cls** (*type of [Tracer](#)*) – type of tracer to create if the global tracer hasn't been initialized
- **params** (*Any*) –

Return type

[Tracer](#)

pytorch_pfn_extras.profiler.load_chrome_trace_as_json

```
pytorch_pfn_extras.profiler.load_chrome_trace_as_json(filename)
```

Parameters

filename (*str*) –

Return type

List[Dict[str, Any]]

pytorch_pfn_extras.profiler.record

`pytorch_pfn_extras.profiler.record(tag, metric=None, use_cuda=False, enable=True, device='cpu', trace=False)`

Parameters

- **tag** (*Optional[str]*) –
- **metric** (*Optional[str]*) –
- **use_cuda** (*bool*) –
- **enable** (*bool*) –
- **device** (*DeviceLike*) –
- **trace** (*Union[Tracer, bool]*) –

Return type

Generator[_ReportNotification, None, None]

pytorch_pfn_extras.profiler.record_function

`pytorch_pfn_extras.profiler.record_function(tag, use_cuda=False, enable=True, device='cpu', trace=False)`

Parameters

- **tag** (*Optional[str]*) –
- **use_cuda** (*bool*) –
- **enable** (*bool*) –
- **device** (*DeviceLike*) –
- **trace** (*Union[Tracer, bool]*) –

Return type

Callable[[Callable[[_T]], _T], Callable[[_T]]]

pytorch_pfn_extras.profiler.record_iterable

`pytorch_pfn_extras.profiler.record_iterable(tag, iter, divide_metric=False, use_cuda=False, enable=True, device='cpu', trace=False)`

Parameters

- **tag** (*Optional[str]*) –
- **iter** (*Iterable[_T]*) –
- **divide_metric** (*bool*) –
- **use_cuda** (*bool*) –
- **enable** (*bool*) –
- **device** (*DeviceLike*) –
- **trace** (*Union[Tracer, bool]*) –

Return type*Iterable[_T]***Classes**

| | |
|--|--|
| <code>pytorch_pfn_extras.profiler.ChromeTracer(...)</code> | Tracer object that outputs a timeline in Chrome format. |
| <code>pytorch_pfn_extras.profiler.TimeSummary(*[, ...])</code> | Online summarization of execution times. |
| <code>pytorch_pfn_extras.profiler.TraceableDataset(...)</code> | Utility class to trace a Dataset inside the DataLoader worker threads. |
| <code>pytorch_pfn_extras.profiler.Tracer()</code> | |

pytorch_pfn_extras.profiler.ChromeTracer**class** `pytorch_pfn_extras.profiler.ChromeTracer`(*max_event_count=None, enable=True, append=True*)Bases: *Tracer*

Tracer object that outputs a timeline in Chrome format.

Parameters

- **max_event_count** (*int*) – Limit the amount of events that can be traced, optional.
- **enable** (*bool*) – Sets the tracer in active state. Optional, defaults to True.
- **append** (*bool*) –

Methods

| |
|--|
| <code>__init__</code> ([<i>max_event_count, enable, append</i>]) |
| <code>add_event</code> (<i>name</i>) |
| <code>add_remote_event</code> (<i>name, event</i>) |
| <code>clear</code> () |
| <code>enable</code> (<i>enable_flag</i>) |
| <code>finalize</code> () |
| <code>flush</code> (<i>filename, writer</i>) |
| <code>initialize_writer</code> (<i>filename, writer</i>) |
| <code>load_state_dict</code> (<i>to_load</i>) |
| <code>state_dict</code> () |

__init__(*max_event_count=None, enable=True, append=True*)

Parameters

- **max_event_count** (*Optional[int]*) –
- **enable** (*bool*) –
- **append** (*bool*) –

Return type

None

add_event(*name*)

Parameters

name (*str*) –

Return type

Generator[None, None, None]

add_remote_event(*name, event*)

Parameters

- **name** (*str*) –
- **event** (*Dict[str, Union[str, int, float]]*) –

Return type

None

clear()

Return type

None

enable(*enable_flag*)

Parameters

enable_flag (*bool*) –

Return type

None

finalize()

Return type

None

flush(*filename, writer*)

Parameters

- **filename** (*str*) –
- **writer** (*Writer*) –

Return type

None

initialize_writer(*filename, writer*)

Parameters

- **filename** (*str*) –

- **writer** (`Writer`) –

Return type

None

load_state_dict (*to_load*)

Parameters

to_load (*Dict[str, Any]*) –

Return type

None

state_dict ()

Return type

Dict[str, Any]

pytorch_pfn_extras.profiler.TimeSummary

class pytorch_pfn_extras.profiler.**TimeSummary** (*, *max_queue_size=1000*, *auto_init=True*)

Bases: object

Online summarization of execution times.

TimeSummary computes the average and standard deviation of exeuction times in both cpu and gpu devices.

Parameters

- **max_queue_size** (*int*) – Length limit of the internal queues that keep reported time info until they are summarized.
- **auto_init** (*bool*) – Whether to automatically call *initialize()* when the instance is created.

Methods

__init__ (*[, *max_queue_size*, *auto_init*])

add (name, value)

complete_report (tag, use_cuda, begin_event, ...)

finalize ()

initialize ()

Initializes the worker threads for TimeSummary.

report (tag[, use_cuda])

Context manager to automatically report execution times.

summary ([clear])

synchronize ()

__init__ (*, *max_queue_size=1000*, *auto_init=True*)

Parameters

- **max_queue_size** (*int*) –
- **auto_init** (*bool*) –

Return type

None

add(*name*, *value*)**Parameters**

- **name** (*str*) –
- **value** (*float*) –

Return type

None

complete_report(*tag*, *use_cuda*, *begin_event*, *begin*)**Parameters**

- **tag** (*str*) –
- **use_cuda** (*bool*) –
- **begin_event** (*Optional[Event]*) –
- **begin** (*float*) –

Return type

None

finalize()**Return type**

None

initialize()

Initializes the worker threads for TimeSummary.

Usually you do not have to call it for yourself. However in case you directly use `ppe.time_summary` outside of `pytorch_pfn_extras.training.extensions.ProfileReport`, you have to explicitly call `initialize()` in advance.

Return type

None

report(*tag*, *use_cuda=False*)

Context manager to automatically report execution times.

The start and completion times are obtained automatically, the user only needs to provide a tag to identify the value in the summary values.

Parameters

- **tag** (*str*) – A name to identify the section of code being profiled.
- **use_cuda** (*bool*) – Indicates if GPU time should also be profiled.

Return type*Generator[_ReportNotification, None, None]*

summary(*clear=False*)

Parameters

clear (*bool*) –

Return type

Generator[Tuple[DictSummary, Dict[str, float]], None, None]

synchronize()

Return type

None

pytorch_pfn_extras.profiler.TraceableDataset

class pytorch_pfn_extras.profiler.**TraceableDataset**(*dataset, tag, tracer=None*)

Bases: *Dataset*

Utility class to trace a Dataset inside the DataLoader worker threads.

Parameters

- **dataset** (*torch.utils.data.Dataset*) – dataset where `__getitem__` will be traced.
- **tag** (*str*) – Tag will be used to name the events.
- **tracer** (*Tracer*) – Tracer object, optional. If *None* it defaults to *ppe.profile.get_tracer()*.

Methods

`__init__`(*dataset, tag[, tracer]*)

`__init__`(*dataset, tag, tracer=None*)

Parameters

- **dataset** (*Dataset*) –
- **tag** (*str*) –
- **tracer** (*Optional[Tracer]*) –

Return type

None

pytorch_pfn_extras.profiler.Tracer

class pytorch_pfn_extras.profiler.**Tracer**

Bases: *object*

Methods

`__init__()`

`add_event(name)`

`add_remote_event(name, value)`

`clear()`

`enable(enable_flag)`

`finalize()`

`flush(filename, writer)`

`initialize_writer(filename, writer)`

`load_state_dict(to_load)`

`state_dict()`

add_event(*name*)

Parameters
name (*str*) –

Return type
Generator[None, None, None]

add_remote_event(*name*, *value*)

Parameters

- **name** (*str*) –

- **value** (*Any*) –

Return type

None

clear()

Return type

None

enable(*enable_flag*)

Parameters
enable_flag (*bool*) –

Return type

None

finalize()

Return type

None

flush(*filename*, *writer*)

Parameters

- **filename** (*str*) –
- **writer** (*Writer*) –

Return type

None

initialize_writer(*filename*, *writer*)

Parameters

- **filename** (*str*) –
- **writer** (*Writer*) –

Return type

None

load_state_dict(*to_load*)

Parameters

to_load (*Dict[str, Any]*) –

Return type

None

state_dict()

Return type

Dict[str, Any]

pytorch_pfn_extras.reporting

Functions

| | |
|--|---|
| <code>pytorch_pfn_extras.reporting.get_current_reporter()</code> | Returns the current reporter object. |
| <code>pytorch_pfn_extras.reporting.overload(func)</code> | Decorator for overloaded functions/methods. |
| <code>pytorch_pfn_extras.reporting.report(values)</code> | Reports observed values with the current reporter object. |
| <code>pytorch_pfn_extras.reporting.report_scope(...)</code> | Returns a report scope with the current reporter. |

pytorch_pfn_extras.reporting.get_current_reporter

`pytorch_pfn_extras.reporting.get_current_reporter()`

Returns the current reporter object.

Return type

Reporter

pytorch_pfn_extras.reporting.overload

`pytorch_pfn_extras.reporting.overload(func)`

Decorator for overloaded functions/methods.

In a stub file, place two or more stub definitions for the same function in a row, each decorated with `@overload`. For example:

```
@overload def utf8(value: None) -> None: ... @overload def utf8(value: bytes) -> bytes: ...
@overload def utf8(value: str) -> bytes: ...
```

In a non-stub file (i.e. a regular .py file), do the same but follow it with an implementation. The implementation should *not* be decorated with `@overload`. For example:

```
@overload def utf8(value: None) -> None: ... @overload def utf8(value: bytes) -> bytes: ...
@overload def utf8(value: str) -> bytes: ... def utf8(value):
    # implementation goes here
```

pytorch_pfn_extras.reporting.report

`pytorch_pfn_extras.reporting.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 `Reporter.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 `nn.Module`. Suppose that a module is registered to the current reporter as an observer (for example, the target module of the optimizer is automatically registered to the main reporter). We can report some values from the link as follows:

```
class MyRegressor:
    def __init__(self, predictor):
        super().__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(abs(pred - y)) / len(x)
        loss = F.mean_squared_error(pred, y)

        # Report the mean absolute and squared errors.
        reporter.report({
            'abs_error': abs_error,
            'squared_error': loss,
        }, self)

    return loss
```

If the module is named 'main' in the hierarchy these reported values are named 'main/abs_error' and 'main/squared_error'.

Parameters

- **values** (*dict*) – Dictionary of observed values.
- **observer** (*Optional[Module]*) – 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.

Return type

None

pytorch_pfn_extras.reporting.report_scope`pytorch_pfn_extras.reporting.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.

Parameters

observation (*Dict[str, Union[Tensor, ndarray, floating, float, Callable[[], float]]]*) –

Return type*Generator[None, None, None]***Classes**

| | |
|---|---|
| <code>pytorch_pfn_extras.reporting.DictSummary()</code> | Online summarization of a sequence of dictionaries. |
| <code>pytorch_pfn_extras.reporting.Reporter()</code> | Object to which observed values are reported. |
| <code>pytorch_pfn_extras.reporting.Summary()</code> | Online summarization of a sequence of scalars. |

pytorch_pfn_extras.reporting.DictSummary**class** `pytorch_pfn_extras.reporting.DictSummary`Bases: `object`

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

| | |
|---------------------------------------|--------------------------------------|
| <code>__init__()</code> | |
| <code>add(d)</code> | Adds a dictionary of scalars. |
| <code>compute_mean()</code> | Creates a dictionary of mean values. |
| <code>load_state_dict(to_load)</code> | |
| <code>make_statistics()</code> | Creates a dictionary of statistics. |
| <code>state_dict()</code> | |

`__init__()`

Return type

None

`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. When the value is a tuple, the second element is interpreted as a weight.

Return type

None

`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

`load_state_dict(to_load)`

Parameters

`to_load` (*Dict[str, Any]*) –

Return type

None

`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

`state_dict()`

Return type

Dict[str, Any]

pytorch_pfn_extras.reporting.Reporter

class pytorch_pfn_extras.reporting.Reporter

Bases: object

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 pytorch_pfn_extras.reporting 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:

```
>>> reporter = Reporter()
>>> observation = {}
>>> with reporter:
...     with report_scope(observation):
...         report({'x': 1})
...
>>> observation
{'x': 1}
```

The most important application of Reporter is to report observed values from each link or chain in the training and validation procedures. and some extensions prepare their own Reporter object with the hierarchy of the target module registered as observers. We can use `report()` function inside any `nn.Module` to report the observed values (e.g., training loss, accuracy, activation statistics, etc.).

observation

Dictionary of observed values.

Methods

| | |
|---|---|
| <code>__init__()</code> | |
| <code>add_observer(name, observer)</code> | Registers an observer of values. |
| <code>add_observers(prefix, observers)</code> | Registers multiple observers at once. |
| <code>report(values[, observer])</code> | Reports observed values. |
| <code>scope(observation)</code> | Creates a scope to report observed values to observation. |

`__init__()`

Return type

None

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** (*Module*) – The observer object. Note that the reporter distinguishes the observers by their object ids (i.e., `id(owner)`), rather than the object equality.

Return type

None

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** (*Sequence[Tuple[str, Module]]*) – Iterator of name and observer pairs.

Return type

None

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: If a value is of type `Tensor`, the variable is copied without preserving the computational graph and the new variable object purged from the graph is stored to the observer.

Parameters

- **values** (*dict*) – Dictionary of observed values.

- **observer** (*Optional[Module]*) – 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.

Return type

None

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.

Return type*Generator*[None, None, None]**pytorch_pfn_extras.reporting.Summary****class** pytorch_pfn_extras.reporting.**Summary**

Bases: object

Online summarization of a sequence of scalars.

Summary computes the statistics of given scalars online.

Methods

| | |
|---------------------------------------|--|
| <code>__init__()</code> | |
| <code>add(value[, weight])</code> | Adds a scalar value. |
| <code>compute_mean()</code> | Computes the mean. |
| <code>load_state_dict(to_load)</code> | |
| <code>make_statistics()</code> | Computes and returns the mean and standard deviation values. |
| <code>state_dict()</code> | |

__init__()**Return type**

None

add(*value*, *weight=1*)

Adds a scalar value.

Parameters

- **value** (*Union[[Tensor](#), [ndarray](#), [floating](#), [float](#), [Callable\[\[\[float\]\(#\)\], \[float\]\(#\)\]\]](#)*) – Scalar value to accumulate. It is either a NumPy scalar or a zero-dimensional array (on CPU or GPU).
- **weight** (*Union[[Tensor](#), [ndarray](#), [floating](#), [float](#)]*) – An optional weight for the value. It is a NumPy scalar or a zero-dimensional array (on CPU or GPU). Default is 1 (integer).

Return type

None

compute_mean()

Computes the mean.

Return type*Union[[Tensor](#), [ndarray](#), [floating](#), [float](#)]***load_state_dict(to_load)****Parameters****to_load** (*Dict[str, Any]*) –**Return type**

None

make_statistics()

Computes and returns the mean and standard deviation values.

Returns

Mean and standard deviation values.

Return type

tuple

state_dict()**Return type***Dict[str, Any]***pytorch_pfn_extras.runtime****Classes**

| | |
|---|---|
| <i>pytorch_pfn_extras.runtime.BaseRuntime(...)</i> | A base class for collections of device-specific callback functions. |
| <i>pytorch_pfn_extras.runtime.PyTorchRuntime(...)</i> | A collections of callback functions for the devices that PyTorch supports by default. |

pytorch_pfn_extras.runtime.BaseRuntime

class pytorch_pfn_extras.runtime.**BaseRuntime**(*device_spec, options*)

Bases: object

A base class for collections of device-specific callback functions.

The function attributes of this class will be called from `ppe.to` or `ppe.handler.Handler`.

`ppe.runtime.runtime_registry` stores the runtime classes and dispatches them by feeding the corresponding name string as an input.

Parameters

- **device_spec** (*torch.device* or *str*) – The device that modules and tensors are transferred to.
- **options** (*dict*) – A configuration dictionary that can be used from runtime method.

Methods

| | |
|--|---|
| <code>__init__(device_spec, options)</code> | |
| <code>convert_batch(args)</code> | Transfers the given batch to the specific device. |
| <code>eval_post_step(evaluator, module, batch_idx, ...)</code> | The method called at the end of each evaluation. |
| <code>eval_pre_step(evaluator, module, batch_idx, ...)</code> | The method called at the beginning of each evaluation. |
| <code>execute(code_block, batch)</code> | Method called by the CodeBlocks API to do device dependent execution. |
| <code>initialize_module(module, loader_or_batch[, ...])</code> | Initializes the module at the beginning of training or inference. |
| <code>map(func, iterable[, out_keys, device])</code> | Method called by the user to apply function to iterable efficiently. |
| <code>move_module(module)</code> | Transfers the module to the specific device. |
| <code>move_tensor(tensor)</code> | Transfers the tensor to the specific device. |
| <code>trace(event_name, arg)</code> | Context manager for tracing PPE events in the custom device tools. |
| <code>train_cleanup(module)</code> | A method called only once when completing a training run. |
| <code>train_epoch_begin(module)</code> | Preprocess of each epoch. |
| <code>train_epoch_end(module)</code> | Completion of each epoch. |
| <code>train_post_step(trainer, module, batch_idx, ...)</code> | Postprocess of each step. |
| <code>train_pre_step(trainer, module, batch_idx, batch)</code> | Preprocess of each step. |
| <code>train_validation_begin(module)</code> | The method called before each evaluation. |
| <code>train_validation_end(module)</code> | The method called after each evaluation. |

`__init__(device_spec, options)`

Parameters

- **device_spec** (*Union[str, device]*) –
- **options** (*Dict[str, Any]*) –

Return type

None

convert_batch(args)

Transfers the given batch to the specific device.

Parameters

args (*object*) – A batch data of any type.

Returns

A batch data transferred to the specific device of the same type as input.

Return type

Any

eval_post_step(evaluator, module, batch_idx, batch, outs)

The method called at the end of each evaluation.

Parameters

- **evaluator** ([Evaluator](#)) – An evaluator.
- **module** (*torch.nn.Module*) – A module.
- **batch_idx** (*int*) – The batch index.
- **batch** (*list of torch.Tensor*) – The list of input tensors of this batch.
- **outs** (*Any*) – (list of torch.Tensor): The list of output tensors of this batch.

Return type

None

Returns: None

eval_pre_step(evaluator, module, batch_idx, batch)

The method called at the beginning of each evaluation.

Parameters

- **evaluator** ([Evaluator](#)) – An evaluator.
- **module** (*torch.nn.Module*) – A module.
- **batch_idx** (*int*) – The batch index.
- **batch** (*list of torch.Tensor*) – The list of input tensors of this batch.

Return type

None

Returns: None

execute(code_block, batch)

Method called by the CodeBlocks API to do device dependent execution.

Parameters

- **code_block** ([CodeBlock](#)) – The codeblock requesting execution.
- **batch** (*dict of str, torch.Tensor*) – The input tensors of this batch.

Returns

The results of executing the codeblock on this runtime.

Return type

Any

initialize_module(*module*, *loader_or_batch*, *optimizer=None*)

Initializes the module at the beginning of training or inference.

Parameters

- **module** (*torch.nn.Module*) – A module.
- **loader_or_batch** (*DataLoader* or *torch.Tensor*) – A data loader or a tensor.
- **optimizer** (*Optimizer* or *None*) – An optimizer. This argument is sometimes used to copy LR from the original optimizer to the training model.

Return type

None

Returns: None

map(*func*, *iterable*, *out_keys=None*, *device='cpu'*)

Method called by the user to apply function to iterable efficiently.

Parameters

- **func** (*CodeBlock*) – The function to be executed
- **iterable** (*Iterable[Any]*) – The data
- **out_keys** (*Optional[Set[str]]*) – The output keys that to be moved to the host device
- **device** (*Any*) – The torch device that contains the final outputs

Returns

The result of *func*

Return type

Iterable[Any]

move_module(*module*)

Transfers the module to the specific device.

Before this method is called, `ppe.to` will add this class as an new attribute (“_ppe_runtime”) to the input module.

Parameters

module (*torch.nn.Module*) – A module.

Returns

A module transferred to the specific device.

Return type

Module

move_tensor(*tensor*)

Transfers the tensor to the specific device.

Parameters

tensor (*torch.Tensor*) – A tensor.

Returns

A tensor transferred to the specific device.

Return type

Tensor

classmethod `trace(event_name, arg)`

Context manager for tracing PPE events in the custom device tools.

Parameters

- **event_name** (*Optional[str]*) – The name of the event being traced
- **arg** (*Any*) – Custom argument for the tracer

Return type

Generator[None, None, None]

train_cleanup(*module*)

A method called only once when compleing a training run.

Parameters

module (*torch.nn.Module*) – A module.

Return type

None

Returns: None

train_epoch_begin(*module*)

Preprocess of each epoch.

Parameters

module (*torch.nn.Module*) – A module.

Return type

None

Returns: None

train_epoch_end(*module*)

Completion of each epoch.

Parameters

module (*torch.nn.Module*) – A module.

Return type

None

Returns: None

train_post_step(*trainer, module, batch_idx, batch, outs*)

Postprocess of each step.

This method is called at the end of every steps: the set of (typically one) iterations and an update.

Parameters

- **trainer** (*Trainer*) – A trainer.
- **module** (*torch.nn.Module*) – A module.
- **batch_idx** (*int*) – The batch index.
- **batch** (*list of torch.Tensor*) – The list of input tensors of this batch.
- **outs** (*Any*) – (list of *torch.Tensor*): The list of output tensors of this batch.

Return type

None

Returns: None

train_pre_step(*trainer, module, batch_idx, batch*)

Preprocess of each step.

This method is called at the beginning of every steps: the set of (typically one) iterations and an update.

Parameters

- **trainer** ([Trainer](#)) – A trainer.
- **module** (*torch.nn.Module*) – A module.
- **batch_idx** (*int*) – The batch index.
- **batch** (*list of torch.Tensor*) – The list of input tensors of this batch.

Return type

None

Returns: None

train_validation_begin(*module*)

The method called before each evaluation.

Parameters

module (*torch.nn.Module*) – A module.

Return type

None

Returns: None

train_validation_end(*module*)

The method called after each evaluation.

Parameters

module (*torch.nn.Module*) – A module.

Return type

None

Returns: None

pytorch_pfn_extras.runtime.PyTorchRuntime

class `pytorch_pfn_extras.runtime.PyTorchRuntime`(*device_spec, options*)

Bases: [BaseRuntime](#)

A collections of callback functions for the devices that PyTorch supports by default.

Parameters

- **device_spec** (*torch.device or str*) – The device.
- **options** (*dict, optional*) – The configuration options.
 - **'autocast' (bool or dict):**
If True, `torch.cuda.amp.autocast` is enabled. using `{"enabled": True, "device_type": "cuda"}` as autocast options. Default is False which corresponds to the following options `{"enabled": False, "device_type": "cuda"}` dict type. If dict, Options to pass to `torch.autocast`. Includes `device_type`, `dtype` among others.

- **'grad_scaler' (torch.cuda.amp.GradScaler):**
A gradient scaler that outputs are applied to.

Methods

| | |
|--|---|
| <code>__init__(device_spec, options)</code> | |
| <code>convert_batch(args)</code> | Transfers the given batch to the specific device. |
| <code>eval_post_step(evaluator, module, batch_idx, ...)</code> | The method called at the end of each evaluation. |
| <code>eval_pre_step(evaluator, module, batch_idx, ...)</code> | The method called at the beginning of each evaluation. |
| <code>execute(code_block, batch)</code> | Method called by the CodeBlocks API to do device dependent execution. |
| <code>initialize_module(module, loader_or_batch[, ...])</code> | Initializes the module at the beginning of training or inference. |
| <code>map(func, iterable[, out_keys, device])</code> | Method called by the user to apply function to iterable efficiently. |
| <code>move_module(module)</code> | Transfers the module to the specific device. |
| <code>move_tensor(tensor)</code> | Transfers the tensor to the specific device. |
| <code>trace(event_name, arg)</code> | Context manager for tracing PPE events in the custom device tools. |
| <code>train_cleanup(module)</code> | A method called only once when completing a training run. |
| <code>train_epoch_begin(module)</code> | Preprocess of each epoch. |
| <code>train_epoch_end(module)</code> | Completion of each epoch. |
| <code>train_post_step(trainer, module, batch_idx, ...)</code> | Postprocess of each step. |
| <code>train_pre_step(trainer, module, batch_idx, batch)</code> | Preprocess of each step. |
| <code>train_validation_begin(module)</code> | The method called before each evaluation. |
| <code>train_validation_end(module)</code> | The method called after each evaluation. |

`__init__(device_spec, options)`

Parameters

- **device_spec** (*Union[str, device]*) –
- **options** (*Dict[str, Any]*) –

Return type

None

eval_post_step(*evaluator, module, batch_idx, batch, outs*)

The method called at the end of each evaluation.

Parameters

- **evaluator** (*Evaluator*) – An evaluator.
- **module** (*torch.nn.Module*) – A module.
- **batch_idx** (*int*) – The batch index.
- **batch** (*list of torch.Tensor*) – The list of input tensors of this batch.
- **outs** (*Any*) – (list of torch.Tensor): The list of output tensors of this batch.

Return type

None

Returns: None

eval_pre_step(*evaluator, module, batch_idx, batch*)

The method called at the beginning of each evaluation.

Parameters

- **evaluator** ([Evaluator](#)) – An evaluator.
- **module** (*torch.nn.Module*) – A module.
- **batch_idx** (*int*) – The batch index.
- **batch** (*list of torch.Tensor*) – The list of input tensors of this batch.

Return type

None

Returns: None

execute(*code_block, batch*)

Method called by the CodeBlocks API to do device dependent execution.

Parameters

- **code_block** ([CodeBlock](#)) – The codeblock requesting execution.
- **batch** (*dict of str, torch.Tensor*) – The input tensors of this batch.

Returns

The results of executing the codeblock on this runtime.

Return type*Any***initialize_module**(*module, loader_or_batch, optimizer=None*)

Initializes the module at the beginning of training or inference.

Parameters

- **module** (*torch.nn.Module*) – A module.
- **loader_or_batch** ([DataLoader](#) or *torch.Tensor*) – A data loader or a tensor.
- **optimizer** (*Optimizer or None*) – An optimizer. This argument is sometimes used to copy LR from the original optimizer to the training model.

Return type

None

Returns: None

map(*func, iterable, out_keys=None, device='cpu'*)

Method called by the user to apply function to iterable efficiently.

Parameters

- **func** ([CodeBlock](#)) – The function to be executed
- **iterable** (*Iterable[Any]*) – The data
- **out_keys** (*Optional[Set[str]]*) – The output keys that to be moved to the host device

- **device** (*Any*) – The torch device that contains the final outputs

Returns

The result of *func*

Return type

Iterable[Any]

move_module(*module*)

Transfers the module to the specific device.

Before this method is called, `ppe.to` will add this class as a new attribute (“_ppe_runtime”) to the input module.

Parameters

module (*torch.nn.Module*) – A module.

Returns

A module transferred to the specific device.

Return type

Module

move_tensor(*tensor*)

Transfers the tensor to the specific device.

Parameters

tensor (*torch.Tensor*) – A tensor.

Returns

A tensor transferred to the specific device.

Return type

Tensor

classmethod trace(*event_name, arg*)

Context manager for tracing PPE events in the custom device tools.

Parameters

- **event_name** (*Optional[str]*) – The name of the event being traced
- **arg** (*Any*) – Custom argument for the tracer

Return type

Generator[None, None, None]

train_cleanup(*module*)

A method called only once when completing a training run.

Parameters

module (*torch.nn.Module*) – A module.

Return type

None

Returns: None

train_epoch_begin(*module*)

Preprocess of each epoch.

Parameters

module (*torch.nn.Module*) – A module.

Return type

None

Returns: None

train_epoch_end(*module*)

Completion of each epoch.

Parameters**module** (*torch.nn.Module*) – A module.**Return type**

None

Returns: None

train_post_step(*trainer, module, batch_idx, batch, outs*)

Postprocess of each step.

This method is called at the end of every steps: the set of (typically one) iterations and an update.

Parameters

- **trainer** (*Trainer*) – A trainer.
- **module** (*torch.nn.Module*) – A module.
- **batch_idx** (*int*) – The batch index.
- **batch** (*list of torch.Tensor*) – The list of input tensors of this batch.
- **outs** (*Any*) – (list of *torch.Tensor*): The list of output tensors of this batch.

Return type

None

Returns: None

train_pre_step(*trainer, module, batch_idx, batch*)

Preprocess of each step.

This method is called at the beginning of every steps: the set of (typically one) iterations and an update.

Parameters

- **trainer** (*Trainer*) – A trainer.
- **module** (*torch.nn.Module*) – A module.
- **batch_idx** (*int*) – The batch index.
- **batch** (*list of torch.Tensor*) – The list of input tensors of this batch.

Return type

None

Returns: None

train_validation_begin(*module*)

The method called before each evaluation.

Parameters**module** (*torch.nn.Module*) – A module.**Return type**

None

Returns: None

train_validation_end(*module*)

The method called after each evaluation.

Parameters

module (*torch.nn.Module*) – A module.

Return type

None

Returns: None

pytorch_pfn_extras.testing

pytorch_pfn_extras.torchscript

Functions

pytorch_pfn_extras.torchscript.

find_inplace(*g*)

pytorch_pfn_extras.torchscript.

run_jit_pass(*p*, ...)

pytorch_pfn_extras.torchscript.find_inplace

`pytorch_pfn_extras.torchscript.find_inplace(g)`

Parameters

g (*Graph*) –

Return type

Tuple[Graph, List[Node]]

pytorch_pfn_extras.torchscript.run_jit_pass

`pytorch_pfn_extras.torchscript.run_jit_pass(p, g, *args, **kwargs)`

Parameters

- **p** (*Callable*) –
- **g** (*Graph*) –
- **args** (*Any*) –
- **kwargs** (*Any*) –

Return type

None

pytorch_pfn_extras.training

Functions

| | |
|---|---|
| <code>pytorch_pfn_extras.training. make_extension(...)</code> | Decorator to make given function into an extension. |
|---|---|

pytorch_pfn_extras.training.make_extension

`pytorch_pfn_extras.training.make_extension(trigger=(1, 'iteration'), default_name=None, priority=100, finalizer=<function <lambda>>, initializer=<function <lambda>>, on_error=<function <lambda>>)`

Decorator to make given function into an extension.

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 extensions. Most of the default values of arguments also follow those for this class.

Parameters

- **trigger** (*TriggerLike*) – Default trigger of the extension.
- **default_name** (*Optional[str]*) – Default name of the extension. The name of a given function is used by default.
- **priority** (*int*) – Default priority of the extension.
- **finalizer** (*ExtensionLike*) – Finalizer function of this extension. It is called at the end of the training loop.
- **initializer** (*ExtensionLike*) – Initializer function of this extension. It is called at the beginning of the training loop.
- **on_error** (*Callable[[ExtensionsManagerProtocol, Exception, TracebackType], None]*) – Error handler callback function of this extension. It is called after an error is raised during the training loop.

Return type

Callable[[ExtensionLike], ExtensionLike]

Classes

| | |
|---|--|
| <code>pytorch_pfn_extras.training.AccuracyMetric(...)</code> | A metric for an evaluator to report accuracy. |
| <code>pytorch_pfn_extras.training.DistributedEvaluator(...)</code> | |
| <code>pytorch_pfn_extras.training.Evaluator(...[, ...])</code> | |
| <code>pytorch_pfn_extras.training.Extension()</code> | Base class of extensions. |
| <code>pytorch_pfn_extras.training.ExtensionEntry(...)</code> | Extension and options. |
| <code>pytorch_pfn_extras.training.ExtensionsManager(...)</code> | Manages the extensions and the current status. |
| <code>pytorch_pfn_extras.training.ExtensionsManagerProtocol(...)</code> | |
| <code>pytorch_pfn_extras.training.IgniteExtensionsManager(...)</code> | Manages extensions and the current status in Ignite training loop. |
| <code>pytorch_pfn_extras.training.StateObjectProtocol(...)</code> | |
| <code>pytorch_pfn_extras.training.Trainer(handler, ...)</code> | |

pytorch_pfn_extras.training.AccuracyMetric

class `pytorch_pfn_extras.training.AccuracyMetric(label_key, output_key)`

Bases: `object`

A metric for an evaluator to report accuracy.

Parameters

- **label_key** (*str*) – The key name of label.
- **output_key** (*str*) – The key name of prediction.

Methods

`__init__(label_key, output_key)`

`__call__(batch, out)`

Call self as a function.

Parameters

- **batch** (*Dict[str, Tensor]*) –
- **out** (*Dict[str, Tensor]*) –

Return type

Dict[str, Any]

`__init__(label_key, output_key)`

Parameters

- **label_key** (*str*) –
- **output_key** (*str*) –

Return type

None

pytorch_pfn_extras.training.DistributedEvaluator

class pytorch_pfn_extras.training.DistributedEvaluator(*handler, models, *, progress_bar=False, metrics=None, profile=None*)

Bases: [Evaluator](#)

Methods

`__init__(handler, models, *, progress_bar, ...)`

`run(loader, *, eval_len)`

Executes the evaluation loop.

Parameters

- **handler** ([BaseHandler](#)) –
- **models** (*Union[Module, Mapping[str, Module]]*) –
- **progress_bar** (*bool*) –
- **metrics** (*Optional[Sequence[MetricType]]*) –
- **profile** (*Optional[profile]*) –

`__init__(handler, models, *, progress_bar=False, metrics=None, profile=None)`

Parameters

- **handler** ([BaseHandler](#)) –
- **models** (*Union[Module, Mapping[str, Module]]*) –
- **progress_bar** (*bool*) –
- **metrics** (*Optional[Sequence[MetricType]]*) –
- **profile** (*Optional[profile]*) –

pytorch_pfn_extras.training.Evaluator

```
class pytorch_pfn_extras.training.Evaluator(handler, models, *, progress_bar=False, metrics=None,
                                           profile=None)
```

Bases: object

Methods

```
__init__(handler, models, *, progress_bar, ...)
```

```
run(loader, *, eval_len) Executes the evaluation loop.
```

Parameters

- **handler** (`BaseHandler`) –
- **models** (`Union[Module, Mapping[str, Module]]`) –
- **progress_bar** (`bool`) –
- **metrics** (`Optional[Sequence[MetricType]]`) –
- **profile** (`Optional[profile]`) –

```
__init__(handler, models, *, progress_bar=False, metrics=None, profile=None)
```

Parameters

- **handler** (`BaseHandler`) –
- **models** (`Union[Module, Mapping[str, Module]]`) –
- **progress_bar** (`bool`) –
- **metrics** (`Optional[Sequence[MetricType]]`) –
- **profile** (`Optional[profile]`) –

```
run(loader, *, eval_len=None)
```

Executes the evaluation loop.

Parameters

- **loader** (`torch.utils.data.DataLoader`) – A data loader for evaluation.
- **eval_len** (`int, optional`) – The number of iterations per one evaluation epoch.

Return type

None

pytorch_pfn_extras.training.Extension

class pytorch_pfn_extras.training.**Extension**

Bases: object

Base class of extensions.

An extension is a callable object that takes the manager 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 `ExtensionsManager.extend()`. For a callable not inheriting this class, the default configurations of this class are used unless the user explicitly specifies them in `ExtensionsManager.extend()` method.

trigger

Default value of trigger for this extension. It is set to (1, 'iteration') by default.

Type

TriggerLike

priority

Default priority of the extension. It is set to `PRIORITY_READER` by default.

Type

int

~Extension.name

Name of the extension. It is set to `None` by default. This value will be overwritten when registering an extension to a manager. See `pytorch_pfn_extras.ExtensionsManager.extend()` for details.

Methods

| | |
|---|--|
| <code>__init__()</code> | |
| <hr/> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <hr/> | |
| <code>load_state_dict(to_load)</code> | |
| <hr/> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before final- ization. |
| <hr/> | |
| <code>state_dict()</code> | Serializes the extension state. |
| <hr/> | |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

`__call__(manager)`

Invokes the extension.

Implementations should override this operator. This method is called at iterations which the corresponding trigger accepts.

Parameters

manager (`ExtensionsManager`) – Manager object to call this operator.

Return type

Any

`property default_name: str`

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.

`finalize(manager)`

Finalizes the extension.

This method is called at the end of the training loop.

Parameters

manager (`ExtensionsManagerProtocol`) –

Return type

None

`initialize(manager)`

Initializes up the manager state.

This method is called before entering the training loop. An extension modifying the state of `ExtensionsManager` can override this method to initialize it.

When the manager has been restored from a snapshot, this method has to recover an appropriate part of the state of the manager.

Parameters

manager (`ExtensionsManager`) – Manager object to call this extension.

Return type

None

is_async = False

load_state_dict(*to_load*)

Parameters

to_load (*Dict[str, Any]*) –

Return type

None

name: *Optional[str]* = None

needs_model_state = False

on_error(*manager, exc, tb*)

Handles the error raised during training before finalization.

This method is called when an exception is thrown during the training loop, before finalize. An extension that needs different error handling from finalize, can override this method to handle errors.

Parameters

- **manager** (*ExtensionsManager*) –
- **extension.** (*Manager object to call this*) –
- **exc** (*Exception*) – arbitrary exception thrown during update loop.
- **tb** (*traceback*) – traceback object of the exception

Return type

None

priority: *int* = 100

state_dict()

Serializes the extension state.

It is called when a manager that owns this extension is serialized. It serializes nothing by default.

Return type

Dict[str, Any]

trigger: *TriggerLike* = (1, 'iteration')

pytorch_pfn_extras.training.ExtensionEntry

class pytorch_pfn_extras.training.**ExtensionEntry**(*extension, *, name=None, priority=None, trigger=None, call_before_training=False*)

Bases: object

Extension and options. When name, priority, or trigger is not specified, it is copied from the attributes of the given extension.

Parameters

- **extension** (*ExtensionLike*) – An extension.
- **name** (*Optional[str]*) – Name of extension.
- **priority** (*Optional[int]*) – Invocation priority of the extension.

- **trigger** (*Optional[TriggerLike]*) – Trigger object that determines when to invoke the extension.
- **call_before_training** (*bool*) – Flag to call extension before training.

See also:

`pytorch_pfn_extras.training.ExtensionsManager.extend()`

Methods

`__init__`(extension, *, name, priority, ...)

`load_state_dict`(to_load)

`state_dict`()

`__init__`(extension, *, name=None, priority=None, trigger=None, call_before_training=False)

Parameters

- **extension** (*ExtensionLike*) –
- **name** (*Optional[str]*) –
- **priority** (*Optional[int]*) –
- **trigger** (*Optional[TriggerLike]*) –
- **call_before_training** (*bool*) –

Return type

None

`load_state_dict`(to_load)

Parameters

to_load (*Dict[str, Any]*) –

Return type

None

`state_dict`()

Return type

Dict[str, Any]

pytorch_pfn_extras.training.ExtensionsManager

```
class pytorch_pfn_extras.training.ExtensionsManager(models, optimizers, max_epochs, *,
                                                    iters_per_epoch, extensions=None,
                                                    out_dir='result', stop_trigger=None,
                                                    writer=None, transform_model=<function
                                                    ExtensionsManager.<lambda>>,
                                                    enable_profile=False, enable_trace=False,
                                                    state_objects={})
```

Bases: `_BaseExtensionsManager`

Manages the extensions and the current status.

Parameters

- **models** (dict or *torch.nn.Module*) – Map of string to Module or an actual Module
- **optimizers** (dict or *torch.Optimizer*) – Map of string to Optimizer or an actual Optimizer.
- **max_epochs** (*int*) – Number of epochs in the whole training loop. Ignored if *stop_trigger* is passed as a kwarg.
- **iters_per_epoch** (*int*) – Number of iterations in one epoch.
- **extensions** (*list* or *None*) – List of Extensions to be used.
- **out_dir** (*str*) – Output directory (default: `result`).
- **stop_trigger** (*trigger object*, *optional*) – to determine whether training has concluded. The default is an interval trigger set to *max_epochs*
- **writer** (*writing.Writer object*) – Writer that can be used by extensions to write data to custom filesystems.
- **enable_profile** (*bool*) – Flag to enable/disable profiling of iterations. Default is *False*.
- **enable_trace** (*bool*) – Flag to enable/disable tracing of iterations. Default is *False*.
- **transform_model** (*Callable[[str, Module], Module]*) –
- **state_objects** (*Dict[str, StateObjectProtocol]*) –

Methods

| | |
|--|--|
| <code>__init__(models, optimizers, max_epochs, *, ...)</code> | |
| <code>extend(extension[, name, trigger, priority, ...])</code> | Registers an extension to the manager. |
| <code>finalize()</code> | |
| <code>get_extension(name)</code> | Returns the extension of a given name. |
| <code>load_state_dict(to_load)</code> | |
| <code>needs_model_state([iteration])</code> | |
| <code>needs_state_this_iteration()</code> | |
| <code>run_extensions()</code> | |
| <code>run_iteration(*[, step_optimizers])</code> | Context manager to run an iteration. |
| <code>start_extensions()</code> | |
| <code>state_dict()</code> | |

Attributes

| |
|--------------------|
| elapsed_time |
| epoch |
| epoch_detail |
| is_before_training |
| iteration |
| models |
| optimizers |
| out |
| raw_models |
| stop_trigger |
| updater |

__init__ (*models*, *optimizers*, *max_epochs*, *, *iters_per_epoch*, *extensions=None*, *out_dir='result'*, *stop_trigger=None*, *writer=None*, *transform_model=<function ExtensionsManager.<lambda>>*, *enable_profile=False*, *enable_trace=False*, *state_objects={}*)

Parameters

- **models** (*Union[Module, Dict[str, Module]]*) –
- **optimizers** (*Union[Optimizer, Dict[str, Optimizer]]*) –
- **max_epochs** (*int*) –
- **iters_per_epoch** (*int*) –
- **extensions** (*Optional[Sequence[extension_module.ExtensionLike]]*) –
- **out_dir** (*str*) –
- **stop_trigger** (*trigger_module.TriggerLike*) –
- **writer** (*Optional[Writer]*) –
- **transform_model** (*Callable[[str, Module], Module]*) –
- **enable_profile** (*bool*) –
- **enable_trace** (*bool*) –
- **state_objects** (*Dict[str, StateObjectProtocol]*) –

Return type

None

finalize()

Return type

None

run_iteration(***, *step_optimizers=None*)

Context manager to run an iteration.

This manager can additionally run a step in the specified optimizers names.

Parameters

- **step_optimizers** (*list* or *None*) – names of the optimizers
- **step** (to call *zero_grad* and) –

Return type

Generator[None, None, None]

pytorch_pfn_extras.training.ExtensionsManagerProtocol

class pytorch_pfn_extras.training.**ExtensionsManagerProtocol**(**args*, ***kwargs*)

Bases: Protocol

Methods

__init__(**args*, ***kwargs*)

get_extension(name)

Attributes

elapsed_time

epoch

epoch_detail

is_before_training

iteration

models

observation

optimizers

out

raw_models

reporter

stop_trigger

writer

`__init__`(**args*, ***kwargs*)**property** `elapsed_time`: float**property** `epoch`: int**property** `epoch_detail`: float**get_extension**(*name*)**Parameters****name** (*str*) –**Return type***Extension***property** `is_before_training`: bool**property** `iteration`: int**property** `models`: Mapping[str, Module]**property** `observation`: reporting.Observation**property** `optimizers`: Mapping[str, Optimizer]

```
property out: str
property raw_models: Mapping[str, Module]
property reporter: reporting.Reporter
property stop_trigger: bool
property writer: Optional[writing.Writer]
```

`pytorch_pfn_extras.training.IgniteExtensionsManager`

```
class pytorch_pfn_extras.training.IgniteExtensionsManager(engine, models, optimizers, max_epochs,
                                                         *, extensions=None, out_dir='result',
                                                         writer=None, enable_profile=False,
                                                         enable_trace=False, state_objects={})
```

Bases: `_BaseExtensionsManager`

Manages extensions and the current status in Ignite training loop.

Parameters

- **engine** (*ignite.engine.Engine*) – Ignite trainer engine
- **models** (*dict or torch.nn.Module*) – Map of string to Module or an actual Module
- **optimizers** (*dict or torch.Optimizer*) – Map of string to Optimizer or an actual Optimizer.
- **max_epochs** (*int*) – Number of epochs in the whole training loop.
- **extensions** (*list or None*) – List of Extensions to be used.
- **out_dir** (*str*) – Output directory (default: `result`).
- **writer** (*writing.Writer object*) – Writer that can be used by extensions to write data to custom filesystems.
- **enable_profile** (*bool*) – Flag to enable/disable profiling of iterations. Default is *False*.
- **enable_trace** (*bool*) – Flag to enable/disable tracing of iterations. Default is *False*.
- **state_objects** (*Dict[str, StateObjectProtocol]*) –

Methods

| | |
|--|--|
| <code>__init__(engine, models, optimizers, ...[, ...])</code> | |
| <code>extend(extension[, name, trigger, priority, ...])</code> | Registers an extension to the manager. |
| <code>get_extension(name)</code> | Returns the extension of a given name. |
| <code>load_state_dict(to_load)</code> | |
| <code>needs_model_state([iteration])</code> | |
| <code>needs_state_this_iteration()</code> | |
| <code>run_extensions()</code> | |
| <code>set_ignite_handlers()</code> | |
| <code>start_extensions()</code> | |
| <code>state_dict()</code> | |

Attributes

| |
|---------------------------------|
| <code>elapsed_time</code> |
| <code>epoch</code> |
| <code>epoch_detail</code> |
| <code>is_before_training</code> |
| <code>iteration</code> |
| <code>models</code> |
| <code>optimizers</code> |
| <code>out</code> |
| <code>raw_models</code> |
| <code>stop_trigger</code> |
| <code>updater</code> |

`__init__(engine, models, optimizers, max_epochs, *, extensions=None, out_dir='result', writer=None, enable_profile=False, enable_trace=False, state_objects={})`

Parameters

- **engine** (*ignite.engine.Engine*) –
- **models** (*Union[Module, Mapping[str, Module]]*) –
- **optimizers** (*Union[Optimizer, Mapping[str, Optimizer]]*) –
- **max_epochs** (*int*) –
- **extensions** (*Optional[Sequence[extension_module.ExtensionLike]]*) –
- **out_dir** (*str*) –
- **writer** (*Optional[Writer]*) –
- **enable_profile** (*bool*) –
- **enable_trace** (*bool*) –
- **state_objects** (*Dict[str, StateObjectProtocol]*) –

Return type

None

load_state_dict(*to_load*)**Parameters****to_load** (*Dict[str, Any]*) –**Return type**

None

set_ignite_handlers()**Return type**

None

state_dict()**Return type***Dict[str, Any]***pytorch_pfn_extras.training.StateObjectProtocol****class** pytorch_pfn_extras.training.StateObjectProtocol(*args, **kwargs)

Bases: Protocol

Methods

__init__(*args, **kwargs)

load_state_dict(state_dict)

state_dict()

__init__(*args, **kwargs)

load_state_dict(*state_dict*)

Parameters

state_dict (*Dict[str, Any]*) –

Return type

None

state_dict()

Return type

Dict[str, Any]

pytorch_pfn_extras.training.Trainer

class pytorch_pfn_extras.training.**Trainer**(*handler, *, evaluator, models, profile=None, **kwargs*)

Bases: object

Methods

__init__(*handler, *, evaluator, models[, ...]*)

extend(*extension[, name, trigger, priority, ...]*)

get_optimizer(*name*)

is_epoch_last_iter(*idx*)

load_state_dict(*to_load*)

run(*train_loader[, val_loader, train_len, ...]*) Executes the training loop.

set_optimizer(*name, optimizer*)

state_dict()

Attributes

epoch

epoch_detail

evaluator

is_before_training

iteration

manager

models

optimizers

stop_trigger

Parameters

- **handler** (*handler_module.BaseHandler*) –
- **evaluator** (*Optional[Union[Evaluator, Tuple[Evaluator, Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]], Mapping[str, Union[Evaluator, Tuple[Evaluator, Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]]]]]*) –
- **models** (*Union[Module, Mapping[str, Module]]*) –
- **profile** (*Optional[profile]*) –
- **kwargs** (*Any*) –

__init__ (*handler, *, evaluator, models, profile=None, **kwargs*)

Parameters

- **handler** (*handler_module.BaseHandler*) –
- **evaluator** (*Optional[Union[Evaluator, Tuple[Evaluator, Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]], Mapping[str, Union[Evaluator, Tuple[Evaluator, Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]]]]]*) –
- **models** (*Union[Module, Mapping[str, Module]]*) –
- **profile** (*Optional[profile]*) –
- **kwargs** (*Any*) –

property epoch: `int`

property epoch_detail: float

property evaluator: Optional[[Evaluator](#)]

extend(*extension*, *name=None*, *trigger=None*, *priority=None*, *, *call_before_training=False*, **kwargs)

Parameters

- **extension** (Union[*extension.ExtensionLike*, [ExtensionEntry](#)]) –
- **name** (Optional[str]) –
- **trigger** (TriggerLike) –
- **priority** (Optional[int]) –
- **call_before_training** (bool) –
- **kwargs** (Any) –

Return type

None

get_optimizer(*name*)

Parameters

name (str) –

Return type

Optimizer

property is_before_training: bool

is_epoch_last_iter(*idx*)

Parameters

idx (int) –

Return type

bool

property iteration: int

load_state_dict(*to_load*)

Parameters

to_load (Dict[str, Any]) –

Return type

None

property manager: ExtensionsManager

property models: Mapping[str, Module]

property optimizers: Mapping[str, Optimizer]

run(*train_loader*, *val_loader=None*, *, *train_len=None*, *eval_len=None*)

Executes the training loop.

Parameters

- **train_loader** (*torch.utils.data.DataLoader*) – A data loader for training.

- **val_loader** (*torch.utils.data.DataLoader*, *optional*) – A data loader passed to `Evaluator.run()`.
- **train_len** (*int*, *optional*) – The number of iterations per one training epoch. The default value is inferred from the size of training data loader.
- **eval_len** (*int*, *optional*) – The number of iterations per one evaluation epoch, passed to `Evaluator.run()`

Return type

None

See also:

- `pytorch_pfn_extras.training._evaluator.Evaluator()`

set_optimizer(*name*, *optimizer*)**Parameters**

- **name** (*str*) –
- **optimizer** (*Optimizer*) –

Return type

None

state_dict()**Return type***Dict[str, Any]***property stop_trigger:** *Trigger***Modules**

`pytorch_pfn_extras.training.extension`

`pytorch_pfn_extras.training.extensions`

`pytorch_pfn_extras.training.manager`

`pytorch_pfn_extras.training.metrics`

`pytorch_pfn_extras.training.trigger`

`pytorch_pfn_extras.training.triggers`

pytorch_pfn_extras.training.extension**Functions**

| | |
|---|---|
| <code>pytorch_pfn_extras.training.extension. make_extension(...)</code> | Decorator to make given function into an extension. |
|---|---|

pytorch_pfn_extras.training.extension.make_extension

`pytorch_pfn_extras.training.extension.make_extension(trigger=(1, 'iteration'), default_name=None, priority=100, finalizer=<function <lambda>>, initializer=<function <lambda>>, on_error=<function <lambda>>)`

Decorator to make given function into an extension.

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 extensions. Most of the default values of arguments also follow those for this class.

Parameters

- **trigger** (*TriggerLike*) – Default trigger of the extension.
- **default_name** (*Optional[str]*) – Default name of the extension. The name of a given function is used by default.
- **priority** (*int*) – Default priority of the extension.
- **finalizer** (*ExtensionLike*) – Finalizer function of this extension. It is called at the end of the training loop.
- **initializer** (*ExtensionLike*) – Initializer function of this extension. It is called at the beginning of the training loop.
- **on_error** (*Callable[[ExtensionsManagerProtocol, Exception, TracebackType], None]*) – Error handler callback function of this extension. It is called after an error is raised during the training loop.

Return type

Callable[[ExtensionLike], ExtensionLike]

Classes

| | |
|--|---------------------------|
| <code>pytorch_pfn_extras.training.extension. Extension()</code> | Base class of extensions. |
| <code>pytorch_pfn_extras.training.extension. ExtensionEntry(...)</code> | Extension and options. |
| <code>pytorch_pfn_extras.training.extension. ExtensionsManagerProtocol(...)</code> | |

pytorch_pfn_extras.training.extension.Extension**class** pytorch_pfn_extras.training.extension.**Extension**

Bases: object

Base class of extensions.

An extension is a callable object that takes the manager 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 `ExtensionsManager.extend()`. For a callable not inheriting this class, the default configurations of this class are used unless the user explicitly specifies them in `ExtensionsManager.extend()` method.

trigger

Default value of trigger for this extension. It is set to (1, 'iteration') by default.

Type

TriggerLike

priority

Default priority of the extension. It is set to `PRIORITY_READER` by default.

Type

int

~Extension.name

Name of the extension. It is set to `None` by default. This value will be overwritten when registering an extension to a manager. See `pytorch_pfn_extras.ExtensionsManager.extend()` for details.

Methods

| | |
|---|--|
| <code>__init__()</code> | |
| <hr/> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <hr/> | |
| <code>load_state_dict(to_load)</code> | |
| <hr/> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before final- ization. |
| <hr/> | |
| <code>state_dict()</code> | Serializes the extension state. |
| <hr/> | |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

`__call__(manager)`

Invokes the extension.

Implementations should override this operator. This method is called at iterations which the corresponding trigger accepts.

Parameters

manager (`ExtensionsManager`) – Manager object to call this operator.

Return type

Any

`property default_name: str`

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.

`finalize(manager)`

Finalizes the extension.

This method is called at the end of the training loop.

Parameters

manager (`ExtensionsManagerProtocol`) –

Return type

None

`initialize(manager)`

Initializes up the manager state.

This method is called before entering the training loop. An extension modifying the state of `ExtensionsManager` can override this method to initialize it.

When the manager has been restored from a snapshot, this method has to recover an appropriate part of the state of the manager.

Parameters

manager (`ExtensionsManager`) – Manager object to call this extension.

Return type

None

is_async = False

load_state_dict(*to_load*)

Parameters

to_load (*Dict[str, Any]*) –

Return type

None

name: *Optional[str]* = None

needs_model_state = False

on_error(*manager, exc, tb*)

Handles the error raised during training before finalization.

This method is called when an exception is thrown during the training loop, before finalize. An extension that needs different error handling from finalize, can override this method to handle errors.

Parameters

- **manager** (*ExtensionsManager*) –
- **extension.** (*Manager object to call this*) –
- **exc** (*Exception*) – arbitrary exception thrown during update loop.
- **tb** (*traceback*) – traceback object of the exception

Return type

None

priority: *int* = 100

state_dict()

Serializes the extension state.

It is called when a manager that owns this extension is serialized. It serializes nothing by default.

Return type

Dict[str, Any]

trigger: *TriggerLike* = (1, 'iteration')

pytorch_pfn_extras.training.extension.ExtensionEntry

```
class pytorch_pfn_extras.training.extension.ExtensionEntry(extension, *, name=None,  
                                                         priority=None, trigger=None,  
                                                         call_before_training=False)
```

Bases: object

Extension and options. When name, priority, or trigger is not specified, it is copied from the attributes of the given extension.

Parameters

- **extension** (*ExtensionLike*) – An extension.
- **name** (*Optional[str]*) – Name of extension.
- **priority** (*Optional[int]*) – Invocation priority of the extension.

- **trigger** (*Optional[TriggerLike]*) – Trigger object that determines when to invoke the extension.
- **call_before_training** (*bool*) – Flag to call extension before training.

See also:

`pytorch_pfn_extras.training.ExtensionsManager.extend()`

Methods

`__init__`(extension, *, name, priority, ...)

`load_state_dict`(to_load)

`state_dict`()

`__init__`(extension, *, name=None, priority=None, trigger=None, call_before_training=False)

Parameters

- **extension** (*ExtensionLike*) –
- **name** (*Optional[str]*) –
- **priority** (*Optional[int]*) –
- **trigger** (*Optional[TriggerLike]*) –
- **call_before_training** (*bool*) –

Return type

None

`load_state_dict`(to_load)

Parameters

to_load (*Dict[str, Any]*) –

Return type

None

`state_dict`()

Return type

Dict[str, Any]

`pytorch_pfn_extras.training.extension.ExtensionsManagerProtocol`

`class pytorch_pfn_extras.training.extension.ExtensionsManagerProtocol(*args, **kwargs)`

Bases: `Protocol`

Methods

`__init__(*args, **kwargs)`

`get_extension(name)`

Attributes

`elapsed_time`

`epoch`

`epoch_detail`

`is_before_training`

`iteration`

`models`

`observation`

`optimizers`

`out`

`raw_models`

`reporter`

`stop_trigger`

`writer`

`__init__(*args, **kwargs)``property elapsed_time: float``property epoch: int``property epoch_detail: float``get_extension(name)`

Parameters

name (*str*) –

Return type

Extension

```

property is_before_training: bool
property iteration: int
property models: Mapping[str, Module]
property observation: reporting.Observation
property optimizers: Mapping[str, Optimizer]
property out: str
property raw_models: Mapping[str, Module]
property reporter: reporting.Reporter
property stop_trigger: bool
property writer: Optional[writing.Writer]

```

pytorch_pfn_extras.training.extensions

Functions

| | |
|--|---|
| <code>pytorch_pfn_extras.training.extensions.observe_lr(...)</code> | Returns an extension to record the learning rate. |
| <code>pytorch_pfn_extras.training.extensions.observe_value(...)</code> | Returns an extension to continuously record a value. |
| <code>pytorch_pfn_extras.training.extensions.snapshot([...])</code> | Returns a trainer extension to take snapshots of the trainer. |
| <code>pytorch_pfn_extras.training.extensions.snapshot_object(...)</code> | Returns an extension to take snapshots of a given object. |

pytorch_pfn_extras.training.extensions.observe_lr

`pytorch_pfn_extras.training.extensions.observe_lr(optimizer, param_group=0, observation_key='lr')`

Returns an extension to record the learning rate.

Parameters

- **optimizer** (*Optimizer*) – Optimizer whose learning rate is recorded.
- **param_group** (*int*) – Param group of the optimizer to observe
- **observation_key** (*str*) – Key of observation to record.

Returns

The extension function.

Return type

Any

This extension is triggered each epoch by default. To change this, use the `trigger` argument with the `ExtensionsManager.extend()` method.

pytorch_pfn_extras.training.extensions.observe_value

`pytorch_pfn_extras.training.extensions.observe_value(observation_key, target_func)`

Returns an 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:`~pytorch_pfn_extras.training.ExtensionsManager` object.

Returns

The extension function.

Return type

Callable[[*ExtensionsManagerProtocol*], None]

This extension is triggered each epoch by default. To change this, use the `trigger` argument with the `ExtensionsManager.extend()` method.

pytorch_pfn_extras.training.extensions.snapshot

`pytorch_pfn_extras.training.extensions.snapshot(savefun=None, filename='snapshot_iter_{.iteration}', *, target=None, condition=None, writer=None, snapshot_on_error=False, n_retains=-1, autoload=False, saver_rank=None, snapshot_mode=SnapshotMode.DEFAULT)`

Returns a trainer extension to take snapshots of the trainer.

This extension serializes the manager 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 manager.

The default priority is -100, which is lower than that of most built-in extensions.

Parameters

- **savefun** (*Optional*[*Any*]) – Function to save the manager. It takes two arguments: the output file path and the manager object. It is `torch.save()` by default. If `writer` is specified, this argument must be `None`.
- **filename** (*str*) – Name of the file into which the manager is serialized. It can be a format string, where the manager object is passed to the `str.format()` method.
- **target** (*Optional*[*Any*]) – Object to serialize. If it is not specified, it will be the manager object.
- **condition** (*Optional*[*Any*]) – Condition object. It must be a callable object that returns boolean without any arguments. If it returns `True`, the snapshot will be done. If not, it will be skipped. The default is a function that always returns `True`.
- **writer** (*Optional*[*Writer*]) – Writer object. It must be a callable object. See below for the list of built-in writers. If `savefun` is other than `None`, this argument must be `None`. In that case, a [SimpleWriter](#) object instantiated with specified `savefun` argument will be used.

- **snapshot_on_error** (*bool*) – Whether to take a snapshot in case training loop has been failed.
- **n_retains** (*int*) – Number of snapshot files to retain through the cleanup. Must be a positive integer for any cleanup to take place. Automatic deletion of old snapshots only works when the filename is string.
- **autoload** (*bool*) – With this enabled, the extension automatically finds the latest snapshot and loads the data to the target. Automatic loading only works when the filename is a string. It is assumed that snapshots are generated by `torch.save()`.
- **saver_rank** (*int*) – If defined, the snapshot will be taken by only one rank when running in distributed mode and restored by all.
- **snapshot_mode** (`SnapshotMode`) – If `SnapshotModel.DEFAULT` is specified, it provides a snapshot feature that operates in single-process mode. However, if `saver_rank` is specified, it provides a snapshot feature that operates in a distributed execution environment.

If `SnapshotModel.DISTRIBUTED` is specified, it provides a snapshot feature that operates in a distributed execution environment. `saver_rank` must be specified simultaneously. In this mode, only the specified `saver_rank` will create a snapshot.

If `SnapshotModel.SHARDED` is specified, it provides a snapshot feature that operates in a distributed execution environment. `saver_rank` must be specified simultaneously. In this mode, all ranks create a snapshot. It creates an appropriate snapshot when the `state_dict` holds a sharded value (e.g. `FullyShardedDataParallel`).

Returns

Snapshot extension object.

Return type

`_Snapshot`

Using asynchronous writers

By specifying `writer` argument, writing operations can be made asynchronous, hiding I/O overhead of snapshots.

```
>>> from pytorch_pfn_extras.training import extensions
>>> from pytorch_pfn_extras import writing
>>> writer = writing.ProcessWriter()
>>> manager.extend(extensions.snapshot(writer=writer), trigger=(1, 'epoch'))
```

To change the format, you can pass a saving function as `savefun` argument of the writer.

```
>>> from pytorch_pfn_extras.training import extensions
>>> from pytorch_pfn_extras import writing
>>> writer = writing.ProcessWriter(
...     savefun=torch.save)
>>> manager.extend(extensions.snapshot(writer=writer), trigger=(1, 'epoch'))
```

This is the list of built-in snapshot writers.

- `pytorch_pfn_extras.writing.SimpleWriter`
- `pytorch_pfn_extras.writing.ThreadWriter`

- `pytorch_pfn_extras.writing.ProcessWriter`
- `pytorch_pfn_extras.writing.ThreadQueueWriter`
- `pytorch_pfn_extras.writing.ProcessQueueWriter`

See also:

- `pytorch_pfn_extras.training.extensions.snapshot_object()`

`pytorch_pfn_extras.training.extensions.snapshot_object`

`pytorch_pfn_extras.training.extensions.snapshot_object(target, filename, savefun=None, *, condition=None, writer=None, snapshot_on_error=False, n_retains=-1, autoload=False)`

Returns an 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 manager.

The default priority is lower than that of most built-in extensions.

Parameters

- **target** (*Any*) – Object to serialize.
- **filename** (*str*) – Name of the file into which the object is serialized. It can be a format string, where the manager object is passed to the `str.format()` method. For example, `'snapshot_{.iteration}'` is converted to `'snapshot_10000'` at the 10,000th iteration.
- **savefun** (*Optional[Any]*) – Function to save the object. It takes two arguments: the output file path and the object to serialize.
- **condition** – Condition object. It must be a callable object that returns boolean without any arguments. If it returns `True`, the snapshot will be done. If not, it will be skipped. The default is a function that always returns `True`.
- **writer** – Writer object. It must be a callable object. See below for the list of built-in writers. If `savefun` is other than `None`, this argument must be `None`. In that case, a [SimpleWriter](#) object instantiated with specified `savefun` argument will be used.
- **snapshot_on_error** (*bool*) – Whether to take a snapshot in case training loop has failed.
- **n_retains** (*int*) – Number of snapshot files to retain through the cleanup. Must be a positive integer for any cleanup to take place. Automatic deletion of old snapshots only works when the filename is string.
- **autoload** (*bool*) – With this enabled, the extension automatically finds the latest snapshot and loads the data to the target. Automatic loading only works when the filename is a string.
- **saver_rank** (*int*) – If defined, the snapshot will be taken by only one rank when running in distributed mode and restored by all.
- **kwargs** (*Any*) –

Returns

Snapshot extension object.

Return type

_Snapshot

See also:

- `pytorch_pfn_extras.training.extensions.snapshot()`

Classes

| | |
|--|---|
| <code>pytorch_pfn_extras.training.extensions.AverageAccumulate(...)</code> | |
| <code>pytorch_pfn_extras.training.extensions.BestValue(...)</code> | Extension traces the best value of a specific key in the observation. |
| <code>pytorch_pfn_extras.training.extensions.DistributedEvaluator(...)</code> | An extension to evaluate models on a validation set in a distributed training setup. |
| <code>pytorch_pfn_extras.training.extensions.Evaluator(...)</code> | An extension to evaluate models on a validation set. |
| <code>pytorch_pfn_extras.training.extensions.FailOnNonNumber(*)</code> | An extension to raise <code>RuntimeError</code> if parameters and its gradients contain NaN or Inf. |
| <code>pytorch_pfn_extras.training.extensions.IgniteEvaluator(...)</code> | |
| <code>pytorch_pfn_extras.training.extensions.LRScheduler(...)</code> | Trainer extension to adjust the learning rate using PyTorch's learning rate scheduler. |
| <code>pytorch_pfn_extras.training.extensions.LogReport([...])</code> | An extension to output the accumulated results to a log file. |
| <code>pytorch_pfn_extras.training.extensions.MaxAccumulate(...)</code> | |
| <code>pytorch_pfn_extras.training.extensions.MaxValue(key)</code> | Extension traces the maximum value of a specific key in the observation. |
| <code>pytorch_pfn_extras.training.extensions.MicroAverage(...)</code> | Calculates micro-average ratio. |
| <code>pytorch_pfn_extras.training.extensions.MinAccumulate(...)</code> | |
| <code>pytorch_pfn_extras.training.extensions.MinValue(key)</code> | Extension traces the maximum value of a specific key in the observation. |
| <code>pytorch_pfn_extras.training.extensions.ParameterStatistics(links)</code> | An extension to report parameter statistics. |
| <code>pytorch_pfn_extras.training.extensions.PlotReport(y_keys)</code> | An extension to output plots. |
| <code>pytorch_pfn_extras.training.extensions.PrintReport([...])</code> | An extension to print the accumulated results. |
| <code>pytorch_pfn_extras.training.extensions.PrintReportCLI</code> | alias of <code>PrintReport</code> |
| <code>pytorch_pfn_extras.training.extensions.ProfileReport([...])</code> | Writes the profile results to a file. |
| <code>pytorch_pfn_extras.training.extensions.ProgressBar([...])</code> | An extension to print a progress bar and recent training status. |
| <code>pytorch_pfn_extras.training.extensions.ProgressBarCLI</code> | alias of <code>ProgressBar</code> |
| <code>pytorch_pfn_extras.training.extensions.Slack(channel)</code> | An extension to communicate with Slack. |
| <code>pytorch_pfn_extras.training.extensions.SlackWebhook(url)</code> | An extension to communicate with Slack using Incoming Webhook. |
| <code>pytorch_pfn_extras.training.extensions.SnapshotMode(value)</code> | An enumeration. |
| <code>pytorch_pfn_extras.training.extensions.StandardDeviationAccumulate(...)</code> | |
| <code>pytorch_pfn_extras.training.extensions.TimelineTrace([...])</code> | Writes the profile timeline to a file. |
| <code>pytorch_pfn_extras.training.extensions.UnbiasedStandardDeviationAccumulate(...)</code> | |
| <code>pytorch_pfn_extras.training.extensions.VariableStatisticsPlot(targets)</code> | An extension to plot statistics for TensorBoard. |

pytorch_pfn_extras.training.extensions.AverageAccumulate

class pytorch_pfn_extras.training.extensions.**AverageAccumulate**(*conversion_key_pair*, *trigger*=(1, 'epoch'), *distributed*=False)

Bases: *AccumulateBase*

Methods

| | |
|--|---|
| <code>__init__(conversion_key_pair[, trigger, ...])</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

Parameters

- **conversion_key_pair** (*Tuple*[*str*, *str*]) –
- **trigger** (*Optional*[*Union*[*Trigger*, *Callable*[[*ExtensionsManagerProtocol*], *bool*], *Tuple*[*float*, *str*]]]) –
- **distributed** (*bool*) –

pytorch_pfn_extras.training.extensions.BestValue**class** pytorch_pfn_extras.training.extensions.**BestValue**(*key*, *compare*, *trigger*=(1, 'epoch'))Bases: *Extension*

Extension traces the best value of a specific key in the observation.

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** (*TriggerLike*) – 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 *BestValueTrigger*.

Methods

| | |
|---|---|
| <i>__init__</i> (key, compare[, trigger]) | |
| <i>finalize</i> (manager) | Finalizes the extension. |
| <i>initialize</i> (manager) | Initializes up the manager state. |
| <i>load_state_dict</i> (to_load) | |
| <i>on_error</i> (manager, exc, tb) | Handles the error raised during training before finalization. |
| <i>state_dict</i> () | Serializes the extension state. |

Attributes

| | |
|--------------------------|--|
| <i>best_epoch</i> | Returns the epoch count that the current best value is observed. |
| <i>best_iteration</i> | Returns the iteration count that the current best value is observed. |
| <i>best_value</i> | Returns the current best value. |
| <i>default_name</i> | |
| <i>is_async</i> | |
| <i>name</i> | |
| <i>needs_model_state</i> | |
| <i>priority</i> | |
| <i>trigger</i> | |

__call__(*manager*)

Invokes the extension.

Implementations should override this operator. This method is called at iterations which the corresponding trigger accepts.

Parameters

manager (*ExtensionsManager*) – Manager object to call this operator.

Return type

None

__init__(*key, compare, trigger=(1, 'epoch')*)

Parameters

- **key** (*str*) –
- **compare** (*Callable[[float, float], bool]*) –
- **trigger** (*TriggerLike*) –

Return type

None

property best_epoch: int

Returns the epoch count that the current best value is observed.

If no value has been observed yet, it raises a RuntimeError.

property best_iteration: int

Returns the iteration count that the current best value is observed.

If no value has been observed yet, it raises a RuntimeError.

property best_value: float

Returns the current best value.

If no value has been observed yet, it raises a RuntimeError.

default_name = 'best_value'

load_state_dict(*to_load*)

Parameters

to_load (*Dict[str, Any]*) –

Return type

None

state_dict()

Serializes the extension state.

It is called when a manager that owns this extension is serialized. It serializes nothing by default.

Return type

Dict[str, Any]

pytorch_pfn_extras.training.extensions.DistributedEvaluator

```
class pytorch_pfn_extras.training.extensions.DistributedEvaluator(self, iterator, target,
                                                                eval_func=None, *,
                                                                progress_bar=False)
```

Bases: [Evaluator](#)

An extension to evaluate models on a validation set in a distributed training setup.

In case torch.distributed is used to parallelize training iterations, it is efficient to also run evaluation in parallel by splitting the validation set to each worker process and conduct evaluation separately followed by aggregation of results of each worker, which can be achieved by :class:`~DistributedEvaluator`.

This extension basically behaves similarly to [Evaluator](#), but adds an aggregation step in [Evaluator.evaluate\(\)](#). A summary of evaluation (DictSummary) in each worker process is collected in “all-gather” manner and then accumulated. Therefore all the worker processes must attend the evaluation, i.e., make sure all the processes have a [Evaluator](#) extension object configured in the [ExtensionManager](#) with the same trigger. All the worker process will get identical evaluation result returned by [Evaluator.evaluate\(\)](#) and reported to an observation.

It is necessary to pass a DataLoader with an appropriate sampler which properly splits the validation dataset to each MPI worker process. PyTorch DistributedSampler implements this, but it allows sampler repetition in order to make the number of samples assigned to each process identical. For evaluation purpose it distorts the evaluation result, hence it is recommended to use [DistributedValidationSampler](#) instead.

Methods

| | |
|---|---|
| <code>__init__(iterator, target[, eval_hook, ...])</code> | |
| <code>add_metric(metric_fn)</code> | Adds a custom metric to the evaluator. |
| <code>eval_func(*args, **kwargs)</code> | |
| <code>evaluate()</code> | Evaluates the model and returns a result dictionary. |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>get_all_iterators()</code> | Returns a dictionary of all iterators. |
| <code>get_all_targets()</code> | Returns a dictionary of all target links. |
| <code>get_iterator(name)</code> | Returns the iterator of the given name. |
| <code>get_target(name)</code> | Returns the target link of the given name. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| |
|-------------------|
| default_name |
| is_async |
| name |
| needs_model_state |
| priority |
| trigger |

Parameters

- **iterator** (*Union*[*DataLoader*[Any], *Dict*[str, *DataLoader*[Any]]]) –
- **target** (*Union*[*Module*, *Dict*[str, *Module*]]) –
- **eval_hook** (*Optional*[*Callable*[[*Evaluator*], *None*]]) –
- **eval_func** (*Optional*[*Callable*[[...], Any]]) –
- **kwargs** (Any) –

__init__(*iterator*, *target*, *eval_hook*=None, *eval_func*=None, **kwargs)

Parameters

- **iterator** (*Union*[*DataLoader*[Any], *Dict*[str, *DataLoader*[Any]]]) –
- **target** (*Union*[*Module*, *Dict*[str, *Module*]]) –
- **eval_hook** (*Optional*[*Callable*[[*Evaluator*], *None*]]) –
- **eval_func** (*Optional*[*Callable*[[...], Any]]) –
- **kwargs** (Any) –

Return type

None

pytorch_pfn_extras.training.extensions.Evaluator

class pytorch_pfn_extras.training.extensions.**Evaluator**(*self*, *iterator*, *target*, *eval_func*=None, *, *progress_bar*=False)

Bases: *Extension*

An 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 manager, 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

This extension is called at the end of each epoch by default.

Parameters

- **iterator** (`Union[DataLoader[Any], Dict[str, DataLoader[Any]]]`) – 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** (`Union[Module, Dict[str, Module]]`) – `torch.nn.Module` object or a dictionary of links to evaluate. If this is just a layer object, the link is registered by the name 'main'.
- **eval_func** (`Optional[Callable[[...], Any]]`) – Evaluation function called at each iteration. The target link to evaluate as a callable is used by default.
- **progress_bar** – Boolean flag to show a progress bar while training, which is similar to `ProgressBar`. (default: `False`)
- **metrics** – List of callables that are called every batch to calculate metrics such as accuracy, roc_auc or others. The signature of the callable is: `def metric_fn(batch, output, last_iteration)` (default: `[]`)
- **eval_hook** (`Optional[Callable[[Evaluator], None]]`) –
- **kwargs** (`Any`) –

Warning: The argument `progress_bar` is experimental. The interface can change in the future.

`eval_hook`

Function to prepare for each evaluation process.

`eval_func`

Evaluation function called at each iteration.

Parameters

- **args** (`Any`) –
- **kwargs** (`Any`) –

Return type

`Any`

Methods

| | |
|---|---|
| <code>__init__(iterator, target[, eval_hook, ...])</code> | |
| <code>add_metric(metric_fn)</code> | Adds a custom metric to the evaluator. |
| <code>eval_func(*args, **kwargs)</code> | |
| <code>evaluate()</code> | Evaluates the model and returns a result dictionary. |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>get_all_iterators()</code> | Returns a dictionary of all iterators. |
| <code>get_all_targets()</code> | Returns a dictionary of all target links. |
| <code>get_iterator(name)</code> | Returns the iterator of the given name. |
| <code>get_target(name)</code> | Returns the target link of the given name. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| |
|--------------------------------|
| <code>default_name</code> |
| <code>is_async</code> |
| <code>name</code> |
| <code>needs_model_state</code> |
| <code>priority</code> |
| <code>trigger</code> |

`__call__(manager=None)`

Executes the evaluator extension.

Unlike usual extensions, this extension can be executed without passing a manager object. This extension reports the performance on validation dataset using the `report()` function. Thus, users can use this extension independently from any manager by manually configuring a `Reporter` object.

Parameters

manager (`ExtensionsManager`) – Manager object that invokes this extension.

Returns

Result dictionary that contains mean statistics of values reported by the evaluation function.

Return type

dict

`__init__(iterator, target, eval_hook=None, eval_func=None, **kwargs)`

Parameters

- **iterator** (*Union*[*DataLoader*[*Any*], *Dict*[*str*, *DataLoader*[*Any*]]]) –
- **target** (*Union*[*Module*, *Dict*[*str*, *Module*]]) –
- **eval_hook** (*Optional*[*Callable*[[*Evaluator*], *None*]]) –
- **eval_func** (*Optional*[*Callable*[[...], *Any*]]) –
- **kwargs** (*Any*) –

Return type

None

add_metric(*metric_fn*)

Adds a custom metric to the evaluator.

The metric is a callable that is executed every batch with the following signature: *def metric_fn(batch, output, last_iteration)*

Batch is the input batch passed to the model. Output is the result of evaluating batch, last_iteration is a boolean flag that indicates if its the last batch in the evaluation.

Parameters

metric_fn (*Callable*[[*Any*, *Any*, *Any*], *None*]) –

Return type

None

default_name = 'validation'

eval_func(*args, **kwargs)**Parameters**

- **args** (*Any*) –
- **kwargs** (*Any*) –

Return type*Any***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.

Users can override this method to customize the evaluation routine.

Returns

Result dictionary. This dictionary is further reported via **report()** without specifying any observer.

Return type

dict

get_all_iterators()

Returns a dictionary of all iterators.

Return type*Dict*[*str*, *DataLoader*[*Any*]]

get_all_targets()

Returns a dictionary of all target links.

Return type

Dict[str, Module]

get_iterator(name)

Returns the iterator of the given name.

Parameters

name (*str*) –

Return type

DataLoader[Any]

get_target(name)

Returns the target link of the given name.

Parameters

name (*str*) –

Return type

Module

priority: `int = 300`

trigger: `TriggerLike = (1, 'epoch')`

pytorch_pfn_extras.training.extensions.FailOnNonNumber

class pytorch_pfn_extras.training.extensions.**FailOnNonNumber**(*, *check_grad=True*)

Bases: *Extension*

An extension to raise `RuntimeError` if parameters and its gradients contain NaN or Inf.

Although parameters including non-number such as NaN and Inf are unnecessary in most cases the training loop 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.

Parameters

check_grad (*bool*) – Set to False to skip checking gradients.

Methods

| | |
|--|---|
| <code>__init__</code> (*[, check_grad]) | |
| <code>finalize</code> (manager) | Finalizes the extension. |
| <code>initialize</code> (manager) | Initializes up the manager state. |
| <code>load_state_dict</code> (to_load) | |
| <code>on_error</code> (manager, exc, tb) | Handles the error raised during training before finalization. |
| <code>state_dict</code> () | Serializes the extension state. |

Attributes

| | |
|---------------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code><i>needs_model_state</i></code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

`__call__`(*manager*)

Invokes the extension.

Implementations should override this operator. This method is called at iterations which the corresponding trigger accepts.

Parameters**`manager`** ([ExtensionsManager](#)) – Manager object to call this operator.**Return type**

None

`__init__`(***, *check_grad=True*)**Parameters****`check_grad`** (*bool*) –**`needs_model_state`** = `True`**`pytorch_pfn_extras.training.extensions.IgniteEvaluator`****class** `pytorch_pfn_extras.training.extensions.IgniteEvaluator`(*evaluator, iterator, target, **kwargs*)Bases: [Evaluator](#)

Methods

| | |
|--|---|
| <code>__init__(evaluator, iterator, target, **kwargs)</code> | |
| <code>add_metric(metric_fn)</code> | Adds a custom metric to the evaluator. |
| <code>eval_func(*args, **kwargs)</code> | |
| <code>evaluate()</code> | Evaluates the model and returns a result dictionary. |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>get_all_iterators()</code> | Returns a dictionary of all iterators. |
| <code>get_all_targets()</code> | Returns a dictionary of all target links. |
| <code>get_iterator(name)</code> | Returns the iterator of the given name. |
| <code>get_target(name)</code> | Returns the target link of the given name. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>set_evaluator_handlers()</code> | |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| |
|--------------------------------|
| <code>default_name</code> |
| <code>is_async</code> |
| <code>name</code> |
| <code>needs_model_state</code> |
| <code>priority</code> |
| <code>trigger</code> |

Parameters

- **evaluator** (*Engine*) –
- **iterator** (*Union[DataLoader[Any], Dict[str, DataLoader[Any]]]*) –
- **target** (*Union[Module, Dict[str, Module]]*) –
- **kwargs** (*Any*) –

`__init__(evaluator, iterator, target, **kwargs)`

Parameters

- **evaluator** (*Engine*) –
- **iterator** (*Union[DataLoader[Any], Dict[str, DataLoader[Any]]]*) –

- **target** (*Union[Module, Dict[str, Module]]*) –
- **kwargs** (*Any*) –

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.

Users can override this method to customize the evaluation routine.

Returns

Result dictionary. This dictionary is further reported via `report()` without specifying any observer.

Return type

dict

set_evaluator_handlers()**Return type**

None

pytorch_pfn_extras.training.extensions.LRScheduler

```
class pytorch_pfn_extras.training.extensions.LRScheduler(scheduler, *, stepper=<function
    _default_stepper>, trigger=(1, 'epoch'),
    wait_for_first_optimizer_step=False,
    is_async=True)
```

Bases: [Extension](#)

Trainer extension to adjust the learning rate using PyTorch's learning rate scheduler.

This extension calls `step()` method of the given LR scheduler. (*torch.optim.lr_scheduler.**). When using *ReduceLROnPlateau*, the latest reported *val/loss* value will be used. This behavior can be customized by passing a custom *stepper* function.

Parameters

- **scheduler** (*_LRScheduler* or *ReduceLROnPlateau*) – Any instance of *torch.optim.lr_scheduler.**.
- **stepper** (*callable*) – Function that performs the step on the scheduler.
- **trigger** (*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) – Frequency to call this extension.
- **wait_for_first_optimizer_step** (*bool*) – Wait until *optimizer.step()* is called before invoking *scheduler.step()*. This can address the issue where *optimizer.step()* is not called from the first iteration when using *GradScaler*.
- **is_async** (*bool*) –

Methods

| | |
|--|---|
| <code>__init__(scheduler, *[, stepper, trigger, ...])</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(state)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>state_dict()</code> | Serializes the extension state. |
| <code>step_by_value(key)</code> | |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

`__call__(manager)`

Invokes the extension.

Implementations should override this operator. This method is called at iterations which the corresponding trigger accepts.

Parameters

manager (`ExtensionsManager`) – Manager object to call this operator.

Return type

None

`__init__(scheduler, *, stepper=<function _default_stepper>, trigger=(1, 'epoch'), wait_for_first_optimizer_step=False, is_async=True)`

Parameters

- **scheduler** (*Any*) –
- **stepper** (*Any*) –
- **trigger** (`Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]`) –
- **wait_for_first_optimizer_step** (*bool*) –
- **is_async** (*bool*) –

Return type

None

load_state_dict(*state*)**Parameters****state** (*Dict[str, Any]*) –**Return type**

None

state_dict()

Serializes the extension state.

It is called when a manager that owns this extension is serialized. It serializes nothing by default.

Return type*Dict[str, Any]***static step_by_value**(*key*)**Parameters****key** (*Optional[str]*) –**Return type***Any***pytorch_pfn_extras.training.extensions.LogReport**

```
class pytorch_pfn_extras.training.extensions.LogReport(keys=None, trigger=(1, 'epoch'),  
                                                    postprocess=None, filename=None,  
                                                    append=False, format=None, **kwargs)
```

Bases: [Extension](#)

An extension to output the accumulated results to a log file.

This extension accumulates the observations of the manager 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 `ExtensionsManager.elapsed_time`.

Parameters

- **keys** (*iterable of str*) – Keys of values to accumulate. If this is None, all the values are accumulated and output to the log file.
- **trigger** (*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) – 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** (*Optional[Callable[[Mapping[str, Any]], None]]*) – 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.
- **filename** (*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. (default: `'log'`)
- **append** (*bool, optional*) – If the file is JSON Lines or YAML, contents will be appended instead of rewriting the file every call.
- **format** (*str, optional*) – accepted values are `'json'`, `'json-lines'` and `'yaml'`.
- **writer** (*writer object, optional*) – must be callable. object to dump the log to. If specified, it needs to have a correct *savefun* defined. The writer can override the save location in the `pytorch_pfn_extras.training.ExtensionsManager` object
- **kwargs** (*Any*) –

Note: Enabling `append=True` reduces size of snapshots (and thus reduces the time needed to take snapshots). Note that extensions relying on the logs from past iterations may behave differently; for example, when resuming from a snapshot, `PrintReport` will not show logs of iterations already done.

Methods

| | |
|--|---|
| <code>__init__([keys, trigger, postprocess, ...])</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>state_dict()</code> | Serializes the extension state. |
| <code>to_dataframe()</code> | |

Attributes

| | |
|---------------------------------------|---|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>log</code> <code>name</code> | The current list of observation dictionaries. |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

`__call__`(*manager*)

Invokes the extension.

Implementations should override this operator. This method is called at iterations which the corresponding trigger accepts.

Parameters

manager ([ExtensionsManager](#)) – Manager object to call this operator.

Return type

None

`__init__`(*keys=None, trigger=(1, 'epoch'), postprocess=None, filename=None, append=False, format=None, **kwargs*)

Parameters

- **keys** (*Optional[Iterable[str]]*) –
- **trigger** (*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) –
- **postprocess** (*Optional[Callable[[Mapping[str, Any]], None]]*) –
- **filename** (*Optional[str]*) –
- **append** (*bool*) –
- **format** (*Optional[str]*) –
- **kwargs** (*Any*) –

`finalize`(*manager*)

Finalizes the extension.

This method is called at the end of the training loop.

Parameters

manager ([ExtensionsManagerProtocol](#)) –

Return type

None

load_state_dict(*to_load*)

Parameters

to_load (*Dict[str, Any]*) –

Return type

None

property log: **List**[**Mapping**[**str**, **Any**]]

The current list of observation dictionaries.

state_dict()

Serializes the extension state.

It is called when a manager that owns this extension is serialized. It serializes nothing by default.

Return type

Dict[*str*, *Any*]

to_dataframe()

Return type

pandas.DataFrame

pytorch_pfn_extras.training.extensions.MaxAccumulate

class `pytorch_pfn_extras.training.extensions.MaxAccumulate`(*conversion_key_pair*, *trigger*=(1, 'epoch'), *distributed*=False)

Bases: *AccumulateBase*

Methods

| | |
|--|---|
| <code>__init__(conversion_key_pair[, trigger, ...])</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

Parameters

- **conversion_key_pair** (*Tuple*[*str*, *str*]) –
- **trigger** (*Optional*[*Union*[*Trigger*, *Callable*[[*ExtensionsManagerProtocol*], *bool*], *Tuple*[*float*, *str*]]]) –
- **distributed** (*bool*) –

pytorch_pfn_extras.training.extensions.MaxValue

class pytorch_pfn_extras.training.extensions.**MaxValue**(*key*, *trigger*=(1, 'epoch'))

Bases: *BestValue*

Extension traces the maximum value of a specific key in the observation.

Parameters

- **key** (*str*) – Key of value.
- **trigger** (*TriggerLike*) – Trigger that decides the comparison interval between current maximum value and new value. This must be a tuple in the form of <int>, 'epoch' or <int>, 'iteration' which is passed to *BestValueTrigger*.

Methods

| | |
|---|---|
| <code>__init__(key[, trigger])</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| | |
|--|--|
| <code>best_epoch</code> | Returns the epoch count that the current best value is observed. |
| <code>best_iteration</code> | Returns the iteration count that the current best value is observed. |
| <code>best_value</code> <code>default_name</code> | Returns the current best value. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

`__init__(key, trigger=(1, 'epoch'))`

Parameters

- **key** (*str*) –
- **trigger** (*TriggerLike*) –

`default_name = 'max_value'`

pytorch_pfn_extras.training.extensions.MicroAverage

class pytorch_pfn_extras.training.extensions.**MicroAverage**(*numerator_key, denominator_key, result_key, trigger=(1, 'epoch')*)

Bases: [Extension](#)

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(torch.nn.Module):
...     def __call__(self, x, y):
...         loss = torch.nn.CrossEntropyLoss(x, y)
...         correct = (x.data.argmax(axis=1) == y.data).sum()
...         total = len(y.data)
...         reporting.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** (*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) – 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

| | |
|--|---|
| <code>__init__(numerator_key, denominator_key, ...)</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

__call__(*manager*)

Invokes the extension.

Implementations should override this operator. This method is called at iterations which the corresponding trigger accepts.

Parameters

manager ([ExtensionsManager](#)) – Manager object to call this operator.

Return type

None

__init__(*numerator_key*, *denominator_key*, *result_key*, *trigger*=(1, 'epoch'))

Parameters

- **numerator_key** (*str*) –
- **denominator_key** (*str*) –
- **result_key** (*str*) –
- **trigger** (*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) –

Return type

None

load_state_dict(*to_load*)

Parameters

to_load (*Dict[str, Any]*) –

Return type

None

priority: `int = 200`

state_dict()

Serializes the extension state.

It is called when a manager that owns this extension is serialized. It serializes nothing by default.

Return type

Dict[str, Any]

pytorch_pfn_extras.training.extensions.MinAccumulate

class `pytorch_pfn_extras.training.extensions.MinAccumulate`(*conversion_key_pair*, *trigger*=(1, 'epoch'), *distributed*=False)

Bases: [AccumulateBase](#)

Methods

| | |
|--|---|
| <code>__init__(conversion_key_pair[, trigger, ...])</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

Parameters

- **conversion_key_pair** (*Tuple*[*str*, *str*]) –
- **trigger** (*Optional*[*Union*[*Trigger*, *Callable*[[*ExtensionsManagerProtocol*], *bool*], *Tuple*[*float*, *str*]]]) –
- **distributed** (*bool*) –

pytorch_pfn_extras.training.extensions.MinValue

class `pytorch_pfn_extras.training.extensions.MinValue`(*key*, *trigger*=(1, 'epoch'))

Bases: [`BestValue`](#)

Extension traces the maximum value of a specific key in the observation.

Parameters

- **key** (*str*) – Key of value.
- **trigger** (*TriggerLike*) – Trigger that decides the comparison interval between current maximum value and new value. This must be a tuple in the form of `<int>`, `'epoch'` or `<int>`, `'iteration'` which is passed to `BestValueTrigger`.

Methods

| | |
|---|---|
| <code>__init__(key[, trigger])</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--|
| <code>best_epoch</code> | Returns the epoch count that the current best value is observed. |
| <code>best_iteration</code> | Returns the iteration count that the current best value is observed. |
| <code>best_value</code> | Returns the current best value. |
| <code>default_name</code> | |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

`__init__(key, trigger=(1, 'epoch'))`

Parameters

- **key** (*str*) –
- **trigger** (*TriggerLike*) –

`default_name = 'min_value'`

pytorch_pfn_extras.training.extensions.ParameterStatistics

```
class pytorch_pfn_extras.training.extensions.ParameterStatistics(links, statistics='default',
                                                                report_params=True,
                                                                report_grads=True,
                                                                prefix=None, trigger=(1,
                                                                'epoch'),
                                                                skip_nan_params=False)
```

Bases: [Extension](#)

An extension to report parameter statistics.

Statistics are collected and reported for a given `Module` or an iterable of `Modules`. If a link contains child modules, the statistics are reported separately for each child.

Any function that takes a one-dimensional `torch.Tensor` 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** (*instance or iterable of Module*) – Module(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 or 'default'*) – 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. If the special value `'default'` is specified, the default statistics functions will be used.
- **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** (*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) – 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.

Note: The default statistic functions are as follows:

- `'mean'` (`xp.mean(x)`)
 - `'std'` (`xp.std(x)`)
 - `'min'` (`xp.min(x)`)
 - `'max'` (`xp.max(x)`)
 - `'zeros'` (`xp.count_nonzero(x == 0)`)
 - `'percentile'` (`xp.percentile(x, (0.13, 2.28, 15.87, 50, 84.13, 97.72, 99.87))`)
-

Methods

| | |
|--|---|
| <code>__init__(links[, statistics, report_params, ...])</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>register_statistics(name, function)</code> | Register a function to compute a certain statistic. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| |
|----------------------------------|
| <code>default_name</code> |
| <code>default_statistics</code> |
| <code>is_async</code> |
| <code>name</code> |
| <code>needs_model_state</code> |
| <code>priority</code> |
| <code>report_key_template</code> |
| <code>trigger</code> |

`__call__(manager)`

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

manager (`ExtensionsManager`) – Associated manager that invoked this extension.

Return type

None

`__init__(links, statistics='default', report_params=True, report_grads=True, prefix=None, trigger=(1, 'epoch'), skip_nan_params=False)`

Parameters

- **links** (*Any*) –
- **statistics** (*Any*) –

- **report_params** (*bool*) –
- **report_grads** (*bool*) –
- **prefix** (*Optional[str]*) –
- **trigger** (*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) –
- **skip_nan_params** (*bool*) –

default_name = 'parameter_statistics'

default_statistics = {'max': <function <lambda>>, 'mean': <function <lambda>>, 'min': <function <lambda>>, 'std': <function <lambda>>, 'zeros': <function <lambda>>}

priority: int = 300

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** (*Any*) – 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.

Return type

None

report_key_template = '{prefix}{param_name}/{attr_name}/{function_name}'

pytorch_pfn_extras.training.extensions.PlotReport

```
class pytorch_pfn_extras.training.extensions.PlotReport(y_keys, x_key='iteration', trigger=(1, 'epoch'), postprocess=None, filename='plot.png', marker='x', grid=True)
```

Bases: [Extension](#)

An extension to output plots.

This extension accumulates the observations of the manager 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 `manager.run_iteration`. For example:

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

manager.extend(
    extensions.PlotReport(['main/loss', 'validation/main/loss'],
                          'epoch', filename='loss.png'))
with manager.run_iteration():
    pass
```

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** (*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) – 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** (*Any*) – 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.
- **filename** (*str*) – Name of the figure file under the output directory. It can be a format string. For historical reasons `file_name` is also accepted as an alias of this argument.
- **marker** (*str*) – The marker used to plot the graph. Default is 'x'. If None is given, it draws with no markers.
- **grid** (*bool*) – If True, set the axis grid on. The default value is True.
- **writer** (*writer object, optional*) – must be callable. object to dump the log to. If specified, it needs to have a correct `savefun` defined. The writer can override the save location in the `pytorch_pfn_extras.training.ExtensionsManager` object
- **kwargs** (*Any*) –

Methods

| | |
|--|---|
| <code>__init__(y_keys[, x_key, trigger, ...])</code> | |
| <code>available()</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

`__call__(manager)`

Invokes the extension.

Implementations should override this operator. This method is called at iterations which the corresponding trigger accepts.

Parameters

manager (`ExtensionsManager`) – Manager object to call this operator.

Return type

None

`__init__(y_keys, x_key='iteration', trigger=(1, 'epoch'), postprocess=None, filename=None, marker='x', grid=True, **kwargs)`

Parameters

- **y_keys** (`Union[Iterable[str], str]`) –
- **x_key** (`str`) –
- **trigger** (`Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]`) –
- **postprocess** (`Optional[Any]`) –
- **filename** (`Optional[str]`) –

- **marker** (*str*) –
- **grid** (*bool*) –
- **kwargs** (*Any*) –

static available()

Return type
bool

finalize(*manager*)

Finalizes the extension.

This method is called at the end of the training loop.

Parameters

manager ([ExtensionsManagerProtocol](#)) –

Return type
None

load_state_dict(*to_load*)

Parameters

to_load (*Dict[str, Any]*) –

Return type
None

state_dict()

Serializes the extension state.

It is called when a manager that owns this extension is serialized. It serializes nothing by default.

Return type

Dict[str, Any]

pytorch_pfn_extras.training.extensions.PrintReport

```
class pytorch_pfn_extras.training.extensions.PrintReport(entries=None, log_report='LogReport',
                                                         out=<_io.TextIOWrapper
                                                         name='<stdout>' mode='w'
                                                         encoding='utf-8'>)
```

Bases: [Extension](#)

An 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 or None*) – List of keys of observations to print. If *None* is passed, automatically infer keys from reported dict.
- **log_report** (*str or LogReport*) – Log report to accumulate the observations. This is either the name of a [LogReport](#) extensions registered to the manager, or a [LogReport](#) instance to use internally.
- **out** ([IO\[Any\]](#)) – Stream to print the bar. Standard output is used by default.

Methods

| | |
|---|---|
| <code>__init__([entries, log_report, out])</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>get_log_report(manager)</code> | |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

`__call__(manager)`

Invokes the extension.

Implementations should override this operator. This method is called at iterations which the corresponding trigger accepts.

Parameters

manager (`ExtensionsManager`) – Manager object to call this operator.

Return type

None

`__init__(entries=None, log_report='LogReport', out=<_io.TextIOWrapper name='<stdout>' mode='w' encoding='utf-8'>)`

Parameters

- **entries** (`Optional[Sequence[str]]`) –
- **log_report** (`Union[str, LogReport]`) –
- **out** (`IO[Any]`) –

Return type

None

get_log_report(*manager*)

Parameters

manager ([ExtensionsManagerProtocol](#)) –

Return type

[LogReport](#)

initialize(*manager*)

Initializes up the manager state.

This method is called before entering the training loop. An extension modifying the state of `ExtensionsManager` can override this method to initialize it.

When the manager has been restored from a snapshot, this method has to recover an appropriate part of the state of the manager.

Parameters

manager ([ExtensionsManager](#)) – Manager object to call this extension.

Return type

None

load_state_dict(*to_load*)

Parameters

to_load (*Dict[str, Any]*) –

Return type

None

state_dict()

Serializes the extension state.

It is called when a manager that owns this extension is serialized. It serializes nothing by default.

Return type

Dict[str, Any]

pytorch_pfn_extras.training.extensions.PrintReportCLI

`pytorch_pfn_extras.training.extensions.PrintReportCLI`

alias of [PrintReport](#)

pytorch_pfn_extras.training.extensions.ProfileReport

```
class pytorch_pfn_extras.training.extensions.ProfileReport(store_keys=None, report_keys=None,
                                                         trigger=(1, 'epoch'), filename=None,
                                                         append=False, format=None,
                                                         **kwargs)
```

Bases: [Extension](#)

Writes the profile results to a file.

Times are reported by using the `pytorch_pfn_extras.profiler.TimeSummary.report()` context manager.

Parameters

- **store_keys** (*iterable of strs*) – Keys of values to write to the profiler report file.
- **report_keys** (*iterable of strs*) – Keys of values that will be reported.
- **trigger** (*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) – 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.
- **filename** (*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.
- **append** (*bool, options1*) – If the file is JSON Lines or YAML, contents will be appended instead of rewriting the file every call.
- **format** (*str, optional*) – accepted values are 'json', 'json-lines' and 'yaml'.
- **writer** (*writer object, optional*) – must be callable. object to dump the log to. If specified, it needs to have a correct *savefun* defined. The writer can override the save location in the `pytorch_pfn_extras.training.ExtensionsManager` object
- **entries** (*list*) – list of str
- **kwargs** (*Any*) –

Returns

header string templates (str): template string for print values.

Return type

header (str)

Methods

| | |
|---|---|
| <code>__init__</code> ([store_keys, report_keys, trigger, ...]) | |
| <code>finalize</code> (manager) | Finalizes the extension. |
| <code>initialize</code> (manager) | Initializes up the manager state. |
| <code>load_state_dict</code> (to_load) | |
| <code>on_error</code> (manager, exc, tb) | Handles the error raised during training before finalization. |
| <code>state_dict</code> () | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

`__call__`(*manager*)

Invokes the extension.

Implementations should override this operator. This method is called at iterations which the corresponding trigger accepts.

Parameters

manager ([ExtensionsManager](#)) – Manager object to call this operator.

Return type

None

`__init__`(*store_keys=None, report_keys=None, trigger=(1, 'epoch'), filename=None, append=False, format=None, **kwargs*)

Parameters

- **store_keys** (*Optional[Iterable[str]]*) –
- **report_keys** (*Optional[Iterable[str]]*) –
- **trigger** (*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) –
- **filename** (*Optional[str]*) –
- **append** (*bool*) –
- **format** (*Optional[str]*) –
- **kwargs** (*Any*) –

`finalize`(*manager*)

Finalizes the extension.

This method is called at the end of the training loop.

Parameters

manager ([ExtensionsManagerProtocol](#)) –

Return type

None

`load_state_dict`(*to_load*)

Parameters

to_load (*Dict[str, Any]*) –

Return type

None

state_dict()

Serializes the extension state.

It is called when a manager that owns this extension is serialized. It serializes nothing by default.

Return type*Dict[str, Any]***pytorch_pfn_extras.training.extensions.ProgressBar**

```
class pytorch_pfn_extras.training.extensions.ProgressBar(training_length=None,
                                                         update_interval=100, bar_length=50,
                                                         out=<_io.TextIOWrapper
                                                         name='<stdout>' mode='w'
                                                         encoding='utf-8'>)
```

Bases: [Extension](#)

An 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 or None*) – 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 manager 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** (*Any*) – Stream to print the bar. Standard output is used by default.

Methods

| | |
|---|---|
| <hr/> <code>__init__</code> ([training_length, update_interval, ...]) | |
| <code>finalize</code> (manager) | Finalizes the extension. |
| <code>initialize</code> (manager) | Initializes up the manager state. |
| <hr/> <code>load_state_dict</code> (to_load) | |
| <code>on_error</code> (manager, exc, tb) | Handles the error raised during training before finalization. |
| <code>state_dict</code> () | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

`__call__(manager)`

Invokes the extension.

Implementations should override this operator. This method is called at iterations which the corresponding trigger accepts.

Parameters

manager ([ExtensionsManager](#)) – Manager object to call this operator.

Return type

None

`__init__(training_length=None, update_interval=100, bar_length=50, out=<_io.TextIOWrapper name='<stdout>' mode='w' encoding='utf-8'>)`

Parameters

- **training_length** (*Optional*[Any]) –
- **update_interval** (*int*) –
- **bar_length** (*int*) –
- **out** (*Any*) –

`finalize(manager)`

Finalizes the extension.

This method is called at the end of the training loop.

Parameters

manager ([ExtensionsManagerProtocol](#)) –

Return type

None

pytorch_pfn_extras.training.extensions.ProgressBarCLI

pytorch_pfn_extras.training.extensions.ProgressBarCLI

alias of *ProgressBar*

pytorch_pfn_extras.training.extensions.Slack

```
class pytorch_pfn_extras.training.extensions.Slack(channel, msg=None, *, start_msg='{default}',
                                                    end_msg='{default}', error_msg='{default}',
                                                    thread=True, filenames=None,
                                                    upload_trigger=None, context=None,
                                                    token=None)
```

Bases: *_SlackBase*

An extension to communicate with Slack.

Example

```
>>> ppe.training.extensions.Slack(
...     channel="experiment-progress",
...     msg="Epoch #{manager.epoch}: loss = {val/loss}",
...     end_msg="{default} \n <@username> Check out the result!",
...     # Upload files at the end of the experiment.
...     filenames=["result/statistics.png"],
...     upload_trigger=(max_epoch, 'epoch'),
... )
```

This extension posts a message when:

- `start_msg`: The training has started
- `msg`: The extension is triggered, usually at the end of each epoch
- `end_msg`: The training has finished
- `error_msg`: An exception has raised during the training

These messages can be specified as a format string, a callable that returns a string, or `None` to disable posting on that event.

When using a format string, the following variables are available for use:

- `manager`: an *ExtensionsManager* object
- `default`: the default message string
- `context`: an arbitrary object passed to this extension
- `error`: an *Exception* object (for `error_msg` only)
- All reported values (`manager.observations`)

When using a callable, it should take (*ExtensionsManager*, *context*) or (*ExtensionsManager*, *Exception*, *context*) (for `error_msg`) and return a string.

This extension can upload files along with the message when triggered. `filenames` can be a list of filenames (the same formatting rule as `msg` apply), or a callable taking (`ExtensionsManager`, `context`) and returning a list of filenames.

To use this extension, you must create a Slack app, then specify the token via an environment variable `SLACK_BOT_TOKEN` or `token` option.

Parameters

- **channel** (*str*) – The channel where messages and files will be sent. This can be a channel name or a channel ID.
- **msg** (*str*, *callable*, or *None*) – A message to be sent when triggered. It can be a string to be formatted using `.format` or a callable that returns a string.
- **start_msg** (*str*, *callable*, or *None*) – A message to be sent at the beginning of the experiment.
- **end_msg** (*str*, *callable*, or *None*) – A message to be sent at the completion of the experiment.
- **error_msg** (*str*, *callable*, or *None*) – A message to be sent when an exception is raised during the experiment.
- **thread** (*bool*) – When True, subsequent messages will be posted as a thread of the original message. Default is True.
- **filenames** (*list of str or callable*) – A list of files that will be uploaded. These are string templates that can take values in the same way as `msg`, or a callable that returns a list of filenames.
- **upload_trigger** (*trigger or None*) – Used to upload files at certain events. If not specified, files will be uploaded in every call.
- **context** (*Any*) – Any arbitrary user object you will need when generating a message.
- **token** (*str*) – Slack bot token. If *None*, the environment variable `SLACK_BOT_TOKEN` will be used. Optional, default is *None*.

Methods

| | |
|--|---|
| <code>__init__(channel[, msg, start_msg, end_msg, ...])</code> | |
| <code>default_end_msg(context)</code> | |
| <code>default_error_msg(exc, context)</code> | |
| <code>default_msg(context)</code> | |
| <code>default_start_msg(context)</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| | |
|-------------------|--------------------------------|
| default_name | Default name of the extension. |
| is_async | |
| name | |
| needs_model_state | |
| priority | |
| <i>trigger</i> | |

```
__init__(channel, msg=None, *, start_msg='{default}', end_msg='{default}', error_msg='{default}',  
          thread=True, filenames=None, upload_trigger=None, context=None, token=None)
```

Parameters

- **channel** (*str*) –
- **msg** (*Optional[Union[[str](#), Callable[[[ExtensionsManagerProtocol](#), Any], [str](#)]]]*) –
- **start_msg** (*Optional[Union[[str](#), Callable[[[ExtensionsManagerProtocol](#), Any], [str](#)]]]*) –
- **end_msg** (*Optional[Union[[str](#), Callable[[[ExtensionsManagerProtocol](#), Any], [str](#)]]]*) –
- **error_msg** (*Optional[Union[[str](#), Callable[[[ExtensionsManagerProtocol](#), Any, [Exception](#)], [str](#)]]]*) –
- **thread** (*bool*) –
- **filenames** (*Optional[Union[Sequence[[str](#)], Callable[[[ExtensionsManagerProtocol](#), Any], Sequence[[str](#)]]]]*) –
- **upload_trigger** (*Optional[Union[[Trigger](#), Callable[[[ExtensionsManagerProtocol](#), [bool](#)], Tuple[[float](#), [str](#)]]]*) –
- **context** (*Optional[Any]*) –
- **token** (*Optional[[str](#)]*) –

Return type

None

```
trigger: Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool],  
Tuple[float, str]]] = (1, 'epoch')
```

pytorch_pfn_extras.training.extensions.SlackWebhook

```
class pytorch_pfn_extras.training.extensions.SlackWebhook(url, msg=None, *, start_msg='{default}',
                                                         end_msg='{default}',
                                                         error_msg='{default}', context=None)
```

Bases: `_SlackBase`

An extension to communicate with Slack using Incoming Webhook.

Example

```
>>> ppe.training.extensions.SlackWebhook(
...     url="https://hooks.slack.com/services/Txxxxx....",
...     msg="Epoch #{manager.epoch}: loss = {val/loss}",
...     end_msg="{default} \n <@username> Check out the result!",
... )
```

This extension posts a message when:

- `start_msg`: The training has started
- `msg`: The extension is triggered, usually at the end of each epoch
- `end_msg`: The training has finished
- `error_msg`: An exception has raised during the training

These messages can be specified as a format string, a callable that returns a string, or `None` to disable posting on that event.

When using a format string, the following variables are available for use:

- `manager`: an `ExtensionsManager` object
- `default`: the default message string
- `context`: an arbitrary object passed to this extension
- `error`: an `Exception` object (for `error_msg` only)
- All reported values (`manager.observations`)

When using a callable, it should take *(ExtensionsManager, context)* or *(ExtensionsManager, Exception, context)* (for `error_msg`) and return a string.

Parameters

- **url** (*str*) – Incoming webhook URL to send messages.
- **msg** (*str, callable, or None*) – A message to be sent when triggered. It can be a string to be formatted using `.format` or a callable that returns a string.
- **start_msg** (*str, callable, or None*) – A message to be sent at the beginning of the experiment.
- **end_msg** (*str, callable, or None*) – A message to be sent at the completion of the experiment.
- **error_msg** (*str, callable, or None*) – A message to be sent when an exception is raised during the experiment.
- **context** (*object*) – Any arbitrary user object you will need when generating a message.

Methods

| | |
|--|---|
| <code>__init__(url[, msg, start_msg, end_msg, ...])</code> | |
| <code>default_end_msg(context)</code> | |
| <code>default_error_msg(exc, context)</code> | |
| <code>default_msg(context)</code> | |
| <code>default_start_msg(context)</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

`__init__(url, msg=None, *, start_msg='{default}', end_msg='{default}', error_msg='{default}', context=None)`

Parameters

- **url** (*str*) –
- **msg** (*Optional[Union[[str](#), Callable[[[ExtensionsManagerProtocol](#), Any], [str](#)]]]*) –
- **start_msg** (*Optional[Union[[str](#), Callable[[[ExtensionsManagerProtocol](#), Any], [str](#)]]]*) –
- **end_msg** (*Optional[Union[[str](#), Callable[[[ExtensionsManagerProtocol](#), Any], [str](#)]]]*) –
- **error_msg** (*Optional[Union[[str](#), Callable[[[ExtensionsManagerProtocol](#), Any, [Exception](#)], [str](#)]]]*) –
- **context** (*Optional[Any]*) –

Return type
None

pytorch_pfn_extras.training.extensions.SnapshotMode

class pytorch_pfn_extras.training.extensions.SnapshotMode(*value*)
Bases: Enum
An enumeration.

Attributes

| |
|------------------------|
| <i>DEFAULT</i> |
| <i>DISTRIBUTED</i> |
| <i>SHARDED</i> |
| DEFAULT = 1 |
| DISTRIBUTED = 2 |
| SHARDED = 3 |

pytorch_pfn_extras.training.extensions.StandardDeviationAccumulate

class pytorch_pfn_extras.training.extensions.StandardDeviationAccumulate(*conversion_key_pair*,
trigger=(1, 'epoch'),
distributed=False)
Bases: *AccumulateBase*

Methods

| | |
|---|--|
| <i>__init__</i> (conversion_key_pair[, trigger, ...]) | |
| <i>finalize</i> (manager) | Finalizes the extension. |
| <i>initialize</i> (manager) | Initializes up the manager state. |
| <i>load_state_dict</i> (to_load) | |
| <i>on_error</i> (manager, exc, tb) | Handles the error raised during training before final- ization. |
| <i>state_dict</i> () | Serializes the extension state. |

Attributes

| | |
|-------------------|--------------------------------|
| default_name | Default name of the extension. |
| is_async | |
| name | |
| needs_model_state | |
| priority | |
| trigger | |

Parameters

- **conversion_key_pair** (*Tuple*[*str*, *str*]) –
- **trigger** (*Optional*[*Union*[*Trigger*, *Callable*[[*ExtensionsManagerProtocol*], *bool*], *Tuple*[*float*, *str*]]]) –
- **distributed** (*bool*) –

pytorch_pfn_extras.training.extensions.TimelineTrace

```
class pytorch_pfn_extras.training.extensions.TimelineTrace(trigger=(1, 'epoch'), filename=None,
                                                           enable=None, disable=None,
                                                           tracer=None, **kwargs)
```

Bases: [Extension](#)

Writes the profile timeline to a file.

Times are reported by using the `pytorch_pfn_extras.profiler.TimeSummary.report()` context manager.

Parameters

- **trigger** (*Optional*[*Union*[*Trigger*, *Callable*[[*ExtensionsManagerProtocol*], *bool*], *Tuple*[*float*, *str*]]]) – Trigger that decides when to output the trace. 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`.
- **enable** (*Optional*[*Union*[*Trigger*, *Callable*[[*ExtensionsManagerProtocol*], *bool*], *Tuple*[*float*, *str*]]]) – Trigger that enables the tracing. Note that since the extensions are executed at the end of an iteration the tracer will be enabled from the iteration after the trigger is fired. If it is a tuple in the form <int>, 'epoch' or <int>, 'iteration', it is passed to `IntervalTrigger`.
- **disable** (*Optional*[*Union*[*Trigger*, *Callable*[[*ExtensionsManagerProtocol*], *bool*], *Tuple*[*float*, *str*]]]) – Trigger that disables the tracing. Note that since the extensions are executed at the end of an iteration the tracer will be disabled from the iteration after the trigger is fired. If it is a tuple in the form <int>, 'epoch' or <int>, 'iteration', it is passed to `IntervalTrigger`.
- **filename** (*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.

- **writer** (*writer object, optional*) – must be callable. object to dump the log to. If specified, it needs to have a correct *savefun* defined. The writer can override the save location in the `pytorch_pfn_extras.training.ExtensionsManager` object
- **tracer** (*tracer object, optional*) – object with the tracing logic. If not specified, the default tracer in the thread local storage will be used.
- **entries** (*list*) – list of str
- **kwargs** (*Any*) –

Returns

header string templates (str): template string for print values.

Return type

header (str)

Methods

| | |
|--|---|
| <code>__init__</code> ([trigger, filename, enable, ...]) | |
| <code>finalize</code> (manager) | Finalizes the extension. |
| <code>initialize</code> (manager) | Initializes up the manager state. |
| <code>load_state_dict</code> (to_load) | |
| <code>on_error</code> (manager, exc, tb) | Handles the error raised during training before finalization. |
| <code>state_dict</code> () | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

`__call__`(manager)

Invokes the extension.

Implementations should override this operator. This method is called at iterations which the corresponding trigger accepts.

Parameters

manager (`ExtensionsManager`) – Manager object to call this operator.

Return type

None

__init__(*trigger=(1, 'epoch'), filename=None, enable=None, disable=None, tracer=None, **kwargs*)**Parameters**

- **trigger** (*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) –
- **filename** (*Optional[str]*) –
- **enable** (*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) –
- **disable** (*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) –
- **tracer** (*Optional[Tracer]*) –
- **kwargs** (*Any*) –

finalize(*manager*)

Finalizes the extension.

This method is called at the end of the training loop.

Parameters**manager** (*ExtensionsManagerProtocol*) –**Return type**

None

initialize(*manager*)

Initializes up the manager state.

This method is called before entering the training loop. An extension modifying the state of `ExtensionsManager` can override this method to initialize it.

When the manager has been restored from a snapshot, this method has to recover an appropriate part of the state of the manager.

Parameters**manager** (*ExtensionsManager*) – Manager object to call this extension.**Return type**

None

load_state_dict(*to_load*)**Parameters****to_load** (*Dict[str, Any]*) –**Return type**

None

state_dict()

Serializes the extension state.

It is called when a manager that owns this extension is serialized. It serializes nothing by default.

Return type*Dict[str, Any]*

pytorch_pfn_extras.training.extensions.UnbiasedStandardDeviationAccumulate

```
class pytorch_pfn_extras.training.extensions.UnbiasedStandardDeviationAccumulate(conversion_key_pair,
                                                                                   trigger=(1,
                                                                                   'epoch'),
                                                                                   distributed=False)
```

Bases: *AccumulateBase*

Methods

| | |
|--|---|
| <code>__init__(conversion_key_pair[, trigger, ...])</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

Parameters

- **conversion_key_pair** (*Tuple*[*str*, *str*]) –
- **trigger** (*Optional*[*Union*[*Trigger*, *Callable*[[*ExtensionsManagerProtocol*], *bool*], *Tuple*[*float*, *str*]]]) –
- **distributed** (*bool*) –

pytorch_pfn_extras.training.extensions.VariableStatisticsPlot

```
class pytorch_pfn_extras.training.extensions.VariableStatisticsPlot(targets,
                                                                    max_sample_size=1000,
                                                                    report_data=True,
                                                                    report_grad=True,
                                                                    plot_mean=True,
                                                                    plot_std=True,
                                                                    percentile_sigmas=(0, 0.13,
                                                                2.28, 15.87, 50, 84.13,
                                                                97.72, 99.87, 100),
                                                                    trigger=(1, 'epoch'),
                                                                    filename='statistics.png',
                                                                    figsize=None,
                                                                    marker=None, grid=True)
```

Bases: [Extension](#)

An extension to plot statistics for Tensors.

This extension collects statistics for a single `torch.Tensor`, a list of `torch.Tensors` or similarly a single or a list of `torch.nn.Modules` containing one or more `torch.Tensors`. In case multiple `torch.Tensors` are found, the means are computed. The collected statistics are plotted and saved as an image in the directory specified by the `Manager`.

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.

:param targets (`torch.Tensor`: or list of either): Parameters for which statistics are collected. :param `torch.nn.Module`: or list of either): Parameters for which statistics are collected. :param `max_sample_size`: Maximum number of running samples. :type `max_sample_size`: `int` :param `report_data`: If `True`, data (e.g. weights) statistics are plotted. If

False, they are neither computed nor plotted.

Parameters

- **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** (*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) – 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`.
- **filename** (*str*) – Name of the output image file under the output directory. For historical reasons `file_name` is also accepted as an alias of this argument.

- **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 the plots or not.
- **writer** (*writer object, optional*) – must be callable. object to dump the log to. If specified, it needs to have a correct *savefun* defined. The writer can override the save location in the `pytorch_pfn_extras.training.ExtensionsManager` object
- **targets** (*Any*) –
- **max_sample_size** (*int*) –
- **report_data** (*bool*) –
- **kwargs** (*Any*) –

Methods

| | |
|--|---|
| <code>__init__(targets[, max_sample_size, ...])</code> | |
| <code>available()</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>save_plot_using_module(plt, manager)</code> | |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

`__call__(manager)`
Invokes the extension.

Implementations should override this operator. This method is called at iterations which the corresponding trigger accepts.

Parameters

manager ([ExtensionsManager](#)) – Manager object to call this operator.

Return type

None

```
__init__(targets, max_sample_size=1000, report_data=True, report_grad=True, plot_mean=True,
         plot_std=True, percentile_sigmas=(0, 0.13, 2.28, 15.87, 50, 84.13, 97.72, 99.87, 100), trigger=(1,
         'epoch'), filename=None, figsize=None, marker=None, grid=True, **kwargs)
```

Parameters

- **targets** (*Any*) –
- **max_sample_size** (*int*) –
- **report_data** (*bool*) –
- **report_grad** (*bool*) –
- **plot_mean** (*bool*) –
- **plot_std** (*bool*) –
- **percentile_sigmas** (*Union*[*float*, *Tuple*[*float*, ...]]) –
- **trigger** (*Optional*[*Union*[[Trigger](#), *Callable*[[[ExtensionsManagerProtocol](#)], *bool*], *Tuple*[*float*, *str*]]]) –
- **filename** (*Optional*[*str*]) –
- **figsize** (*Optional*[*Tuple*[*int*, ...]]) –
- **marker** (*Optional*[*str*]) –
- **grid** (*bool*) –
- **kwargs** (*Any*) –

static available()

Return type

bool

finalize(manager)

Finalizes the extension.

This method is called at the end of the training loop.

Parameters

manager ([ExtensionsManagerProtocol](#)) –

Return type

None

save_plot_using_module(plt, manager)

Parameters

- **plt** (*Any*) –
- **manager** ([ExtensionsManagerProtocol](#)) –

Return type

None

Modules

`pytorch_pfn_extras.training.extensions.
accumulate`

`pytorch_pfn_extras.training.extensions.
best_value`

`pytorch_pfn_extras.training.extensions.
evaluator`

`pytorch_pfn_extras.training.extensions.
fail_on_non_number`

`pytorch_pfn_extras.training.extensions.
log_report`

`pytorch_pfn_extras.training.extensions.
lr_scheduler`

`pytorch_pfn_extras.training.extensions.
micro_average`

`pytorch_pfn_extras.training.extensions.
parameter_statistics`

`pytorch_pfn_extras.training.extensions.
plot_report`

`pytorch_pfn_extras.training.extensions.
print_report`

`pytorch_pfn_extras.training.extensions.
profile_report`

`pytorch_pfn_extras.training.extensions.
progress_bar`

`pytorch_pfn_extras.training.extensions.
slack`

`pytorch_pfn_extras.training.extensions.
snapshot_writers`

`pytorch_pfn_extras.training.extensions.
timeline_trace`

`pytorch_pfn_extras.training.extensions.
util`

`pytorch_pfn_extras.training.extensions.
value_observation`

`pytorch_pfn_extras.training.extensions.
variable_statistics_plot`

pytorch_pfn_extras.training.extensions.accumulate**Classes**

`pytorch_pfn_extras.training.extensions.
accumulate.AverageAccumulate(...)`

`pytorch_pfn_extras.training.extensions.
accumulate.MaxAccumulate(...)`

`pytorch_pfn_extras.training.extensions.
accumulate.MinAccumulate(...)`

`pytorch_pfn_extras.training.extensions.
accumulate.StandardDeviationAccumulate(...)`

`pytorch_pfn_extras.training.
extensions.accumulate.
UnbiasedStandardDeviationAccumulate(...)`

pytorch_pfn_extras.training.extensions.accumulate.AverageAccumulate

class pytorch_pfn_extras.training.extensions.accumulate.**AverageAccumulate**(*conversion_key_pair*,
trigger=(1, 'epoch'),
distributed=False)

Bases: *AccumulateBase*

Methods

| | |
|--|--|
| <code>__init__(conversion_key_pair[, trigger, ...])</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before final- ization. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

Parameters

- **conversion_key_pair** (*Tuple[str, str]*) –
- **trigger** (*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) –
- **distributed** (*bool*) –

pytorch_pfn_extras.training.extensions.accumulate.MaxAccumulate

class pytorch_pfn_extras.training.extensions.accumulate.**MaxAccumulate**(*conversion_key_pair*,
trigger=(1, 'epoch'),
distributed=False)

Bases: *AccumulateBase*

Methods

| | |
|--|---|
| <code>__init__(conversion_key_pair[, trigger, ...])</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

Parameters

- **conversion_key_pair** (*Tuple[str, str]*) –
- **trigger** (*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) –
- **distributed** (*bool*) –

pytorch_pfn_extras.training.extensions.accumulate.MinAccumulate

```
class pytorch_pfn_extras.training.extensions.accumulate.MinAccumulate(conversion_key_pair,  
                                                                    trigger=(1, 'epoch'),  
                                                                    distributed=False)
```

Bases: *AccumulateBase*

Methods

| | |
|--|---|
| <hr/> | |
| <code>__init__(conversion_key_pair[, trigger, ...])</code> | |
| <hr/> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <hr/> | |
| <code>load_state_dict(to_load)</code> | |
| <hr/> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <hr/> | |
| <code>state_dict()</code> | Serializes the extension state. |
| <hr/> | |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <hr/> | |
| <code>is_async</code> | |
| <hr/> | |
| <code>name</code> | |
| <hr/> | |
| <code>needs_model_state</code> | |
| <hr/> | |
| <code>priority</code> | |
| <hr/> | |
| <code>trigger</code> | |
| <hr/> | |

Parameters

- **conversion_key_pair** (*Tuple*[*str*, *str*]) –
- **trigger** (*Optional*[*Union*[*Trigger*, *Callable*[[*ExtensionsManagerProtocol*], *bool*], *Tuple*[*float*, *str*]]]) –
- **distributed** (*bool*) –

pytorch_pfn_extras.training.extensions.accumulate.StandardDeviationAccumulate

```
class pytorch_pfn_extras.training.extensions.accumulate.StandardDeviationAccumulate(conversion_key_pair,
                                                                                   trigger=(1,
                                                                                   'epoch'),
                                                                                   dis-
                                                                                   tributed=False)
```

Bases: *AccumulateBase*

Methods

| | |
|--|---|
| <code>__init__(conversion_key_pair[, trigger, ...])</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

Parameters

- **conversion_key_pair** (*Tuple[str, str]*) –
- **trigger** (*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) –
- **distributed** (*bool*) –

pytorch_pfn_extras.training.extensions.accumulate.UnbiasedStandardDeviationAccumulate

```
class pytorch_pfn_extras.training.extensions.accumulate.UnbiasedStandardDeviationAccumulate(conversion_key,
trig-
ger=(1,
'epoch'),
dis-
tributed=False
```

Bases: *AccumulateBase*

Methods

| | |
|--|--|
| <code>__init__(conversion_key_pair[, trigger, ...])</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before final- ization. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

Parameters

- **conversion_key_pair** (*Tuple*[*str*, *str*]) –
- **trigger** (*Optional*[*Union*[*Trigger*, *Callable*[[*ExtensionsManagerProtocol*], *bool*], *Tuple*[*float*, *str*]]]) –
- **distributed** (*bool*) –

Modules

```
pytorch_pfn_extras.training.extensions.
accumulate.average_accumulate
pytorch_pfn_extras.training.extensions.
accumulate.max_accumulate
pytorch_pfn_extras.training.extensions.
accumulate.min_accumulate
pytorch_pfn_extras.training.extensions.
accumulate.standard_deviation_accumulate
pytorch_pfn_extras.training.
extensions.accumulate.
unbiased_standard_deviation_accumulate
```

pytorch_pfn_extras.training.extensions.accumulate.average_accumulate

Functions

```
pytorch_pfn_extras.training.extensions.
accumulate.average_accumulate.
all_gather_object(obj)
```

pytorch_pfn_extras.training.extensions.accumulate.average_accumulate.all_gather_object

`pytorch_pfn_extras.training.extensions.accumulate.average_accumulate.all_gather_object(obj)`

Parameters

obj (*T*) –

Return type

List[Optional[T]]

Classes

```
pytorch_pfn_extras.training.extensions.
accumulate.average_accumulate.
AccumulateBase(...)
pytorch_pfn_extras.training.extensions.
accumulate.average_accumulate.
AverageAccumulate(...)
pytorch_pfn_extras.training.extensions.
accumulate.average_accumulate.
AverageSummary()
pytorch_pfn_extras.training.extensions.
accumulate.average_accumulate.SummaryBase()
```

pytorch_pfn_extras.training.extensions.accumulate.average_accumulate.AccumulateBase

```
class pytorch_pfn_extras.training.extensions.accumulate.average_accumulate.AccumulateBase(conversion_key_L
trig-
ger=(1,
'epoch'),
dis-
tributed=False)
```

Bases: ABC, [Extension](#)

Methods

| | |
|--|--|
| <code>__init__(conversion_key_pair[, trigger, ...])</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before final- ization. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

Parameters

- **conversion_key_pair** (*Tuple*[*str*, *str*]) –
- **trigger** (*Optional*[*Union*[[Trigger](#), *Callable*[[[ExtensionsManagerProtocol](#)], *bool*], *Tuple*[*float*, *str*]]]) –
- **distributed** (*bool*) –

`__call__(manager)`

Call self as a function.

Parameters

manager ([ExtensionsManagerProtocol](#)) –

Return type
None

`__init__(conversion_key_pair, trigger=(1, 'epoch'), distributed=False)`

Parameters

- **conversion_key_pair** (*Tuple[str, str]*) –
- **trigger** (*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) –
- **distributed** (*bool*) –

Return type
None

`load_state_dict(to_load)`

Parameters
to_load (*Dict[str, Any]*) –

Return type
None

priority: `int = 200`

state_dict()
Serializes the extension state.

It is called when a manager that owns this extension is serialized. It serializes nothing by default.

Return type
Dict[str, Any]

pytorch_pfn_extras.training.extensions.accumulate.average_accumulate.AverageAccumulate

`class pytorch_pfn_extras.training.extensions.accumulate.average_accumulate.AverageAccumulate` (*conversion_k...*
trig-
ger=(1,
'epoch'),
dis-
tributed=False)

Bases: *AccumulateBase*

Methods

| | |
|--|--|
| <code>__init__(conversion_key_pair[, trigger, ...])</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before final- ization. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

Parameters

- **conversion_key_pair** (*Tuple*[*str*, *str*]) –
- **trigger** (*Optional*[*Union*[*Trigger*, *Callable*[[*ExtensionsManagerProtocol*], *bool*], *Tuple*[*float*, *str*]]]) –
- **distributed** (*bool*) –

pytorch_pfn_extras.training.extensions.accumulate.average_accumulate.AverageSummary

class `pytorch_pfn_extras.training.extensions.accumulate.average_accumulate.AverageSummary`

Bases: [*SummaryBase*](#)

Methods

| |
|---------------------------------------|
| <code>__init__()</code> |
| <code>add(value[, weight])</code> |
| <code>compute_accumulate()</code> |
| <code>compute_average()</code> |
| <code>load_state_dict(to_load)</code> |
| <code>state_dict()</code> |

`__init__()`

Return type

None

add(*value*, *weight=1*)

Parameters

- **value** (*Union[Tensor, ndarray, floating, float, Callable[[], float]]*) –
- **weight** (*Union[Tensor, ndarray, floating, float]*) –

Return type

None

compute_accumulate()

Return type

Union[Tensor, ndarray, floating, float]

compute_average()

Return type

Union[Tensor, ndarray, floating, float]

load_state_dict(*to_load*)

Parameters

to_load (*Dict[str, Any]*) –

Return type

None

state_dict()

Return type

Dict[str, Any]

pytorch_pfn_extras.training.extensions.accumulate.average_accumulate.SummaryBase

class pytorch_pfn_extras.training.extensions.accumulate.average_accumulate.**SummaryBase**

Bases: ABC

Methods

__init__()

add(*value*[, *weight*])

compute_accumulate()

load_state_dict(*to_load*)

state_dict()

__init__()

Return type

None

abstract add(*value*, *weight=1*)

Parameters

- **value** (*Union[[Tensor](#), [ndarray](#), [floating](#), [float](#), [Callable\[\[\[float\]\(#\)\], \[float\]\(#\)\]\]](#)) –*
- **weight** (*Union[[Tensor](#), [ndarray](#), [floating](#), [float](#)]*) –

Return type

None

abstract compute_accumulate()

Return type

Union[[Tensor](#), [ndarray](#), [floating](#), [float](#)]

abstract load_state_dict(*to_load*)

Parameters

to_load (*Dict[str, Any]*) –

Return type

None

abstract state_dict()

Return type

Dict[str, Any]

`pytorch_pfn_extras.training.extensions.accumulate.max_accumulate`

Functions

```
pytorch_pfn_extras.training.  
extensions.accumulate.max_accumulate.  
all_gather_object(obj)
```

`pytorch_pfn_extras.training.extensions.accumulate.max_accumulate.all_gather_object`

`pytorch_pfn_extras.training.extensions.accumulate.max_accumulate.all_gather_object(obj)`

Parameters

obj (*T*) –

Return type

List[Optional[T]]

Classes

```
pytorch_pfn_extras.training.
extensions.accumulate.max_accumulate.
AccumulateBase(...)
pytorch_pfn_extras.training.extensions.
accumulate.max_accumulate.MaxAccumulate(...)
pytorch_pfn_extras.training.extensions.
accumulate.max_accumulate.MaxSummary()
pytorch_pfn_extras.training.extensions.
accumulate.max_accumulate.SummaryBase()
```

pytorch_pfn_extras.training.extensions.accumulate.max_accumulate.AccumulateBase

```
class pytorch_pfn_extras.training.extensions.accumulate.max_accumulate.AccumulateBase(conversion_key_pair,
                                                                                       trigger=1,
                                                                                       'epoch'),
                                                                                       distributed=False)
```

Bases: ABC, *Extension*

Methods

```
__init__(conversion_key_pair[, trigger, ...])
```

| | |
|-----------------------------|-----------------------------------|
| <i>finalize</i> (manager) | Finalizes the extension. |
| <i>initialize</i> (manager) | Initializes up the manager state. |

```
load_state_dict(to_load)
```

| | |
|------------------------------------|---|
| <i>on_error</i> (manager, exc, tb) | Handles the error raised during training before finalization. |
|------------------------------------|---|

```
state_dict()
```

```
Serializes the extension state.
```

Attributes

| | |
|---------------------|--------------------------------|
| <i>default_name</i> | Default name of the extension. |
|---------------------|--------------------------------|

```
is_async
```

```
name
```

```
needs_model_state
```

```
priority
```

```
trigger
```

Parameters

- **conversion_key_pair** (*Tuple*[*str*, *str*]) –
- **trigger** (*Optional*[*Union*[*Trigger*, *Callable*[[*ExtensionsManagerProtocol*], *bool*], *Tuple*[*float*, *str*]]]) –
- **distributed** (*bool*) –

__call__ (*manager*)

Call self as a function.

Parameters**manager** (*ExtensionsManagerProtocol*) –**Return type**

None

__init__ (*conversion_key_pair*, *trigger*=(1, 'epoch'), *distributed*=False)**Parameters**

- **conversion_key_pair** (*Tuple*[*str*, *str*]) –
- **trigger** (*Optional*[*Union*[*Trigger*, *Callable*[[*ExtensionsManagerProtocol*], *bool*], *Tuple*[*float*, *str*]]]) –
- **distributed** (*bool*) –

Return type

None

load_state_dict (*to_load*)**Parameters****to_load** (*Dict*[*str*, *Any*]) –**Return type**

None

priority: *int* = 200**state_dict** ()

Serializes the extension state.

It is called when a manager that owns this extension is serialized. It serializes nothing by default.

Return type*Dict*[*str*, *Any*]**pytorch_pfn_extras.training.extensions.accumulate.max_accumulate.MaxAccumulate**

```
class pytorch_pfn_extras.training.extensions.accumulate.max_accumulate.MaxAccumulate(conversion_key_pair,  
                                                                                       trigger=(1,  
                                                                                       'epoch'),  
                                                                                       distributed=False)
```

Bases: *AccumulateBase*

Methods

| | |
|--|---|
| <code>__init__(conversion_key_pair[, trigger, ...])</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

Parameters

- **conversion_key_pair** (*Tuple*[*str*, *str*]) –
- **trigger** (*Optional*[*Union*[*Trigger*, *Callable*[[*ExtensionsManagerProtocol*], *bool*], *Tuple*[*float*, *str*]]]) –
- **distributed** (*bool*) –

pytorch_pfn_extras.training.extensions.accumulate.max_accumulate.MaxSummary

class pytorch_pfn_extras.training.extensions.accumulate.max_accumulate.**MaxSummary**

Bases: [SummaryBase](#)

Methods

`__init__()`

`add(value[, weight])`

`compute_accumulate()`

`compute_max()`

`load_state_dict(to_load)`

`state_dict()`

`__init__()`**Return type**

None

`add(value, weight=1)`**Parameters**

- **value** (*Union[[Tensor](#), [ndarray](#), [floating](#), [float](#), [Callable\[\[\[float\]\(#\)\], \[float\]\(#\)\]\]](#)*) –
- **weight** (*Union[[Tensor](#), [ndarray](#), [floating](#), [float](#)]*) –

Return type

None

`compute_accumulate()`**Return type***Union[[Tensor](#), [ndarray](#), [floating](#), [float](#)]*`compute_max()`**Return type***Union[[Tensor](#), [ndarray](#), [floating](#), [float](#)]*`load_state_dict(to_load)`**Parameters****to_load** (*Dict[str, Any]*) –**Return type**

None

`state_dict()`**Return type***Dict[str, Any]*

pytorch_pfn_extras.training.extensions.accumulate.max_accumulate.SummaryBase**class** pytorch_pfn_extras.training.extensions.accumulate.max_accumulate.**SummaryBase**

Bases: ABC

Methods*__init__()**add*(value[, weight])*compute_accumulate()**load_state_dict*(to_load)*state_dict()***__init__()****Return type**

None

abstract add(value, weight=1)**Parameters**

- **value** (*Union[[Tensor](#), [ndarray](#), [floating](#), [float](#), [Callable\[\[\[float\]\(#\)\], \[float\]\(#\)\]\]](#)*) –
- **weight** (*Union[[Tensor](#), [ndarray](#), [floating](#), [float](#)]*) –

Return type

None

abstract compute_accumulate()**Return type***Union[[Tensor](#), [ndarray](#), [floating](#), [float](#)]***abstract load_state_dict**(to_load)**Parameters****to_load** (*Dict[[str](#), [Any](#)]*) –**Return type**

None

abstract state_dict()**Return type***Dict[[str](#), [Any](#)]*

pytorch_pfn_extras.training.extensions.accumulate.min_accumulate**Functions**

```
pytorch_pfn_extras.training.  
extensions.accumulate.min_accumulate.  
all_gather_object(obj)
```

pytorch_pfn_extras.training.extensions.accumulate.min_accumulate.all_gather_object

```
pytorch_pfn_extras.training.extensions.accumulate.min_accumulate.all_gather_object(obj)
```

Parameters

obj (*T*) –

Return type

List[Optional[T]]

Classes

```
pytorch_pfn_extras.training.  
extensions.accumulate.min_accumulate.  
AccumulateBase(...)
```

```
pytorch_pfn_extras.training.extensions.  
accumulate.min_accumulate.MinAccumulate(...)
```

```
pytorch_pfn_extras.training.extensions.  
accumulate.min_accumulate.MinSummary()
```

```
pytorch_pfn_extras.training.extensions.  
accumulate.min_accumulate.SummaryBase()
```

pytorch_pfn_extras.training.extensions.accumulate.min_accumulate.AccumulateBase

```
class pytorch_pfn_extras.training.extensions.accumulate.min_accumulate.AccumulateBase(conversion_key_pair,  
trig-  
ger=(1,  
'epoch'),  
dis-  
tributed=False)
```

Bases: ABC, *Extension*

Methods

| | |
|--|---|
| <code>__init__(conversion_key_pair[, trigger, ...])</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

Parameters

- **conversion_key_pair** (*Tuple*[*str*, *str*]) –
- **trigger** (*Optional*[*Union*[*Trigger*, *Callable*[[*ExtensionsManagerProtocol*], *bool*], *Tuple*[*float*, *str*]]]) –
- **distributed** (*bool*) –

`__call__(manager)`

Call self as a function.

Parameters

manager (*ExtensionsManagerProtocol*) –

Return type

None

`__init__(conversion_key_pair, trigger=(1, 'epoch'), distributed=False)`

Parameters

- **conversion_key_pair** (*Tuple*[*str*, *str*]) –
- **trigger** (*Optional*[*Union*[*Trigger*, *Callable*[[*ExtensionsManagerProtocol*], *bool*], *Tuple*[*float*, *str*]]]) –
- **distributed** (*bool*) –

Return type

None

load_state_dict(*to_load*)**Parameters****to_load** (*Dict[str, Any]*) –**Return type**

None

priority: `int = 200`**state_dict**()

Serializes the extension state.

It is called when a manager that owns this extension is serialized. It serializes nothing by default.

Return type*Dict[str, Any]***pytorch_pfn_extras.training.extensions.accumulate.min_accumulate.MinAccumulate**

```
class pytorch_pfn_extras.training.extensions.accumulate.min_accumulate.MinAccumulate(conversion_key_pair,  
                                                                                       trigger=(1,  
                                                                                       'epoch'),  
                                                                                       distributed=False)
```

Bases: [*AccumulateBase*](#)**Methods**

| | |
|--|---|
| <hr/> | |
| <code>__init__(conversion_key_pair[, trigger, ...])</code> | |
| <hr/> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <hr/> | |
| <code>load_state_dict(to_load)</code> | |
| <hr/> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <hr/> | |
| <code>state_dict()</code> | Serializes the extension state. |
| <hr/> | |

Attributes

| | |
|-------------------|--------------------------------|
| default_name | Default name of the extension. |
| is_async | |
| name | |
| needs_model_state | |
| priority | |
| trigger | |

Parameters

- **conversion_key_pair** (*Tuple*[*str*, *str*]) –
- **trigger** (*Optional*[*Union*[*Trigger*, *Callable*[[*ExtensionsManagerProtocol*], *bool*], *Tuple*[*float*, *str*]]]) –
- **distributed** (*bool*) –

pytorch_pfn_extras.training.extensions.accumulate.min_accumulate.MinSummary

class pytorch_pfn_extras.training.extensions.accumulate.min_accumulate.MinSummary

Bases: *SummaryBase*

Methods

| |
|----------------------------------|
| <i>__init__</i> () |
| <i>add</i> (value[, weight]) |
| <i>compute_accumulate</i> () |
| <i>compute_min</i> () |
| <i>load_state_dict</i> (to_load) |
| <i>state_dict</i> () |

__init__()

Return type
None

add(*value*, *weight=1*)

Parameters

- **value** (*Union[[Tensor](#), [ndarray](#), [floating](#), [float](#), [Callable\[\[\[Tensor\]\(#\), \[float\]\(#\)\]\]](#)]*) –
- **weight** (*Union[[Tensor](#), [ndarray](#), [floating](#), [float](#)]*) –

Return type

None

compute_accumulate()

Return type

Union[[Tensor](#), [ndarray](#), [floating](#), [float](#)]

compute_min()

Return type

Union[[Tensor](#), [ndarray](#), [floating](#), [float](#)]

load_state_dict(*to_load*)

Parameters

to_load (*Dict[str, Any]*) –

Return type

None

state_dict()

Return type

Dict[str, Any]

pytorch_pfn_extras.training.extensions.accumulate.min_accumulate.SummaryBase

class `pytorch_pfn_extras.training.extensions.accumulate.min_accumulate.SummaryBase`

Bases: `ABC`

Methods

`__init__()`

`add(value[, weight])`

`compute_accumulate()`

`load_state_dict(to_load)`

`state_dict()`

__init__()

Return type

None

abstract add(*value*, *weight=1*)

Parameters

- **value** (*Union*[*Tensor*, *ndarray*, *floating*, *float*, *Callable*[[*float*]]]) –
- **weight** (*Union*[*Tensor*, *ndarray*, *floating*, *float*]) –

Return type

None

abstract compute_accumulate()

Return type

Union[*Tensor*, *ndarray*, *floating*, *float*]

abstract load_state_dict(*to_load*)

Parameters

to_load (*Dict*[*str*, *Any*]) –

Return type

None

abstract state_dict()

Return type

Dict[*str*, *Any*]

pytorch_pfn_extras.training.extensions.accumulate.standard_deviation_accumulate

Functions

```
pytorch_pfn_extras.training.extensions.
accumulate.standard_deviation_accumulate.
all_gather_object(obj)
```

pytorch_pfn_extras.training.extensions.accumulate.standard_deviation_accumulate.all_gather_object

pytorch_pfn_extras.training.extensions.accumulate.standard_deviation_accumulate.all_gather_object(*obj*)

Parameters

obj (*T*) –

Return type

List[*Optional*[*T*]]

Classes

| |
|---|
| <code>pytorch_pfn_extras.training.extensions.accumulate.standard_deviation_accumulate.AccumulateBase(...)</code> |
| <code>pytorch_pfn_extras.training.extensions.accumulate.standard_deviation_accumulate.StandardDeviationAccumulate(...)</code> |
| <code>pytorch_pfn_extras.training.extensions.accumulate.standard_deviation_accumulate.StandardDeviationSummary()</code> |
| <code>pytorch_pfn_extras.training.extensions.accumulate.standard_deviation_accumulate.SummaryBase()</code> |

pytorch_pfn_extras.training.extensions.accumulate.standard_deviation_accumulate.AccumulateBase

`class pytorch_pfn_extras.training.extensions.accumulate.standard_deviation_accumulate.AccumulateBase(conversion_key_pair, trigger, generator, epoch_trigger, distributed_trigger)`

Bases: `ABC`, `Extension`

Methods

| | |
|--|---|
| <code>__init__(conversion_key_pair[, trigger, ...])</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

Parameters

- **conversion_key_pair** (`Tuple[str, str]`) –
- **trigger** (`Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]`) –
- **distributed** (`bool`) –

__call__ (*manager*)

Call self as a function.

Parameters

manager (`ExtensionsManagerProtocol`) –

Return type

None

__init__ (*conversion_key_pair*, *trigger*=(1, 'epoch'), *distributed*=False)

Parameters

- **conversion_key_pair** (`Tuple[str, str]`) –
- **trigger** (`Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]`) –
- **distributed** (`bool`) –

Return type

None

load_state_dict (*to_load*)

Parameters

to_load (`Dict[str, Any]`) –

Return type

None

priority: `int` = 200

state_dict ()

Serializes the extension state.

It is called when a manager that owns this extension is serialized. It serializes nothing by default.

Return type
Dict[str, Any]

`pytorch_pfn_extras.training.extensions.accumulate.standard_deviation_accumulate.StandardDeviationAccumulate`

`class pytorch_pfn_extras.training.extensions.accumulate.standard_deviation_accumulate.StandardDeviationAccumulate`

Bases: *AccumulateBase*

Methods

| | |
|--|--|
| <code>__init__(conversion_key_pair[, trigger, ...])</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before final- ization. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

Parameters

- `conversion_key_pair` (*Tuple[str, str]*) –
- `trigger` (*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) –
- `distributed` (*bool*) –

pytorch_pfn_extras.training.extensions.accumulate.standard_deviation_accumulate.StandardDeviationSummary

class pytorch_pfn_extras.training.extensions.accumulate.standard_deviation_accumulate.
StandardDeviationSummary

Bases: *SummaryBase*

Methods

__init__()

add(value[, weight])

compute_accumulate()

compute_mean()

compute_standard_deviation()

load_state_dict(to_load)

state_dict()

__init__()

Return type

None

add(value, weight=1)

Parameters

- **value** (*Union[[Tensor](#), [ndarray](#), [floating](#), [float](#), [Callable\[\[\[float\]\(#\)\], \[float\]\(#\)\]\]](#)*) –
- **weight** (*Union[[Tensor](#), [ndarray](#), [floating](#), [float](#)]*) –

Return type

None

compute_accumulate()

Return type

Union[[Tensor](#), [ndarray](#), [floating](#), [float](#)]

compute_mean()

Return type

Union[[Tensor](#), [ndarray](#), [floating](#), [float](#)]

compute_standard_deviation()

Return type

Union[[Tensor](#), [ndarray](#), [floating](#), [float](#)]

load_state_dict(*to_load*)

Parameters

to_load (*Dict[str, Any]*) –

Return type

None

state_dict()

Return type

Dict[str, Any]

pytorch_pfn_extras.training.extensions.accumulate.standard_deviation_accumulate.SummaryBase

class pytorch_pfn_extras.training.extensions.accumulate.standard_deviation_accumulate.
SummaryBase

Bases: ABC

Methods

__init__()

add(value[, weight])

compute_accumulate()

load_state_dict(to_load)

state_dict()

__init__()

Return type

None

abstract add(*value*, *weight=1*)

Parameters

- **value** (*Union[Tensor, ndarray, floating, float, Callable[[], float]]*) –
- **weight** (*Union[Tensor, ndarray, floating, float]*) –

Return type

None

abstract compute_accumulate()

Return type

Union[Tensor, ndarray, floating, float]

abstract `load_state_dict(to_load)`

Parameters

`to_load` (*Dict[str, Any]*) –

Return type

None

abstract `state_dict()`

Return type

Dict[str, Any]

`pytorch_pfn_extras.training.extensions.accumulate.unbiased_standard_deviation_accumulate`

Functions

```
pytorch_pfn_extras.training.
extensions.accumulate.
unbiased_standard_deviation_accumulate.
all_gather_object(obj)
```

`pytorch_pfn_extras.training.extensions.accumulate.unbiased_standard_deviation_accumulate.all_gather_object`

`pytorch_pfn_extras.training.extensions.accumulate.unbiased_standard_deviation_accumulate.all_gather_obj`

Parameters

`obj` (*T*) –

Return type

List[Optional[T]]

Classes

```
pytorch_pfn_extras.training.
extensions.accumulate.
unbiased_standard_deviation_accumulate.
AccumulateBase(...)
```

```
pytorch_pfn_extras.training.
extensions.accumulate.
unbiased_standard_deviation_accumulate.
SummaryBase()
```

```
pytorch_pfn_extras.training.
extensions.accumulate.
unbiased_standard_deviation_accumulate.
UnbiasedStandardDeviationAccumulate(...)
```

```
pytorch_pfn_extras.training.
extensions.accumulate.
unbiased_standard_deviation_accumulate.
UnbiasedStandardDeviationSummary()
```

`pytorch_pfn_extras.training.extensions.accumulate.unbiased_standard_deviation_accumulate.AccumulateBase`

`class pytorch_pfn_extras.training.extensions.accumulate.unbiased_standard_deviation_accumulate.Accumula`

Bases: `ABC`, `Extension`

Methods

| | |
|--|---|
| <code>__init__(conversion_key_pair[, trigger, ...])</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

Parameters

- **conversion_key_pair** (`Tuple[str, str]`) –
- **trigger** (`Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]`) –
- **distributed** (`bool`) –

`__call__(manager)`

Call self as a function.

Parameters

manager (`ExtensionsManagerProtocol`) –

Return type

None

`__init__(conversion_key_pair, trigger=(1, 'epoch'), distributed=False)`**Parameters**

- **conversion_key_pair** (`Tuple[str, str]`) –
- **trigger** (`Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]`) –
- **distributed** (`bool`) –

Return type

None

`load_state_dict(to_load)`**Parameters**`to_load` (`Dict[str, Any]`) –**Return type**

None

`priority: int = 200``state_dict()`

Serializes the extension state.

It is called when a manager that owns this extension is serialized. It serializes nothing by default.

Return type`Dict[str, Any]`**pytorch_pfn_extras.training.extensions.accumulate.unbiased_standard_deviation_accumulate.SummaryBase**

class `pytorch_pfn_extras.training.extensions.accumulate.unbiased_standard_deviation_accumulate.SummaryBase`

Bases: `ABC`**Methods**`__init__()``add(value[, weight])``compute_accumulate()``load_state_dict(to_load)``state_dict()`

__init__()

Return type

None

abstract add(*value*, *weight=1*)

Parameters

- **value** (*Union[[Tensor](#), [ndarray](#), [floating](#), [float](#), [Callable\[\[\[float\]\(#\)\], \[float\]\(#\)\]\]](#)) –*
- **weight** (*Union[[Tensor](#), [ndarray](#), [floating](#), [float](#)]*) –

Return type

None

abstract compute_accumulate()

Return type

Union[[Tensor](#), [ndarray](#), [floating](#), [float](#)]

abstract load_state_dict(*to_load*)

Parameters

to_load (*Dict[[str](#), [Any](#)]*) –

Return type

None

abstract state_dict()

Return type

Dict[[str](#), [Any](#)]

[pytorch_pfn_extras.training.extensions.accumulate.unbiased_standard_deviation_accumulate.UnbiasedStandardDeviationAccumulate](#)

`class pytorch_pfn_extras.training.extensions.accumulate.unbiased_standard_deviation_accumulate.UnbiasedStandardDeviationAccumulate`

Bases: [AccumulateBase](#)

Methods

| | |
|--|---|
| <hr/> | |
| <code>__init__(conversion_key_pair[, trigger, ...])</code> | |
| <hr/> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <hr/> | |
| <code>load_state_dict(to_load)</code> | |
| <hr/> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <hr/> | |
| <code>state_dict()</code> | Serializes the extension state. |
| <hr/> | |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

Parameters

- **conversion_key_pair** (*Tuple*[*str*, *str*]) –
- **trigger** (*Optional*[*Union*[*Trigger*, *Callable*[[*ExtensionsManagerProtocol*], *bool*], *Tuple*[*float*, *str*]]]) –
- **distributed** (*bool*) –

`pytorch_pfn_extras.training.extensions.accumulate.unbiased_standard_deviation_accumulate.UnbiasedStandardDeviationSummary`

`class pytorch_pfn_extras.training.extensions.accumulate.unbiased_standard_deviation_accumulate.UnbiasedStandardDeviationSummary`

Bases: *SummaryBase*

Methods

| |
|--|
| <code>__init__()</code> |
| <code>add(value[, weight])</code> |
| <code>compute_accumulate()</code> |
| <code>compute_mean()</code> |
| <code>compute_unbiased_standard_deviation()</code> |
| <code>load_state_dict(to_load)</code> |
| <code>state_dict()</code> |
| <code>__init__()</code> |

Return type

None

add(*value*, *weight=1*)**Parameters**

- **value** (*Union[[Tensor](#), [ndarray](#), [floating](#), [float](#), [Callable\[\[\[float\]\(#\)\], \[float\]\(#\)\]\]](#)) –*
- **weight** (*Union[[Tensor](#), [ndarray](#), [floating](#), [float](#)]*) –

Return type

None

compute_accumulate()**Return type***Union[[Tensor](#), [ndarray](#), [floating](#), [float](#)]***compute_mean**()**Return type***Union[[Tensor](#), [ndarray](#), [floating](#), [float](#)]***compute_unbiased_standard_deviation**()**Return type***Union[[Tensor](#), [ndarray](#), [floating](#), [float](#)]***load_state_dict**(*to_load*)**Parameters****to_load** (*Dict[str, Any]*) –**Return type**

None

state_dict()**Return type***Dict[str, Any]***pytorch_pfn_extras.training.extensions.best_value****Classes**

| | |
|---|--|
| <code>pytorch_pfn_extras.training.extensions.best_value.BestValue(...)</code> | Extension traces the best value of a specific key in the observation. |
| <code>pytorch_pfn_extras.training.extensions.best_value.ExtensionsManagerProtocol(...)</code> | |
| <code>pytorch_pfn_extras.training.extensions.best_value.MaxValue(key)</code> | Extension traces the maximum value of a specific key in the observation. |
| <code>pytorch_pfn_extras.training.extensions.best_value.MinValue(key)</code> | Extension traces the maximum value of a specific key in the observation. |

pytorch_pfn_extras.training.extensions.best_value.BestValue

```
class pytorch_pfn_extras.training.extensions.best_value.BestValue(key, compare, trigger=(1, 'epoch'))
```

Bases: [*Extension*](#)

Extension traces the best value of a specific key in the observation.

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** (*TriggerLike*) – 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 `BestValueTrigger`.

Methods

| | |
|--|---|
| <code>__init__(key, compare[, trigger])</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--|
| <code>best_epoch</code> | Returns the epoch count that the current best value is observed. |
| <code>best_iteration</code> | Returns the iteration count that the current best value is observed. |
| <code>best_value</code> | Returns the current best value. |
| <code>default_name</code> | |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

__call__(*manager*)

Invokes the extension.

Implementations should override this operator. This method is called at iterations which the corresponding trigger accepts.

Parameters

manager (*ExtensionsManager*) – Manager object to call this operator.

Return type

None

__init__(*key, compare, trigger=(1, 'epoch')*)

Parameters

- **key** (*str*) –
- **compare** (*Callable[[float, float], bool]*) –
- **trigger** (*TriggerLike*) –

Return type

None

property best_epoch: int

Returns the epoch count that the current best value is observed.

If no value has been observed yet, it raises a `RuntimeError`.

property best_iteration: int

Returns the iteration count that the current best value is observed.

If no value has been observed yet, it raises a `RuntimeError`.

property best_value: float

Returns the current best value.

If no value has been observed yet, it raises a `RuntimeError`.

default_name = 'best_value'

load_state_dict(*to_load*)

Parameters

to_load (*Dict[str, Any]*) –

Return type

None

state_dict()

Serializes the extension state.

It is called when a manager that owns this extension is serialized. It serializes nothing by default.

Return type

Dict[str, Any]

pytorch_pfn_extras.training.extensions.best_value.ExtensionsManagerProtocol

```
class pytorch_pfn_extras.training.extensions.best_value.ExtensionsManagerProtocol(*args,
                                                                                    **kwargs)
```

Bases: Protocol

Methods

```
__init__(*args, **kwargs)
```

```
get_extension(name)
```

Attributes

```
elapsed_time
```

```
epoch
```

```
epoch_detail
```

```
is_before_training
```

```
iteration
```

```
models
```

```
observation
```

```
optimizers
```

```
out
```

```
raw_models
```

```
reporter
```

```
stop_trigger
```

```
writer
```

```
__init__(*args, **kwargs)
```

```
property elapsed_time: float
```

```
property epoch: int
```

```
property epoch_detail: float
```

`get_extension(name)`

Parameters

name (*str*) –

Return type

Extension

property `is_before_training`: `bool`

property `iteration`: `int`

property `models`: `Mapping[str, Module]`

property `observation`: `reporting.Observation`

property `optimizers`: `Mapping[str, Optimizer]`

property `out`: `str`

property `raw_models`: `Mapping[str, Module]`

property `reporter`: *reporting.Reporter*

property `stop_trigger`: `bool`

property `writer`: `Optional[writing.Writer]`

`pytorch_pfn_extras.training.extensions.best_value.MaxValue`

`class pytorch_pfn_extras.training.extensions.best_value.MaxValue(key, trigger=(1, 'epoch'))`

Bases: *BestValue*

Extension traces the maximum value of a specific key in the observation.

Parameters

- **key** (*str*) – Key of value.
- **trigger** (*TriggerLike*) – Trigger that decides the comparison interval between current maximum value and new value. This must be a tuple in the form of `<int>, 'epoch'` or `<int>, 'iteration'` which is passed to `BestValueTrigger`.

Methods

| | |
|---|---|
| <code>__init__(key[, trigger])</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| | |
|--|--|
| <code>best_epoch</code> | Returns the epoch count that the current best value is observed. |
| <code>best_iteration</code> | Returns the iteration count that the current best value is observed. |
| <code>best_value</code> <i>default_name</i> | Returns the current best value. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

`__init__(key, trigger=(1, 'epoch'))`

Parameters

- **key** (*str*) –
- **trigger** (*TriggerLike*) –

`default_name = 'max_value'`

pytorch_pfn_extras.training.extensions.best_value.MinValue

class `pytorch_pfn_extras.training.extensions.best_value.MinValue`(*key, trigger=(1, 'epoch')*)

Bases: *BestValue*

Extension traces the maximum value of a specific key in the observation.

Parameters

- **key** (*str*) – Key of value.
- **trigger** (*TriggerLike*) – Trigger that decides the comparison interval between current maximum value and new value. This must be a tuple in the form of `<int>, 'epoch'` or `<int>, 'iteration'` which is passed to `BestValueTrigger`.

Methods

| | |
|---|---|
| <code>__init__(key[, trigger])</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--|
| <code>best_epoch</code> | Returns the epoch count that the current best value is observed. |
| <code>best_iteration</code> | Returns the iteration count that the current best value is observed. |
| <code>best_value</code> | Returns the current best value. |
| <code>default_name</code> | |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

`__init__(key, trigger=(1, 'epoch'))`

Parameters

- **key** (*str*) –
- **trigger** (*TriggerLike*) –

`default_name = 'min_value'`

pytorch_pfn_extras.training.extensions.evaluator**Classes**

| | |
|--|--|
| <code>pytorch_pfn_extras.training.extensions.evaluator.DistributedEvaluator(...)</code> | An extension to evaluate models on a validation set in a distributed training setup. |
| <code>pytorch_pfn_extras.training.extensions.evaluator.Evaluator(...)</code> | An extension to evaluate models on a validation set. |
| <code>pytorch_pfn_extras.training.extensions.evaluator.ExtensionsManagerProtocol(...)</code> | |
| <code>pytorch_pfn_extras.training.extensions.evaluator.IgniteEvaluator(...)</code> | |
| <code>pytorch_pfn_extras.training.extensions.evaluator.IterationStatus(size)</code> | |
| <code>pytorch_pfn_extras.training.extensions.evaluator.TextIO(...)</code> | Typed version of the return of <code>open()</code> in text mode. |

pytorch_pfn_extras.training.extensions.evaluator.DistributedEvaluator

```
class pytorch_pfn_extras.training.extensions.evaluator.DistributedEvaluator(self, iterator,
                                                                    target,
                                                                    eval_func=None,
                                                                    *,
                                                                    progress_bar=False)
```

Bases: [*Evaluator*](#)

An extension to evaluate models on a validation set in a distributed training setup.

In case `torch.distributed` is used to parallelize training iterations, it is efficient to also run evaluation in parallel by splitting the validation set to each worker process and conduct evaluation separately followed by aggregation of results of each worker, which can be achieved by `:class:`~DistributedEvaluator``.

This extension basically behaves similarly to [*Evaluator*](#), but adds an aggregation step in [*Evaluator.evaluate\(\)*](#). A summary of evaluation (`DictSummary`) in each worker process is collected in “all-gather” manner and then accumulated. Therefore all the worker processes must attend the evaluation, i.e., make sure all the processes have a [*Evaluator*](#) extension object configured in the `ExtensionManager` with the same trigger. All the worker process will get identical evaluation result returned by [*Evaluator.evaluate\(\)*](#) and reported to an observation.

It is necessary to pass a `DataLoader` with an appropriate sampler which properly splits the validation dataset to each MPI worker process. PyTorch `DistributedSampler` implements this, but it allows sampler repetition in order to make the number of samples assigned to each process identical. For evaluation purpose it distorts the evaluation result, hence it is recommended to use `DistributedValidationSampler` instead.

Methods

| | |
|---|---|
| <code>__init__(iterator, target[, eval_hook, ...])</code> | |
| <code>add_metric(metric_fn)</code> | Adds a custom metric to the evaluator. |
| <code>eval_func(*args, **kwargs)</code> | |
| <code>evaluate()</code> | Evaluates the model and returns a result dictionary. |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>get_all_iterators()</code> | Returns a dictionary of all iterators. |
| <code>get_all_targets()</code> | Returns a dictionary of all target links. |
| <code>get_iterator(name)</code> | Returns the iterator of the given name. |
| <code>get_target(name)</code> | Returns the target link of the given name. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| |
|--------------------------------|
| <code>default_name</code> |
| <code>is_async</code> |
| <code>name</code> |
| <code>needs_model_state</code> |
| <code>priority</code> |
| <code>trigger</code> |

Parameters

- **iterator** (`Union[DataLoader[Any], Dict[str, DataLoader[Any]]]`) –
- **target** (`Union[Module, Dict[str, Module]]`) –
- **eval_hook** (`Optional[Callable[[Evaluator], None]]`) –
- **eval_func** (`Optional[Callable[..., Any]]`) –
- **kwargs** (`Any`) –

`__init__(iterator, target, eval_hook=None, eval_func=None, **kwargs)`

Parameters

- **iterator** (`Union[DataLoader[Any], Dict[str, DataLoader[Any]]]`) –
- **target** (`Union[Module, Dict[str, Module]]`) –

- **eval_hook** (*Optional*[Callable[[[Evaluator](#)], None]]) –
- **eval_func** (*Optional*[Callable[[...], Any]]) –
- **kwargs** (Any) –

Return type

None

pytorch_pfn_extras.training.extensions.evaluator.Evaluator

```
class pytorch_pfn_extras.training.extensions.evaluator.Evaluator(self, iterator, target,
                                                                eval_func=None, *,
                                                                progress_bar=False)
```

Bases: [Extension](#)

An 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 manager, 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

This extension is called at the end of each epoch by default.

Parameters

- **iterator** (*Union*[[DataLoader](#)[Any], *Dict*[str, [DataLoader](#)[Any]]]) – 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** (*Union*[*Module*, *Dict*[str, *Module*]]) – `torch.nn.Module` object or a dictionary of links to evaluate. If this is just a layer object, the link is registered by the name 'main'.
- **eval_func** (*Optional*[Callable[[...], Any]]) – Evaluation function called at each iteration. The target link to evaluate as a callable is used by default.
- **progress_bar** – Boolean flag to show a progress bar while training, which is similar to [ProgressBar](#). (default: False)
- **metrics** – List of callables that are called every batch to calculate metrics such as accuracy, roc_auc or others The signature of the callable is: `def metric_fn(batch, output, last_iteration)` (default: [])
- **eval_hook** (*Optional*[Callable[[[Evaluator](#)], None]]) –

- **kwargs** (*Any*) –

Warning: The argument `progress_bar` is experimental. The interface can change in the future.

eval_hook

Function to prepare for each evaluation process.

eval_func

Evaluation function called at each iteration.

Parameters

- **args** (*Any*) –
- **kwargs** (*Any*) –

Return type

Any

Methods

| | |
|---|---|
| <code>__init__(iterator, target[, eval_hook, ...])</code> | |
| <code>add_metric(metric_fn)</code> | Adds a custom metric to the evaluator. |
| <code>eval_func(*args, **kwargs)</code> | |
| <code>evaluate()</code> | Evaluates the model and returns a result dictionary. |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>get_all_iterators()</code> | Returns a dictionary of all iterators. |
| <code>get_all_targets()</code> | Returns a dictionary of all target links. |
| <code>get_iterator(name)</code> | Returns the iterator of the given name. |
| <code>get_target(name)</code> | Returns the target link of the given name. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

default_name

is_async

name

needs_model_state

priority

trigger

__call__(*manager=None*)

Executes the evaluator extension.

Unlike usual extensions, this extension can be executed without passing a manager object. This extension reports the performance on validation dataset using the `report()` function. Thus, users can use this extension independently from any manager by manually configuring a `Reporter` object.

Parameters

manager (`ExtensionsManager`) – Manager object that invokes this extension.

Returns

Result dictionary that contains mean statistics of values reported by the evaluation function.

Return type

dict

__init__(*iterator, target, eval_hook=None, eval_func=None, **kwargs*)

Parameters

- **iterator** (`Union[DataLoader[Any], Dict[str, DataLoader[Any]]]`) –
- **target** (`Union[Module, Dict[str, Module]]`) –
- **eval_hook** (`Optional[Callable[[Evaluator], None]]`) –
- **eval_func** (`Optional[Callable[[...], Any]]`) –
- **kwargs** (`Any`) –

Return type

None

add_metric(*metric_fn*)

Adds a custom metric to the evaluator.

The metric is a callable that is executed every batch with the following signature: `def metric_fn(batch, output, last_iteration)`

Batch is the input batch passed to the model. Output is the result of evaluating batch, last_iteration is a boolean flag that indicates if its the last batch in the evaluation.

Parameters

metric_fn (`Callable[[Any, Any, Any], None]`) –

Return type

None

default_name = 'validation'**eval_func**(*args, **kwargs)**Parameters**

- **args** (*Any*) –
- **kwargs** (*Any*) –

Return type*Any***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.

Users can override this method to customize the evaluation routine.

Returns

Result dictionary. This dictionary is further reported via `report()` without specifying any observer.

Return type

dict

get_all_iterators()

Returns a dictionary of all iterators.

Return type*Dict*[str, [DataLoader](#)[*Any*]]**get_all_targets()**

Returns a dictionary of all target links.

Return type*Dict*[str, *Module*]**get_iterator**(*name*)

Returns the iterator of the given name.

Parameters**name** (*str*) –**Return type**[DataLoader](#)[*Any*]**get_target**(*name*)

Returns the target link of the given name.

Parameters**name** (*str*) –**Return type***Module***priority**: int = 300**trigger**: `TriggerLike` = (1, 'epoch')

pytorch_pfn_extras.training.extensions.evaluator.ExtensionsManagerProtocol

```
class pytorch_pfn_extras.training.extensions.evaluator.ExtensionsManagerProtocol(*args,
                                                                                **kwargs)
```

Bases: Protocol

Methods

```
__init__(*args, **kwargs)
```

```
get_extension(name)
```

Attributes

```
elapsed_time
```

```
epoch
```

```
epoch_detail
```

```
is_before_training
```

```
iteration
```

```
models
```

```
observation
```

```
optimizers
```

```
out
```

```
raw_models
```

```
reporter
```

```
stop_trigger
```

```
writer
```

```
__init__(*args, **kwargs)
```

```
property elapsed_time: float
```

```
property epoch: int
```

```
property epoch_detail: float
```

`get_extension(name)`

Parameters

name (*str*) –

Return type

Extension

`property is_before_training:` `bool`

`property iteration:` `int`

`property models:` `Mapping[str, Module]`

`property observation:` `reporting.Observation`

`property optimizers:` `Mapping[str, Optimizer]`

`property out:` `str`

`property raw_models:` `Mapping[str, Module]`

`property reporter:` `reporting.Reporter`

`property stop_trigger:` `bool`

`property writer:` `Optional[writing.Writer]`

`pytorch_pfn_extras.training.extensions.evaluator.IgniteEvaluator`

`class pytorch_pfn_extras.training.extensions.evaluator.IgniteEvaluator` (*evaluator, iterator, target, **kwargs*)

Bases: `Evaluator`

Methods

| | |
|--|---|
| <code>__init__(evaluator, iterator, target, **kwargs)</code> | |
| <code>add_metric(metric_fn)</code> | Adds a custom metric to the evaluator. |
| <code>eval_func(*args, **kwargs)</code> | |
| <code>evaluate()</code> | Evaluates the model and returns a result dictionary. |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>get_all_iterators()</code> | Returns a dictionary of all iterators. |
| <code>get_all_targets()</code> | Returns a dictionary of all target links. |
| <code>get_iterator(name)</code> | Returns the iterator of the given name. |
| <code>get_target(name)</code> | Returns the target link of the given name. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>set_evaluator_handlers()</code> | |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

default_name

is_async

name

needs_model_state

priority

trigger

Parameters

- **evaluator** (*Engine*) –
- **iterator** (*Union[DataLoader[Any], Dict[str, DataLoader[Any]]]*) –
- **target** (*Union[Module, Dict[str, Module]]*) –
- **kwargs** (*Any*) –

__init__(*evaluator, iterator, target, **kwargs*)**Parameters**

- **evaluator** (*Engine*) –
- **iterator** (*Union[DataLoader[Any], Dict[str, DataLoader[Any]]]*) –
- **target** (*Union[Module, Dict[str, Module]]*) –
- **kwargs** (*Any*) –

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.

Users can override this method to customize the evaluation routine.

Returns

Result dictionary. This dictionary is further reported via `report()` without specifying any observer.

Return type

dict

set_evaluator_handlers()**Return type**

None

pytorch_pfn_extras.training.extensions.evaluator.IterationStatus**class** pytorch_pfn_extras.training.extensions.evaluator.**IterationStatus**(*size*)

Bases: object

Methods

`__init__`(*size*)

Attributes

`epoch_detail`

Parameters**size** (*int*) –`__init__`(*size*)**Parameters****size** (*int*) –**Return type**

None

property epoch_detail: float**pytorch_pfn_extras.training.extensions.evaluator.TextIO****class** pytorch_pfn_extras.training.extensions.evaluator.**TextIO**(*args, **kws)Bases: [IO](#)[str]

Typed version of the return of open() in text mode.

Methods

`__init__()`

`close()`

`fileno()`

`flush()`

`isatty()`

`read([n])`

`readable()`

`readline([limit])`

`readlines([hint])`

`seek(offset[, whence])`

`seekable()`

`tell()`

`truncate([size])`

`writable()`

`write(s)`

`writelines(lines)`

Attributes

| |
|-----------------------|
| <i>buffer</i> |
| closed |
| <i>encoding</i> |
| <i>errors</i> |
| <i>line_buffering</i> |
| mode |
| name |
| <i>newlines</i> |

```
abstract property buffer: BinaryIO
abstract property encoding: str
abstract property errors: Optional[str]
abstract property line_buffering: bool
abstract property newlines: Any
```

pytorch_pfn_extras.training.extensions.fail_on_non_number

Classes

| | |
|---|--|
| <i>pytorch_pfn_extras.training.extensions.fail_on_non_number.ExtensionsManagerProtocol(...)</i> | |
| <i>pytorch_pfn_extras.training.extensions.fail_on_non_number.FailOnNonNumber(*)</i> | An extension to raise RuntimeError if parameters and its gradients contain NaN or Inf. |

pytorch_pfn_extras.training.extensions.fail_on_non_number.ExtensionsManagerProtocol

```
class pytorch_pfn_extras.training.extensions.fail_on_non_number.ExtensionsManagerProtocol(*args,
                                                                                          **kwargs)
    Bases: Protocol
```


Methods

`__init__(*args, **kwargs)`

`get_extension(name)`

Attributes

`elapsed_time`

`epoch`

`epoch_detail`

`is_before_training`

`iteration`

`models`

`observation`

`optimizers`

`out`

`raw_models`

`reporter`

`stop_trigger`

`writer`

`__init__(*args, **kwargs)``property elapsed_time: float``property epoch: int``property epoch_detail: float``get_extension(name)`

Parameters

name (*str*) –

Return type

Extension

```
property is_before_training: bool
property iteration: int
property models: Mapping[str, Module]
property observation: reporting.Observation
property optimizers: Mapping[str, Optimizer]
property out: str
property raw_models: Mapping[str, Module]
property reporter: reporting.Reporter
property stop_trigger: bool
property writer: Optional[writing.Writer]
```

pytorch_pfn_extras.training.extensions.fail_on_non_number.FailOnNonNumber

```
class pytorch_pfn_extras.training.extensions.fail_on_non_number.FailOnNonNumber(*,
                                                                                check_grad=True)
```

Bases: *Extension*

An extension to raise `RuntimeError` if parameters and its gradients contain NaN or Inf.

Although parameters including non-number such as NaN and Inf are unnecessary in most cases the training loop 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.

Parameters

check_grad (*bool*) – Set to False to skip checking gradients.

Methods

| | |
|--|---|
| <code>__init__</code> (*[, check_grad]) | |
| <code>finalize</code> (manager) | Finalizes the extension. |
| <code>initialize</code> (manager) | Initializes up the manager state. |
| <code>load_state_dict</code> (to_load) | |
| <code>on_error</code> (manager, exc, tb) | Handles the error raised during training before finalization. |
| <code>state_dict</code> () | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

`__call__`(*manager*)

Invokes the extension.

Implementations should override this operator. This method is called at iterations which the corresponding trigger accepts.

Parameters

manager ([ExtensionsManager](#)) – Manager object to call this operator.

Return type

None

`__init__`(***, *check_grad=True*)

Parameters

check_grad (*bool*) –

needs_model_state = True

pytorch_pfn_extras.training.extensions.log_report

Classes

| | |
|---|---|
| <code>pytorch_pfn_extras.training.extensions.log_report.ExtensionsManagerProtocol(...)</code> | |
| <code>pytorch_pfn_extras.training.extensions.log_report.LogReport([...])</code> | An extension to output the accumulated results to a log file. |
| <code>pytorch_pfn_extras.training.extensions.log_report.LogWriterSaveFunc(...)</code> | |

pytorch_pfn_extras.training.extensions.log_report.ExtensionsManagerProtocol

```
class pytorch_pfn_extras.training.extensions.log_report.ExtensionsManagerProtocol(*args,  
                                                                                **kwargs)
```

Bases: Protocol

Methods

```
__init__(*args, **kwargs)
```

```
get_extension(name)
```

Attributes

```
elapsed_time
```

```
epoch
```

```
epoch_detail
```

```
is_before_training
```

```
iteration
```

```
models
```

```
observation
```

```
optimizers
```

```
out
```

```
raw_models
```

```
reporter
```

```
stop_trigger
```

```
writer
```

```
__init__(*args, **kwargs)
```

```
property elapsed_time: float
```

```
property epoch: int
```

```
property epoch_detail: float
```

`get_extension(name)`

Parameters

`name` (*str*) –

Return type

Extension

`property is_before_training: bool`

`property iteration: int`

`property models: Mapping[str, Module]`

`property observation: reporting.Observation`

`property optimizers: Mapping[str, Optimizer]`

`property out: str`

`property raw_models: Mapping[str, Module]`

`property reporter: reporting.Reporter`

`property stop_trigger: bool`

`property writer: Optional[writing.Writer]`

`pytorch_pfn_extras.training.extensions.log_report.LogReport`

```
class pytorch_pfn_extras.training.extensions.log_report.LogReport(keys=None, trigger=(1,
                                          'epoch'), postprocess=None,
                                          filename=None,
                                          append=False, format=None,
                                          **kwargs)
```

Bases: [Extension](#)

An extension to output the accumulated results to a log file.

This extension accumulates the observations of the manager 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 `ExtensionsManager.elapsed_time`.

Parameters

- `keys` (*iterable of str*) – Keys of values to accumulate. If this is None, all the values are accumulated and output to the log file.

- **trigger** (*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) – 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** (*Optional[Callable[[Mapping[str, Any]], None]]*) – 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.
- **filename** (*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. (default: `'log'`)
- **append** (*bool, optional*) – If the file is JSON Lines or YAML, contents will be appended instead of rewriting the file every call.
- **format** (*str, optional*) – accepted values are `'json'`, `'json-lines'` and `'yaml'`.
- **writer** (*writer object, optional*) – must be callable. object to dump the log to. If specified, it needs to have a correct `savefun` defined. The writer can override the save location in the `pytorch_pfn_extras.training.ExtensionsManager` object
- **kwargs** (*Any*) –

Note: Enabling `append=True` reduces size of snapshots (and thus reduces the time needed to take snapshots). Note that extensions relying on the logs from past iterations may behave differently; for example, when resuming from a snapshot, `PrintReport` will not show logs of iterations already done.

Methods

| | |
|--|---|
| <code>__init__([keys, trigger, postprocess, ...])</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>state_dict()</code> | Serializes the extension state. |
| <code>to_dataframe()</code> | |

Attributes

| | |
|--------------------------------|---|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>log</code> | The current list of observation dictionaries. |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

`__call__`(*manager*)

Invokes the extension.

Implementations should override this operator. This method is called at iterations which the corresponding trigger accepts.

Parameters

manager ([ExtensionsManager](#)) – Manager object to call this operator.

Return type

None

`__init__`(*keys=None, trigger=(1, 'epoch'), postprocess=None, filename=None, append=False, format=None, **kwargs*)

Parameters

- **keys** (*Optional[Iterable[str]]*) –
- **trigger** (*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) –
- **postprocess** (*Optional[Callable[[Mapping[str, Any]], None]]*) –
- **filename** (*Optional[str]*) –
- **append** (*bool*) –
- **format** (*Optional[str]*) –
- **kwargs** (*Any*) –

`finalize`(*manager*)

Finalizes the extension.

This method is called at the end of the training loop.

Parameters

manager ([ExtensionsManagerProtocol](#)) –

Return type

None

load_state_dict(*to_load*)

Parameters

to_load (*Dict[str, Any]*) –

Return type

None

property log: **List**[**Mapping**[**str**, **Any**]]

The current list of observation dictionaries.

state_dict()

Serializes the extension state.

It is called when a manager that owns this extension is serialized. It serializes nothing by default.

Return type

Dict[str, Any]

to_dataframe()

Return type

pandas.DataFrame

pytorch_pfn_extras.training.extensions.log_report.LogWriterSaveFunc

class pytorch_pfn_extras.training.extensions.log_report.**LogWriterSaveFunc**(*format, append*)

Bases: object

Methods

__init__(*format, append*)

Parameters

- **format** (*str*) –
- **append** (*bool*) –

__call__(*target, file_o*)

Call self as a function.

Parameters

- **target** (*Dict[str, Any]*) –
- **file_o** (*Any*) –

Return type

None

__init__(*format, append*)

Parameters

- **format** (*str*) –
- **append** (*bool*) –

Return type
None

pytorch_pfn_extras.training.extensions.lr_scheduler

Classes

| | |
|--|---|
| <hr/> | |
| <code>pytorch_pfn_extras.training.extensions. lr_scheduler.ExtensionsManagerProtocol(...)</code> | |
| <code>pytorch_pfn_extras.training.extensions. lr_scheduler.LRScheduler(...)</code> | Trainer extension to adjust the learning rate using Py- Torch's learning rate scheduler. |
| <code>pytorch_pfn_extras.training.extensions. lr_scheduler.ReduceLROnPlateau(...)</code> | Reduce learning rate when a metric has stopped improv- ing. |
| <hr/> | |

pytorch_pfn_extras.training.extensions.lr_scheduler.ExtensionsManagerProtocol

class pytorch_pfn_extras.training.extensions.lr_scheduler.**ExtensionsManagerProtocol**(*args,
**kwargs)

Bases: Protocol

Methods

| |
|--|
| <code>__init__(*args, **kwargs)</code> |
| <code>get_extension(name)</code> |
| <hr/> |

Attributes

elapsed_time

epoch

epoch_detail

is_before_training

iteration

models

observation

optimizers

out

raw_models

reporter

stop_trigger

writer

__init__(*args, **kwargs)**property** elapsed_time: float**property** epoch: int**property** epoch_detail: float**get_extension**(name)**Parameters****name** (str) –**Return type***Extension***property** is_before_training: bool**property** iteration: int**property** models: Mapping[str, Module]**property** observation: reporting.Observation**property** optimizers: Mapping[str, Optimizer]

```

property out: str

property raw_models: Mapping[str, Module]

property reporter: reporting.Reporter

property stop_trigger: bool

property writer: Optional[writing.Writer]

```

pytorch_pfn_extras.training.extensions.lr_scheduler.LRScheduler

```

class pytorch_pfn_extras.training.extensions.lr_scheduler.LRScheduler(scheduler, *,
                             stepper=<function
                                _default_stepper>,
                             trigger=(1, 'epoch'),
                             wait_for_first_optimizer_step=False,
                             is_async=True)

```

Bases: [Extension](#)

Trainer extension to adjust the learning rate using PyTorch's learning rate scheduler.

This extension calls *step()* method of the given LR scheduler. (*torch.optim.lr_scheduler.**). When using *ReduceLROnPlateau*, the latest reported *val/loss* value will be used. This behavior can be customized by passing a custom *stepper* function.

Parameters

- **scheduler** (*_LRScheduler* or *ReduceLROnPlateau*) – Any instance of *torch.optim.lr_scheduler.**.
- **stepper** (*callable*) – Function that performs the step on the scheduler.
- **trigger** (*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) – Frequency to call this extension.
- **wait_for_first_optimizer_step** (*bool*) – Wait until *optimizer.step()* is called before invoking *scheduler.step()*. This can address the issue where *optimizer.step()* is not called from the first iteration when using *GradScaler*.
- **is_async** (*bool*) –

Methods

| | |
|--|---|
| <code>__init__(scheduler, *, stepper, trigger, ...)</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(state)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>state_dict()</code> | Serializes the extension state. |
| <code>step_by_value(key)</code> | |

Attributes

| | |
|-------------------|--------------------------------|
| default_name | Default name of the extension. |
| is_async | |
| name | |
| needs_model_state | |
| priority | |
| trigger | |

__call__(*manager*)

Invokes the extension.

Implementations should override this operator. This method is called at iterations which the corresponding trigger accepts.

Parameters

manager ([ExtensionsManager](#)) – Manager object to call this operator.

Return type

None

__init__(*scheduler*, *, *stepper*=<function _default_stepper>, *trigger*=(1, 'epoch'),
wait_for_first_optimizer_step=False, *is_async*=True)

Parameters

- **scheduler** (*Any*) –
- **stepper** (*Any*) –
- **trigger** (*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) –
- **wait_for_first_optimizer_step** (*bool*) –
- **is_async** (*bool*) –

Return type

None

load_state_dict(*state*)

Parameters

state (*Dict[str, Any]*) –

Return type

None

state_dict()

Serializes the extension state.

It is called when a manager that owns this extension is serialized. It serializes nothing by default.

Return type

Dict[str, Any]

static step_by_value(*key*)

Parameters

key (*Optional[str]*) –

Return type

Any

pytorch_pfn_extras.training.extensions.lr_scheduler.ReduceLROnPlateau

```
class pytorch_pfn_extras.training.extensions.lr_scheduler.ReduceLROnPlateau(optimizer,
                                                                            mode='min',
                                                                            factor=0.1,
                                                                            patience=10,
                                                                            thresh-
                                                                            old=0.0001,
                                                                            thresh-
                                                                            old_mode='rel',
                                                                            cooldown=0,
                                                                            min_lr=0,
                                                                            eps=1e-08,
                                                                            verbose=False)
```

Bases: object

Reduce learning rate when a metric has stopped improving. Models often benefit from reducing the learning rate by a factor of 2-10 once learning stagnates. This scheduler reads a metrics quantity and if no improvement is seen for a ‘patience’ number of epochs, the learning rate is reduced.

Parameters

- **optimizer** (*Optimizer*) – Wrapped optimizer.
- **mode** (*str*) – One of *min*, *max*. In *min* mode, lr will be reduced when the quantity monitored has stopped decreasing; in *max* mode it will be reduced when the quantity monitored has stopped increasing. Default: ‘min’.
- **factor** (*float*) – Factor by which the learning rate will be reduced. $\text{new_lr} = \text{lr} * \text{factor}$. Default: 0.1.
- **patience** (*int*) – Number of epochs with no improvement after which learning rate will be reduced. For example, if *patience* = 2, then we will ignore the first 2 epochs with no improvement, and will only decrease the LR after the 3rd epoch if the loss still hasn’t improved then. Default: 10.
- **threshold** (*float*) – Threshold for measuring the new optimum, to only focus on significant changes. Default: 1e-4.
- **threshold_mode** (*str*) – One of *rel*, *abs*. In *rel* mode, $\text{dynamic_threshold} = \text{best} * (1 + \text{threshold})$ in ‘max’ mode or $\text{best} * (1 - \text{threshold})$ in *min* mode. In *abs* mode, $\text{dynamic_threshold} = \text{best} + \text{threshold}$ in *max* mode or $\text{best} - \text{threshold}$ in *min* mode. Default: ‘rel’.
- **cooldown** (*int*) – Number of epochs to wait before resuming normal operation after lr has been reduced. Default: 0.
- **min_lr** (*float or list*) – A scalar or a list of scalars. A lower bound on the learning rate of all param groups or each group respectively. Default: 0.

- **eps** (*float*) – Minimal decay applied to lr. If the difference between new and old lr is smaller than eps, the update is ignored. Default: 1e-8.
- **verbose** (*bool*) – If True, prints a message to stdout for each update. Default: False.

Example

```
>>> # xdoctest: +SKIP
>>> optimizer = torch.optim.SGD(model.parameters(), lr=0.1, momentum=0.9)
>>> scheduler = ReduceLROnPlateau(optimizer, 'min')
>>> for epoch in range(10):
>>>     train(...)
>>>     val_loss = validate(...)
>>>     # Note that step should be called after validate()
>>>     scheduler.step(val_loss)
```

Methods

`__init__(optimizer[, mode, factor, ...])`

`is_better(a, best)`

`load_state_dict(state_dict)`

`state_dict()`

`step(metrics[, epoch])`

Attributes

`in_cooldown`

`__init__(optimizer, mode='min', factor=0.1, patience=10, threshold=0.0001, threshold_mode='rel',
cooldown=0, min_lr=0, eps=1e-08, verbose=False)``property in_cooldown``is_better(a, best)``load_state_dict(state_dict)``state_dict()``step(metrics, epoch=None)`

pytorch_pfn_extras.training.extensions.micro_average**Classes**

| | |
|---|---------------------------------|
| <code>pytorch_pfn_extras.training.extensions. micro_average.ExtensionsManagerProtocol(...)</code> | |
| <code>pytorch_pfn_extras.training.extensions. micro_average.MicroAverage(...)</code> | Calculates micro-average ratio. |

pytorch_pfn_extras.training.extensions.micro_average.ExtensionsManagerProtocol

```
class pytorch_pfn_extras.training.extensions.micro_average.ExtensionsManagerProtocol(*args,  
                                                                                    **kwargs)
```

Bases: Protocol

Methods

| |
|--|
| <code>__init__(*args, **kwargs)</code> |
| <code>get_extension(name)</code> |

Attributes

elapsed_time

epoch

epoch_detail

is_before_training

iteration

models

observation

optimizers

out

raw_models

reporter

stop_trigger

writer

__init__(*args, **kwargs)**property** elapsed_time: float**property** epoch: int**property** epoch_detail: float**get_extension**(name)**Parameters****name** (str) –**Return type***Extension***property** is_before_training: bool**property** iteration: int**property** models: Mapping[str, Module]**property** observation: reporting.Observation**property** optimizers: Mapping[str, Optimizer]


```

property out: str

property raw_models: Mapping[str, Module]

property reporter: reporting.Reporter

property stop_trigger: bool

property writer: Optional[writing.Writer]

```

pytorch_pfn_extras.training.extensions.micro_average.MicroAverage

```

class pytorch_pfn_extras.training.extensions.micro_average.MicroAverage(
    numerator_key,
    denominator_key,
    result_key, trigger=(1,
    'epoch'))

```

Bases: *Extension*

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(torch.nn.Module):
...     def __call__(self, x, y):
...         loss = torch.nn.CrossEntropyLoss(x, y)
...         correct = (x.data.argmax(axis=1) == y.data).sum()
...         total = len(y.data)
...         reporting.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** (*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) – 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

| | |
|--|---|
| <code>__init__(numerator_key, denominator_key, ...)</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

`__call__(manager)`

Invokes the extension.

Implementations should override this operator. This method is called at iterations which the corresponding trigger accepts.

Parameters

manager (*ExtensionsManager*) – Manager object to call this operator.

Return type

None

`__init__(numerator_key, denominator_key, result_key, trigger=(1, 'epoch'))`

Parameters

- **numerator_key** (*str*) –
- **denominator_key** (*str*) –
- **result_key** (*str*) –

- **trigger**(*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) –

Return type

None

load_state_dict(*to_load*)

Parameters

to_load(*Dict[str, Any]*) –

Return type

None

priority: `int = 200`

state_dict()

Serializes the extension state.

It is called when a manager that owns this extension is serialized. It serializes nothing by default.

Return type

Dict[str, Any]

pytorch_pfn_extras.training.extensions.parameter_statistics

Classes

pytorch_pfn_extras.training.

extensions.parameter_statistics.

ExtensionsManagerProtocol(...)

pytorch_pfn_extras.training.

extensions.parameter_statistics.

ParameterStatistics(links)

An extension to report parameter statistics.

pytorch_pfn_extras.training.extensions.parameter_statistics.ExtensionsManagerProtocol

class `pytorch_pfn_extras.training.extensions.parameter_statistics.ExtensionsManagerProtocol`(*args, **kwargs)

Bases: `Protocol`

Methods

__init__(*args, **kwargs)

get_extension(name)

Attributes

elapsed_time

epoch

epoch_detail

is_before_training

iteration

models

observation

optimizers

out

raw_models

reporter

stop_trigger

writer

__init__(*args, **kwargs)**property** elapsed_time: float**property** epoch: int**property** epoch_detail: float**get_extension**(name)**Parameters****name** (str) –**Return type***Extension***property** is_before_training: bool**property** iteration: int**property** models: Mapping[str, Module]**property** observation: reporting.Observation**property** optimizers: Mapping[str, Optimizer]

```

property out: str

property raw_models: Mapping[str, Module]

property reporter: reporting.Reporter

property stop_trigger: bool

property writer: Optional[writing.Writer]

```

pytorch_pfn_extras.training.extensions.parameter_statistics.ParameterStatistics

```

class pytorch_pfn_extras.training.extensions.parameter_statistics.ParameterStatistics(links,
                                                                                      statis-
                                                                                      tics='default',
                                                                                      re-
                                                                                      port_params=True,
                                                                                      re-
                                                                                      port_grads=True,
                                                                                      pre-
                                                                                      fix=None,
                                                                                      trig-
                                                                                      ger=(1,
                                                                                      'epoch'),
                                                                                      skip_nan_params=False)

```

Bases: *Extension*

An extension to report parameter statistics.

Statistics are collected and reported for a given `Module` or an iterable of `Modules`. If a link contains child modules, the statistics are reported separately for each child.

Any function that takes a one-dimensional `torch.Tensor` 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** (*instance or iterable of Module*) – Module(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 or 'default'*) – 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. If the special value `'default'` is specified, the default statistics functions will be used.
- **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** (*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) – 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.

Note: The default statistic functions are as follows:

- 'mean' (`xp.mean(x)`)
 - 'std' (`xp.std(x)`)
 - 'min' (`xp.min(x)`)
 - 'max' (`xp.max(x)`)
 - 'zeros' (`xp.count_nonzero(x == 0)`)
 - 'percentile' (`xp.percentile(x, (0.13, 2.28, 15.87, 50, 84.13, 97.72, 99.87))`)
-

Methods

| | |
|--|---|
| <code>__init__(links[, statistics, report_params, ...])</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>register_statistics(name, function)</code> | Register a function to compute a certain statistic. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| |
|----------------------------------|
| <code>default_name</code> |
| <code>default_statistics</code> |
| <code>is_async</code> |
| <code>name</code> |
| <code>needs_model_state</code> |
| <code>priority</code> |
| <code>report_key_template</code> |
| <code>trigger</code> |

__call__(*manager*)

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

manager ([ExtensionsManager](#)) – Associated manager that invoked this extension.

Return type

None

__init__(*links*, *statistics*='default', *report_params*=True, *report_grads*=True, *prefix*=None, *trigger*=(1, 'epoch'), *skip_nan_params*=False)

Parameters

- **links** (*Any*) –
- **statistics** (*Any*) –
- **report_params** (*bool*) –
- **report_grads** (*bool*) –
- **prefix** (*Optional*[*str*]) –
- **trigger** (*Optional*[*Union*[[Trigger](#), *Callable*[[[ExtensionsManagerProtocol](#)], *bool*], *Tuple*[*float*, *str*]]]) –
- **skip_nan_params** (*bool*) –

default_name = 'parameter_statistics'

default_statistics = {'max': <function <lambda>>, 'mean': <function <lambda>>, 'min': <function <lambda>>, 'std': <function <lambda>>, 'zeros': <function <lambda>>}

priority: *int* = 300

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** (*Any*) – 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.

Return type

None

report_key_template = '{prefix}{param_name}/{attr_name}/{function_name}'

pytorch_pfn_extras.training.extensions.plot_report**Functions**

`pytorch_pfn_extras.training.extensions.
plot_report.matplotlib_savefun(...)`

pytorch_pfn_extras.training.extensions.plot_report.matplotlib_savefun`pytorch_pfn_extras.training.extensions.plot_report.matplotlib_savefun(target, file_o)`**Parameters**

- **target** (*Tuple*[Any, Any, Any]) –
- **file_o** (Any) –

Return type

None

Classes

`pytorch_pfn_extras.training.extensions.
plot_report.ExtensionsManagerProtocol(...)`

| | |
|---|-------------------------------|
| <code>pytorch_pfn_extras.training.extensions. plot_report.PlotReport(y_keys)</code> | An extension to output plots. |
|---|-------------------------------|

pytorch_pfn_extras.training.extensions.plot_report.ExtensionsManagerProtocol`class pytorch_pfn_extras.training.extensions.plot_report.ExtensionsManagerProtocol(*args,
**kwargs)`

Bases: Protocol

Methods

`__init__(*args, **kwargs)`

`get_extension(name)`

Attributes

elapsed_time

epoch

epoch_detail

is_before_training

iteration

models

observation

optimizers

out

raw_models

reporter

stop_trigger

writer

`__init__`(**args*, ***kwargs*)**property** `elapsed_time`: float**property** `epoch`: int**property** `epoch_detail`: float**get_extension**(*name*)**Parameters****name** (*str*) –**Return type***Extension***property** `is_before_training`: bool**property** `iteration`: int**property** `models`: Mapping[str, Module]**property** `observation`: reporting.Observation**property** `optimizers`: Mapping[str, Optimizer]

```
property out: str
property raw_models: Mapping[str, Module]
property reporter: reporting.Reporter
property stop_trigger: bool
property writer: Optional[writing.Writer]
```

pytorch_pfn_extras.training.extensions.plot_report.PlotReport

```
class pytorch_pfn_extras.training.extensions.plot_report.PlotReport(y_keys, x_key='iteration',
                                                                    trigger=(1, 'epoch'),
                                                                    postprocess=None,
                                                                    filename='plot.png',
                                                                    marker='x', grid=True)
```

Bases: *Extension*

An extension to output plots.

This extension accumulates the observations of the manager 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 `manager.run_iteration`. For example:

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

manager.extend(
    extensions.PlotReport(['main/loss', 'validation/main/loss'],
                          'epoch', filename='loss.png'))
with manager.run_iteration():
    pass
```

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** (*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) – 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** (*Any*) – 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.
- **filename** (*str*) – Name of the figure file under the output directory. It can be a format string. For historical reasons `file_name` is also accepted as an alias of this argument.
- **marker** (*str*) – The marker used to plot the graph. Default is 'x'. If None is given, it draws with no markers.
- **grid** (*bool*) – If True, set the axis grid on. The default value is True.
- **writer** (*writer object, optional*) – must be callable. object to dump the log to. If specified, it needs to have a correct `savefun` defined. The writer can override the save location in the `pytorch_pfn_extras.training.ExtensionsManager` object
- **kwargs** (*Any*) –

Methods

| | |
|--|---|
| <code>__init__(y_keys[, x_key, trigger, ...])</code> | |
| <code>available()</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

__call__(*manager*)

Invokes the extension.

Implementations should override this operator. This method is called at iterations which the corresponding trigger accepts.

Parameters

manager ([ExtensionsManager](#)) – Manager object to call this operator.

Return type

None

__init__(*y_keys*, *x_key*='iteration', *trigger*=(1, 'epoch'), *postprocess*=None, *filename*=None, *marker*='x', *grid*=True, ***kwargs*)

Parameters

- **y_keys** ([Union](#)[[Iterable](#)[*str*], *str*]) –
- **x_key** (*str*) –
- **trigger** ([Optional](#)[[Union](#)[[Trigger](#), [Callable](#)[[\[ExtensionsManagerProtocol\]](#), *bool*], [Tuple](#)[*float*, *str*]]) –
- **postprocess** ([Optional](#)[*Any*]) –
- **filename** ([Optional](#)[*str*]) –
- **marker** (*str*) –
- **grid** (*bool*) –
- **kwargs** (*Any*) –

static available()

Return type

bool

finalize(*manager*)

Finalizes the extension.

This method is called at the end of the training loop.

Parameters

manager ([ExtensionsManagerProtocol](#)) –

Return type

None

load_state_dict(*to_load*)

Parameters

to_load ([Dict](#)[*str*, *Any*]) –

Return type

None

state_dict()

Serializes the extension state.

It is called when a manager that owns this extension is serialized. It serializes nothing by default.

Return type

[Dict](#)[*str*, *Any*]

pytorch_pfn_extras.training.extensions.print_report**Functions**

| | |
|---|--|
| <code>pytorch_pfn_extras.training.extensions.print_report.create_header_and_templates(entries)</code> | Construct header and templates from <i>entries</i> |
| <code>pytorch_pfn_extras.training.extensions.print_report.deepcopy(x)</code> | Deep copy operation on arbitrary Python objects. |
| <code>pytorch_pfn_extras.training.extensions.print_report.filter_and_sort_entries(...)</code> | |

pytorch_pfn_extras.training.extensions.print_report.create_header_and_templates

`pytorch_pfn_extras.training.extensions.print_report.create_header_and_templates(entries)`

Construct header and templates from *entries*

Parameters

entries (*list*) – list of str

Returns

header string templates (str): template string for print values.

Return type

header (str)

pytorch_pfn_extras.training.extensions.print_report.deepcopy

`pytorch_pfn_extras.training.extensions.print_report.deepcopy(x, memo=None, _nil=[])`

Deep copy operation on arbitrary Python objects.

See the module's `__doc__` string for more info.

pytorch_pfn_extras.training.extensions.print_report.filter_and_sort_entries

`pytorch_pfn_extras.training.extensions.print_report.filter_and_sort_entries(all_entries, unit='epoch')`

Parameters

- **all_entries** (*List[str]*) –
- **unit** (*str*) –

Return type

List[str]

Classes

| | |
|--|--|
| <code>pytorch_pfn_extras.training.extensions. print_report.ExtensionsManagerProtocol(...)</code> | |
| <code>pytorch_pfn_extras.training.extensions. print_report.IO(...)</code> | Generic base class for TextIO and BinaryIO. |
| <code>pytorch_pfn_extras.training.extensions. print_report.PrintReport(...)</code> | An extension to print the accumulated results. |

`pytorch_pfn_extras.training.extensions.print_report.ExtensionsManagerProtocol`

class `pytorch_pfn_extras.training.extensions.print_report.ExtensionsManagerProtocol(*args,
**kwargs)`

Bases: `Protocol`

Methods

| |
|--|
| <code>__init__(*args, **kwargs)</code> |
| <code>get_extension(name)</code> |

Attributes

elapsed_time

epoch

epoch_detail

is_before_training

iteration

models

observation

optimizers

out

raw_models

reporter

stop_trigger

writer

`__init__`(**args*, ***kwargs*)**property** `elapsed_time`: float**property** `epoch`: int**property** `epoch_detail`: float**get_extension**(*name*)**Parameters****name** (*str*) –**Return type***Extension***property** `is_before_training`: bool**property** `iteration`: int**property** `models`: Mapping[str, Module]**property** `observation`: reporting.Observation**property** `optimizers`: Mapping[str, Optimizer]

```
property out: str
property raw_models: Mapping[str, Module]
property reporter: reporting.Reporter
property stop_trigger: bool
property writer: Optional[writing.Writer]
```

`pytorch_pfn_extras.training.extensions.print_report.IO`

```
class pytorch_pfn_extras.training.extensions.print_report.IO(*args, **kws)
```

Bases: `Generic`

Generic base class for `TextIO` and `BinaryIO`.

This is an abstract, generic version of the return of `open()`.

NOTE: This does not distinguish between the different possible classes (text vs. binary, read vs. write vs. read/write, append-only, unbuffered). The `TextIO` and `BinaryIO` subclasses below capture the distinctions between text vs. binary, which is pervasive in the interface; however we currently do not offer a way to track the other distinctions in the type system.

Methods

__init__()

close()

fileno()

flush()

isatty()

read([n])

readable()

readline([limit])

readlines([hint])

seek(offset[, whence])

seekable()

tell()

truncate([size])

writable()

write(s)

writelines(lines)

Attributes

closed

mode

name

abstract close()**Return type**

None

abstract property closed: bool

abstract `fileno()`

Return type
 int

abstract `flush()`

Return type
 None

abstract `isatty()`

Return type
 bool

abstract `property mode: str`

abstract `property name: str`

abstract `read(n=-1)`

Parameters
 n (int) –

Return type
 AnyStr

abstract `readable()`

Return type
 bool

abstract `readline(limit=-1)`

Parameters
 limit (int) –

Return type
 AnyStr

abstract `readlines(hint=-1)`

Parameters
 hint (int) –

Return type
 List

abstract `seek(offset, whence=0)`

Parameters

- *offset* (int) –
- *whence* (int) –

Return type
 int

abstract `seekable()`

Return type
 bool

```

abstract tell()

    Return type
    int

abstract truncate(size=None)

    Parameters
    size (Optional[int]) –

    Return type
    int

abstract writable()

    Return type
    bool

abstract write(s)

    Parameters
    s (AnyStr) –

    Return type
    int

abstract writelines(lines)

    Parameters
    lines (List) –

    Return type
    None

```

pytorch_pfn_extras.training.extensions.print_report.PrintReport

```

class pytorch_pfn_extras.training.extensions.print_report.PrintReport(entries=None,
                                                                    log_report='LogReport',
                                                                    out=<_io.TextIOWrapper
                                                                    name='<stdout>'
                                                                    mode='w'
                                                                    encoding='utf-8'>)

```

Bases: [Extension](#)

An 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 or None*) – List of keys of observations to print. If *None* is passed, automatically infer keys from reported dict.
- **log_report** (*str or LogReport*) – Log report to accumulate the observations. This is either the name of a `LogReport` extensions registered to the manager, or a `LogReport` instance to use internally.
- **out** (*IO[Any]*) – Stream to print the bar. Standard output is used by default.

Methods

| | |
|---|---|
| <code>__init__([entries, log_report, out])</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>get_log_report(manager)</code> | |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

`__call__(manager)`

Invokes the extension.

Implementations should override this operator. This method is called at iterations which the corresponding trigger accepts.

Parameters

manager (`ExtensionsManager`) – Manager object to call this operator.

Return type

None

`__init__(entries=None, log_report='LogReport', out=<_io.TextIOWrapper name='<stdout>' mode='w' encoding='utf-8'>)`

Parameters

- **entries** (`Optional[Sequence[str]]`) –
- **log_report** (`Union[str, LogReport]`) –
- **out** (`IO[Any]`) –

Return type

None

get_log_report(*manager*)

Parameters

manager ([ExtensionsManagerProtocol](#)) –

Return type

[LogReport](#)

initialize(*manager*)

Initializes up the manager state.

This method is called before entering the training loop. An extension modifying the state of `ExtensionsManager` can override this method to initialize it.

When the manager has been restored from a snapshot, this method has to recover an appropriate part of the state of the manager.

Parameters

manager ([ExtensionsManager](#)) – Manager object to call this extension.

Return type

None

load_state_dict(*to_load*)

Parameters

to_load (*Dict[str, Any]*) –

Return type

None

state_dict()

Serializes the extension state.

It is called when a manager that owns this extension is serialized. It serializes nothing by default.

Return type

Dict[str, Any]

pytorch_pfn_extras.training.extensions.profile_report

Functions

`pytorch_pfn_extras.training.extensions.
profile_report.get_time_summary()`

pytorch_pfn_extras.training.extensions.profile_report.get_time_summary

`pytorch_pfn_extras.training.extensions.profile_report.get_time_summary()`

Return type

[TimeSummary](#)

Classes

```
pytorch_pfn_extras.training.  
extensions.profile_report.  
ExtensionsManagerProtocol(...)
```

```
pytorch_pfn_extras.training.extensions.  
profile_report.OrderedDict
```

Dictionary that remembers insertion order

```
pytorch_pfn_extras.training.extensions.  
profile_report.ProfileReport(...)
```

pytorch_pfn_extras.training.extensions.profile_report.ExtensionsManagerProtocol

```
class pytorch_pfn_extras.training.extensions.profile_report.ExtensionsManagerProtocol(*args,  
                                                                                       **kwargs)
```

Bases: Protocol

Methods

```
__init__(*args, **kwargs)
```

```
get_extension(name)
```

Attributes

elapsed_time

epoch

epoch_detail

is_before_training

iteration

models

observation

optimizers

out

raw_models

reporter

stop_trigger

writer

`__init__`(**args*, ***kwargs*)**property** `elapsed_time`: float**property** `epoch`: int**property** `epoch_detail`: float**get_extension**(*name*)**Parameters****name** (*str*) –**Return type***Extension***property** `is_before_training`: bool**property** `iteration`: int**property** `models`: Mapping[str, Module]**property** `observation`: reporting.Observation**property** `optimizers`: Mapping[str, Optimizer]

```

property out: str
property raw_models: Mapping[str, Module]
property reporter: reporting.Reporter
property stop_trigger: bool
property writer: Optional[writing.Writer]

```

pytorch_pfn_extras.training.extensions.profile_report.OrderedDict

class pytorch_pfn_extras.training.extensions.profile_report.OrderedDict

Bases: dict

Dictionary that remembers insertion order

Methods

| | |
|------------------------------------|--|
| <i>__init__</i> (*args, **kwargs) | |
| <i>clear</i> () | |
| <i>copy</i> () | |
| <i>fromkeys</i> ([value]) | Create a new ordered dictionary with keys from iterable and values set to value. |
| <i>get</i> (key[, default]) | Return the value for key if key is in the dictionary, else default. |
| <i>items</i> () | |
| <i>keys</i> () | |
| <i>move_to_end</i> (key[, last]) | Move an existing element to the end (or beginning if last is false). |
| <i>pop</i> (k[,d]) | value. |
| <i>popitem</i> ([last]) | Remove and return a (key, value) pair from the dictionary. |
| <i>setdefault</i> (key[, default]) | Insert key with a value of default if key is not in the dictionary. |
| <i>update</i> ([E,]**F) | If E is present and has a .keys() method, then does: for k in E: D[k] = E[k] If E is present and lacks a .keys() method, then does: for k, v in E: D[k] = v In either case, this is followed by: for k in F: D[k] = F[k] |
| <i>values</i> () | |
| <hr/> | |
| <i>__init__</i> (*args, **kwargs) | |
| clear () | → None. Remove all items from od. |
| copy () | → a shallow copy of od |

fromkeys(*value=None*)

Create a new ordered dictionary with keys from iterable and values set to value.

items() → a set-like object providing a view on D's items

keys() → a set-like object providing a view on D's keys

move_to_end(*key, last=True*)

Move an existing element to the end (or beginning if last is false).

Raise KeyError if the element does not exist.

pop(*k[, d]*) → *v*, remove specified key and return the corresponding value. If key is not found, *d* is returned if given, otherwise KeyError is raised.

popitem(*last=True*)

Remove and return a (key, value) pair from the dictionary.

Pairs are returned in LIFO order if last is true or FIFO order if false.

setdefault(*key, default=None*)

Insert key with a value of default if key is not in the dictionary.

Return the value for key if key is in the dictionary, else default.

update(*[E], **F*) → None. Update D from dict/iterable E and F.

If E is present and has a .keys() method, then does: for k in E: D[k] = E[k] If E is present and lacks a .keys() method, then does: for k, v in E: D[k] = v In either case, this is followed by: for k in F: D[k] = F[k]

values() → an object providing a view on D's values

pytorch_pfn_extras.training.extensions.profile_report.ProfileReport

```
class pytorch_pfn_extras.training.extensions.profile_report.ProfileReport(store_keys=None,
                                                                           report_keys=None,
                                                                           trigger=(1, 'epoch'),
                                                                           filename=None,
                                                                           append=False,
                                                                           format=None,
                                                                           **kwargs)
```

Bases: [Extension](#)

Writes the profile results to a file.

Times are reported by using the [pytorch_pfn_extras.profiler.TimeSummary.report\(\)](#) context manager.

Parameters

- **store_keys** (*iterable of strs*) – Keys of values to write to the profiler report file.
- **report_keys** (*iterable of strs*) – Keys of values that will be reported.
- **trigger** (*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) – 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.

- **filename** (*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.
- **append** (*bool*, *options1*) – If the file is JSON Lines or YAML, contents will be appended instead of rewriting the file every call.
- **format** (*str*, *optional*) – accepted values are 'json', 'json-lines' and 'yaml'.
- **writer** (*writer object*, *optional*) – must be callable. object to dump the log to. If specified, it needs to have a correct *savefun* defined. The writer can override the save location in the [pytorch_pfn_extras.training.ExtensionsManager](#) object
- **entries** (*list*) – list of str
- **kwargs** (*Any*) –

Returns

header string templates (*str*): template string for print values.

Return type

header (*str*)

Methods

| | |
|---|---|
| <code>__init__</code> ([store_keys, report_keys, trigger, ...]) | |
| <code>finalize</code> (manager) | Finalizes the extension. |
| <code>initialize</code> (manager) | Initializes up the manager state. |
| <code>load_state_dict</code> (to_load) | |
| <code>on_error</code> (manager, exc, tb) | Handles the error raised during training before finalization. |
| <code>state_dict</code> () | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

`__call__`(*manager*)
Invokes the extension.

Implementations should override this operator. This method is called at iterations which the corresponding trigger accepts.

Parameters

manager ([ExtensionsManager](#)) – Manager object to call this operator.

Return type

None

__init__(*store_keys=None, report_keys=None, trigger=(1, 'epoch'), filename=None, append=False, format=None, **kwargs*)

Parameters

- **store_keys** (*Optional[Iterable[str]]*) –
- **report_keys** (*Optional[Iterable[str]]*) –
- **trigger** (*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) –
- **filename** (*Optional[str]*) –
- **append** (*bool*) –
- **format** (*Optional[str]*) –
- **kwargs** (*Any*) –

finalize(*manager*)

Finalizes the extension.

This method is called at the end of the training loop.

Parameters

manager ([ExtensionsManagerProtocol](#)) –

Return type

None

load_state_dict(*to_load*)

Parameters

to_load (*Dict[str, Any]*) –

Return type

None

state_dict()

Serializes the extension state.

It is called when a manager that owns this extension is serialized. It serializes nothing by default.

Return type

Dict[str, Any]

pytorch_pfn_extras.training.extensions.progress_bar**Classes**

*pytorch_pfn_extras.training.extensions.
progress_bar.ExtensionsManagerProtocol(...)*

| | |
|--|--|
| <i>pytorch_pfn_extras.training.extensions. progress_bar.ProgressBar(...)</i> | An extension to print a progress bar and recent training status. |
|--|--|

pytorch_pfn_extras.training.extensions.progress_bar.ExtensionsManagerProtocol

```
class pytorch_pfn_extras.training.extensions.progress_bar.ExtensionsManagerProtocol(*args,  
                                                                                    **kwargs)
```

Bases: Protocol

Methods

__init__(*args, **kwargs)

get_extension(name)

Attributes

elapsed_time

epoch

epoch_detail

is_before_training

iteration

models

observation

optimizers

out

raw_models

reporter

stop_trigger

writer

`__init__`(**args, **kwargs*)**property** `elapsed_time`: float**property** `epoch`: int**property** `epoch_detail`: float**get_extension**(*name*)**Parameters****name** (*str*) –**Return type***Extension***property** `is_before_training`: bool**property** `iteration`: int**property** `models`: Mapping[str, Module]**property** `observation`: reporting.Observation**property** `optimizers`: Mapping[str, Optimizer]

```
property out: str

property raw_models: Mapping[str, Module]

property reporter: reporting.Reporter

property stop_trigger: bool

property writer: Optional[writing.Writer]
```

pytorch_pfn_extras.training.extensions.progress_bar.ProgressBar

```
class pytorch_pfn_extras.training.extensions.progress_bar.ProgressBar(training_length=None,
                                                                    update_interval=100,
                                                                    bar_length=50,
                                                                    out=<_io.TextIOWrapper
                                                                    name='<stdout>'
                                                                    mode='w'
                                                                    encoding='utf-8'>)
```

Bases: *Extension*

An 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* or *None*) – 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 manager 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** (*Any*) – Stream to print the bar. Standard output is used by default.

Methods

| | |
|--|---|
| <hr/> | |
| <code>__init__([training_length, update_interval, ...])</code> | |
| <hr/> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <hr/> | |
| <code>load_state_dict(to_load)</code> | |
| <hr/> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <hr/> | |
| <code>state_dict()</code> | Serializes the extension state. |
| <hr/> | |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

`__call__(manager)`

Invokes the extension.

Implementations should override this operator. This method is called at iterations which the corresponding trigger accepts.

Parameters

manager ([ExtensionsManager](#)) – Manager object to call this operator.

Return type

None

`__init__(training_length=None, update_interval=100, bar_length=50, out=<_io.TextIOWrapper name='<stdout>' mode='w' encoding='utf-8'>)`

Parameters

- **training_length** (*Optional*[Any]) –
- **update_interval** (*int*) –
- **bar_length** (*int*) –
- **out** (*Any*) –

`finalize(manager)`

Finalizes the extension.

This method is called at the end of the training loop.

Parameters

manager ([ExtensionsManagerProtocol](#)) –

Return type

None

pytorch_pfn_extras.training.extensions.slack**Classes**

| | |
|--|--|
| <code>pytorch_pfn_extras.training.extensions.slack.ExtensionsManagerProtocol(...)</code> | |
| <code>pytorch_pfn_extras.training.extensions.slack.Slack(channel)</code> | An extension to communicate with Slack. |
| <code>pytorch_pfn_extras.training.extensions.slack.SlackWebhook(url)</code> | An extension to communicate with Slack using Incoming Webhook. |

pytorch_pfn_extras.training.extensions.slack.ExtensionsManagerProtocol

```
class pytorch_pfn_extras.training.extensions.slack.ExtensionsManagerProtocol(*args,  
                                                                              **kwargs)
```

Bases: Protocol

Methods

```
__init__(*args, **kwargs)
```

```
get_extension(name)
```

Attributes

elapsed_time

epoch

epoch_detail

is_before_training

iteration

models

observation

optimizers

out

raw_models

reporter

stop_trigger

writer

`__init__`(**args*, ***kwargs*)**property** `elapsed_time`: float**property** `epoch`: int**property** `epoch_detail`: float**get_extension**(*name*)**Parameters****name** (*str*) –**Return type***Extension***property** `is_before_training`: bool**property** `iteration`: int**property** `models`: Mapping[str, Module]**property** `observation`: reporting.Observation**property** `optimizers`: Mapping[str, Optimizer]

```
property out: str
property raw_models: Mapping[str, Module]
property reporter: reporting.Reporter
property stop_trigger: bool
property writer: Optional[writing.Writer]
```

pytorch_pfn_extras.training.extensions.slack.Slack

```
class pytorch_pfn_extras.training.extensions.slack.Slack(channel, msg=None, *,
                                                         start_msg='{default}',
                                                         end_msg='{default}',
                                                         error_msg='{default}', thread=True,
                                                         filenames=None, upload_trigger=None,
                                                         context=None, token=None)
```

Bases: `_SlackBase`

An extension to communicate with Slack.

Example

```
>>> ppe.training.extensions.Slack(
...     channel="experiment-progress",
...     msg="Epoch #{manager.epoch}: loss = {val/loss}",
...     end_msg="{default} \n <@username> Check out the result!",
...     # Upload files at the end of the experiment.
...     filenames=["result/statistics.png"],
...     upload_trigger=(max_epoch, 'epoch'),
... )
```

This extension posts a message when:

- `start_msg`: The training has started
- `msg`: The extension is triggered, usually at the end of each epoch
- `end_msg`: The training has finished
- `error_msg`: An exception has raised during the training

These messages can be specified as a format string, a callable that returns a string, or `None` to disable posting on that event.

When using a format string, the following variables are available for use:

- `manager`: an `ExtensionsManager` object
- `default`: the default message string
- `context`: an arbitrary object passed to this extension
- `error`: an `Exception` object (for `error_msg` only)
- All reported values (`manager.observations`)

When using a callable, it should take *(ExtensionsManager, context)* or *(ExtensionsManager, Exception, context)* (for `error_msg`) and return a string.

This extension can upload files along with the message when triggered. `filenames` can be a list of filenames (the same formatting rule as `msg` apply), or a callable taking *(ExtensionsManager, context)* and returning a list of filenames.

To use this extension, you must create a Slack app, then specify the token via an environment variable `SLACK_BOT_TOKEN` or `token` option.

Parameters

- **channel** (*str*) – The channel where messages and files will be sent. This can be a channel name or a channel ID.
- **msg** (*str, callable, or None*) – A message to be sent when triggered. It can be a string to be formatted using `.format` or a callable that returns a string.
- **start_msg** (*str, callable, or None*) – A message to be sent at the beginning of the experiment.
- **end_msg** (*str, callable, or None*) – A message to be sent at the completion of the experiment.
- **error_msg** (*str, callable, or None*) – A message to be sent when an exception is raised during the experiment.
- **thread** (*bool*) – When True, subsequent messages will be posted as a thread of the original message. Default is True.
- **filenames** (*list of str or callable*) – A list of files that will be uploaded. These are string templates that can take values in the same way as `msg`, or a callable that returns a list of filenames.
- **upload_trigger** (*trigger or None*) – Used to upload files at certain events. If not specified, files will be uploaded in every call.
- **context** (*Any*) – Any arbitrary user object you will need when generating a message.
- **token** (*str*) – Slack bot token. If None, the environment variable `SLACK_BOT_TOKEN` will be used. Optional, default is None.

Methods

| | |
|--|---|
| <code>__init__(channel[, msg, start_msg, end_msg, ...])</code> | |
| <code>default_end_msg(context)</code> | |
| <code>default_error_msg(exc, context)</code> | |
| <code>default_msg(context)</code> | |
| <code>default_start_msg(context)</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

`__init__(channel, msg=None, *, start_msg='{default}', end_msg='{default}', error_msg='{default}', thread=True, filenames=None, upload_trigger=None, context=None, token=None)`

Parameters

- **channel** (*str*) –
- **msg** (*Optional[Union[[str](#), Callable[[[ExtensionsManagerProtocol](#), Any], [str](#)]]]*) –
- **start_msg** (*Optional[Union[[str](#), Callable[[[ExtensionsManagerProtocol](#), Any], [str](#)]]]*) –
- **end_msg** (*Optional[Union[[str](#), Callable[[[ExtensionsManagerProtocol](#), Any], [str](#)]]]*) –
- **error_msg** (*Optional[Union[[str](#), Callable[[[ExtensionsManagerProtocol](#), Any, [Exception](#)], [str](#)]]]*) –
- **thread** (*bool*) –

- **filenames** *(Optional[Union[Sequence[str], Callable[[ExtensionsManagerProtocol, Any], Sequence[str]]]])* –
- **upload_trigger** *(Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]])* –
- **context** *(Optional[Any])* –
- **token** *(Optional[str])* –

Return type

None

trigger: Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]] = (1, 'epoch')

pytorch_pfn_extras.training.extensions.slack.SlackWebhook

```
class pytorch_pfn_extras.training.extensions.slack.SlackWebhook(url, msg=None, *,
                                                                start_msg='{default}',
                                                                end_msg='{default}',
                                                                error_msg='{default}',
                                                                context=None)
```

Bases: _SlackBase

An extension to communicate with Slack using Incoming Webhook.

Example

```
>>> ppe.training.extensions.SlackWebhook(
...     url="https://hooks.slack.com/services/Txxxxx.....",
...     msg="Epoch #{manager.epoch}: loss = {val/loss}",
...     end_msg="{default} \n <@username> Check out the result!",
... )
```

This extension posts a message when:

- **start_msg**: The training has started
- **msg**: The extension is triggered, usually at the end of each epoch
- **end_msg**: The training has finished
- **error_msg**: An exception has raised during the training

These messages can be specified as a format string, a callable that returns a string, or None to disable posting on that event.

When using a format string, the following variables are available for use:

- **manager**: an ExtensionsManager object
- **default**: the default message string
- **context**: an arbitrary object passed to this extension
- **error**: an Exception object (for **error_msg** only)
- All reported values (**manager.observations**)

When using a callable, it should take *(ExtensionsManager, context)* or *(ExtensionsManager, Exception, context)* (for `error_msg`) and return a string.

Parameters

- **url** (*str*) – Incoming webhook URL to send messages.
- **msg** (*str, callable, or None*) – A message to be sent when triggered. It can be a string to be formatted using `.format` or a callable that returns a string.
- **start_msg** (*str, callable, or None*) – A message to be sent at the beginning of the experiment.
- **end_msg** (*str, callable, or None*) – A message to be sent at the completion of the experiment.
- **error_msg** (*str, callable, or None*) – A message to be sent when an exception is raised during the experiment.
- **context** (*object*) – Any arbitrary user object you will need when generating a message.

Methods

| | |
|--|---|
| <code>__init__(url[, msg, start_msg, end_msg, ...])</code> | |
| <code>default_end_msg(context)</code> | |
| <code>default_error_msg(exc, context)</code> | |
| <code>default_msg(context)</code> | |
| <code>default_start_msg(context)</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

```
__init__(url, msg=None, *, start_msg='{default}', end_msg='{default}', error_msg='{default}',
         context=None)
```

Parameters

- **url** (*str*) –
- **msg** (*Optional[Union[*str*, Callable[[ExtensionsManagerProtocol, Any], *str*]]]*) –
- **start_msg** (*Optional[Union[*str*, Callable[[ExtensionsManagerProtocol, Any], *str*]]]*) –
- **end_msg** (*Optional[Union[*str*, Callable[[ExtensionsManagerProtocol, Any], *str*]]]*) –
- **error_msg** (*Optional[Union[*str*, Callable[[ExtensionsManagerProtocol, Any, Exception], *str*]]]*) –
- **context** (*Optional[Any]*) –

Return type

None

pytorch_pfn_extras.training.extensions.snapshot_writers

Classes

| | |
|--|---|
| <code>pytorch_pfn_extras.training.extensions.snapshot_writers.ProcessQueueWriter(...)</code> | Snapshot writer that uses process queue. |
| <code>pytorch_pfn_extras.training.extensions.snapshot_writers.ProcessWriter(...)</code> | Snapshot writer that uses a separate process. |
| <code>pytorch_pfn_extras.training.extensions.snapshot_writers.QueueWriter(...)</code> | Base class of queue snapshot writers. |
| <code>pytorch_pfn_extras.training.extensions.snapshot_writers.SimpleWriter(...)</code> | The most simple snapshot writer. |
| <code>pytorch_pfn_extras.training.extensions.snapshot_writers.StandardWriter(...)</code> | Base class of snapshot writers which use thread or process. |
| <code>pytorch_pfn_extras.training.extensions.snapshot_writers.TensorBoardWriter(...)</code> | Writer that sends statistics to TensorBoard. |
| <code>pytorch_pfn_extras.training.extensions.snapshot_writers.ThreadQueueWriter(...)</code> | Snapshot writer that uses a thread queue. |
| <code>pytorch_pfn_extras.training.extensions.snapshot_writers.ThreadWriter(...)</code> | Snapshot writer that uses a separate thread. |
| <code>pytorch_pfn_extras.training.extensions.snapshot_writers.Writer(...)</code> | Base class of snapshot writers. |

pytorch_pfn_extras.training.extensions.snapshot_writers.ProcessQueueWriter

```
class pytorch_pfn_extras.training.extensions.snapshot_writers.ProcessQueueWriter(savefun=<function save>,
                                                                              fs=None,
                                                                              out_dir="",
                                                                              task=None)
```

Bases: [QueueWriter](#)[Process]

Snapshot writer that uses process queue.

This class creates a process and a queue by `multiprocessing` module. The process will be a consumer of this queue, and the main process will be a producer of this queue.

Note: Forking a new process from MPI process might be danger. Consider using [ThreadQueueWriter](#) instead of `ProcessQueueWriter` if you are using MPI.

See also:

- [pytorch_pfn_extras.training.extensions.snapshot\(\)](#)

Methods

`__init__([savefun, fs, out_dir, task])`

`consume(q)`

`create_consumer(q)`

`create_queue()`

`create_task(savefun)`

`finalize()` Finalizes the writer.

`initialize(out_dir)`

`save(filename, out_dir, target, savefun, ...)`

```
__init__(savefun=<function save>, fs=None, out_dir="", task=None)
```

Parameters

- **savefun** (`Callable[[...], None]`) –
- **fs** (`Optional[Any]`) –
- **out_dir** (`str`) –
- **task** (`Optional[Callable[[...], None]]`) –

Return type

None

create_consumer(*q*)

Parameters

q (`queue.Queue[_QueUnit]`) –

Return type

`Process`

create_queue()

Return type

`queue.Queue[_QueUnit]`

pytorch_pfn_extras.training.extensions.snapshot_writers.ProcessWriter

```
class pytorch_pfn_extras.training.extensions.snapshot_writers.ProcessWriter(
    savefun=<function save>, fs=None,
    out_dir="",
    **kws)
```

Bases: `StandardWriter`[`Process`]

Snapshot writer that uses a separate process.

This class creates a new process that invokes the actual saving function.

Note: Forking a new process from a MPI process might be danger. Consider using `ThreadWriter` instead of `ProcessWriter` if you are using MPI.

See also:

- `pytorch_pfn_extras.training.extensions.snapshot()`

Methods

`__init__`(*savefun*, *fs*, *out_dir*)

`create_worker`(*filename*, *out_dir*, *target*, *) Creates a worker for the snapshot.

`finalize`() Finalizes the writer.

`initialize`(*out_dir*)

`save`(*filename*, *out_dir*, *target*, *savefun*, ...)

`__init__`(*savefun*=<function save>, *fs*=None, *out_dir*="", ***kws*)

Parameters

- **savefun** (`Callable[[...], None]`) –
- **fs** (`Optional[Any]`) –
- **out_dir** (`str`) –
- **kws** (`Any`) –

Return type

None

create_worker(*filename*, *out_dir*, *target*, *, *savefun*=None, *append*=False, ***savefun_kwargs*)

Creates a worker for the snapshot.

This method creates a thread or a process to take a snapshot. The created worker must have `start()` and `join()` methods. If the worker has an `exitcode` attribute (e.g., `multiprocessing.Process`), the value will be tested.

Parameters

- **filename** (*str*) –
- **out_dir** (*str*) –
- **target** (*Union[Sequence[Any], Mapping[str, Any]]*) –
- **savefun** (*Optional[Callable[[...], None]]*) –
- **append** (*bool*) –
- **savefun_kwargs** (*Any*) –

Return type*Process***pytorch_pfn_extras.training.extensions.snapshot_writers.QueueWriter**

```
class pytorch_pfn_extras.training.extensions.snapshot_writers.QueueWriter(savefun=<function  
save>, fs=None,  
out_dir="",  
task=None)
```

Bases: [Writer](#), [Generic\[_Worker\]](#)

Base class of queue snapshot writers.

This class is a base class of snapshot writers that use a queue. A Queue is created when this class is constructed, and every time when `__call__` is invoked, a snapshot task is put into the queue.

Parameters

- **savefun** – Callable object which is passed to the [create_task\(\)](#) if the task is None. It takes three arguments: the output file path, the serialized dictionary object, and the optional keyword arguments.
- **fs** – FileSystem abstracting interface to implement all the operations. optional, defaults to None
- **out_dir** – str. Specifies the directory this writer will use. It takes precedence over the one specified in `__call__` optional, defaults to ''
- **task** – Callable object. Its `__call__` must have a same interface to `Writer.__call__`. This object is directly put into the queue.

See also:

- [pytorch_pfn_extras.training.extensions.snapshot\(\)](#)

Methods

| | |
|--|-----------------------|
| <code>__init__([savefun, fs, out_dir, task])</code> | |
| <code>consume(q)</code> | |
| <code>create_consumer(q)</code> | |
| <code>create_queue()</code> | |
| <code>create_task(savefun)</code> | |
| <code>finalize()</code> | Finalizes the writer. |
| <code>initialize(out_dir)</code> | |
| <code>save(filename, out_dir, target, savefun, ...)</code> | |

`__call__(filename, out_dir, target, *, savefun=None, append=False)`

Does the actual writing to the file.

This method is invoked by a Snapshot object every time it takes a snapshot.

Parameters

- **filename** (*str*) – Name of the file into which the serialized target is saved. It is a concrete file name, i.e. not a pre-formatted template string.
- **out_dir** (*str*) – Output directory. Corresponds to :py:attr:`ExtensionsManager.out`
<pytorch_pfn_extras.training.ExtensionsManager.out>`.
- **target** (*dict*) – Serialized object which will be saved.
- **savefun** (*callable*) – A callable that accepts a two positional arguments (an object to be serialized, file path) like *torch.save*.
- **append** (*bool*) – Mode used to open the file. True to use the append mode, False to use the write mode (truncates the file if it already exists).

Return type

None

`__init__(savefun=<function save>, fs=None, out_dir="", task=None)`

Parameters

- **savefun** (*Callable[[...], None]*) –
- **fs** (*Optional[Any]*) –
- **out_dir** (*str*) –
- **task** (*Optional[Callable[[...], None]]*) –

Return type

None

`consume(q)`

Parameters

q (*queue.Queue[_QueUnit]*) –

Return type

None

create_consumer(*q*)

Parameters

q (*queue.Queue[_QueUnit]*) –

Return type

_Worker

create_queue()

Return type

queue.Queue[_QueUnit]

create_task(*savefun*)

Parameters

savefun (*Callable[[...], None]*) –

Return type

Callable[[...], None]

finalize()

Finalizes the writer.

Calling this method on already-finalized Writer does nothing.

Return type

None

pytorch_pfn_extras.training.extensions.snapshot_writers.SimpleWriter

```
class pytorch_pfn_extras.training.extensions.snapshot_writers.SimpleWriter(savefun=<function  
                                save>, fs=None,  
                                out_dir="",  
                                **kwargs)
```

Bases: *Writer*

The most simple snapshot writer.

This class just passes the arguments to the actual saving function.

Parameters

- **savefun** (*Callable[[...], None]*) – Callable object. It takes three arguments: the output file path, the serialized dictionary object, and the optional keyword arguments.
- **fs** (*Any*) – FileSystem abstracting interface to implement all the operations. optional, defaults to None
- **out_dir** (*str*) – str. Specifies the directory this writer will use. It takes precedence over the one specified in `__call__` optional, defaults to ''
- **kwargs** (*Any*) – Keyword arguments for the `savefun`.

See also:

- `pytorch_pfn_extras.training.extensions.snapshot()`

Methods

| | |
|--|-----------------------|
| <code>__init__([savefun, fs, out_dir])</code> | |
| <code>finalize()</code> | Finalizes the writer. |
| <code>initialize(out_dir)</code> | |
| <code>save(filename, out_dir, target, savefun, ...)</code> | |

`__call__(filename, out_dir, target, *, savefun=None, append=False)`

Does the actual writing to the file.

This method is invoked by a `Snapshot` object every time it takes a snapshot.

Parameters

- **filename** (*str*) – Name of the file into which the serialized target is saved. It is a concrete file name, i.e. not a pre-formatted template string.
- **out_dir** (*str*) – Output directory. Corresponds to `:py:attr:`ExtensionsManager.out`` `<pytorch_pfn_extras.training.ExtensionsManager.out>`.
- **target** (*dict*) – Serialized object which will be saved.
- **savefun** (*callable*) – A callable that accepts a two positional arguments (an object to be serialized, file path) like `torch.save`.
- **append** (*bool*) – Mode used to open the file. True to use the append mode, False to use the write mode (truncates the file if it already exists).

Return type

None

`__init__(savefun=<function save>, fs=None, out_dir="", **kwds)`

Parameters

- **savefun** (*Callable[[...], None]*) –
- **fs** (*Optional[Any]*) –
- **out_dir** (*str*) –
- **kwds** (*Any*) –

Return type

None

pytorch_pfn_extras.training.extensions.snapshot_writers.StandardWriter

```
class pytorch_pfn_extras.training.extensions.snapshot_writers.StandardWriter(savefun=<function
    save>,
    fs=None,
    out_dir="",
    **kwargs)
```

Bases: [Writer](#), [Generic\[_Worker\]](#)

Base class of snapshot writers which use thread or process.

This class creates a new thread or a process every time when `__call__` is invoked.

Parameters

- **savefun** – Callable object. It takes three arguments: the output file path, the serialized dictionary object, and the optional keyword arguments.
- **fs** – FileSystem abstracting interface to implement all the operations. optional, defaults to None
- **out_dir** – str. Specifies the directory this writer will use. It takes precedence over the one specified in `__call__` optional, defaults to ''
- **kwargs** – Keyword arguments for the `savefun`.

See also:

- [pytorch_pfn_extras.training.extensions.snapshot\(\)](#)

Methods

| | |
|--|------------------------------------|
| <code>__init__([savefun, fs, out_dir])</code> | |
| <code>create_worker(filename, out_dir, target, *)</code> | Creates a worker for the snapshot. |
| <code>finalize()</code> | Finalizes the writer. |
| <code>initialize(out_dir)</code> | |
| <code>save(filename, out_dir, target, savefun, ...)</code> | |

`__call__(filename, out_dir, target, *, savefun=None, append=False)`

Does the actual writing to the file.

This method is invoked by a `Snapshot` object every time it takes a snapshot.

Parameters

- **filename** (*str*) – Name of the file into which the serialized target is saved. It is a concrete file name, i.e. not a pre-formatted template string.
- **out_dir** (*str*) – Output directory. Corresponds to `:py:attr:`ExtensionsManager.out`` `<pytorch_pfn_extras.training.ExtensionsManager.out>`.
- **target** (*dict*) – Serialized object which will be saved.
- **savefun** (*callable*) – A callable that accepts a two positional arguments (an object to be serialized, file path) like `torch.save`.

- **append** (*bool*) – Mode used to open the file. True to use the append mode, False to use the write mode (truncates the file if it already exists).

Return type

None

__init__(*savefun=<function save>, fs=None, out_dir="", **kws*)

Parameters

- **savefun** (*Callable[[...], None]*) –
- **fs** (*Optional[Any]*) –
- **out_dir** (*str*) –
- **kws** (*Any*) –

Return type

None

create_worker(*filename, out_dir, target, *, savefun=None, append=False, **savefun_kwargs*)

Creates a worker for the snapshot.

This method creates a thread or a process to take a snapshot. The created worker must have `start()` and `join()` methods. If the worker has an `exitcode` attribute (e.g., `multiprocessing.Process`), the value will be tested.

Parameters

- **filename** (*str*) –
- **out_dir** (*str*) –
- **target** (*Union[Sequence[Any], Mapping[str, Any]]*) –
- **savefun** (*Optional[Callable[[...], None]]*) –
- **append** (*bool*) –
- **savefun_kwargs** (*Any*) –

Return type*_Worker*

finalize()

Finalizes the writer.

Calling this method on already-finalized Writer does nothing.

Return type

None

pytorch_pfn_extras.training.extensions.snapshot_writers.TensorBoardWriter

```
class pytorch_pfn_extras.training.extensions.snapshot_writers.TensorBoardWriter(savefun=None,
                                                                              fs=None,
                                                                              out_dir="",
                                                                              stats=None,
                                                                              **kws)
```

Bases: `object`

Writer that sends statistics to TensorBoard.

This class contains a `torch.utils.tensorboard.SummaryWriter` object that is used to send the collected statistics to TensorBoard. A list of stats can be specified to report only the desired ones.

Parameters

- **savefun** (*Optional*[*Callable*[[...], None]]) – Ignored.
- **fs** (*Any*) – Ignored.
- **out_dir** (*str*) – Passed as `log_dir` argument to `SummaryWriter`.
- **stats** (*list*) – List of statistic keys.
- **kwds** (*Any*) – Passed as an additional arguments to `SummaryWriter`.

Methods

`__init__`(`savefun`, `fs`, `out_dir`, `stats`)

`finalize`()

`__call__`(`filename`, `out_dir`, `target`, *, `savefun=None`, `append=False`)

Sends the statistics to the TensorBoard.

Parameters

- **filename** (*str*) – Ignored.
- **out_dir** (*str*) – Ignored.
- **target** (*dict or list*) – The statistics of the iteration. If given as a list, only the last element (assumed to be a dict containing the latest iteration statistics) is reported.
- **savefun** (*Optional*[*Callable*[[...], None]]) – Ignored.
- **append** (*bool*) – Ignored.

Return type

None

`__init__`(`savefun=None`, `fs=None`, `out_dir=""`, `stats=None`, `**kwds`)

Parameters

- **savefun** (*Optional*[*Callable*[[...], None]]) –
- **fs** (*Optional*[*Any*]) –
- **out_dir** (*str*) –
- **stats** (*Optional*[*KeysView*[*str*]]) –
- **kwds** (*Any*) –

Return type

None

`finalize`()

Return type

None

pytorch_pfn_extras.training.extensions.snapshot_writers.ThreadQueueWriter

```
class pytorch_pfn_extras.training.extensions.snapshot_writers.ThreadQueueWriter(savefun=<function
                                                    save>,
                                                    fs=None,
                                                    out_dir="",
                                                    task=None)
```

Bases: [QueueWriter](#)[Thread]

Snapshot writer that uses a thread queue.

This class creates a thread and a queue by `threading` and `queue` modules respectively. The thread will be a consumer of the queue, and the main thread will be a producer of the queue.

See also:

- [pytorch_pfn_extras.training.extensions.snapshot\(\)](#)

Methods

| | |
|---|-----------------------|
| <code>__init__</code> ([savefun, fs, out_dir, task]) | |
| <code>consume</code> (q) | |
| <code>create_consumer</code> (q) | |
| <code>create_queue</code> () | |
| <code>create_task</code> (savefun) | |
| <code>finalize</code> () | Finalizes the writer. |
| <code>initialize</code> (out_dir) | |
| <code>save</code> (filename, out_dir, target, savefun, ...) | |

```
__init__(savefun=<function save>, fs=None, out_dir="", task=None)
```

Parameters

- **savefun** (`Callable[[...], None]`) –
- **fs** (`Optional[Any]`) –
- **out_dir** (`str`) –
- **task** (`Optional[Callable[[...], None]]`) –

Return type

None

```
create_consumer(q)
```

Parameters

- **q** (`queue.Queue[_QueUnit]`) –

Return type*Thread***create_queue()****Return type***queue.Queue[_QueUnit]***pytorch_pfn_extras.training.extensions.snapshot_writers.ThreadWriter**

```
class pytorch_pfn_extras.training.extensions.snapshot_writers.ThreadWriter(savefun=<function
save>, fs=None,
out_dir="",
**kws)
```

Bases: *StandardWriter*[*Thread*]

Snapshot writer that uses a separate thread.

This class creates a new thread that invokes the actual saving function.

See also:

- *pytorch_pfn_extras.training.extensions.snapshot()*

Methods

__init__([savefun, fs, out_dir])

create_worker(filename, out_dir, target, *)
Creates a worker for the snapshot.

finalize()
Finalizes the writer.

initialize(out_dir)

save(filename, out_dir, target, savefun, ...)

__init__(savefun=<function save>, fs=None, out_dir="", **kws)**Parameters**

- **savefun** (*Callable*[[...], None]) –
- **fs** (*Optional*[Any]) –
- **out_dir** (*str*) –
- **kws** (*Any*) –

Return type

None

create_worker(filename, out_dir, target, *, savefun=None, append=False, **savefun_kwargs)

Creates a worker for the snapshot.

This method creates a thread or a process to take a snapshot. The created worker must have `start()` and `join()` methods. If the worker has an `exitcode` attribute (e.g., `multiprocessing.Process`), the value will be tested.

Parameters

- **filename** (*str*) –
- **out_dir** (*str*) –
- **target** (*Union[Sequence[Any], Mapping[str, Any]]*) –
- **savefun** (*Optional[Callable[[...], None]]*) –
- **append** (*bool*) –
- **savefun_kwargs** (*Any*) –

Return type*Thread***pytorch_pfn_extras.training.extensions.snapshot_writers.Writer**

class pytorch_pfn_extras.training.extensions.snapshot_writers.**Writer**(*fs=None, out_dir=""*)

Bases: object

Base class of snapshot writers.

Snapshot invokes `__call__` of this class every time when taking a snapshot. This class determines how the actual saving function will be invoked.

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.

See also:

- [`pytorch_pfn_extras.training.extensions.snapshot\(\)`](#)

Methods

`__init__`(*[fs, out_dir]*)

`finalize`() Finalizes the writer.

`initialize`(*out_dir*)

`save`(*filename, out_dir, target, savefun, ...*)

Parameters

- **fs** (*Any*) –
- **out_dir** (*str*) –

`__call__`(*filename, out_dir, target, *, savefun=None, append=False*)

Does the actual writing to the file.

This method is invoked by a Snapshot object every time it takes a snapshot.

Parameters

- **filename** (*str*) – Name of the file into which the serialized target is saved. It is a concrete file name, i.e. not a pre-formatted template string.
- **out_dir** (*str*) – Output directory. Corresponds to :py:attr:`ExtensionsManager.out`
<pytorch_pfn_extras.training.ExtensionsManager.out>`.
- **target** (*dict*) – Serialized object which will be saved.
- **savefun** (*callable*) – A callable that accepts a two positional arguments (an object to be serialized, file path) like *torch.save*.
- **append** (*bool*) – Mode used to open the file. True to use the append mode, False to use the write mode (truncates the file if it already exists).

Return type

None

__init__(*fs=None, out_dir=""*)**Parameters**

- **fs** (*Optional[Any]*) –
- **out_dir** (*str*) –

Return type

None

finalize()

Finalizes the writer.

Calling this method on already-finalized Writer does nothing.

Return type

None

initialize(*out_dir*)**Parameters****out_dir** (*str*) –**Return type**

None

save(*filename, out_dir, target, savefun, append, **savefun_kwargs*)**Parameters**

- **filename** (*str*) –
- **out_dir** (*str*) –
- **target** (*Union[Sequence[Any], Mapping[str, Any]]*) –
- **savefun** (*Callable[[...], None]*) –
- **append** (*bool*) –
- **savefun_kwargs** (*Any*) –

Return type

None

pytorch_pfn_extras.training.extensions.timeline_trace

Functions

| | |
|---|---------------------------------|
| <code>pytorch_pfn_extras.training.extensions. timeline_trace.get_tracer([...])</code> | Gets the current global tracer. |
|---|---------------------------------|

pytorch_pfn_extras.training.extensions.timeline_trace.get_tracer

`pytorch_pfn_extras.training.extensions.timeline_trace.get_tracer(tracer_cls=<class 'pytorch_pfn_extras.profiler._tracing.ChromeTracer'>, *params)`

Gets the current global tracer.

Parameters

- **tracer_cls** (type of `Tracer`) – type of tracer to create if the global tracer hasn't been initialized
- **params** (*Any*) –

Return type

`Tracer`

Classes

| | |
|---|--|
| <code>pytorch_pfn_extras.training. extensions.timeline_trace. ExtensionsManagerProtocol(...)</code> | |
| <code>pytorch_pfn_extras.training.extensions. timeline_trace.TimelineTrace([...])</code> | Writes the profile timeline to a file. |
| <code>pytorch_pfn_extras.training.extensions. timeline_trace.Tracer()</code> | |

pytorch_pfn_extras.training.extensions.timeline_trace.ExtensionsManagerProtocol

`class pytorch_pfn_extras.training.extensions.timeline_trace.ExtensionsManagerProtocol(*args, **kwargs)`

Bases: `Protocol`

Methods

`__init__(*args, **kwargs)`

`get_extension(name)`

Attributes

`elapsed_time`

`epoch`

`epoch_detail`

`is_before_training`

`iteration`

`models`

`observation`

`optimizers`

`out`

`raw_models`

`reporter`

`stop_trigger`

`writer`

`__init__(*args, **kwargs)``property elapsed_time: float``property epoch: int``property epoch_detail: float``get_extension(name)`

Parameters

name (*str*) –

Return type

Extension

```

property is_before_training: bool
property iteration: int
property models: Mapping[str, Module]
property observation: reporting.Observation
property optimizers: Mapping[str, Optimizer]
property out: str
property raw_models: Mapping[str, Module]
property reporter: reporting.Reporter
property stop_trigger: bool
property writer: Optional[writing.Writer]

```

`pytorch_pfn_extras.training.extensions.timeline_trace.TimelineTrace`

```

class pytorch_pfn_extras.training.extensions.timeline_trace.TimelineTrace(trigger=(1, 'epoch'),
                                                                           filename=None,
                                                                           enable=None,
                                                                           disable=None,
                                                                           tracer=None,
                                                                           **kwargs)

```

Bases: [Extension](#)

Writes the profile timeline to a file.

Times are reported by using the `pytorch_pfn_extras.profiler.TimeSummary.report()` context manager.

Parameters

- **trigger** (`Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]`) – Trigger that decides when to output the trace. 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`.
- **enable** (`Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]`) – Trigger that enables the tracing. Note that since the extensions are executed at the end of an iteration the tracer will be enabled from the iteration after the trigger is fired. If it is a tuple in the form `<int>`, `'epoch'` or `<int>`, `'iteration'`, it is passed to `IntervalTrigger`.
- **disable** (`Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]`) – Trigger that disables the tracing. Note that since the extensions are executed at the end of an iteration the tracer will be disabled from the iteration after the trigger is fired. If it is a tuple in the form `<int>`, `'epoch'` or `<int>`, `'iteration'`, it is passed to `IntervalTrigger`.
- **filename** (`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.

- **writer** (*writer object, optional*) – must be callable. object to dump the log to. If specified, it needs to have a correct *savefun* defined. The writer can override the save location in the `pytorch_pfn_extras.training.ExtensionsManager` object
- **tracer** (*tracer object, optional*) – object with the tracing logic. If not specified, the default tracer in the thread local storage with be used.
- **entries** (*list*) – list of str
- **kwargs** (*Any*) –

Returns

header string templates (str): template string for print values.

Return type

header (str)

Methods

| | |
|--|---|
| <code>__init__</code> ([trigger, filename, enable, ...]) | |
| <code>finalize</code> (manager) | Finalizes the extension. |
| <code>initialize</code> (manager) | Initializes up the manager state. |
| <code>load_state_dict</code> (to_load) | |
| <code>on_error</code> (manager, exc, tb) | Handles the error raised during training before finalization. |
| <code>state_dict</code> () | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

`__call__` (*manager*)

Invokes the extension.

Implementations should override this operator. This method is called at iterations which the corresponding trigger accepts.

Parameters

manager (`ExtensionsManager`) – Manager object to call this operator.

Return type

None

__init__(*trigger=(1, 'epoch'), filename=None, enable=None, disable=None, tracer=None, **kwargs*)

Parameters

- **trigger** (*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) –
- **filename** (*Optional[str]*) –
- **enable** (*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) –
- **disable** (*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) –
- **tracer** (*Optional[Tracer]*) –
- **kwargs** (*Any*) –

finalize(*manager*)

Finalizes the extension.

This method is called at the end of the training loop.

Parameters

manager (*ExtensionsManagerProtocol*) –

Return type

None

initialize(*manager*)

Initializes up the manager state.

This method is called before entering the training loop. An extension modifying the state of `ExtensionsManager` can override this method to initialize it.

When the manager has been restored from a snapshot, this method has to recover an appropriate part of the state of the manager.

Parameters

manager (*ExtensionsManager*) – Manager object to call this extension.

Return type

None

load_state_dict(*to_load*)

Parameters

to_load (*Dict[str, Any]*) –

Return type

None

state_dict()

Serializes the extension state.

It is called when a manager that owns this extension is serialized. It serializes nothing by default.

Return type

Dict[str, Any]

pytorch_pfn_extras.training.extensions.timeline_trace.Tracer**class** pytorch_pfn_extras.training.extensions.timeline_trace.Tracer

Bases: object

Methods

`__init__()`

`add_event(name)`

`add_remote_event(name, value)`

`clear()`

`enable(enable_flag)`

`finalize()`

`flush(filename, writer)`

`initialize_writer(filename, writer)`

`load_state_dict(to_load)`

`state_dict()`

add_event(*name*)**Parameters****name** (*str*) –**Return type***Generator*[None, None, None]**add_remote_event**(*name*, *value*)**Parameters**

- **name** (*str*) –
- **value** (*Any*) –

Return type

None

clear()**Return type**

None

enable(*enable_flag*)**Parameters****enable_flag** (*bool*) –

Return type
None

finalize()

Return type
None

flush(*filename*, *writer*)

Parameters

- **filename** (*str*) –
- **writer** (*Writer*) –

Return type
None

initialize_writer(*filename*, *writer*)

Parameters

- **filename** (*str*) –
- **writer** (*Writer*) –

Return type
None

load_state_dict(*to_load*)

Parameters
to_load (*Dict[str, Any]*) –

Return type
None

state_dict()

Return type
Dict[str, Any]

pytorch_pfn_extras.training.extensions.util

Classes

| | |
|---|---|
| <code>pytorch_pfn_extras.training.extensions.util.ExtensionsManagerProtocol(...)</code> | |
| <code>pytorch_pfn_extras.training.extensions.util.ProgressBar([out])</code> | |
| <code>pytorch_pfn_extras.training.extensions.util.TextIO(...)</code> | Typed version of the return of open() in text mode. |

pytorch_pfn_extras.training.extensions.util.ExtensionsManagerProtocol

```
class pytorch_pfn_extras.training.extensions.util.ExtensionsManagerProtocol(*args,  
                                                                            **kwargs)
```

Bases: Protocol

Methods

```
__init__(*args, **kwargs)
```

```
get_extension(name)
```

Attributes

```
elapsed_time
```

```
epoch
```

```
epoch_detail
```

```
is_before_training
```

```
iteration
```

```
models
```

```
observation
```

```
optimizers
```

```
out
```

```
raw_models
```

```
reporter
```

```
stop_trigger
```

```
writer
```

```
__init__(*args, **kwargs)
```

```
property elapsed_time: float
```

```
property epoch: int
```

```
property epoch_detail: float
```

`get_extension(name)`

Parameters

`name` (*str*) –

Return type

Extension

property `is_before_training`: `bool`

property `iteration`: `int`

property `models`: `Mapping[str, Module]`

property `observation`: `reporting.Observation`

property `optimizers`: `Mapping[str, Optimizer]`

property `out`: `str`

property `raw_models`: `Mapping[str, Module]`

property `reporter`: *reporting.Reporter*

property `stop_trigger`: `bool`

property `writer`: `Optional[writing.Writer]`

pytorch_pfn_extras.training.extensions.util.ProgressBar

`class pytorch_pfn_extras.training.extensions.util.ProgressBar(out=None)`

Bases: `object`

Methods

__init__([*out*])

close()

erase_console()

flush()

get_lines()

move_cursor_up(*n*)

update([*manager*])

update_speed(*iteration*, *epoch_detail*)

Parameters

`out` (*Optional[TextIO]*) –

__init__(*out=None*)

Parameters

out (*Optional*[[TextIO](#)]) –

Return type

None

close()

Return type

None

erase_console()

Return type

None

flush()

Return type

None

get_lines()

Return type

Sequence[*str*]

move_cursor_up(*n*)

Parameters

n (*int*) –

Return type

None

update(*manager=None*)

Parameters

manager (*Optional*[[ExtensionsManagerProtocol](#)]) –

Return type

None

update_speed(*iteration, epoch_detail*)

Parameters

- **iteration** (*int*) –
- **epoch_detail** (*float*) –

Return type

Tuple[*float*, *float*]

pytorch_pfn_extras.training.extensions.util.TextIO**class** pytorch_pfn_extras.training.extensions.util.**TextIO**(*args, **kws)Bases: [IO](#)[str]Typed version of the return of `open()` in text mode.**Methods**

`__init__()`

`close()`

`fileno()`

`flush()`

`isatty()`

`read([n])`

`readable()`

`readline([limit])`

`readlines([hint])`

`seek(offset[, whence])`

`seekable()`

`tell()`

`truncate([size])`

`writable()`

`write(s)`

`writelines(lines)`

Attributes

| |
|-----------------------|
| <i>buffer</i> |
| closed |
| <i>encoding</i> |
| <i>errors</i> |
| <i>line_buffering</i> |
| mode |
| name |
| <i>newlines</i> |

```
abstract property buffer: BinaryIO
abstract property encoding: str
abstract property errors: Optional[str]
abstract property line_buffering: bool
abstract property newlines: Any
```

pytorch_pfn_extras.training.extensions.value_observation

Functions

| | |
|--|--|
| <i>pytorch_pfn_extras.training.extensions.value_observation.observe_lr(...)</i> | Returns an extension to record the learning rate. |
| <i>pytorch_pfn_extras.training.extensions.value_observation.observe_value(...)</i> | Returns an extension to continuously record a value. |

pytorch_pfn_extras.training.extensions.value_observation.observe_lr

`pytorch_pfn_extras.training.extensions.value_observation.observe_lr(optimizer, param_group=0, observation_key='lr')`

Returns an extension to record the learning rate.

Parameters

- **optimizer** (*Optimizer*) – Optimizer whose learning rate is recorded.
- **param_group** (*int*) – Param group of the optimizer to observe
- **observation_key** (*str*) – Key of observation to record.

Returns

The extension function.

Return type

Any

This extension is triggered each epoch by default. To change this, use the `trigger` argument with the `ExtensionsManager.extend()` method.

pytorch_pfn_extras.training.extensions.value_observation.observe_value

`pytorch_pfn_extras.training.extensions.value_observation.observe_value(observation_key, target_func)`

Returns an 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:`~pytorch_pfn_extras.training.ExtensionsManager` object.

Returns

The extension function.

Return type

Callable[[*ExtensionsManagerProtocol*], None]

This extension is triggered each epoch by default. To change this, use the `trigger` argument with the `ExtensionsManager.extend()` method.

Classes

```
pytorch_pfn_extras.training.  
extensions.value_observation.  
ExtensionsManagerProtocol(...)
```

pytorch_pfn_extras.training.extensions.value_observation.ExtensionsManagerProtocol

```
class pytorch_pfn_extras.training.extensions.value_observation.ExtensionsManagerProtocol(*args,  
                                                                                       **kwargs)
```

Bases: `Protocol`

Methods

```
__init__(*args, **kwargs)
```

```
get_extension(name)
```

Attributes

elapsed_time

epoch

epoch_detail

is_before_training

iteration

models

observation

optimizers

out

raw_models

reporter

stop_trigger

writer

__init__(*args, **kwargs)**property** elapsed_time: float**property** epoch: int**property** epoch_detail: float**get_extension**(name)**Parameters****name** (str) –**Return type***Extension***property** is_before_training: bool**property** iteration: int**property** models: Mapping[str, Module]**property** observation: reporting.Observation**property** optimizers: Mapping[str, Optimizer]

```

property out: str
property raw_models: Mapping[str, Module]
property reporter: reporting.Reporter
property stop_trigger: bool
property writer: Optional[writing.Writer]

```

pytorch_pfn_extras.training.extensions.variable_statistics_plot

Functions

```

pytorch_pfn_extras.training.
extensions.variable_statistics_plot.
matplotlib_savefun(...)

```

```

pytorch_pfn_extras.training.extensions.
variable_statistics_plot.percentile(a, ...)

```

pytorch_pfn_extras.training.extensions.variable_statistics_plot.matplotlib_savefun

```

pytorch_pfn_extras.training.extensions.variable_statistics_plot.matplotlib_savefun(target,
                                                                                   file_o)

```

Parameters

- **target** (*Tuple*[*Any*, *Any*]) –
- **file_o** (*Any*) –

Return type

None

pytorch_pfn_extras.training.extensions.variable_statistics_plot.percentile

```

pytorch_pfn_extras.training.extensions.variable_statistics_plot.percentile(a, q, axis)

```

Parameters

- **a** (*Tensor*) –
- **q** (*Union*[*float*, *Tuple*[*float*, ...]]) –
- **axis** (*int*) –

Return type

Any

Classes

| | |
|--|--|
| <code>pytorch_pfn_extras.training. extensions.variable_statistics_plot. ExtensionsManagerProtocol(...)</code> | |
| <code>pytorch_pfn_extras.training.extensions. variable_statistics_plot.Reservoir(...)</code> | Reservoir sample with a fixed sized buffer. |
| <code>pytorch_pfn_extras.training.extensions. variable_statistics_plot.Statistician(...)</code> | Helper to compute basic NumPy-like statistics. |
| <code>pytorch_pfn_extras.training. extensions.variable_statistics_plot. VariableStatisticsPlot(targets)</code> | An extension to plot statistics for Tensors. |

`pytorch_pfn_extras.training.extensions.variable_statistics_plot.ExtensionsManagerProtocol`

class `pytorch_pfn_extras.training.extensions.variable_statistics_plot.ExtensionsManagerProtocol(*args, **kwargs)`

Bases: `Protocol`

Methods

| |
|--|
| <code>__init__(*args, **kwargs)</code> |
| <code>get_extension(name)</code> |

Attributes

elapsed_time

epoch

epoch_detail

is_before_training

iteration

models

observation

optimizers

out

raw_models

reporter

stop_trigger

writer

`__init__`(**args*, ***kwargs*)**property** `elapsed_time`: float**property** `epoch`: int**property** `epoch_detail`: float**get_extension**(*name*)**Parameters****name** (*str*) –**Return type***Extension***property** `is_before_training`: bool**property** `iteration`: int**property** `models`: Mapping[str, Module]**property** `observation`: reporting.Observation**property** `optimizers`: Mapping[str, Optimizer]

```
property out: str
property raw_models: Mapping[str, Module]
property reporter: reporting.Reporter
property stop_trigger: bool
property writer: Optional[writing.Writer]
```

pytorch_pfn_extras.training.extensions.variable_statistics_plot.Reservoir

```
class pytorch_pfn_extras.training.extensions.variable_statistics_plot.Reservoir(size,
                                                                              data_shape,
                                                                              dtype=<class
                                                                              'numpy.float32'>)
```

Bases: object

Reservoir sample with a fixed sized buffer.

Methods

```
__init__(size, data_shape[, dtype])
```

```
add(x[, idx])
```

```
get_data()
```

Parameters

- **size** (*int*) –
- **data_shape** (*Tuple[int, ...]*) –
- **dtype** (*Any*) –

```
__init__(size, data_shape, dtype=<class 'numpy.float32'>)
```

Parameters

- **size** (*int*) –
- **data_shape** (*Tuple[int, ...]*) –
- **dtype** (*Any*) –

Return type

None

```
add(x, idx=None)
```

Parameters

- **x** (*Any*) –
- **idx** (*Optional[Any]*) –

Return type

None

get_data()**Return type***Tuple[Any, Any]***pytorch_pfn_extras.training.extensions.variable_statistics_plot.Statistician**

class pytorch_pfn_extras.training.extensions.variable_statistics_plot.**Statistician**(*collect_mean*,
collect_std,
percentile_sigmas)

Bases: object

Helper to compute basic NumPy-like statistics.

Methods

__init__(collect_mean, collect_std, ...)

collect(x, axis)

Parameters

- **collect_mean** (*bool*) –
- **collect_std** (*bool*) –
- **percentile_sigmas** (*Union[float, Tuple[float, ...]]*) –

__call__(x, axis=0, dtype=None)

Call self as a function.

Parameters

- **x** (*Any*) –
- **axis** (*Any*) –
- **dtype** (*Optional[Any]*) –

Return type*Dict[str, Any]**__init__*(collect_mean, collect_std, percentile_sigmas)**Parameters**

- **collect_mean** (*bool*) –
- **collect_std** (*bool*) –
- **percentile_sigmas** (*Union[float, Tuple[float, ...]]*) –

Return type

None

collect(*x*, *axis*)**Parameters**

- **x** (*Any*) –
- **axis** (*int*) –

Return type*Dict*[*str*, *Any*]**pytorch_pfn_extras.training.extensions.variable_statistics_plot.VariableStatisticsPlot**

```
class pytorch_pfn_extras.training.extensions.variable_statistics_plot.VariableStatisticsPlot(targets,  
max_sample_size=int,  
replot=bool,  
port_data=bool,  
replot_grad=bool,  
plot_mean=bool,  
plot_std=bool,  
percentile_sigma=list,  
0.13,  
2.28,  
15.87,  
50,  
84.13,  
97.72,  
99.87,  
100),  
trigger=int,  
'epoch'),  
filename=str,  
figure_size=tuple,  
marker=str,  
grid=bool)
```

Bases: [Extension](#)

An extension to plot statistics for Tensors.

This extension collects statistics for a single `torch.Tensor`, a list of `torch.Tensors` or similarly a single or a list of `torch.nn.Modules` containing one or more `torch.Tensors`. In case multiple `torch.Tensors` are found, the means are computed. The collected statistics are plotted and saved as an image in the directory specified by the `Manager`.

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.

:param targets (`torch.Tensor`: or list of either): Parameters for which statistics are collected. :param `torch.nn.Module`: or list of either): Parameters for which statistics are collected. :param `max_sample_size`: Maximum number of running samples. :type `max_sample_size`: `int` :param `report_data`: If `True`, data (e.g. weights) statistics are plotted. If

`False`, they are neither computed nor plotted.

Parameters

- **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** (*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) – 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`.
- **filename** (*str*) – Name of the output image file under the output directory. For historical reasons `file_name` is also accepted as an alias of this argument.
- **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.
- **writer** (*writer object, optional*) – must be callable. object to dump the log to. If specified, it needs to have a correct `savefun` defined. The writer can override the save location in the `pytorch_pfn_extras.training.ExtensionsManager` object
- **targets** (*Any*) –
- **max_sample_size** (*int*) –
- **report_data** (*bool*) –
- **kwargs** (*Any*) –

Methods

| | |
|--|---|
| <code>__init__(targets[, max_sample_size, ...])</code> | |
| <code>available()</code> | |
| <code>finalize(manager)</code> | Finalizes the extension. |
| <code>initialize(manager)</code> | Initializes up the manager state. |
| <code>load_state_dict(to_load)</code> | |
| <code>on_error(manager, exc, tb)</code> | Handles the error raised during training before finalization. |
| <code>save_plot_using_module(plt, manager)</code> | |
| <code>state_dict()</code> | Serializes the extension state. |

Attributes

| | |
|--------------------------------|--------------------------------|
| <code>default_name</code> | Default name of the extension. |
| <code>is_async</code> | |
| <code>name</code> | |
| <code>needs_model_state</code> | |
| <code>priority</code> | |
| <code>trigger</code> | |

`__call__(manager)`

Invokes the extension.

Implementations should override this operator. This method is called at iterations which the corresponding trigger accepts.

Parameters

manager (`ExtensionsManager`) – Manager object to call this operator.

Return type

None

`__init__(targets, max_sample_size=1000, report_data=True, report_grad=True, plot_mean=True, plot_std=True, percentile_sigmas=(0, 0.13, 2.28, 15.87, 50, 84.13, 97.72, 99.87, 100), trigger=(1, 'epoch'), filename=None, figsize=None, marker=None, grid=True, **kwargs)`

Parameters

- **targets** (*Any*) –
- **max_sample_size** (*int*) –
- **report_data** (*bool*) –
- **report_grad** (*bool*) –

- **plot_mean** (*bool*) –
- **plot_std** (*bool*) –
- **percentile_sigmas** (*Union[float, Tuple[float, ...]]*) –
- **trigger** (*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) –
- **filename** (*Optional[str]*) –
- **figsize** (*Optional[Tuple[int, ...]]*) –
- **marker** (*Optional[str]*) –
- **grid** (*bool*) –
- **kwargs** (*Any*) –

static available()

Return type
bool

finalize(*manager*)

Finalizes the extension.

This method is called at the end of the training loop.

Parameters

manager (*ExtensionsManagerProtocol*) –

Return type
None

save_plot_using_module(*plt, manager*)

Parameters

- **plt** (*Any*) –
- **manager** (*ExtensionsManagerProtocol*) –

Return type
None

pytorch_pfn_extras.training.manager

Functions

pytorch_pfn_extras.training.manager.

default_transform_model(*n, x*)

pytorch_pfn_extras.training.manager.

record(*tag*)

Classes

| | |
|---|--|
| <code>pytorch_pfn_extras.training.manager.ExtensionsManager(...)</code> | Manages the extensions and the current status. |
| <code>pytorch_pfn_extras.training.manager.IgniteExtensionsManager(...)</code> | Manages extensions and the current status in Ignite training loop. |
| <code>pytorch_pfn_extras.training.manager.StateObjectProtocol(...)</code> | |

pytorch_pfn_extras.training.metrics

Classes

| | |
|--|---|
| <code>pytorch_pfn_extras.training.metrics.AccuracyMetric(...)</code> | A metric for an evaluator to report accuracy. |
|--|---|

pytorch_pfn_extras.training.metrics.AccuracyMetric

class `pytorch_pfn_extras.training.metrics.AccuracyMetric(label_key, output_key)`

Bases: `object`

A metric for an evaluator to report accuracy.

Parameters

- **label_key** (*str*) – The key name of label.
- **output_key** (*str*) – The key name of prediction.

Methods

| |
|--|
| <code>__init__(label_key, output_key)</code> |
|--|

`__call__(batch, out)`

Call self as a function.

Parameters

- **batch** (*Dict[str, Tensor]*) –
- **out** (*Dict[str, Tensor]*) –

Return type

Dict[str, Any]

`__init__(label_key, output_key)`

Parameters

- **label_key** (*str*) –
- **output_key** (*str*) –

Return type

None

pytorch_pfn_extras.training.trigger**Functions**

| | |
|---|------------------------|
| <code>pytorch_pfn_extras.training.trigger.get_trigger(trigger)</code> | Gets a trigger object. |
|---|------------------------|

pytorch_pfn_extras.training.trigger.get_trigger`pytorch_pfn_extras.training.trigger.get_trigger(trigger)`

Gets a trigger object.

Trigger object is a callable that accepts a [ExtensionsManager](#) 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 `extend()` method of a manager, 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 creates a `IntervalTrigger`.

Parameters

trigger (*Optional[Union[Trigger, Callable[[ExtensionsManagerProtocol], bool], Tuple[float, str]]]*) – Trigger object. It can be either an already built trigger object (i.e., a callable object that accepts a manager 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.

Return type[Trigger](#)**Classes**

| | |
|---|------------------------------------|
| <code>pytorch_pfn_extras.training.trigger.IntervalTrigger(...)</code> | Trigger based on a fixed interval. |
| <code>pytorch_pfn_extras.training.trigger.Trigger()</code> | Base class for triggers. |

pytorch_pfn_extras.training.trigger.IntervalTrigger

class pytorch_pfn_extras.training.trigger.IntervalTrigger(*period, unit*)

Bases: *Trigger*

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 manager.

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

`__init__`(*period, unit*)

`get_training_length`()

`load_state_dict`(*state*)

| | |
|--|--|
| <code>may_fire</code> (<i>iteration, epoch_length</i>) | Flags if the trigger may fire at the current iteration |
| <code>state_dict</code> () | |

`__call__`(*manager*)

Decides whether the extension should be called on this iteration.

Parameters

manager (*ExtensionsManager*) – Manager object that this trigger is associated with. The iteration related information in this manager is used to determine if the trigger should fire.

Returns

True if the corresponding extension should be invoked in this iteration.

Return type

bool

`__init__`(*period, unit*)

Parameters

- **period** (*float*) –
- **unit** (*UnitLiteral*) –

get_training_length()

Return type

Tuple[float, str]

may_fire(*iteration*, *epoch_length*)

Flags if the trigger may fire at the current iteration

This must not alter the trigger state

Parameters

- **iteration** (*int*) –
- **epoch_length** (*int*) –

Return type

bool

pytorch_pfn_extras.training.trigger.Trigger

class pytorch_pfn_extras.training.trigger.Trigger

Bases: object

Base class for triggers.

Methods

__init__()

load_state_dict(state)

| | |
|--|--|
| <i>may_fire</i> (iteration, epoch_len) | Flags if the trigger may fire at the current iteration |
| <i>state_dict</i> () | |

__call__(manager)

Call self as a function.

Parameters

manager (*ExtensionsManagerProtocol*) –

Return type

bool

load_state_dict(state)

Parameters

state (*Dict*[str, Any]) –

Return type

None

may_fire(*iteration*, *epoch_len*)

Flags if the trigger may fire at the current iteration

This must not alter the trigger state

Parameters

- **iteration** (*int*) –
- **epoch_len** (*int*) –

Return type

bool

state_dict()**Return type***Dict*[*str*, *Any*]**pytorch_pfn_extras.training.triggers****Classes**

| | |
|--|--|
| <code>pytorch_pfn_extras.training.triggers.BestValueTrigger(...)</code> | Trigger invoked when specific value becomes best. |
| <code>pytorch_pfn_extras.training.triggers.EarlyStoppingTrigger(self)</code> | Trigger for Early Stopping |
| <code>pytorch_pfn_extras.training.triggers.FunctionTrigger(fn)</code> | |
| <code>pytorch_pfn_extras.training.triggers.IntervalTrigger(...)</code> | Trigger based on a fixed interval. |
| <code>pytorch_pfn_extras.training.triggers.ManualScheduleTrigger(...)</code> | Trigger invoked at specified point(s) of iterations or epochs. |
| <code>pytorch_pfn_extras.training.triggers.MaxValueTrigger(key)</code> | Trigger invoked when specific value becomes maximum. |
| <code>pytorch_pfn_extras.training.triggers.MinValueTrigger(key)</code> | Trigger invoked when specific value becomes minimum. |
| <code>pytorch_pfn_extras.training.triggers.OnceTrigger([...])</code> | Trigger based on the starting point of the iteration. |
| <code>pytorch_pfn_extras.training.triggers.TimeTrigger(period)</code> | Trigger based on a fixed time interval. |

pytorch_pfn_extras.training.triggers.BestValueTrigger**class** `pytorch_pfn_extras.training.triggers.BestValueTrigger`(*key*, *compare*, *trigger*=(1, 'epoch'))Bases: *Trigger*

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** (*TriggerLike*) – 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

`__init__(key, compare[, trigger])`

`load_state_dict(to_load)`

| | |
|--|--|
| <code>may_fire(iteration, epoch_length)</code> | Flags if the trigger may fire at the current iteration |
| <code>state_dict()</code> | |

`__call__(manager)`

Decides whether the extension should be called on this iteration.

Parameters

manager (`ExtensionsManager`) – Manager object that this trigger is associated with. The observation of this manager is used to determine if the trigger should fire.

Returns

True if the corresponding extension should be invoked in this iteration.

Return type

bool

`__init__(key, compare, trigger=(1, 'epoch'))`

Parameters

- **key** (`str`) –
- **compare** (`Callable[[float, float], bool]`) –
- **trigger** (`TriggerLike`) –

Return type

None

`load_state_dict(to_load)`

Parameters

to_load (`Dict[str, Any]`) –

Return type

None

`may_fire(iteration, epoch_length)`

Flags if the trigger may fire at the current iteration

This must not alter the trigger state

Parameters

- **iteration** (`int`) –
- **epoch_length** (`int`) –

Return type

bool

`state_dict()`

Return type

`Dict[str, Any]`

pytorch_pfn_extras.training.triggers.EarlyStoppingTrigger

```
class pytorch_pfn_extras.training.triggers.EarlyStoppingTrigger(self, check_trigger=(1, 'epoch'),
                                                                monitor='main/loss',
                                                                patience=3, mode='auto',
                                                                verbose=False,
                                                                max_trigger=(100, 'epoch'))
```

Bases: [Trigger](#)

Trigger for Early Stopping

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 `patience`), this trigger fires.

Parameters

- **monitor** (*str*) – The metric you want to monitor
- **check_trigger** (*TriggerLike*) – 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`.
- **patience** (*int*) – Counts to let the trigger be patient. The trigger will not fire until the condition is met for successive `patience` 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** (*Tuple[int, UnitLiteral]*) – Upper bound of the number of training loops

Methods

`__init__`(`[check_trigger, monitor, patience, ...]`)

`get_training_length`()

`load_state_dict`(`state`)

`may_fire`(`iteration, epoch_length`) Flags if the trigger may fire at the current iteration
`state_dict`()

`__call__`(*manager*)

Decides whether the training loop should be stopped.

Parameters

manager ([ExtensionsManager](#)) – Manager object that this trigger is associated with. The observation of this manager is used to determine if the trigger should fire.

Returns

True if the training loop should be stopped.

Return type

bool

__init__(*check_trigger*=(1, 'epoch'), *monitor*='main/loss', *patience*=3, *mode*='auto', *verbose*=False, *max_trigger*=(100, 'epoch'))

Parameters

- **check_trigger** (*TriggerLike*) –
- **monitor** (*str*) –
- **patience** (*int*) –
- **mode** (*str*) –
- **verbose** (*bool*) –
- **max_trigger** (*Tuple[int, UnitLiteral]*) –

Return type

None

get_training_length()

Return type*Tuple*[float, str]

may_fire(*iteration*, *epoch_length*)

Flags if the trigger may fire at the current iteration

This must not alter the trigger state

Parameters

- **iteration** (*int*) –
- **epoch_length** (*int*) –

Return type

bool

pytorch_pfn_extras.training.triggers.FunctionTrigger

class pytorch_pfn_extras.training.triggers.**FunctionTrigger**(*fn*, *args*=None, *kwargs*=None, *trigger*=(1, 'iteration'))

Bases: *Trigger*

Methods

__init__(*fn*[, *args*, *kwargs*, *trigger*])

load_state_dict(*to_load*)

may_fire(*iteration*, *epoch_len*) Flags if the trigger may fire at the current iteration

state_dict()

Parameters

- **fn** (*Callable*[[...], bool]) –
- **args** (*Optional*[*Sequence*[Any]]) –
- **kwargs** (*Optional*[*Mapping*[str, Any]]) –
- **trigger** (*TriggerLike*) –

__call__(*manager*)

Call self as a function.

Parameters

manager (*ExtensionsManagerProtocol*) –

Return type

bool

__init__(*fn*, *args=None*, *kwargs=None*, *trigger=(1, 'iteration')*)

Parameters

- **fn** (*Callable*[[...], bool]) –
- **args** (*Optional*[*Sequence*[Any]]) –
- **kwargs** (*Optional*[*Mapping*[str, Any]]) –
- **trigger** (*TriggerLike*) –

Return type

None

load_state_dict(*to_load*)

Parameters

to_load (*Dict*[str, Any]) –

Return type

None

may_fire(*iteration*, *epoch_len*)

Flags if the trigger may fire at the current iteration

This must not alter the trigger state

Parameters

- **iteration** (*int*) –
- **epoch_len** (*int*) –

Return type

bool

state_dict()

Return type

Dict[str, Any]

pytorch_pfn_extras.training.triggers.IntervalTrigger

class pytorch_pfn_extras.training.triggers.IntervalTrigger(*period, unit*)

Bases: [Trigger](#)

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 manager.

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

`__init__`(*period, unit*)

`get_training_length`()

`load_state_dict`(*state*)

| | |
|--|--|
| <code>may_fire</code> (<i>iteration, epoch_length</i>) | Flags if the trigger may fire at the current iteration |
|--|--|

`state_dict`()

`__call__`(*manager*)

Decides whether the extension should be called on this iteration.

Parameters

manager ([ExtensionsManager](#)) – Manager object that this trigger is associated with. The iteration related information in this manager is used to determine if the trigger should fire.

Returns

True if the corresponding extension should be invoked in this iteration.

Return type

bool

`__init__`(*period, unit*)

Parameters

- **period** (*float*) –
- **unit** (*UnitLiteral*) –

get_training_length()

Return type
Tuple[float, str]

may_fire(*iteration, epoch_length*)

Flags if the trigger may fire at the current iteration

This must not alter the trigger state

Parameters

- **iteration** (*int*) –
- **epoch_length** (*int*) –

Return type
bool

pytorch_pfn_extras.training.triggers.ManualScheduleTrigger

class pytorch_pfn_extras.training.triggers.**ManualScheduleTrigger**(*points, unit*)

Bases: *Trigger*

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 manager.

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

__init__(points, unit)

load_state_dict(state)

| | |
|---|--|
| <i>may_fire</i> (iteration, epoch_length) | Flags if the trigger may fire at the current iteration |
| state_dict() | |

__call__(manager)

Decides whether the extension should be called on this iteration.

Parameters

manager (*ExtensionsManager*) – Manager object that this trigger is associated with. The iteration information in this manager is used to determine if the trigger should fire.

Returns

True if the corresponding extension should be invoked in this iteration.

Return type

bool

`__init__(points, unit)`

Parameters

- **points** (*Union[float, Sequence[float]]*) –
- **unit** (*UnitLiteral*) –

`may_fire(iteration, epoch_length)`

Flags if the trigger may fire at the current iteration

This must not alter the trigger state

Parameters

- **iteration** (*int*) –
- **epoch_length** (*int*) –

Return type

bool

pytorch_pfn_extras.training.triggers.MaxValueTrigger

class pytorch_pfn_extras.training.triggers.MaxValueTrigger(*key, trigger=(1, 'epoch')*)

Bases: [BestValueTrigger](#)

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** (*TriggerLike*) – 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

`__init__(key[, trigger])`

`load_state_dict(to_load)`

| | |
|--|--|
| <code>may_fire(iteration, epoch_length)</code> | Flags if the trigger may fire at the current iteration |
| <code>state_dict()</code> | |

```
__init__(key, trigger=(1, 'epoch'))
```

Parameters

- **key** (*str*) –
- **trigger** (*TriggerLike*) –

pytorch_pfn_extras.training.triggers.MinValueTrigger

```
class pytorch_pfn_extras.training.triggers.MinValueTrigger(key, trigger=(1, 'epoch'))
```

Bases: *BestValueTrigger*

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** (*TriggerLike*) – 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

```
__init__(key[, trigger])
```

```
load_state_dict(to_load)
```

| | |
|--|--|
| <pre>may_fire(iteration, epoch_length)</pre> | Flags if the trigger may fire at the current iteration |
| <pre>state_dict()</pre> | |

```
__init__(key, trigger=(1, 'epoch'))
```

Parameters

- **key** (*str*) –
- **trigger** (*TriggerLike*) –

pytorch_pfn_extras.training.triggers.OnceTrigger

```
class pytorch_pfn_extras.training.triggers.OnceTrigger(call_on_resume=False)
```

Bases: *Trigger*

Trigger based on the starting point of the iteration.

This trigger accepts only once at starting point of the iteration. There are two ways to specify the starting point: only starting point in whole iteration or called again when training resumed.

Parameters

call_on_resume (*bool*) – Whether the extension is called again or not when restored from a snapshot. It is set to `False` by default.

finished

Flag that indicates whether or not this trigger will

Type

`bool`

fire in the future. This flag is used to determine if the extension should be initialized after resume.

Methods

`__init__`(`call_on_resume`)

`load_state_dict`(`to_load`)

| | |
|--|--|
| <code>may_fire</code> (<code>iteration</code> , <code>epoch_length</code>) | Flags if the trigger may fire at the current iteration |
|--|--|

`state_dict`()

Attributes

`finished`

`__call__`(*manager*)

Call self as a function.

Parameters

manager (`ExtensionsManagerProtocol`) –

Return type

`bool`

`__init__`(`call_on_resume=False`)

Parameters

call_on_resume (*bool*) –

Return type

`None`

property finished: `bool`

`load_state_dict`(`to_load`)

Parameters

to_load (`Dict[str, Any]`) –

Return type

`None`

may_fire(*iteration*, *epoch_length*)

Flags if the trigger may fire at the current iteration

This must not alter the trigger state

Parameters

- **iteration** (*int*) –
- **epoch_length** (*int*) –

Return type

bool

state_dict()

Return type

Dict[str, *Any*]

pytorch_pfn_extras.training.triggers.TimeTrigger

class pytorch_pfn_extras.training.triggers.**TimeTrigger**(*period*)

Bases: *Trigger*

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

__init__(*period*)

load_state_dict(*to_load*)

may_fire(*iteration*, *epoch_len*)

Flags if the trigger may fire at the current iteration

state_dict()

__call__(*manager*)

Call self as a function.

Parameters

manager (*ExtensionsManagerProtocol*) –

Return type

bool

__init__(*period*)

Parameters

period (*float*) –

Return type

None

load_state_dict(*to_load*)

Parameters

to_load (*Dict[str, Any]*) –

Return type

None

state_dict()

Return type

Dict[str, Any]

Modules

*pytorch_pfn_extras.training.triggers.
early_stopping_trigger*

*pytorch_pfn_extras.training.triggers.
function_trigger*

*pytorch_pfn_extras.training.triggers.
interval_trigger*

*pytorch_pfn_extras.training.triggers.
manual_schedule_trigger*

*pytorch_pfn_extras.training.triggers.
minmax_value_trigger*

*pytorch_pfn_extras.training.triggers.
once_trigger*

*pytorch_pfn_extras.training.triggers.
time_trigger*

pytorch_pfn_extras.training.triggers.early_stopping_trigger

Classes

| | |
|---|----------------------------|
| <i>pytorch_pfn_extras.training. triggers.early_stopping_trigger. EarlyStoppingTrigger(self)</i> | Trigger for Early Stopping |
|---|----------------------------|

*pytorch_pfn_extras.training.
triggers.early_stopping_trigger.
ExtensionsManagerProtocol(...)*

pytorch_pfn_extras.training.triggers.early_stopping_trigger.EarlyStoppingTrigger

```
class pytorch_pfn_extras.training.triggers.early_stopping_trigger.EarlyStoppingTrigger(self,
                                                                                       check_trigger=(1,
                                                                                       'epoch'),
                                                                                       mon-
                                                                                       i-
                                                                                       tor='main/loss',
                                                                                       pa-
                                                                                       tience=3,
                                                                                       mode='auto',
                                                                                       ver-
                                                                                       bose=False,
                                                                                       max_trigger=(100,
                                                                                       'epoch'))
```

Bases: [Trigger](#)

Trigger for Early Stopping

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 `patience`), this trigger fires.

Parameters

- **monitor** (*str*) – The metric you want to monitor
- **check_trigger** (*TriggerLike*) – 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`.
- **patience** (*int*) – Counts to let the trigger be patient. The trigger will not fire until the condition is met for successive `patience` 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** (*Tuple[int, UnitLiteral]*) – Upper bound of the number of training loops

Methods

`__init__`([`check_trigger`, `monitor`, `patience`, ...])

`get_training_length`()

`load_state_dict`(`state`)

`may_fire`(`iteration`, `epoch_length`) Flags if the trigger may fire at the current iteration
`state_dict`()

`__call__`(*manager*)

Decides whether the training loop should be stopped.

Parameters

manager (`ExtensionsManager`) – Manager object that this trigger is associated with.
The observation of this manager is used to determine if the trigger should fire.

Returns

True if the training loop should be stopped.

Return type

bool

```
__init__(check_trigger=(1, 'epoch'), monitor='main/loss', patience=3, mode='auto', verbose=False,  
max_trigger=(100, 'epoch'))
```

Parameters

- **check_trigger** (`TriggerLike`) –
- **monitor** (`str`) –
- **patience** (`int`) –
- **mode** (`str`) –
- **verbose** (`bool`) –
- **max_trigger** (`Tuple[int, UnitLiteral]`) –

Return type

None

```
get_training_length()
```

Return type

`Tuple[float, str]`

```
may_fire(iteration, epoch_length)
```

Flags if the trigger may fire at the current iteration

This must not alter the trigger state

Parameters

- **iteration** (`int`) –
- **epoch_length** (`int`) –

Return type

bool

`pytorch_pfn_extras.training.triggers.early_stopping_trigger.ExtensionsManagerProtocol`

```
class pytorch_pfn_extras.training.triggers.early_stopping_trigger.ExtensionsManagerProtocol(*args,  
                                                                                          **kwargs)
```

Bases: `Protocol`

Methods

`__init__(*args, **kwargs)`

`get_extension(name)`

Attributes

`elapsed_time`

`epoch`

`epoch_detail`

`is_before_training`

`iteration`

`models`

`observation`

`optimizers`

`out`

`raw_models`

`reporter`

`stop_trigger`

`writer`

`__init__(*args, **kwargs)``property elapsed_time: float``property epoch: int``property epoch_detail: float``get_extension(name)`

Parameters

name (*str*) –

Return type

Extension

```

property is_before_training: bool
property iteration: int
property models: Mapping[str, Module]
property observation: reporting.Observation
property optimizers: Mapping[str, Optimizer]
property out: str
property raw_models: Mapping[str, Module]
property reporter: reporting.Reporter
property stop_trigger: bool
property writer: Optional[writing.Writer]

```

pytorch_pfn_extras.training.triggers.function_trigger

Classes

```

pytorch_pfn_extras.training.
triggers.function_trigger.
ExtensionsManagerProtocol(...)
pytorch_pfn_extras.training.triggers.
function_trigger.FunctionTrigger(fn)

```

pytorch_pfn_extras.training.triggers.function_trigger.ExtensionsManagerProtocol

```

class pytorch_pfn_extras.training.triggers.function_trigger.ExtensionsManagerProtocol(*args,
                                                                                       **kwargs)

```

Bases: Protocol

Methods

```

__init__(*args, **kwargs)

```

```

get_extension(name)

```

Attributes

elapsed_time

epoch

epoch_detail

is_before_training

iteration

models

observation

optimizers

out

raw_models

reporter

stop_trigger

writer

__init__(*args, **kwargs)**property** elapsed_time: float**property** epoch: int**property** epoch_detail: float**get_extension**(name)**Parameters****name** (str) –**Return type***Extension***property** is_before_training: bool**property** iteration: int**property** models: Mapping[str, Module]**property** observation: reporting.Observation**property** optimizers: Mapping[str, Optimizer]


```

property out: str

property raw_models: Mapping[str, Module]

property reporter: reporting.Reporter

property stop_trigger: bool

property writer: Optional[writing.Writer]

```

pytorch_pfn_extras.training.triggers.function_trigger.FunctionTrigger

```

class pytorch_pfn_extras.training.triggers.function_trigger.FunctionTrigger(fn, args=None,
                                                                              kwargs=None,
                                                                              trigger=(1,
                                                                              'iteration'))

```

Bases: *Trigger*

Methods

```
__init__(fn[, args, kwargs, trigger])
```

```
load_state_dict(to_load)
```

| | |
|---|--|
| <i>may_fire</i> (<i>iteration</i> , <i>epoch_len</i>) | Flags if the trigger may fire at the current iteration |
| <i>state_dict</i> () | |

Parameters

- **fn** (*Callable*[*...*], *bool*) –
- **args** (*Optional*[*Sequence*[*Any*]]) –
- **kwargs** (*Optional*[*Mapping*[*str*, *Any*]]) –
- **trigger** (*TriggerLike*) –

```
__call__(manager)
```

Call self as a function.

Parameters

manager (*ExtensionsManagerProtocol*) –

Return type

bool

```
__init__(fn, args=None, kwargs=None, trigger=(1, 'iteration'))
```

Parameters

- **fn** (*Callable*[*...*], *bool*) –
- **args** (*Optional*[*Sequence*[*Any*]]) –
- **kwargs** (*Optional*[*Mapping*[*str*, *Any*]]) –

- **trigger** (*TriggerLike*) –

Return type

None

load_state_dict(*to_load*)

Parameters

to_load (*Dict[str, Any]*) –

Return type

None

may_fire(*iteration, epoch_len*)

Flags if the trigger may fire at the current iteration

This must not alter the trigger state

Parameters

- **iteration** (*int*) –

- **epoch_len** (*int*) –

Return type

bool

state_dict()

Return type

Dict[str, Any]

pytorch_pfn_extras.training.triggers.interval_trigger

Classes

*pytorch_pfn_extras.training.
triggers.interval_trigger.
ExtensionsManagerProtocol(...)*

*pytorch_pfn_extras.training.triggers.
interval_trigger.IntervalTrigger(...)*

Trigger based on a fixed interval.

pytorch_pfn_extras.training.triggers.interval_trigger.ExtensionsManagerProtocol

class `pytorch_pfn_extras.training.triggers.interval_trigger.ExtensionsManagerProtocol(*args, **kwargs)`

Bases: `Protocol`

Methods

`__init__(*args, **kwargs)`

`get_extension(name)`

Attributes

`elapsed_time`

`epoch`

`epoch_detail`

`is_before_training`

`iteration`

`models`

`observation`

`optimizers`

`out`

`raw_models`

`reporter`

`stop_trigger`

`writer`

`__init__(*args, **kwargs)``property elapsed_time: float``property epoch: int``property epoch_detail: float``get_extension(name)`

Parameters

name (*str*) –

Return type

Extension

```
property is_before_training: bool
property iteration: int
property models: Mapping[str, Module]
property observation: reporting.Observation
property optimizers: Mapping[str, Optimizer]
property out: str
property raw_models: Mapping[str, Module]
property reporter: reporting.Reporter
property stop_trigger: bool
property writer: Optional[writing.Writer]
```

`pytorch_pfn_extras.training.triggers.interval_trigger.IntervalTrigger`

class `pytorch_pfn_extras.training.triggers.interval_trigger.IntervalTrigger`(*period, unit*)

Bases: *Trigger*

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 manager.

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

`__init__`(period, unit)

`get_training_length`()

`load_state_dict`(state)

`may_fire`(iteration, epoch_length) Flags if the trigger may fire at the current iteration
`state_dict`()

__call__(*manager*)

Decides whether the extension should be called on this iteration.

Parameters

manager ([ExtensionsManager](#)) – Manager object that this trigger is associated with. The iteration related information in this manager is used to determine if the trigger should fire.

Returns

True if the corresponding extension should be invoked in this iteration.

Return type

bool

__init__(*period, unit*)

Parameters

- **period** (*float*) –
- **unit** (*UnitLiteral*) –

get_training_length()

Return type

Tuple[float, str]

may_fire(*iteration, epoch_length*)

Flags if the trigger may fire at the current iteration

This must not alter the trigger state

Parameters

- **iteration** (*int*) –
- **epoch_length** (*int*) –

Return type

bool

pytorch_pfn_extras.training.triggers.manual_schedule_trigger

Classes

pytorch_pfn_extras.training.
triggers.manual_schedule_trigger.
ExtensionsManagerProtocol(...)

pytorch_pfn_extras.training.
triggers.manual_schedule_trigger.
ManualScheduleTrigger(...)

Trigger invoked at specified point(s) of iterations or epochs.

pytorch_pfn_extras.training.triggers.manual_schedule_trigger.ExtensionsManagerProtocol

```
class pytorch_pfn_extras.training.triggers.manual_schedule_trigger.ExtensionsManagerProtocol(*args,  
                                                                                          **kwargs)
```

Bases: Protocol

Methods

```
__init__(*args, **kwargs)
```

```
get_extension(name)
```

Attributes

```
elapsed_time
```

```
epoch
```

```
epoch_detail
```

```
is_before_training
```

```
iteration
```

```
models
```

```
observation
```

```
optimizers
```

```
out
```

```
raw_models
```

```
reporter
```

```
stop_trigger
```

```
writer
```

```
__init__(*args, **kwargs)
```

```
property elapsed_time: float
```

```
property epoch: int
```

```
property epoch_detail: float
```

`get_extension(name)`

Parameters

name (*str*) –

Return type

Extension

`property is_before_training: bool`

`property iteration: int`

`property models: Mapping[str, Module]`

`property observation: reporting.Observation`

`property optimizers: Mapping[str, Optimizer]`

`property out: str`

`property raw_models: Mapping[str, Module]`

`property reporter: reporting.Reporter`

`property stop_trigger: bool`

`property writer: Optional[writing.Writer]`

`pytorch_pfn_extras.training.triggers.manual_schedule_trigger.ManualScheduleTrigger`

`class pytorch_pfn_extras.training.triggers.manual_schedule_trigger.ManualScheduleTrigger(points, unit)`

Bases: *Trigger*

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 manager.

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

| | |
|--|--|
| <code>__init__(points, unit)</code> | |
| <code>load_state_dict(state)</code> | |
| <code>may_fire(iteration, epoch_length)</code> | Flags if the trigger may fire at the current iteration |
| <code>state_dict()</code> | |

`__call__(manager)`

Decides whether the extension should be called on this iteration.

Parameters

manager (`ExtensionsManager`) – Manager object that this trigger is associated with. The iteration information in this manager is used to determine if the trigger should fire.

Returns

True if the corresponding extension should be invoked in this iteration.

Return type

bool

`__init__(points, unit)`

Parameters

- **points** (`Union[float, Sequence[float]]`) –
- **unit** (`UnitLiteral`) –

`may_fire(iteration, epoch_length)`

Flags if the trigger may fire at the current iteration

This must not alter the trigger state

Parameters

- **iteration** (`int`) –
- **epoch_length** (`int`) –

Return type

bool

`pytorch_pfn_extras.training.triggers.minmax_value_trigger`

Classes

| | |
|---|--|
| <code>pytorch_pfn_extras.training.triggers.minmax_value_trigger.BestValueTrigger(...)</code> | Trigger invoked when specific value becomes best. |
| <code>pytorch_pfn_extras.training.triggers.minmax_value_trigger.ExtensionsManagerProtocol(...)</code> | |
| <code>pytorch_pfn_extras.training.triggers.minmax_value_trigger.MaxValueTrigger(key)</code> | Trigger invoked when specific value becomes maximum. |
| <code>pytorch_pfn_extras.training.triggers.minmax_value_trigger.MinValueTrigger(key)</code> | Trigger invoked when specific value becomes minimum. |

pytorch_pfn_extras.training.triggers.minmax_value_trigger.BestValueTrigger

class pytorch_pfn_extras.training.triggers.minmax_value_trigger.**BestValueTrigger**(*key*, *compare*, *trigger*=(1, 'epoch'))

Bases: *Trigger*

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** (*TriggerLike*) – 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

| | |
|--|--|
| <code>__init__(key, compare[, trigger])</code> | |
| <code>load_state_dict(to_load)</code> | |
| <code>may_fire(iteration, epoch_length)</code> | Flags if the trigger may fire at the current iteration |
| <code>state_dict()</code> | |

`__call__(manager)`

Decides whether the extension should be called on this iteration.

Parameters

manager (*ExtensionsManager*) – Manager object that this trigger is associated with. The observation of this manager is used to determine if the trigger should fire.

Returns

True if the corresponding extension should be invoked in this iteration.

Return type

bool

`__init__(key, compare, trigger=(1, 'epoch'))`**Parameters**

- **key** (*str*) –
- **compare** (*Callable[[float, float], bool]*) –
- **trigger** (*TriggerLike*) –

Return type

None

`load_state_dict(to_load)`**Parameters**`to_load` (*Dict[str, Any]*) –**Return type**

None

`may_fire(iteration, epoch_length)`

Flags if the trigger may fire at the current iteration

This must not alter the trigger state

Parameters

- **iteration** (*int*) –
- **epoch_length** (*int*) –

Return type

bool

`state_dict()`**Return type***Dict[str, Any]***pytorch_pfn_extras.training.triggers.minmax_value_trigger.ExtensionsManagerProtocol**

```
class pytorch_pfn_extras.training.triggers.minmax_value_trigger.ExtensionsManagerProtocol(*args,  
                                                                                          **kwargs)
```

Bases: Protocol

Methods

`__init__(*args, **kwargs)`

`get_extension(name)`

Attributes

elapsed_time

epoch

epoch_detail

is_before_training

iteration

models

observation

optimizers

out

raw_models

reporter

stop_trigger

writer

`__init__`(**args*, ***kwargs*)**property** `elapsed_time`: float**property** `epoch`: int**property** `epoch_detail`: float**get_extension**(*name*)**Parameters****name** (*str*) –**Return type***Extension***property** `is_before_training`: bool**property** `iteration`: int**property** `models`: Mapping[str, Module]**property** `observation`: reporting.Observation**property** `optimizers`: Mapping[str, Optimizer]

```
property out: str
property raw_models: Mapping[str, Module]
property reporter: reporting.Reporter
property stop_trigger: bool
property writer: Optional[writing.Writer]
```

pytorch_pfn_extras.training.triggers.minmax_value_trigger.MaxValueTrigger

```
class pytorch_pfn_extras.training.triggers.minmax_value_trigger.MaxValueTrigger(key,
                                                                                 trigger=(1,
                                                                                 'epoch'))
```

Bases: *BestValueTrigger*

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** (*TriggerLike*) – 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

```
__init__(key[, trigger])
```

```
load_state_dict(to_load)
```

```
may_fire(iteration, epoch_length)           Flags if the trigger may fire at the current iteration
state_dict()
```

```
__init__(key, trigger=(1, 'epoch'))
```

Parameters

- **key** (*str*) –
- **trigger** (*TriggerLike*) –

pytorch_pfn_extras.training.triggers.minmax_value_trigger.MinValueTrigger

```
class pytorch_pfn_extras.training.triggers.minmax_value_trigger.MinValueTrigger(key,
                                                                                trigger=(1,
                                                                                'epoch'))
```

Bases: *BestValueTrigger*

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** (*TriggerLike*) – 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

| | |
|--|--|
| <code>__init__(key[, trigger])</code> | |
| <code>load_state_dict(to_load)</code> | |
| <code>may_fire(iteration, epoch_length)</code> | Flags if the trigger may fire at the current iteration |
| <code>state_dict()</code> | |

```
__init__(key, trigger=(1, 'epoch'))
```

Parameters

- **key** (*str*) –
- **trigger** (*TriggerLike*) –

pytorch_pfn_extras.training.triggers.once_trigger**Classes**

| | |
|--|---|
| <code>pytorch_pfn_extras.training.triggers. once_trigger.ExtensionsManagerProtocol(...)</code> | |
| <code>pytorch_pfn_extras.training.triggers. once_trigger.OnceTrigger(...)</code> | Trigger based on the starting point of the iteration. |

pytorch_pfn_extras.training.triggers.once_trigger.ExtensionsManagerProtocol

```
class pytorch_pfn_extras.training.triggers.once_trigger.ExtensionsManagerProtocol(*args,  
                                                                                  **kwargs)
```

Bases: Protocol

Methods

```
__init__(*args, **kwargs)
```

```
get_extension(name)
```

Attributes

```
elapsed_time
```

```
epoch
```

```
epoch_detail
```

```
is_before_training
```

```
iteration
```

```
models
```

```
observation
```

```
optimizers
```

```
out
```

```
raw_models
```

```
reporter
```

```
stop_trigger
```

```
writer
```

```
__init__(*args, **kwargs)
```

```
property elapsed_time: float
```

```
property epoch: int
```

```
property epoch_detail: float
```

`get_extension(name)`

Parameters

name (*str*) –

Return type

Extension

property is_before_training: `bool`

property iteration: `int`

property models: `Mapping[str, Module]`

property observation: `reporting.Observation`

property optimizers: `Mapping[str, Optimizer]`

property out: `str`

property raw_models: `Mapping[str, Module]`

property reporter: `reporting.Reporter`

property stop_trigger: `bool`

property writer: `Optional[writing.Writer]`

`pytorch_pfn_extras.training.triggers.once_trigger.OnceTrigger`

`class pytorch_pfn_extras.training.triggers.once_trigger.OnceTrigger(call_on_resume=False)`

Bases: `Trigger`

Trigger based on the starting point of the iteration.

This trigger accepts only once at starting point of the iteration. There are two ways to specify the starting point: only starting point in whole iteration or called again when training resumed.

Parameters

call_on_resume (*bool*) – Whether the extension is called again or not when restored from a snapshot. It is set to `False` by default.

finished

Flag that indicates whether or not this trigger will

Type

`bool`

fire in the future. This flag is used to determine if the extension should be initialized after resume.

Methods

| | |
|--|--|
| <code>__init__</code> (<code>call_on_resume</code>) | |
| <code>load_state_dict</code> (<code>to_load</code>) | |
| <code>may_fire</code> (<code>iteration</code> , <code>epoch_length</code>) | Flags if the trigger may fire at the current iteration |
| <code>state_dict</code> () | |

Attributes

| |
|-----------------------|
| <code>finished</code> |
|-----------------------|

`__call__`(*manager*)
Call self as a function.

Parameters
`manager` (`ExtensionsManagerProtocol`) –

Return type
`bool`

`__init__`(*call_on_resume=False*)

Parameters
`call_on_resume` (*bool*) –

Return type
`None`

property finished: `bool`

`load_state_dict`(*to_load*)

Parameters
`to_load` (`Dict[str, Any]`) –

Return type
`None`

`may_fire`(*iteration*, *epoch_length*)
Flags if the trigger may fire at the current iteration
This must not alter the trigger state

Parameters

- `iteration` (*int*) –
- `epoch_length` (*int*) –

Return type
`bool`

`state_dict()`

Return type
Dict[str, Any]

`pytorch_pfn_extras.training.triggers.time_trigger`

Classes

| | |
|---|---|
| <code>pytorch_pfn_extras.training.triggers.time_trigger.ExtensionsManagerProtocol(...)</code> | |
| <code>pytorch_pfn_extras.training.triggers.time_trigger.TimeTrigger(period)</code> | Trigger based on a fixed time interval. |

`pytorch_pfn_extras.training.triggers.time_trigger.ExtensionsManagerProtocol`

`class pytorch_pfn_extras.training.triggers.time_trigger.ExtensionsManagerProtocol(*args, **kwargs)`

Bases: Protocol

Methods

| |
|--|
| <code>__init__(*args, **kwargs)</code> |
| <code>get_extension(name)</code> |

Attributes

elapsed_time

epoch

epoch_detail

is_before_training

iteration

models

observation

optimizers

out

raw_models

reporter

stop_trigger

writer

__init__(*args, **kwargs)**property** elapsed_time: float**property** epoch: int**property** epoch_detail: float**get_extension**(name)**Parameters****name** (str) –**Return type***Extension***property** is_before_training: bool**property** iteration: int**property** models: Mapping[str, Module]**property** observation: reporting.Observation**property** optimizers: Mapping[str, Optimizer]

```

property out: str

property raw_models: Mapping[str, Module]

property reporter: reporting.Reporter

property stop_trigger: bool

property writer: Optional[writing.Writer]

```

pytorch_pfn_extras.training.triggers.time_trigger.TimeTrigger

class pytorch_pfn_extras.training.triggers.time_trigger.TimeTrigger(*period*)

Bases: *Trigger*

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

__init__(*period*)

load_state_dict(*to_load*)

| | |
|---|--|
| <i>may_fire</i> (<i>iteration</i> , <i>epoch_len</i>) | Flags if the trigger may fire at the current iteration |
|---|--|

state_dict()

__call__(*manager*)

Call self as a function.

Parameters

manager (*ExtensionsManagerProtocol*) –

Return type

bool

__init__(*period*)

Parameters

period (*float*) –

Return type

None

load_state_dict(*to_load*)

Parameters

to_load (*Dict[str, Any]*) –

Return type

None

state_dict()

Return type

Dict[str, Any]

pytorch_pfn_extras.utils

Modules

pytorch_pfn_extras.utils.checkpoint

pytorch_pfn_extras.utils.comparer

pytorch_pfn_extras.utils.checkpoint

Functions

*pytorch_pfn_extras.utils.checkpoint.
checkpoint(...)*

pytorch_pfn_extras.utils.checkpoint.checkpoint

pytorch_pfn_extras.utils.checkpoint.**checkpoint** (*function*, **args*, ***kwargs*)

Parameters

- **function** (*Module*) –
- **args** (*Any*) –
- **kwargs** (*Any*) –

Return type

Any

pytorch_pfn_extras.utils.comparer

Functions

| | |
|---|------------------------------------|
| <i>pytorch_pfn_extras.utils.comparer. get_default_comparer(...)</i> | Creates default comparer function. |
| <i>pytorch_pfn_extras.utils.comparer. intermediate_value(...)</i> | |

pytorch_pfn_extras.utils.comparer.get_default_comparer

`pytorch_pfn_extras.utils.comparer.get_default_comparer(rtol=0.0001, atol=0, equal_nan=True)`

Creates default comparer function.

The created function will compare the outputs by using `torch.testing.assert_allclose` with specified options.

Parameters

- **rtol** (*float*) – Relative tolerance.
- **atol** (*float*) – Absolute tolerance.
- **equal_nan** (*bool*) – If True, NaNs will be ignored.

Return type

Callable[[*str*, *str*, *str*, *Any*, *Any*], *None*]

pytorch_pfn_extras.utils.comparer.intermediate_value

`pytorch_pfn_extras.utils.comparer.intermediate_value(name, value)`

Parameters

- **name** (*str*) –
- **value** (*Tensor*) –

Return type

None

Classes

| | |
|---|---|
| <code>pytorch_pfn_extras.utils.comparer.Comparer(*)</code> | A class for comparison of iteration outputs and model parameters. |
| <code>pytorch_pfn_extras.utils.comparer.ModelComparer(engines)</code> | A class for comparison of iteration model parameters. |
| <code>pytorch_pfn_extras.utils.comparer.OutputsComparer(engines)</code> | A class for comparison of iteration outputs. |

pytorch_pfn_extras.utils.comparer.Comparer

```
class pytorch_pfn_extras.utils.comparer.Comparer(*, trigger=None, compare_fn=<function
get_default_comparer:<locals>.compare_fn>,
concurrency=None, outputs=True, params=False,
baseline=None)
```

Bases: `object`

A class for comparison of iteration outputs and model parameters.

This class is mainly used to compare results between different devices.

Parameters

- **trigger** (*Trigger*) – Trigger object that determines when to compare values.

- **compare_fn** (*function*) – Comparison function. Default is `get_default_comparer()`.
- **concurrency** (*int, optional*) – The upper bound limit on the number of workers that run concurrently. If `None`, inferred from the size of `engines`.
- **outputs** (*tuple of str or bool*) – A set of keys of output dict to compare.
- **params** (*tuple of str or bool*) – A set of keys of model parameters to compare.
- **baseline** (*str, optional*) – The baseline engine that is assumed to be correct.

Examples

```
>>> trainer_cpu = ppe.engine.create_trainer(
    model, optimizer, 1, device='cpu')
>>> trainer_gpu = ppe.engine.create_trainer(
    model, optimizer, 1, device='cuda:0')
>>> comp = ppe.utils.comparer.Comparer()
>>> comp.add_engine("cpu", engine_cpu, train_1, eval_1)
>>> comp.add_engine("gpu", engine_gpu, train_2, eval_2)
>>> comp.compare()
```

Methods

| | |
|---|---|
| <code>__init__</code> (*[, trigger, compare_fn, ...]) | A class for comparison of iteration outputs and model parameters. |
| <code>add_dump</code> (name, dir) | Add an engine to compare variables. |
| <code>add_engine</code> (name, engine, *args, **kwargs) | Add an engine to compare variables. |
| <code>compare</code> () | Compares outputs. |
| <code>dump</code> (engine, dir, *args, **kwargs) | Add an engine to compare variables. |

`__init__`(*[, trigger=None, compare_fn=<function get_default_comparer.<locals>.compare_fn>, concurrency=None, outputs=True, params=False, baseline=None)

A class for comparison of iteration outputs and model parameters.

This class is mainly used to compare results between different devices.

Parameters

- **trigger** (`Trigger`) – Trigger object that determines when to compare values.
- **compare_fn** (*function*) – Comparison function. Default is `get_default_comparer()`.
- **concurrency** (*int, optional*) – The upper bound limit on the number of workers that run concurrently. If `None`, inferred from the size of `engines`.
- **outputs** (*tuple of str or bool*) – A set of keys of output dict to compare.
- **params** (*tuple of str or bool*) – A set of keys of model parameters to compare.
- **baseline** (*str, optional*) – The baseline engine that is assumed to be correct.

Return type

`None`

Examples

```
>>> trainer_cpu = ppe.engine.create_trainer(
    model, optimizer, 1, device='cpu')
>>> trainer_gpu = ppe.engine.create_trainer(
    model, optimizer, 1, device='cuda:0')
>>> comp = ppe.utils.comparer.Comparer()
>>> comp.add_engine("cpu", engine_cpu, train_1, eval_1)
>>> comp.add_engine("gpu", engine_gpu, train_2, eval_2)
>>> comp.compare()
```

add_dump(*name*, *dir*)

Add an engine to compare variables.

Parameters

- **name** (*str*) – The name of dump.
- **dir** (*str*) – The directory that the results are saved to.

Return type

None

add_engine(*name*, *engine*, **args*, ***kwargs*)

Add an engine to compare variables.

Parameters

- **name** (*str*) – Engine name.
- **engine** (*Trainer or Evaluator*) – An engine to compare variables.
- ****kwargs** (**args and*) – Arguments passed to `engine.run`.
- **args** (*Any*) –
- **kwargs** (*Any*) –

Return type

None

compare()

Compares outputs.

Return type

None

dump(*engine*, *dir*, **args*, ***kwargs*)

Add an engine to compare variables.

Parameters

- **engine** (*Trainer or Evaluator*) – An engine to compare variables.
- **dir** (*str*) – Name of the directory that the results are saved to.
- ****kwargs** (**args and*) – Arguments passed to `engine.run`.
- **args** (*Any*) –
- **kwargs** (*Any*) –

Return type

None

pytorch_pfn_extras.utils.comparer.ModelComparer

```
class pytorch_pfn_extras.utils.comparer.ModelComparer(engines, to_compare_keys=None, *,
                                                    compare_fn=<function
                                                    get_default_comparer.<locals>.compare_fn>,
                                                    concurrency=None)
```

Bases: `_ComparerBase`

A class for comparison of iteration model parameters.

This class is mainly used to compare results between different devices.

Parameters

- **engines** (*dict of Engines*) – Trainers or Evaluators to compare outputs.
- **to_compare_keys** (*tuple of str, optional*) – A set of keys of model parameters to compare.
- **compare_fn** (*function*) – Comparison function. Default is `get_default_comparer()`.
- **concurrency** (*int, optional*) – The upper bound limit on the number of workers that run concurrently. If None, inferred from the size of **engines**.

Examples

```
>>> trainer_cpu = ppe.engine.create_trainer(
    model, optimizer, 1, device='cpu')
>>> trainer_gpu = ppe.engine.create_trainer(
    model, optimizer, 1, device='cuda:0')
>>> comp = ppe.utils.comparer.ModelComparer(
    {"cpu": trainer_cpu, "gpu": trainer_gpu})
>>> comp.compare({"cpu": loader, "gpu": loader})
```

Methods

| | |
|---|---|
| <code>__init__(engines[, to_compare_keys, ...])</code> | A class for comparison of iteration model parameters. |
| <code>compare(loaders[, n_iters])</code> | Compares outputs. |
| <code>compare_targets(name, engine, batch_idx, target)</code> | |
| <code>run_engine(engine, loaders)</code> | |

```
__init__(engines, to_compare_keys=None, *, compare_fn=<function
get_default_comparer.<locals>.compare_fn>, concurrency=None)
```

A class for comparison of iteration model parameters.

This class is mainly used to compare results between different devices.

Parameters

- **engines** (*dict of Engines*) – Trainers or Evaluators to compare outputs.

- **to_compare_keys** (*tuple of str, optional*) – A set of keys of model parameters to compare.
- **compare_fn** (*function*) – Comparison function. Default is `get_default_comparer()`.
- **concurrency** (*int, optional*) – The upper bound limit on the number of workers that run concurrently. If None, inferred from the size of **engines**.

Examples

```
>>> trainer_cpu = ppe.engine.create_trainer(
    model, optimizer, 1, device='cpu')
>>> trainer_gpu = ppe.engine.create_trainer(
    model, optimizer, 1, device='cuda:0')
>>> comp = ppe.utils.comparer.ModelComparer(
    {"cpu": trainer_cpu, "gpu": trainer_gpu})
>>> comp.compare({"cpu": loader, "gpu": loader})
```

pytorch_pfn_extras.utils.comparer.OutputsComparer

```
class pytorch_pfn_extras.utils.comparer.OutputsComparer(engines, to_compare_keys=None, *,
    compare_fn=<function
    get_default_comparer.<locals>.compare_fn>,
    concurrency=None)
```

Bases: `_ComparerBase`

A class for comparison of iteration outputs.

This class is mainly used to compare results between different devices.

Parameters

- **engines** (*dict of Engines*) – Trainers or Evaluators to compare outputs.
- **to_compare_keys** (*tuple of str, optional*) – A set of keys of output dict to compare.
- **compare_fn** (*function*) – Comparison function. Default is `get_default_comparer()`.
- **concurrency** (*int, optional*) – The upper bound limit on the number of workers that run concurrently. If None, inferred from the size of **engines**.

Examples

```
>>> trainer_cpu = ppe.engine.create_trainer(
    model, optimizer, 1, device='cpu')
>>> trainer_gpu = ppe.engine.create_trainer(
    model, optimizer, 1, device='cuda:0')
>>> comp = ppe.utils.comparer.OutputsComparer(
    {"cpu": trainer_cpu, "gpu": trainer_gpu})
>>> comp.compare({"cpu": loader, "gpu": loader})
```

Methods

| | |
|---|--|
| <code>__init__(engines[, to_compare_keys, ...])</code> | A class for comparison of iteration outputs. |
| <code>compare(loaders[, n_iters])</code> | Compares outputs. |
| <code>compare_targets(name, engine, batch_idx, target)</code> | |
| <code>run_engine(engine, loaders)</code> | |

`__init__(engines, to_compare_keys=None, *, compare_fn=<function
get_default_comparer.<locals>.compare_fn>, concurrency=None)`

A class for comparison of iteration outputs.

This class is mainly used to compare results between different devices.

Parameters

- **engines** (*dict of Engines*) – Trainers or Evaluators to compare outputs.
- **to_compare_keys** (*tuple of str, optional*) – A set of keys of output dict to compare.
- **compare_fn** (*function*) – Comparison function. Default is `get_default_comparer()`.
- **concurrency** (*int, optional*) – The upper bound limit on the number of workers that run concurrently. If `None`, inferred from the size of **engines**.

Return type

`None`

Examples

```
>>> trainer_cpu = ppe.engine.create_trainer(
    model, optimizer, 1, device='cpu')
>>> trainer_gpu = ppe.engine.create_trainer(
    model, optimizer, 1, device='cuda:0')
>>> comp = ppe.utils.comparer.OutputsComparer(
    {"cpu": trainer_cpu, "gpu": trainer_gpu})
>>> comp.compare({"cpu": loader, "gpu": loader})
```

Classes

| | |
|---|---|
| <code>pytorch_pfn_extras.writing.ProcessQueueWriter(...)</code> | Snapshot writer that uses process queue. |
| <code>pytorch_pfn_extras.writing.ProcessWriter(...)</code> | Snapshot writer that uses a separate process. |
| <code>pytorch_pfn_extras.writing.QueueWriter(...)</code> | Base class of queue snapshot writers. |
| <code>pytorch_pfn_extras.writing.SimpleWriter(...)</code> | The most simple snapshot writer. |
| <code>pytorch_pfn_extras.writing.StandardWriter(...)</code> | Base class of snapshot writers which use thread or process. |
| <code>pytorch_pfn_extras.writing.TensorBoardWriter(...)</code> | Writer that sends statistics to TensorBoard. |
| <code>pytorch_pfn_extras.writing.ThreadQueueWriter(...)</code> | Snapshot writer that uses a thread queue. |
| <code>pytorch_pfn_extras.writing.ThreadWriter(...)</code> | Snapshot writer that uses a separate thread. |
| <code>pytorch_pfn_extras.writing.Writer(fs, out_dir)</code> | Base class of snapshot writers. |

pytorch_pfn_extras.writing.ProcessQueueWriter

class `pytorch_pfn_extras.writing.ProcessQueueWriter`(*savefun=<function save>, fs=None, out_dir="", task=None*)

Bases: `QueueWriter[Process]`

Snapshot writer that uses process queue.

This class creates a process and a queue by multiprocessing module. The process will be a consumer of this queue, and the main process will be a producer of this queue.

Note: Forking a new process from MPI process might be danger. Consider using `ThreadQueueWriter` instead of `ProcessQueueWriter` if you are using MPI.

See also:

- `pytorch_pfn_extras.training.extensions.snapshot()`

Methods

| | |
|------------------------------|---|
| <code>__init__</code> | ([savefun, fs, out_dir, task]) |
| <code>consume</code> | (q) |
| <code>create_consumer</code> | (q) |
| <code>create_queue</code> | () |
| <code>create_task</code> | (savefun) |
| <code>finalize</code> | Finalizes the writer. |
| <code>initialize</code> | (out_dir) |
| <code>save</code> | (filename, out_dir, target, savefun, ...) |

`__init__` (savefun=<function save>, fs=None, out_dir="", task=None)

Parameters

- **savefun** (Callable[[...], None]) –
- **fs** (Optional[Any]) –
- **out_dir** (str) –
- **task** (Optional[Callable[[...], None]]) –

Return type

None

`create_consumer`(q)

Parameters

q (queue.Queue[_QueUnit]) –

Return type

Process

`create_queue`()

Return type

queue.Queue[_QueUnit]

pytorch_pfn_extras.writing.ProcessWriter

class pytorch_pfn_extras.writing.**ProcessWriter**(savefun=<function save>, fs=None, out_dir="",
**kwds)

Bases: [StandardWriter](#)[Process]

Snapshot writer that uses a separate process.

This class creates a new process that invokes the actual saving function.

Note: Forking a new process from a MPI process might be danger. Consider using [ThreadWriter](#) instead of ProcessWriter if you are using MPI.

See also:

- `pytorch_pfn_extras.training.extensions.snapshot()`

Methods

| | |
|--|------------------------------------|
| <code>__init__([savefun, fs, out_dir])</code> | |
| <code>create_worker(filename, out_dir, target, *)</code> | Creates a worker for the snapshot. |
| <code>finalize()</code> | Finalizes the writer. |
| <code>initialize(out_dir)</code> | |
| <code>save(filename, out_dir, target, savefun, ...)</code> | |

`__init__(savefun=<function save>, fs=None, out_dir="", **kws)`

Parameters

- **savefun** (`Callable[[...], None]`) –
- **fs** (`Optional[Any]`) –
- **out_dir** (`str`) –
- **kws** (`Any`) –

Return type

`None`

`create_worker(filename, out_dir, target, *, savefun=None, append=False, **savefun_kwargs)`

Creates a worker for the snapshot.

This method creates a thread or a process to take a snapshot. The created worker must have `start()` and `join()` methods. If the worker has an `exitcode` attribute (e.g., `multiprocessing.Process`), the value will be tested.

Parameters

- **filename** (`str`) –
- **out_dir** (`str`) –
- **target** (`Union[Sequence[Any], Mapping[str, Any]]`) –
- **savefun** (`Optional[Callable[[...], None]]`) –
- **append** (`bool`) –
- **savefun_kwargs** (`Any`) –

Return type

`Process`

pytorch_pfn_extras.writing.QueueWriter

```
class pytorch_pfn_extras.writing.QueueWriter(savefun=<function save>, fs=None, out_dir="",
                                              task=None)
```

Bases: [Writer](#), [Generic\[_Worker\]](#)

Base class of queue snapshot writers.

This class is a base class of snapshot writers that use a queue. A Queue is created when this class is constructed, and every time when `__call__` is invoked, a snapshot task is put into the queue.

Parameters

- **savefun** – Callable object which is passed to the [create_task\(\)](#) if the task is `None`. It takes three arguments: the output file path, the serialized dictionary object, and the optional keyword arguments.
- **fs** – `FileSystem` abstracting interface to implement all the operations. optional, defaults to `None`
- **out_dir** – str. Specifies the directory this writer will use. It takes precedence over the one specified in `__call__` optional, defaults to ''
- **task** – Callable object. Its `__call__` must have a same interface to `Writer.__call__`. This object is directly put into the queue.

See also:

- [pytorch_pfn_extras.training.extensions.snapshot\(\)](#)

Methods

`__init__([savefun, fs, out_dir, task])`

`consume(q)`

`create_consumer(q)`

`create_queue()`

`create_task(savefun)`

`finalize()` Finalizes the writer.

`initialize(out_dir)`

`save(filename, out_dir, target, savefun, ...)`

`__call__(filename, out_dir, target, *, savefun=None, append=False)`

Does the actual writing to the file.

This method is invoked by a `Snapshot` object every time it takes a snapshot.

Parameters

- **filename** (*str*) – Name of the file into which the serialized target is saved. It is a concrete file name, i.e. not a pre-formatted template string.

- **out_dir** (*str*) – Output directory. Corresponds to :py:attr:`ExtensionsManager.out`
<pytorch_pfn_extras.training.ExtensionsManager.out>`.
- **target** (*dict*) – Serialized object which will be saved.
- **savefun** (*callable*) – A callable that accepts a two positional arguments (an object to be serialized, file path) like *torch.save*.
- **append** (*bool*) – Mode used to open the file. True to use the append mode, False to use the write mode (truncates the file if it already exists).

Return type

None

```
__init__(savefun=<function save>, fs=None, out_dir="", task=None)
```

Parameters

- **savefun** (*Callable[[...], None]*) –
- **fs** (*Optional[Any]*) –
- **out_dir** (*str*) –
- **task** (*Optional[Callable[[...], None]]*) –

Return type

None

```
consume(q)
```

Parameters

- **q** (*queue.Queue[_QueUnit]*) –

Return type

None

```
create_consumer(q)
```

Parameters

- **q** (*queue.Queue[_QueUnit]*) –

Return type*_Worker*

```
create_queue()
```

Return type*queue.Queue[_QueUnit]*

```
create_task(savefun)
```

Parameters

- **savefun** (*Callable[[...], None]*) –

Return type*Callable[[...], None]*

```
finalize()
```

Finalizes the writer.

Calling this method on already-finalized Writer does nothing.

Return type

None

pytorch_pfn_extras.writing.SimpleWriter

class pytorch_pfn_extras.writing.**SimpleWriter**(*savefun=<function save>, fs=None, out_dir='', **kwds*)

Bases: [Writer](#)

The most simple snapshot writer.

This class just passes the arguments to the actual saving function.

Parameters

- **savefun** (*Callable[[...], None]*) – Callable object. It takes three arguments: the output file path, the serialized dictionary object, and the optional keyword arguments.
- **fs** (*Any*) – FileSystem abstracting interface to implement all the operations. optional, defaults to None
- **out_dir** (*str*) – str. Specifies the directory this writer will use. It takes precedence over the one specified in `__call__` optional, defaults to ''
- **kwds** (*Any*) – Keyword arguments for the `savefun`.

See also:

- [pytorch_pfn_extras.training.extensions.snapshot\(\)](#)

Methods

| | |
|--|-----------------------|
| <code>__init__</code> (<i>savefun, fs, out_dir</i>) | |
| <code>finalize</code> () | Finalizes the writer. |
| <code>initialize</code> (<i>out_dir</i>) | |
| <code>save</code> (<i>filename, out_dir, target, savefun, ...</i>) | |

__call__(*filename, out_dir, target, *, savefun=None, append=False*)

Does the actual writing to the file.

This method is invoked by a `Snapshot` object every time it takes a snapshot.

Parameters

- **filename** (*str*) – Name of the file into which the serialized target is saved. It is a concrete file name, i.e. not a pre-formatted template string.
- **out_dir** (*str*) – Output directory. Corresponds to `:py:attr:`ExtensionsManager.out` <pytorch_pfn_extras.training.ExtensionsManager.out>`.
- **target** (*dict*) – Serialized object which will be saved.
- **savefun** (*callable*) – A callable that accepts a two positional arguments (an object to be serialized, file path) like `torch.save`.
- **append** (*bool*) – Mode used to open the file. True to use the append mode, False to use the write mode (truncates the file if it already exists).

Return type

None


```
__init__(savefun=<function save>, fs=None, out_dir="", **kws)
```

Parameters

- **savefun** (*Callable*[[...], None]) –
- **fs** (*Optional*[Any]) –
- **out_dir** (*str*) –
- **kws** (*Any*) –

Return type

None

pytorch_pfn_extras.writing.StandardWriter

```
class pytorch_pfn_extras.writing.StandardWriter(savefun=<function save>, fs=None, out_dir="",
                                                **kws)
```

Bases: [Writer](#), [Generic](#)[_Worker]

Base class of snapshot writers which use thread or process.

This class creates a new thread or a process every time when `__call__` is invoked.

Parameters

- **savefun** – Callable object. It takes three arguments: the output file path, the serialized dictionary object, and the optional keyword arguments.
- **fs** – *FileSystem* abstracting interface to implement all the operations. optional, defaults to None
- **out_dir** – str. Specifies the directory this writer will use. It takes precedence over the one specified in `__call__` optional, defaults to ''
- **kws** – Keyword arguments for the `savefun`.

See also:

- [pytorch_pfn_extras.training.extensions.snapshot\(\)](#)

Methods

| | |
|---|------------------------------------|
| <code>__init__</code> ([savefun, fs, out_dir]) | |
| <code>create_worker</code> (filename, out_dir, target, *) | Creates a worker for the snapshot. |
| <code>finalize</code> () | Finalizes the writer. |
| <code>initialize</code> (out_dir) | |
| <code>save</code> (filename, out_dir, target, savefun, ...) | |

```
__call__(filename, out_dir, target, *, savefun=None, append=False)
```

Does the actual writing to the file.

This method is invoked by a `Snapshot` object every time it takes a snapshot.

Parameters

- **filename** (*str*) – Name of the file into which the serialized target is saved. It is a concrete file name, i.e. not a pre-formatted template string.
- **out_dir** (*str*) – Output directory. Corresponds to :py:attr:`ExtensionsManager.out`
<pytorch_pfn_extras.training.ExtensionsManager.out>`.
- **target** (*dict*) – Serialized object which will be saved.
- **savefun** (*callable*) – A callable that accepts a two positional arguments (an object to be serialized, file path) like *torch.save*.
- **append** (*bool*) – Mode used to open the file. True to use the append mode, False to use the write mode (truncates the file if it already exists).

Return type

None

```
__init__(savefun=<function save>, fs=None, out_dir="", **kws)
```

Parameters

- **savefun** (*Callable[[...], None]*) –
- **fs** (*Optional[Any]*) –
- **out_dir** (*str*) –
- **kws** (*Any*) –

Return type

None

```
create_worker(filename, out_dir, target, *, savefun=None, append=False, **savefun_kwargs)
```

Creates a worker for the snapshot.

This method creates a thread or a process to take a snapshot. The created worker must have `start()` and `join()` methods. If the worker has an `exitcode` attribute (e.g., `multiprocessing.Process`), the value will be tested.

Parameters

- **filename** (*str*) –
- **out_dir** (*str*) –
- **target** (*Union[Sequence[Any], Mapping[str, Any]]*) –
- **savefun** (*Optional[Callable[[...], None]]*) –
- **append** (*bool*) –
- **savefun_kwargs** (*Any*) –

Return type*_Worker*

```
finalize()
```

Finalizes the writer.

Calling this method on already-finalized Writer does nothing.

Return type

None

pytorch_pfn_extras.writing.TensorBoardWriter

```
class pytorch_pfn_extras.writing.TensorBoardWriter(savefun=None, fs=None, out_dir="", stats=None,
                                                    **kwds)
```

Bases: object

Writer that sends statistics to TensorBoard.

This class contains a `torch.utils.tensorboard.SummaryWriter` object that is used to send the collected statistics to TensorBoard. A list of stats can be specified to report only the desired ones.

Parameters

- **savefun** (*Optional*[*Callable*[[...], None]]) – Ignored.
- **fs** (*Any*) – Ignored.
- **out_dir** (*str*) – Passed as `log_dir` argument to `SummaryWriter`.
- **stats** (*list*) – List of statistic keys.
- **kwds** (*Any*) – Passed as an additional arguments to `SummaryWriter`.

Methods

```
__init__([savefun, fs, out_dir, stats])
```

```
finalize()
```

```
__call__(filename, out_dir, target, *, savefun=None, append=False)
```

Sends the statistics to the TensorBoard.

Parameters

- **filename** (*str*) – Ignored.
- **out_dir** (*str*) – Ignored.
- **target** (*dict or list*) – The statistics of the iteration. If given as a list, only the last element (assumed to be a dict containing the latest iteration statistics) is reported.
- **savefun** (*Optional*[*Callable*[[...], None]]) – Ignored.
- **append** (*bool*) – Ignored.

Return type

None

```
__init__(savefun=None, fs=None, out_dir="", stats=None, **kwds)
```

Parameters

- **savefun** (*Optional*[*Callable*[[...], None]]) –
- **fs** (*Optional*[*Any*]) –
- **out_dir** (*str*) –
- **stats** (*Optional*[*KeysView*[*str*]]) –
- **kwds** (*Any*) –

Return type

None

finalize()**Return type**

None

pytorch_pfn_extras.writing.ThreadQueueWriter

```
class pytorch_pfn_extras.writing.ThreadQueueWriter(savefun=<function save>, fs=None, out_dir="", task=None)
```

Bases: [QueueWriter](#)[Thread]

Snapshot writer that uses a thread queue.

This class creates a thread and a queue by `threading` and `queue` modules respectively. The thread will be a consumer of the queue, and the main thread will be a producer of the queue.

See also:

- [pytorch_pfn_extras.training.extensions.snapshot\(\)](#)

Methods

`__init__`([savefun, fs, out_dir, task])

`consume`(q)

`create_consumer`(q)

`create_queue`()

`create_task`(savefun)

`finalize`() Finalizes the writer.

`initialize`(out_dir)

`save`(filename, out_dir, target, savefun, ...)

```
__init__(savefun=<function save>, fs=None, out_dir="", task=None)
```

Parameters

- **savefun** ([Callable](#)[[...], None]) –
- **fs** ([Optional](#)[Any]) –
- **out_dir** ([str](#)) –
- **task** ([Optional](#)[[Callable](#)[[...], None]]) –

Return type

None

create_consumer(*q*)

Parameters

q (*queue.Queue[_QueUnit]*) –

Return type

Thread

create_queue()

Return type

queue.Queue[_QueUnit]

pytorch_pfn_extras.writing.ThreadWriter

class `pytorch_pfn_extras.writing.ThreadWriter`(*savefun=<function save>, fs=None, out_dir="", **kws*)

Bases: *StandardWriter*[*Thread*]

Snapshot writer that uses a separate thread.

This class creates a new thread that invokes the actual saving function.

See also:

- `pytorch_pfn_extras.training.extensions.snapshot()`

Methods

| | |
|--|------------------------------------|
| <code>__init__</code> (<i>savefun, fs, out_dir</i>) | |
| <code>create_worker</code> (<i>filename, out_dir, target, *</i>) | Creates a worker for the snapshot. |
| <code>finalize</code> () | Finalizes the writer. |
| <code>initialize</code> (<i>out_dir</i>) | |
| <code>save</code> (<i>filename, out_dir, target, savefun, ...</i>) | |

`__init__`(*savefun=<function save>, fs=None, out_dir="", **kws*)

Parameters

- **savefun** (*Callable[[...], None]*) –
- **fs** (*Optional[Any]*) –
- **out_dir** (*str*) –
- **kws** (*Any*) –

Return type

None

create_worker(*filename, out_dir, target, *, savefun=None, append=False, **savefun_kwargs*)

Creates a worker for the snapshot.

This method creates a thread or a process to take a snapshot. The created worker must have `start()` and `join()` methods. If the worker has an `exitcode` attribute (e.g., `multiprocessing.Process`), the value will be tested.

Parameters

- **filename** (*str*) –
- **out_dir** (*str*) –
- **target** (*Union[Sequence[Any], Mapping[str, Any]]*) –
- **savefun** (*Optional[Callable[[...], None]]*) –
- **append** (*bool*) –
- **savefun_kwargs** (*Any*) –

Return type*Thread***pytorch_pfn_extras.writing.Writer****class** pytorch_pfn_extras.writing.**Writer**(*fs=None, out_dir=""*)

Bases: object

Base class of snapshot writers.

Snapshot invokes `__call__` of this class every time when taking a snapshot. This class determines how the actual saving function will be invoked.

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.

See also:

- [`pytorch_pfn_extras.training.extensions.snapshot\(\)`](#)

Methods

`__init__`(*[fs, out_dir]*)

`finalize`() Finalizes the writer.

`initialize`(*out_dir*)

`save`(*filename, out_dir, target, savefun, ...*)

Parameters

- **fs** (*Any*) –
- **out_dir** (*str*) –

`__call__`(*filename, out_dir, target, *, savefun=None, append=False*)

Does the actual writing to the file.

This method is invoked by a Snapshot object every time it takes a snapshot.

Parameters

- **filename** (*str*) – Name of the file into which the serialized target is saved. It is a concrete file name, i.e. not a pre-formatted template string.
- **out_dir** (*str*) – Output directory. Corresponds to :py:attr:`ExtensionsManager.out`
<pytorch_pfn_extras.training.ExtensionsManager.out>`.
- **target** (*dict*) – Serialized object which will be saved.
- **savefun** (*callable*) – A callable that accepts a two positional arguments (an object to be serialized, file path) like *torch.save*.
- **append** (*bool*) – Mode used to open the file. True to use the append mode, False to use the write mode (truncates the file if it already exists).

Return type

None

__init__(*fs=None, out_dir=""*)**Parameters**

- **fs** (*Optional[Any]*) –
- **out_dir** (*str*) –

Return type

None

finalize()

Finalizes the writer.

Calling this method on already-finalized Writer does nothing.

Return type

None

initialize(*out_dir*)**Parameters****out_dir** (*str*) –**Return type**

None

save(*filename, out_dir, target, savefun, append, **savefun_kwargs*)**Parameters**

- **filename** (*str*) –
- **out_dir** (*str*) –
- **target** (*Union[Sequence[Any], Mapping[str, Any]]*) –
- **savefun** (*Callable[[...], None]*) –
- **append** (*bool*) –
- **savefun_kwargs** (*Any*) –

Return type

None

2.2 Training Loop

2.2.1 Trainer

| | |
|--|---|
| <code>engine.create_trainer(models, optimizers, ...)</code> | Creates a trainer object. |
| <code>engine.create_evaluator(models, *[, ...])</code> | Creates an evaluator object. |
| <code>handler.BaseLogic([options])</code> | |
| <code>handler.Logic([model_name, options])</code> | A set of methods that defines the training logic. |
| <code>handler.BaseHandler(logic, options, *args, ...)</code> | Base class of Handler. |
| <code>handler.Handler(logic, entry_runtime, options)</code> | A set of callback functions to perform device-specific operations. |
| <code>runtime.BaseRuntime(device_spec, options)</code> | A base class for collections of device-specific callback functions. |
| <code>runtime.PyTorchRuntime(device_spec, options)</code> | A collections of callback functions for the devices that PyTorch supports by default. |

2.2.2 Extensions Manager

| | |
|---|--|
| <code>training.ExtensionsManager(models, ...[, ...])</code> | Manages the extensions and the current status. |
| <code>training.IgniteExtensionsManager(engine, ...)</code> | Manages extensions and the current status in Ignite training loop. |

2.2.3 Extensions

| | |
|--|---|
| <code>training.extension.make_extension([trigger, ...])</code> | Decorator to make given function into an extension. |
| <code>training.extension.Extension()</code> | Base class of extensions. |
| <code>training.extension.ExtensionEntry(extension, *)</code> | Extension and options. |

| | |
|---|--|
| <code>training.extensions.BestValue(key, compare)</code> | Extension traces the best value of a specific key in the observation. |
| <code>training.extensions.Evaluator(self, ..., ...)</code> | An extension to evaluate models on a validation set. |
| <code>training.extensions.LogReport([keys, ...])</code> | An extension to output the accumulated results to a log file. |
| <code>training.extensions.MaxValue(key[, trigger])</code> | Extension traces the maximum value of a specific key in the observation. |
| <code>training.extensions.MicroAverage(...[, trigger])</code> | Calculates micro-average ratio. |
| <code>training.extensions.MinValue(key[, trigger])</code> | Extension traces the maximum value of a specific key in the observation. |
| <code>training.extensions.observe_lr(optimizer[, ...])</code> | Returns an extension to record the learning rate. |
| <code>training.extensions.observe_value(...)</code> | Returns an extension to continuously record a value. |
| <code>training.extensions. ParameterStatistics(links)</code> | An extension to report parameter statistics. |
| <code>training.extensions.PlotReport(y_keys[, ...])</code> | An extension to output plots. |
| <code>training.extensions.PrintReport([entries, ...])</code> | An extension to print the accumulated results. |
| <code>training.extensions.ProgressBar([...])</code> | An extension to print a progress bar and recent training status. |
| <code>training.extensions.ProfileReport([...])</code> | Writes the profile results to a file. |
| <code>training.extensions.snapshot([savefun, ...])</code> | Returns a trainer extension to take snapshots of the trainer. |
| <code>training.extensions.Slack(channel[, msg, ...])</code> | An extension to communicate with Slack. |
| <code>training.extensions.SlackWebhook(url[, msg, ...])</code> | An extension to communicate with Slack using Incoming Webhook. |
| <code>training.extensions. VariableStatisticsPlot(targets)</code> | An extension to plot statistics for Tensors. |

2.2.4 Triggers

| | |
|---|--|
| <code>training.triggers.EarlyStoppingTrigger(self)</code> | Trigger for Early Stopping |
| <code>training.triggers.IntervalTrigger(period, unit)</code> | Trigger based on a fixed interval. |
| <code>training.triggers.ManualScheduleTrigger(...)</code> | Trigger invoked at specified point(s) of iterations or epochs. |
| <code>training.triggers.BestValueTrigger(key, compare)</code> | Trigger invoked when specific value becomes best. |
| <code>training.triggers.MaxValueTrigger(key[, trigger])</code> | Trigger invoked when specific value becomes maximum. |
| <code>training.triggers.MinValueTrigger(key[, trigger])</code> | Trigger invoked when specific value becomes minimum. |
| <code>training.triggers. OnceTrigger([call_on_resume])</code> | Trigger based on the starting point of the iteration. |
| <code>training.triggers.TimeTrigger(period)</code> | Trigger based on a fixed time interval. |

2.2.5 Reporting

| | |
|---|---|
| <code>reporting.Reporter()</code> | Object to which observed values are reported. |
| <code>reporting.report(values[, observer])</code> | Reports observed values with the current reporter object. |
| <code>reporting.report_scope(observation)</code> | Returns a report scope with the current reporter. |

2.2.6 Logging

| | |
|---------------------------------------|--|
| <code>logging.get_logger(name)</code> | Returns a child logger to be used by applications. |
|---------------------------------------|--|

2.2.7 Profiler

| | |
|--|--|
| <code>profiler.TimeSummary.report(tag[, use_cuda])</code> | Context manager to automatically report execution times. |
| <code>profiler.clear_tracer()</code> | Resets the status of the global tracer. |
| <code>profiler.enable_global_trace(enable)</code> | Enable or disable tracing for all the threads. |
| <code>profiler.enable_thread_trace(enable)</code> | Enable or disable tracing for the current thread. |
| <code>profiler.get_tracer([tracer_cls])</code> | Gets the current global tracer. |
| <code>profiler.ChromeTracer([max_event_count, ...])</code> | Tracer object that outputs a timeline in Chrome format. |
| <code>profiler.TraceableDataset(dataset, tag[, tracer])</code> | Utility class to trace a Dataset inside the DataLoader worker threads. |

2.3 Distributed Training

| | |
|---|--|
| <code>nn.parallel.DistributedDataParallel(module)</code> | Module for distributed data parallelism |
| <code>distributed.initialize_omp_environment(*[, ...])</code> | Initialize <i>torch.distributed</i> environments with values taken from OpenMPI. |

2.4 Check Pointing

| | |
|-------------------------------|--|
| <code>utils.checkpoint</code> | |
|-------------------------------|--|

2.5 Lazy Modules

| | |
|--|---|
| <code>nn.Ensure(*[, shape, dtype, broadcastable, ...])</code> | Module to check the shape of a tensor. |
| <code>nn.ensure(tensor[, shape, dtype, ...])</code> | Checks the shape and type of a tensor. |
| <code>nn.LazyLinear(in_features, *args, **kwargs)</code> | Linear module with lazy weight initialization. |
| <code>nn.LazyConv1d(in_channels, *args, **kwargs)</code> | Conv1d module with lazy weight initialization. |
| <code>nn.LazyConv2d(in_channels, *args, **kwargs)</code> | Conv2d module with lazy weight initialization. |
| <code>nn.LazyConv3d(in_channels, *args, **kwargs)</code> | Conv3d module with lazy weight initialization. |
| <code>nn.LazyBatchNorm1d(num_features, *args, **kwargs)</code> | BatchNorm1d module with lazy weight initialization. |
| <code>nn.LazyBatchNorm2d(num_features, *args, **kwargs)</code> | BatchNorm2d module with lazy weight initialization. |
| <code>nn.LazyBatchNorm3d(num_features, *args, **kwargs)</code> | BatchNorm3d module with lazy weight initialization. |

2.6 ONNX

2.6.1 Export

| | |
|--|---|
| <code>onnx.export(model, args, f[, return_output, ...])</code> | Export model into ONNX Graph. |
| <code>onnx.export_testcase(model, args, out_dir, *)</code> | Export model and I/O tensors of the model in protobuf format. |

2.6.2 Annotation

| | |
|--|---|
| <code>onnx.annotate(**attrs)</code> | Annotation parameters to the target function. |
| <code>onnx.apply_annotation(fn, *args, **attrs)</code> | Annotation applier to the target function |
| <code>onnx.scoped_anchor(**attrs)</code> | Add anchor node to the scoped modules |
| <code>onnx.export(model, args, f[, return_output, ...])</code> | Export model into ONNX Graph. |
| <code>onnx.export_testcase(model, args, out_dir, *)</code> | Export model and I/O tensors of the model in protobuf format. |

2.7 Datasets

| | |
|---|---|
| <code>dataset.SharedDataset(sm_size[, cache_type])</code> | Dataset that caches the load samples in shared memory |
| <code>dataset.TabularDataset(*args, **kws)</code> | An abstract class that represents tabular dataset. |
| <code>dataset.ItemNotFoundException</code> | |

2.8 Config

config.Config(config[, types])

config_types.optuna_types(trial)

config_types.load_path_with_optuna_types(...)

2.9 NumPy/CuPy Compatibility

| | |
|--|---|
| <i>from_ndarray</i> (ndarray) | Creates a <i>torch.Tensor</i> from a <i>numpy.ndarray</i> or <i>cupy.ndarray</i> . |
| <i>as_ndarray</i> (tensor) | Creates a <i>numpy.ndarray</i> or <i>cupy.ndarray</i> from <i>torch.Tensor</i> . |
| <i>get_xp</i> (obj) | Returns a module of ndarray implementation (<i>numpy</i> or <i>cupy</i>) for the given <i>obj</i> . |
| <i>as_numpy_dtype</i> (torch_dtype) | Returns NumPy dtype for the given PyTorch dtype. |
| <i>from_numpy_dtype</i> (numpy_dtype) | Returns PyTorch dtype for the given NumPy dtype. |
| <hr/> | |
| <i>cuda.stream</i> (stream) | Context-manager that selects a given stream. |
| <i>cuda.use_torch_mempool_in_cupy</i> () | Use the PyTorch memory pool in CuPy. |
| <i>cuda.use_default_mempool_in_cupy</i> () | Use the default memory pool in CuPy. |

PYTHON MODULE INDEX

p

`pytorch_pfn_extras`, 37
`pytorch_pfn_extras.config`, 42
`pytorch_pfn_extras.config_types`, 43
`pytorch_pfn_extras.cuda`, 44
`pytorch_pfn_extras.dataloaders`, 45
`pytorch_pfn_extras.dataloaders.data_loader`, 48
`pytorch_pfn_extras.dataloaders.utils`, 55
`pytorch_pfn_extras.dataset`, 56
`pytorch_pfn_extras.dataset.shared_dataset`, 61
`pytorch_pfn_extras.dataset.tabular`, 63
`pytorch_pfn_extras.dataset.tabular.delegate_dataset`, 66
`pytorch_pfn_extras.dataset.tabular.tabular_dataset`, 68
`pytorch_pfn_extras.distributed`, 73
`pytorch_pfn_extras.engine`, 75
`pytorch_pfn_extras.handler`, 85
`pytorch_pfn_extras.logging`, 105
`pytorch_pfn_extras.nn`, 105
`pytorch_pfn_extras.nn.modules`, 129
`pytorch_pfn_extras.nn.modules.ensure_shape`, 129
`pytorch_pfn_extras.nn.modules.extended_sequential`, 133
`pytorch_pfn_extras.nn.modules.lazy`, 136
`pytorch_pfn_extras.nn.modules.lazy_batchnorm`, 156
`pytorch_pfn_extras.nn.modules.lazy_conv`, 183
`pytorch_pfn_extras.nn.modules.lazy_linear`, 210
`pytorch_pfn_extras.nn.parallel`, 232
`pytorch_pfn_extras.nn.parallel.distributed`, 238
`pytorch_pfn_extras.onnx`, 249
`pytorch_pfn_extras.onnx.load`, 254
`pytorch_pfn_extras.onnx.pfto_exporter`, 263
`pytorch_pfn_extras.onnx.strip_large_tensor`, 263
`pytorch_pfn_extras.onnx.symbolic_registry`, 264
`pytorch_pfn_extras.onnx.unstrip_tensor`, 264
`pytorch_pfn_extras.ops`, 269
`pytorch_pfn_extras.profiler`, 271
`pytorch_pfn_extras.reporting`, 280
`pytorch_pfn_extras.runtime`, 287
`pytorch_pfn_extras.testing`, 297
`pytorch_pfn_extras.torchscript`, 297
`pytorch_pfn_extras.training`, 298
`pytorch_pfn_extras.training.extension`, 317
`pytorch_pfn_extras.training.extensions`, 323
`pytorch_pfn_extras.training.extensions.accumulate`, 378
`pytorch_pfn_extras.training.extensions.accumulate.average`, 383
`pytorch_pfn_extras.training.extensions.accumulate.max_accuracy`, 388
`pytorch_pfn_extras.training.extensions.accumulate.min_accuracy`, 394
`pytorch_pfn_extras.training.extensions.accumulate.standard_deviation`, 399
`pytorch_pfn_extras.training.extensions.accumulate.unbiased_average`, 405
`pytorch_pfn_extras.training.extensions.best_value`, 410
`pytorch_pfn_extras.training.extensions.evaluator`, 417
`pytorch_pfn_extras.training.extensions.fail_on_non_number`, 428
`pytorch_pfn_extras.training.extensions.log_report`, 431
`pytorch_pfn_extras.training.extensions.lr_scheduler`, 437
`pytorch_pfn_extras.training.extensions.micro_average`, 443
`pytorch_pfn_extras.training.extensions.parameter_statistics`, 447
`pytorch_pfn_extras.training.extensions.plot_report`, 452
`pytorch_pfn_extras.training.extensions.print_report`, 457
`pytorch_pfn_extras.training.extensions.profile_report`, 465
`pytorch_pfn_extras.training.extensions.progress_bar`,

472
pytorch_pfn_extras.training.extensions.slack,
476
pytorch_pfn_extras.training.extensions.snapshot_writers,
483
pytorch_pfn_extras.training.extensions.timeline_trace,
497
pytorch_pfn_extras.training.extensions.util,
503
pytorch_pfn_extras.training.extensions.value_observation,
508
pytorch_pfn_extras.training.extensions.variable_statistics_plot,
511
pytorch_pfn_extras.training.manager, 519
pytorch_pfn_extras.training.metrics, 520
pytorch_pfn_extras.training.trigger, 521
pytorch_pfn_extras.training.triggers, 524
pytorch_pfn_extras.training.triggers.early_stopping_trigger,
535
pytorch_pfn_extras.training.triggers.function_trigger,
539
pytorch_pfn_extras.training.triggers.interval_trigger,
542
pytorch_pfn_extras.training.triggers.manual_schedule_trigger,
545
pytorch_pfn_extras.training.triggers.minmax_value_trigger,
548
pytorch_pfn_extras.training.triggers.once_trigger,
553
pytorch_pfn_extras.training.triggers.time_trigger,
557
pytorch_pfn_extras.utils, 560
pytorch_pfn_extras.utils.checkpoint, 560
pytorch_pfn_extras.utils.comparer, 560
pytorch_pfn_extras.writing, 566

Symbols

`__call__()` (`pytorch_pfn_extras.dataloaders.utils.CollateAsDict` method), 55
`__call__()` (`pytorch_pfn_extras.handler.CodeBlock` method), 95
`__call__()` (`pytorch_pfn_extras.training.AccuracyMetric` method), 299
`__call__()` (`pytorch_pfn_extras.training.Extension` method), 303
`__call__()` (`pytorch_pfn_extras.training.extension.Extension` method), 319
`__call__()` (`pytorch_pfn_extras.training.extensions.BestValue` method), 330
`__call__()` (`pytorch_pfn_extras.training.extensions.Evaluator` method), 335
`__call__()` (`pytorch_pfn_extras.training.extensions.FailOnNonNumber` method), 338
`__call__()` (`pytorch_pfn_extras.training.extensions.LRScheduler` method), 341
`__call__()` (`pytorch_pfn_extras.training.extensions.LogReport` method), 344
`__call__()` (`pytorch_pfn_extras.training.extensions.MicroAverage` method), 348
`__call__()` (`pytorch_pfn_extras.training.extensions.ParameterStatistics` method), 353
`__call__()` (`pytorch_pfn_extras.training.extensions.PlotReport` method), 356
`__call__()` (`pytorch_pfn_extras.training.extensions.PrintReport` method), 358
`__call__()` (`pytorch_pfn_extras.training.extensions.ProfileReport` method), 361
`__call__()` (`pytorch_pfn_extras.training.extensions.ProgressBar` method), 363
`__call__()` (`pytorch_pfn_extras.training.extensions.TimelineTrace` method), 371
`__call__()` (`pytorch_pfn_extras.training.extensions.VariableStatisticsPlot` method), 375
`__call__()` (`pytorch_pfn_extras.training.extensions.accumulate.average_accumulate.AccumulateBase` method), 384
`__call__()` (`pytorch_pfn_extras.training.extensions.accumulate.max_accumulate.AccumulateBase` method), 390
`__call__()` (`pytorch_pfn_extras.training.extensions.accumulate.min_accumulate.AccumulateBase` method), 395
`__call__()` (`pytorch_pfn_extras.training.extensions.accumulate.standard_accumulate.AccumulateBase` method), 401
`__call__()` (`pytorch_pfn_extras.training.extensions.accumulate.unbiased_accumulate.AccumulateBase` method), 406
`__call__()` (`pytorch_pfn_extras.training.extensions.best_value.BestValue` method), 411
`__call__()` (`pytorch_pfn_extras.training.extensions.evaluator.Evaluator` method), 421
`__call__()` (`pytorch_pfn_extras.training.extensions.fail_on_non_number.FailOnNonNumber` method), 431
`__call__()` (`pytorch_pfn_extras.training.extensions.log_report.LogReport` method), 435
`__call__()` (`pytorch_pfn_extras.training.extensions.log_report.LogWriter` method), 436
`__call__()` (`pytorch_pfn_extras.training.extensions.lr_scheduler.LRScheduler` method), 440
`__call__()` (`pytorch_pfn_extras.training.extensions.micro_average.MicroAverage` method), 446
`__call__()` (`pytorch_pfn_extras.training.extensions.parameter_statistics.ParameterStatistics` method), 451
`__call__()` (`pytorch_pfn_extras.training.extensions.plot_report.PlotReport` method), 455
`__call__()` (`pytorch_pfn_extras.training.extensions.print_report.PrintReport` method), 464
`__call__()` (`pytorch_pfn_extras.training.extensions.profile_report.ProfileReport` method), 470
`__call__()` (`pytorch_pfn_extras.training.extensions.progress_bar.ProgressBar` method), 475
`__call__()` (`pytorch_pfn_extras.training.extensions.snapshot_writers.QueueSnapshotWriter` method), 487
`__call__()` (`pytorch_pfn_extras.training.extensions.snapshot_writers.SimpleSnapshotWriter` method), 489
`__call__()` (`pytorch_pfn_extras.training.extensions.snapshot_writers.StandardSnapshotWriter` method), 490
`__call__()` (`pytorch_pfn_extras.training.extensions.snapshot_writers.TensorSnapshotWriter` method), 492
`__call__()` (`pytorch_pfn_extras.training.extensions.snapshot_writers.WriterSnapshotWriter` method), 495
`__call__()` (`pytorch_pfn_extras.training.extensions.timeline_trace.TimelineTrace` method), 500
`__call__()` (`pytorch_pfn_extras.training.extensions.variable_statistics_plot.VariableStatisticsPlot` method), 505

`method`), 515
`__call__` () (pytorch_pfn_extras.training.extensions.variable_invariant_statistic.Plots.SharedDataset
`method`), 518
`__call__` () (pytorch_pfn_extras.training.metrics.AccuracyMonitor.__init__ (pytorch_pfn_extras.dataset.shared_dataset.InfiniteCache
`method`), 520
`__call__` () (pytorch_pfn_extras.training.trigger.IntervalTrigger.__init__ (pytorch_pfn_extras.dataset.shared_dataset.SharedDataset
`method`), 522
`__call__` () (pytorch_pfn_extras.training.trigger.Trigger.__init__ (pytorch_pfn_extras.dataset.tabular.DelegateDataset
`method`), 523
`__call__` () (pytorch_pfn_extras.training.triggers.BestValueTrigger.__init__ (pytorch_pfn_extras.dataset.tabular.delegate_dataset.Delegate
`method`), 525
`__call__` () (pytorch_pfn_extras.training.triggers.EarlyStoppingTrigger.__init__ (pytorch_pfn_extras.distributed.DistributedValidationSampler
`method`), 526
`__call__` () (pytorch_pfn_extras.training.triggers.FunctionTrigger.__init__ (pytorch_pfn_extras.engine.DistributedEvaluator
`method`), 528
`__call__` () (pytorch_pfn_extras.training.triggers.IntervalTrigger.__init__ (pytorch_pfn_extras.engine.Evaluator
`method`), 529
`__call__` () (pytorch_pfn_extras.training.triggers.ManualScheduleTrigger.__init__ (pytorch_pfn_extras.engine.StateObjectProtocol
`method`), 530
`__call__` () (pytorch_pfn_extras.training.triggers.OnceTrigger.__init__ (pytorch_pfn_extras.engine.Trainer
`method`), 533
`__call__` () (pytorch_pfn_extras.training.triggers.TimeTrigger.__init__ (pytorch_pfn_extras.handler.BaseHandler
`method`), 534
`__call__` () (pytorch_pfn_extras.training.triggers.early_stopping_trigger.EarlyStoppingTrigger.__init__ (pytorch_pfn_extras.handler.BaseLogic
`method`), 536
`__call__` () (pytorch_pfn_extras.training.triggers.function_trigger.FunctionTrigger.__init__ (pytorch_pfn_extras.handler.CodeBlock
`method`), 541
`__call__` () (pytorch_pfn_extras.training.triggers.interval_trigger.IntervalTrigger.__init__ (pytorch_pfn_extras.handler.CodeBlockLogic
`method`), 544
`__call__` () (pytorch_pfn_extras.training.triggers.manual_schedule_trigger.ManualScheduleTrigger.__init__ (pytorch_pfn_extras.handler.Handler
`method`), 548
`__call__` () (pytorch_pfn_extras.training.triggers.minmax_value_trigger.MinMaxValueTrigger.__init__ (pytorch_pfn_extras.handler.Logic
`method`), 549
`__call__` () (pytorch_pfn_extras.training.triggers.once_trigger.OnceTrigger.__init__ (pytorch_pfn_extras.nn.Ensure
`method`), 556
`__call__` () (pytorch_pfn_extras.training.triggers.time_trigger.TimeTrigger.__init__ (pytorch_pfn_extras.nn.LazyLinear
`method`), 559
`__call__` () (pytorch_pfn_extras.writing.QueueWriter.__init__ (pytorch_pfn_extras.nn.modules.ensure_shape.Ensure
`method`), 570
`__call__` () (pytorch_pfn_extras.writing.SimpleWriter.__init__ (pytorch_pfn_extras.nn.modules.extended_sequential.TypeVar
`method`), 572
`__call__` () (pytorch_pfn_extras.writing.StandardWriter.__init__ (pytorch_pfn_extras.nn.modules.lazy.LazyInitializationMixin
`method`), 573
`__call__` () (pytorch_pfn_extras.writing.TensorBoardWriter.__init__ (pytorch_pfn_extras.nn.modules.lazy_batchnorm.LazyInitializ
`method`), 575
`__call__` () (pytorch_pfn_extras.writing.Writer.__init__ (pytorch_pfn_extras.nn.modules.lazy_conv.LazyInitializationM
`method`), 578
`__init__` () (pytorch_pfn_extras.config.Config `method`), `__init__` () (pytorch_pfn_extras.nn.modules.lazy_linear.LazyInitialization
42
`__init__` () (pytorch_pfn_extras.dataloaders.DataLoader.__init__ (pytorch_pfn_extras.nn.modules.lazy_linear.LazyLinear
`method`), 47
`__init__` () (pytorch_pfn_extras.dataloaders.dataloader.DataLoader.__init__ (pytorch_pfn_extras.nn.parallel.DistributedDataParallel
`method`), 53
`__init__` () (pytorch_pfn_extras.dataloaders.utils.CollateAsDict.__init__ (pytorch_pfn_extras.nn.parallel.distributed.DistributedDataPa

method), 242

`__init__` () (pytorch_pfn_extras.nn.parallel.distributed.OrderedDict method), 245

`__init__` () (pytorch_pfn_extras.nn.parallel.distributed.TypeVar method), 247

`__init__` () (pytorch_pfn_extras.nn.parallel.distributed.record_function method), 248

`__init__` () (pytorch_pfn_extras.ops.OpDesc method), 270

`__init__` () (pytorch_pfn_extras.profiler.ChromeTracer method), 274

`__init__` () (pytorch_pfn_extras.profiler.TimeSummary method), 276

`__init__` () (pytorch_pfn_extras.profiler.TraceableDataset method), 278

`__init__` () (pytorch_pfn_extras.reporting.DictSummary method), 283

`__init__` () (pytorch_pfn_extras.reporting.Reporter method), 285

`__init__` () (pytorch_pfn_extras.reporting.Summary method), 286

`__init__` () (pytorch_pfn_extras.runtime.BaseRuntime method), 288

`__init__` () (pytorch_pfn_extras.runtime.PyTorchRuntime method), 293

`__init__` () (pytorch_pfn_extras.training.AccuracyMetric method), 299

`__init__` () (pytorch_pfn_extras.training.DistributedEvaluator method), 300

`__init__` () (pytorch_pfn_extras.training.Evaluator method), 301

`__init__` () (pytorch_pfn_extras.training.ExtensionEntry method), 305

`__init__` () (pytorch_pfn_extras.training.ExtensionsManager method), 307

`__init__` () (pytorch_pfn_extras.training.ExtensionsManagerProtocol method), 309

`__init__` () (pytorch_pfn_extras.training.IgniteExtensionsManager method), 311

`__init__` () (pytorch_pfn_extras.training.StateObjectProtocol method), 312

`__init__` () (pytorch_pfn_extras.training.Trainer method), 314

`__init__` () (pytorch_pfn_extras.training.extension.ExtensionEntry method), 321

`__init__` () (pytorch_pfn_extras.training.extension.ExtensionManager method), 322

`__init__` () (pytorch_pfn_extras.training.extensions.BestValue method), 331

`__init__` () (pytorch_pfn_extras.training.extensions.DistributedEvaluator method), 333

`__init__` () (pytorch_pfn_extras.training.extensions.Evaluator method), 335

`__init__` () (pytorch_pfn_extras.training.extensions.FailOnNonNumber method), 338

`__init__` () (pytorch_pfn_extras.training.extensions.IgniteEvaluator method), 339

`__init__` () (pytorch_pfn_extras.training.extensions.LRScheduler method), 341

`__init__` () (pytorch_pfn_extras.training.extensions.LogReport method), 344

`__init__` () (pytorch_pfn_extras.training.extensions.MaxValue method), 347

`__init__` () (pytorch_pfn_extras.training.extensions.MicroAverage method), 349

`__init__` () (pytorch_pfn_extras.training.extensions.MinValue method), 351

`__init__` () (pytorch_pfn_extras.training.extensions.ParameterStatistics method), 353

`__init__` () (pytorch_pfn_extras.training.extensions.PlotReport method), 356

`__init__` () (pytorch_pfn_extras.training.extensions.PrintReport method), 358

`__init__` () (pytorch_pfn_extras.training.extensions.ProfileReport method), 361

`__init__` () (pytorch_pfn_extras.training.extensions.ProgressBar method), 363

`__init__` () (pytorch_pfn_extras.training.extensions.Slack method), 366

`__init__` () (pytorch_pfn_extras.training.extensions.SlackWebhook method), 368

`__init__` () (pytorch_pfn_extras.training.extensions.TimelineTrace method), 372

`__init__` () (pytorch_pfn_extras.training.extensions.VariableStatisticsPlot method), 376

`__init__` () (pytorch_pfn_extras.training.extensions.accumulate.average_method), 385

`__init__` () (pytorch_pfn_extras.training.extensions.accumulate.average_method), 386

`__init__` () (pytorch_pfn_extras.training.extensions.accumulate.average_method), 387

`__init__` () (pytorch_pfn_extras.training.extensions.accumulate.max_accu method), 390

`__init__` () (pytorch_pfn_extras.training.extensions.accumulate.max_accu method), 392

`__init__` () (pytorch_pfn_extras.training.extensions.accumulate.max_accu method), 393

`__init__` () (pytorch_pfn_extras.training.extensions.accumulate.min_accu method), 395

`__init__` () (pytorch_pfn_extras.training.extensions.accumulate.min_accu method), 397

`__init__` () (pytorch_pfn_extras.training.extensions.accumulate.min_accu method), 398

`__init__` () (pytorch_pfn_extras.training.extensions.accumulate.min_accu method), 401

`__init__` () (pytorch_pfn_extras.training.extensions.accumulate.standard method), 403

`__init__` () (pytorch_pfn_extras.training.extensions.accumulate.standard method), 403

`__init__` () (pytorch_pfn_extras.training.extensions.accumulate.standard method), 403

| | |
|---|----------------------------|
| <code>method</code>), 404 | <code>method</code>), 459 |
| <code>__init__</code> () (pytorch_pfn_extras.training.extensions.accumulation_indexing_extensions.AccumulateBase. <code>PrintReport</code> <code>method</code>), 407 | <code>method</code>), 464 |
| <code>__init__</code> () (pytorch_pfn_extras.training.extensions.accumulation_indexing_extensions.SnapshotBase. <code>report</code> . <code>ExtensionsManager</code> <code>method</code>), 407 | <code>method</code>), 467 |
| <code>__init__</code> () (pytorch_pfn_extras.training.extensions.accumulation_indexing_extensions.SnapshotBase. <code>report</code> . <code>ExtensionsManager</code> <code>method</code>), 409 | <code>method</code>), 468 |
| <code>__init__</code> () (pytorch_pfn_extras.training.extensions.best_value_finder. <code>BestValue</code> (pytorch_pfn_extras.training.extensions.profile_report. <code>ProfileReport</code> <code>method</code>), 412 | <code>method</code>), 471 |
| <code>__init__</code> () (pytorch_pfn_extras.training.extensions.best_value_finder. <code>BestValue</code> (pytorch_pfn_extras.training.extensions.progress_bar. <code>ExtensionsManager</code> <code>method</code>), 413 | <code>method</code>), 473 |
| <code>__init__</code> () (pytorch_pfn_extras.training.extensions.best_value_finder. <code>BestValue</code> (pytorch_pfn_extras.training.extensions.progress_bar. <code>ProgressReporter</code> <code>method</code>), 415 | <code>method</code>), 475 |
| <code>__init__</code> () (pytorch_pfn_extras.training.extensions.best_value_finder. <code>BestValue</code> (pytorch_pfn_extras.training.extensions.slack. <code>ExtensionsManager</code> <code>method</code>), 416 | <code>method</code>), 477 |
| <code>__init__</code> () (pytorch_pfn_extras.training.extensions.evaluator. <code>DistributedEvaluator</code> (pytorch_pfn_extras.training.extensions.slack. <code>Slack</code> <code>method</code>), 418 | <code>method</code>), 480 |
| <code>__init__</code> () (pytorch_pfn_extras.training.extensions.evaluator. <code>Evaluator</code> (pytorch_pfn_extras.training.extensions.slack. <code>SlackWebhook</code> <code>method</code>), 421 | <code>method</code>), 483 |
| <code>__init__</code> () (pytorch_pfn_extras.training.extensions.evaluator. <code>ExtensionsManager</code> (pytorch_pfn_extras.training.extensions.snapshot_writers. <code>ProgressReporter</code> <code>method</code>), 423 | <code>method</code>), 484 |
| <code>__init__</code> () (pytorch_pfn_extras.training.extensions.evaluator. <code>Evaluator</code> (pytorch_pfn_extras.training.extensions.snapshot_writers. <code>ProgressReporter</code> <code>method</code>), 425 | <code>method</code>), 485 |
| <code>__init__</code> () (pytorch_pfn_extras.training.extensions.evaluator. <code>Evaluator</code> (pytorch_pfn_extras.training.extensions.snapshot_writers. <code>Queue</code> <code>method</code>), 426 | <code>method</code>), 487 |
| <code>__init__</code> () (pytorch_pfn_extras.training.extensions.fail_on_nan. <code>FailOnNaN</code> (pytorch_pfn_extras.training.extensions.snapshot_writers. <code>SimpleWriter</code> <code>method</code>), 429 | <code>method</code>), 489 |
| <code>__init__</code> () (pytorch_pfn_extras.training.extensions.fail_on_nan. <code>FailOnNaN</code> (pytorch_pfn_extras.training.extensions.snapshot_writers. <code>StandardWriter</code> <code>method</code>), 431 | <code>method</code>), 491 |
| <code>__init__</code> () (pytorch_pfn_extras.training.extensions.log_reporter. <code>LogReporter</code> (pytorch_pfn_extras.training.extensions.snapshot_writers. <code>TensorBoardWriter</code> <code>method</code>), 432 | <code>method</code>), 492 |
| <code>__init__</code> () (pytorch_pfn_extras.training.extensions.log_reporter. <code>LogReporter</code> (pytorch_pfn_extras.training.extensions.snapshot_writers. <code>ThriftWriter</code> <code>method</code>), 435 | <code>method</code>), 493 |
| <code>__init__</code> () (pytorch_pfn_extras.training.extensions.log_reporter. <code>LogReporter</code> (pytorch_pfn_extras.training.extensions.snapshot_writers. <code>ThriftWriter</code> <code>method</code>), 436 | <code>method</code>), 494 |
| <code>__init__</code> () (pytorch_pfn_extras.training.extensions.lr_scheduler. <code>LRScheduler</code> (pytorch_pfn_extras.training.extensions.snapshot_writers. <code>Writer</code> <code>method</code>), 438 | <code>method</code>), 496 |
| <code>__init__</code> () (pytorch_pfn_extras.training.extensions.lr_scheduler. <code>LRScheduler</code> (pytorch_pfn_extras.training.extensions.timeline_trace. <code>ExtensionsManager</code> <code>method</code>), 440 | <code>method</code>), 498 |
| <code>__init__</code> () (pytorch_pfn_extras.training.extensions.lr_scheduler. <code>LRScheduler</code> (pytorch_pfn_extras.training.extensions.timeline_trace. <code>TimelineTrace</code> <code>method</code>), 442 | <code>method</code>), 500 |
| <code>__init__</code> () (pytorch_pfn_extras.training.extensions.micro_optimizer. <code>MicroOptimizer</code> (pytorch_pfn_extras.training.extensions.util. <code>ExtensionsManager</code> <code>method</code>), 444 | <code>method</code>), 504 |
| <code>__init__</code> () (pytorch_pfn_extras.training.extensions.micro_optimizer. <code>MicroOptimizer</code> (pytorch_pfn_extras.training.extensions.util. <code>ProgressBar</code> <code>method</code>), 446 | <code>method</code>), 505 |
| <code>__init__</code> () (pytorch_pfn_extras.training.extensions.parameter_observer. <code>ParameterObserver</code> (pytorch_pfn_extras.training.extensions.value_observation. <code>ExtensionsManager</code> <code>method</code>), 448 | <code>method</code>), 510 |
| <code>__init__</code> () (pytorch_pfn_extras.training.extensions.parameter_observer. <code>ParameterObserver</code> (pytorch_pfn_extras.training.extensions.variable_statistics. <code>PlotReporter</code> <code>method</code>), 451 | <code>method</code>), 513 |
| <code>__init__</code> () (pytorch_pfn_extras.training.extensions.plot_reporter. <code>PlotReporter</code> (pytorch_pfn_extras.training.extensions.variable_statistics. <code>PlotReporter</code> <code>method</code>), 453 | <code>method</code>), 514 |
| <code>__init__</code> () (pytorch_pfn_extras.training.extensions.plot_reporter. <code>PlotReporter</code> (pytorch_pfn_extras.training.extensions.variable_statistics. <code>PlotReporter</code> <code>method</code>), 456 | <code>method</code>), 515 |
| <code>__init__</code> () (pytorch_pfn_extras.training.extensions.print_reporter. <code>PrintReporter</code> (pytorch_pfn_extras.training.extensions.variable_statistics. <code>PlotReporter</code> <code>method</code>), 459 | <code>method</code>), 515 |

method), 518

__init__ () (pytorch_pfn_extras.training.metrics.AccuracyMetric method), 520

__init__ () (pytorch_pfn_extras.training.trigger.IntervalTrigger method), 522

__init__ () (pytorch_pfn_extras.training.triggers.BestValueTrigger method), 525

__init__ () (pytorch_pfn_extras.training.triggers.EarlyStoppingTrigger method), 527

__init__ () (pytorch_pfn_extras.training.triggers.FunctionTrigger method), 528

__init__ () (pytorch_pfn_extras.training.triggers.IntervalTrigger method), 529

__init__ () (pytorch_pfn_extras.training.triggers.ManualScheduleTrigger method), 531

__init__ () (pytorch_pfn_extras.training.triggers.MaxValueTrigger method), 531

__init__ () (pytorch_pfn_extras.training.triggers.MinValueTrigger method), 532

__init__ () (pytorch_pfn_extras.training.triggers.OnceTrigger method), 533

__init__ () (pytorch_pfn_extras.training.triggers.TimeTrigger method), 534

__init__ () (pytorch_pfn_extras.training.triggers.early_stopping_trigger.EarlyStoppingTrigger method), 537

__init__ () (pytorch_pfn_extras.training.triggers.early_stopping_trigger.ExtensionsManagerProtocol method), 538

__init__ () (pytorch_pfn_extras.training.triggers.function_extensions_manager_protocol.FunctionExtensionsManagerProtocol method), 540

__init__ () (pytorch_pfn_extras.training.triggers.function_extensions_manager_protocol.FunctionExtensionsManagerProtocol method), 541

__init__ () (pytorch_pfn_extras.training.triggers.interval_trigger.IntervalTrigger method), 543

__init__ () (pytorch_pfn_extras.training.triggers.interval_trigger.IntervalTrigger method), 545

__init__ () (pytorch_pfn_extras.training.triggers.manual_schedule_trigger.ExtensionsManagerProtocol method), 546

__init__ () (pytorch_pfn_extras.training.triggers.manual_schedule_trigger.ManualScheduleTrigger method), 548

__init__ () (pytorch_pfn_extras.training.triggers.minmax_value_trigger.BestValueTrigger method), 550

__init__ () (pytorch_pfn_extras.training.triggers.minmax_value_trigger.ExtensionsManagerProtocol method), 551

__init__ () (pytorch_pfn_extras.training.triggers.minmax_value_trigger.MaxValueTrigger method), 552

__init__ () (pytorch_pfn_extras.training.triggers.minmax_value_trigger.MinValueTrigger method), 553

__init__ () (pytorch_pfn_extras.training.triggers.once_trigger.ExtensionsManagerProtocol method), 554

__init__ () (pytorch_pfn_extras.training.triggers.once_trigger.OnceTrigger method), 556

__init__ () (pytorch_pfn_extras.training.triggers.time_trigger.ExtensionsManagerProtocol method), 558

__init__ () (pytorch_pfn_extras.training.triggers.time_trigger.TimeTrigger method), 559

__init__ () (pytorch_pfn_extras.utils.comparer.Comparer method), 562

__init__ () (pytorch_pfn_extras.utils.comparer.ModelComparer method), 564

__init__ () (pytorch_pfn_extras.utils.comparer.OutputsComparer method), 566

__init__ () (pytorch_pfn_extras.writing.ProcessQueueWriter method), 568

__init__ () (pytorch_pfn_extras.writing.ProcessWriter method), 569

__init__ () (pytorch_pfn_extras.writing.QueueWriter method), 571

__init__ () (pytorch_pfn_extras.writing.SimpleWriter method), 572

__init__ () (pytorch_pfn_extras.writing.StandardWriter method), 574

__init__ () (pytorch_pfn_extras.writing.TensorBoardWriter method), 575

__init__ () (pytorch_pfn_extras.writing.ThreadQueueWriter method), 576

__init__ () (pytorch_pfn_extras.writing.ThreadWriter method), 577

__init__ () (pytorch_pfn_extras.writing.Writer method), 579

A

absolute_path (pytorch_pfn_extras.onnx.load.Path method), 260

absolute_path (pytorch_pfn_extras.onnx.unstrip_tensor.Path method), 267

accumulate (pytorch_pfn_extras.training.extensions.accumulate.average_accumulate method), 389

accumulate (pytorch_pfn_extras.training.extensions.accumulate.max_accumulate method), 389

accumulate (pytorch_pfn_extras.training.extensions.accumulate.min_accumulate method), 389

accumulate (pytorch_pfn_extras.training.extensions.accumulate.standard_deviation method), 400

accumulate (pytorch_pfn_extras.training.extensions.accumulate.unbiased_standard_deviation method), 400

AccuracyMetric (class in pytorch_pfn_extras.training), 520

AccuracyMetric (class in pytorch_pfn_extras.training.metrics), 520

add () (pytorch_pfn_extras.profiler.TimeSummary method), 284

add () (pytorch_pfn_extras.reporting.DictSummary method), 283

add() (pytorch_pfn_extras.reporting.Summary method), 286
 add_to_cache() (pytorch_pfn_extras.dataset.shared_dataset.InfiniteCache method), 62
 add() (pytorch_pfn_extras.training.extensions.accumulate_average.SummaryBase method), 386
 add() (pytorch_pfn_extras.training.extensions.accumulate_average.SummaryBase method), 388
 add() (pytorch_pfn_extras.training.extensions.accumulate_average.SummaryBase method), 392
 add() (pytorch_pfn_extras.training.extensions.accumulate_average.SummaryBase method), 393
 add() (pytorch_pfn_extras.training.extensions.accumulate.min_accumulate.MinSummary method), 397
 add() (pytorch_pfn_extras.training.extensions.accumulate.min_accumulate.MinSummary method), 399
 add() (pytorch_pfn_extras.training.extensions.accumulate.min_accumulate.MinSummary method), 403
 add() (pytorch_pfn_extras.training.extensions.accumulate.standard_deviation_accumulate.SummaryBase method), 404
 add() (pytorch_pfn_extras.training.extensions.accumulate.unbiased_standard_deviation_accumulate.SummaryBase method), 408
 add() (pytorch_pfn_extras.training.extensions.accumulate.unbiased_standard_deviation_accumulate.SummaryBase method), 410
 add() (pytorch_pfn_extras.training.extensions.variable_statistics_plot.Reservoir method), 514
 add_dump() (pytorch_pfn_extras.utils.comparer.Comparer apply_annotation() method), 563
 add_engine() (pytorch_pfn_extras.utils.comparer.Comparer as_ndarray() method), 563
 add_event() (pytorch_pfn_extras.profiler.ChromeTracer as_numpy_dtype() method), 275
 add_event() (pytorch_pfn_extras.profiler.Tracer as_output() method), 279
 add_event() (pytorch_pfn_extras.training.extensions.timeline_trace.AsDict method), 502
 add_metric() (pytorch_pfn_extras.training.extensions.Evaluator asdict() method), 336
 add_metric() (pytorch_pfn_extras.training.extensions.evaluator.AsDict method), 421
 add_observer() (pytorch_pfn_extras.reporting.Reporter asdict() method), 285
 add_observers() (pytorch_pfn_extras.reporting.Reporter asdict() method), 285
 add_remote_event() (pytorch_pfn_extras.profiler.ChromeTracer asdict() method), 275
 add_remote_event() (pytorch_pfn_extras.profiler.Tracer asdict() method), 279
 add_remote_event() (pytorch_pfn_extras.training.extensions.timeline_trace.AsDict method), 502
 add_to_cache() (pytorch_pfn_extras.dataset.shared_dataset.Cache method), 61
 add_to_cache() (pytorch_pfn_extras.dataset.shared_dataset.InfiniteCache method), 62
 affine() (pytorch_pfn_extras.modules.lazy_batchnorm.LazyBatchNorm1d attribute), 158
 affine() (pytorch_pfn_extras.modules.lazy_batchnorm.LazyBatchNorm2d attribute), 161
 affine() (pytorch_pfn_extras.modules.lazy_batchnorm.LazyBatchNorm3d attribute), 164
 all_gather_object() (in module pytorch_pfn_extras.training.extensions.accumulate.average_accumulate), 387
 all_gather_object() (in module pytorch_pfn_extras.training.extensions.accumulate.min_accumulate), 388
 all_gather_object() (in module pytorch_pfn_extras.training.extensions.accumulate.min_accumulate), 399
 all_gather_object() (in module pytorch_pfn_extras.training.extensions.accumulate.min_accumulate), 403
 all_gather_object() (in module pytorch_pfn_extras.training.extensions.accumulate.min_accumulate), 404
 all_gather_object() (in module pytorch_pfn_extras.training.extensions.accumulate.min_accumulate), 408
 all_gather_object() (in module pytorch_pfn_extras.training.extensions.accumulate.min_accumulate), 410
 annotate() (in module pytorch_pfn_extras.onnx), 249
 apply_annotation() (in module pytorch_pfn_extras.onnx), 250
 as_ndarray() (in module pytorch_pfn_extras), 38
 as_numpy_dtype() (in module pytorch_pfn_extras), 38
 as_output() (in module pytorch_pfn_extras.onnx), 250
 asdict() (pytorch_pfn_extras.dataset.tabular.tabular_dataset.TabularDataset method), 70
 asdict() (pytorch_pfn_extras.dataset.TabularDataset method), 58
 astuple() (pytorch_pfn_extras.dataset.tabular.tabular_dataset.TabularDataset method), 70
 astuple() (pytorch_pfn_extras.dataset.TabularDataset method), 58
 available() (pytorch_pfn_extras.training.extensions.plot_report.PlotReport static method), 456
 available() (pytorch_pfn_extras.training.extensions.PlotReport static method), 357
 available() (pytorch_pfn_extras.training.extensions.variable_statistics_plot.VariableStatisticsPlot static method), 519
 available() (pytorch_pfn_extras.training.extensions.VariableStatisticsPlot static method), 376
 AverageAccumulate (class in pytorch_pfn_extras.training.extensions), 329
 AverageAccumulate (class in pytorch_pfn_extras.training.extensions.accumulate), 378
 AverageAccumulate (class in pytorch_pfn_extras.training.extensions.accumulate.average_accumulate), 387
 AverageSummary (class in pytorch_pfn_extras.training.extensions.accumulate.min_accumulate), 397

- `torch_pfn_extras.training.extensions.accumulate_grad_into_cache_item()` (pytorch_pfn_extras.dataset.shared_dataset.SharedDataset method), 62
- `cache_item()` (pytorch_pfn_extras.dataset.SharedDataset method), 56
- `cast()` (in module `pytorch_pfn_extras.engine`), 75
- `check_worker_number_rationality()` (pytorch_pfn_extras.dataloaders.DataLoader method), 48
- `check_worker_number_rationality()` (pytorch_pfn_extras.dataloaders.dataloader.DataLoader method), 54
- `checkpoint()` (in module `pytorch_pfn_extras.utils.checkpoint`), 560
- `chmod()` (pytorch_pfn_extras.onnx.load.Path method), 260
- `chmod()` (pytorch_pfn_extras.onnx.unstrip_tensor.Path method), 267
- `ChromeTracer` (class in `pytorch_pfn_extras.profiler`), 274
- `clear()` (pytorch_pfn_extras.nn.parallel.distributed.OrderedDict method), 245
- `clear()` (pytorch_pfn_extras.profiler.ChromeTracer method), 275
- `clear()` (pytorch_pfn_extras.profiler.Tracer method), 279
- `clear_report()` (pytorch_pfn_extras.training.extensions.profile_report.OrderedDict method), 468
- `clear()` (pytorch_pfn_extras.training.extensions.timeline_trace.Tracer method), 502
- `clear_tracer()` (in module `pytorch_pfn_extras.profiler`), 271
- `close()` (pytorch_pfn_extras.onnx.load.IO method), 256
- `close()` (pytorch_pfn_extras.training.extensions.print_report.IO method), 461
- `close()` (pytorch_pfn_extras.training.extensions.util.ProgressBar method), 506
- `closed` (pytorch_pfn_extras.onnx.load.IO property), 256
- `closed` (pytorch_pfn_extras.training.extensions.print_report.IO property), 461
- `ClosureLogic` (class in `pytorch_pfn_extras.handler`), 92
- `CodeBlock` (class in `pytorch_pfn_extras.handler`), 94
- `CodeBlockLogic` (class in `pytorch_pfn_extras.handler`), 96
- `CollateAsDict` (class in `pytorch_pfn_extras.dataloaders.utils`), 55
- `collect()` (pytorch_pfn_extras.training.extensions.variable_statistics_plot method), 516
- `compare()` (pytorch_pfn_extras.utils.comparer.Comparer method), 563
- `Comparer` (class in `pytorch_pfn_extras.utils.comparer`), 561
- `compile()` (in module `pytorch_pfn_extras`), 38
- `complete_report()` (pytorch_pfn_extras.training.extensions.accumulate_grad_into_cache_item() method), 62
- B**
- `backprop` (pytorch_pfn_extras.handler.CodeBlock attribute), 95
- `backprop_from` (pytorch_pfn_extras.handler.CodeBlock attribute), 95
- `backprop_to` (pytorch_pfn_extras.handler.CodeBlock attribute), 95
- `BaseHandler` (class in `pytorch_pfn_extras.handler`), 86
- `BaseLogic` (class in `pytorch_pfn_extras.handler`), 90
- `BaseRuntime` (class in `pytorch_pfn_extras.runtime`), 288
- `batch_size` (pytorch_pfn_extras.dataloaders.DataLoader attribute), 48
- `batch_size` (pytorch_pfn_extras.dataloaders.dataloader.DataLoader attribute), 54
- `best_epoch` (pytorch_pfn_extras.training.extensions.best_value property), 412
- `best_epoch` (pytorch_pfn_extras.training.extensions.BestValue property), 331
- `best_iteration` (pytorch_pfn_extras.training.extensions.best_value property), 412
- `best_iteration` (pytorch_pfn_extras.training.extensions.BestValue property), 331
- `best_value` (pytorch_pfn_extras.training.extensions.best_value property), 412
- `best_value` (pytorch_pfn_extras.training.extensions.BestValue property), 331
- `BestValue` (class in `pytorch_pfn_extras.training.extensions`), 330
- `BestValue` (class in `pytorch_pfn_extras.training.extensions.best_value`), 411
- `BestValueTrigger` (class in `pytorch_pfn_extras.training.triggers`), 524
- `BestValueTrigger` (class in `pytorch_pfn_extras.training.triggers.minmax_value_trigger`), 549
- `bias` (pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv1d attribute), 185
- `bias` (pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv2d attribute), 188
- `bias` (pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv3d attribute), 191
- `buffer` (pytorch_pfn_extras.training.extensions.evaluator.TextIO property), 428
- `buffer` (pytorch_pfn_extras.training.extensions.util.TextIO property), 508
- C**
- `Cache` (class in `pytorch_pfn_extras.dataset.shared_dataset`), 61

`torch_pfn_extras.profiler.TimeSummary` (method), 70
`method()`, 277
`compute_accumulate()` (py- `method()`, 58
`torch_pfn_extras.training.extensions.accumulate_and_config (class) pytorch-pfn-extras.config`, 42
`method()`, 387
`compute_accumulate()` (py- `method()`, 487
`torch_pfn_extras.training.extensions.accumulate_and_consume()` (pytorch_pfn_extras.training.extensions.snapshot_writers.QueueWriter), 487
`method()`, 388
`compute_accumulate()` (py- `consume_options()` (py- `method()`, 571
`torch_pfn_extras.training.extensions.accumulate_and_consume_options()` (py- `method()`, 571
`method()`, 392
`compute_accumulate()` (py- `consume_options()` (py- `method()`, 87
`torch_pfn_extras.training.extensions.accumulate_and_consume_options()` (py- `method()`, 87
`method()`, 393
`compute_accumulate()` (py- `consume_options()` (py- `method()`, 90
`torch_pfn_extras.training.extensions.accumulate_and_consume_options()` (py- `method()`, 90
`method()`, 398
`compute_accumulate()` (py- `consume_options()` (py- `method()`, 93
`torch_pfn_extras.training.extensions.accumulate_and_consume_options()` (py- `method()`, 93
`method()`, 399
`compute_accumulate()` (py- `consume_options()` (py- `method()`, 96
`torch_pfn_extras.training.extensions.accumulate_and_consume_options()` (py- `method()`, 96
`method()`, 403
`compute_accumulate()` (py- `consume_options()` (pytorch_pfn_extras.handler.BaseHandler), 99
`torch_pfn_extras.training.extensions.accumulate_and_consume_options()` (pytorch_pfn_extras.handler.BaseHandler), 99
`method()`, 404
`compute_accumulate()` (py- `consume_options()` (pytorch_pfn_extras.handler.BaseHandler), 99
`torch_pfn_extras.training.extensions.accumulate_and_consume_options()` (pytorch_pfn_extras.handler.BaseHandler), 99
`method()`, 408
`compute_accumulate()` (py- `consume_options()` (pytorch_pfn_extras.handler.BaseHandler), 99
`torch_pfn_extras.training.extensions.accumulate_and_consume_options()` (pytorch_pfn_extras.handler.BaseHandler), 99
`method()`, 410
`compute_average()` (py- `convert_batch()` (py- `method()`, 288
`torch_pfn_extras.training.extensions.accumulate_and_convert_batch()` (py- `method()`, 288
`method()`, 392
`compute_max()` (pytorch_pfn_extras.training.extensions.accumulate_and_convert_batch() (py- `method()`, 288
`method()`, 392
`compute_mean()` (pytorch_pfn_extras.reporting.DictSummary), 283
`method()`, 283
`compute_mean()` (pytorch_pfn_extras.reporting.Summary), 287
`method()`, 287
`compute_mean()` (pytorch_pfn_extras.training.extensions.accumulate_and_convert_batch() (py- `method()`, 288
`method()`, 403
`compute_mean()` (pytorch_pfn_extras.training.extensions.accumulate_and_convert_batch() (py- `method()`, 288
`method()`, 410
`compute_min()` (pytorch_pfn_extras.training.extensions.accumulate_and_convert_batch() (py- `method()`, 288
`method()`, 398
`compute_standard_deviation()` (py- `method()`, 493
`torch_pfn_extras.training.extensions.accumulate_and_convert_batch()` (py- `method()`, 288
`method()`, 403
`compute_unbiased_standard_deviation()` (py- `method()`, 568
`torch_pfn_extras.training.extensions.accumulate_and_convert_batch()` (py- `method()`, 288
`method()`, 410
`concat()` (pytorch_pfn_extras.dataset.tabular.tabular_dataset.TabularDataset), 571

| | |
|--|--|
| <code>create_consumer()</code> (pytorch_pfn_extras.writing.ThreadQueueWriter method), 576 | D |
| <code>create_distributed_subset_indices()</code> (in module pytorch_pfn_extras.distributed), 73 | <code>DataLoader</code> (class in pytorch_pfn_extras.dataloaders), 45 |
| <code>create_evaluator()</code> (in module pytorch_pfn_extras.engine), 76 | <code>DataLoader</code> (class in pytorch_pfn_extras.dataloaders.dataloader), 51 |
| <code>create_header_and_templates()</code> (in module pytorch_pfn_extras.training.extensions.print_report), 457 | <code>Dataset</code> (class in pytorch_pfn_extras.dataset.tabular.tabular_dataset), 68 |
| <code>create_queue()</code> (pytorch_pfn_extras.training.extensions.snapshot_writers.ProcessQueueWriter method), 485 | <code>dataset</code> (pytorch_pfn_extras.dataloaders.DataLoader attribute), 48 |
| <code>create_queue()</code> (pytorch_pfn_extras.training.extensions.snapshot_writers.QueueWriter method), 488 | <code>dataset</code> (pytorch_pfn_extras.dataloaders.dataloader.DataLoader attribute), 54 |
| <code>create_queue()</code> (pytorch_pfn_extras.training.extensions.snapshot_writers.ThreadQueueWriter method), 494 | <code>deepcopy()</code> (in module pytorch_pfn_extras.training.extensions.print_report), 457 |
| <code>create_queue()</code> (pytorch_pfn_extras.writing.ProcessQueueWriter method), 568 | <code>DEFAULT</code> (pytorch_pfn_extras.training.extensions.SnapshotMode attribute), 369 |
| <code>create_queue()</code> (pytorch_pfn_extras.writing.QueueWriter method), 571 | <code>default_collate()</code> (in module pytorch_pfn_extras.dataloaders.dataloader), 49 |
| <code>create_queue()</code> (pytorch_pfn_extras.writing.ThreadQueueWriter method), 577 | <code>default_convert()</code> (in module pytorch_pfn_extras.dataloaders.dataloader), 50 |
| <code>create_task()</code> (pytorch_pfn_extras.training.extensions.snapshot_writers.QueueWriter method), 488 | <code>default_name</code> (pytorch_pfn_extras.training.Extension property), 303 |
| <code>create_task()</code> (pytorch_pfn_extras.writing.QueueWriter method), 571 | <code>default_name</code> (pytorch_pfn_extras.training.extension.Extension property), 319 |
| <code>create_trainer()</code> (in module pytorch_pfn_extras.engine), 76 | <code>default_name</code> (pytorch_pfn_extras.training.extensions.best_value.BestValue attribute), 412 |
| <code>create_worker()</code> (pytorch_pfn_extras.training.extensions.snapshot_writers.ProcessWriter method), 486 | <code>default_name</code> (pytorch_pfn_extras.training.extensions.best_value.MaxValue attribute), 415 |
| <code>create_worker()</code> (pytorch_pfn_extras.training.extensions.snapshot_writers.StandardWriter method), 491 | <code>default_name</code> (pytorch_pfn_extras.training.extensions.best_value.MinValue attribute), 416 |
| <code>create_worker()</code> (pytorch_pfn_extras.training.extensions.snapshot_writers.ThreadWriter method), 494 | <code>default_name</code> (pytorch_pfn_extras.training.extensions.BestValue attribute), 331 |
| <code>create_worker()</code> (pytorch_pfn_extras.writing.ProcessWriter method), 569 | <code>default_name</code> (pytorch_pfn_extras.training.extensions.Evaluator attribute), 336 |
| <code>create_worker()</code> (pytorch_pfn_extras.writing.StandardWriter method), 574 | <code>default_name</code> (pytorch_pfn_extras.training.extensions.evaluator.Evaluator attribute), 422 |
| <code>create_worker()</code> (pytorch_pfn_extras.writing.ThreadWriter method), 577 | <code>default_name</code> (pytorch_pfn_extras.training.extensions.MaxValue attribute), 347 |
| <code>customize_type()</code> (in module pytorch_pfn_extras.config), 42 | <code>default_name</code> (pytorch_pfn_extras.training.extensions.MinValue attribute), 351 |
| <code>cwd()</code> (pytorch_pfn_extras.onnx.load.Path class method), 260 | <code>default_name</code> (pytorch_pfn_extras.training.extensions.parameter_statistics attribute), 451 |
| <code>cwd()</code> (pytorch_pfn_extras.onnx.unstrip_tensor.Path class method), 267 | <code>default_name</code> (pytorch_pfn_extras.training.extensions.ParameterStatistics attribute), 354 |
| | <code>default_statistics</code> (pytorch_pfn_extras.training.extensions.parameter_statistics.ParameterStatistics attribute), 451 |
| | <code>default_statistics</code> (pytorch_pfn_extras.training.extensions.ParameterStatistics attribute), 354 |

`default_transform_model()` (in module `pytorch_pfn_extras.engine`), 77
`DelegatedDataset` (class in `pytorch_pfn_extras.dataset.tabular`), 64
`DelegatedDataset` (class in `pytorch_pfn_extras.dataset.tabular.delegate_dataset`), 66
`DictSummary` (class in `pytorch_pfn_extras.reporting`), 282
`dilation` (`pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv1d` attribute), 185
`dilation` (`pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv2d` attribute), 188
`dilation` (`pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv3d` attribute), 191
`DISTRIBUTED` (`pytorch_pfn_extras.training.extensions.SnapshotMode` attribute), 369
`DistributedDataParallel` (class in `pytorch_pfn_extras.nn.parallel`), 233
`DistributedDataParallel` (class in `pytorch_pfn_extras.nn.parallel.distributed`), 239
`DistributedEvaluator` (class in `pytorch_pfn_extras.engine`), 78
`DistributedEvaluator` (class in `pytorch_pfn_extras.training`), 300
`DistributedEvaluator` (class in `pytorch_pfn_extras.training.extensions`), 332
`DistributedEvaluator` (class in `pytorch_pfn_extras.training.extensions.evaluator`), 417
`DistributedValidationSampler` (class in `pytorch_pfn_extras.distributed`), 74
`drop_last` (`pytorch_pfn_extras.dataloaders.DataLoader` attribute), 48
`drop_last` (`pytorch_pfn_extras.dataloaders.dataloader.DataLoader` attribute), 54
`dump()` (`pytorch_pfn_extras.utils.comparer.Comparer` method), 563

E

`EarlyStoppingTrigger` (class in `pytorch_pfn_extras.training.triggers`), 526
`EarlyStoppingTrigger` (class in `pytorch_pfn_extras.training.triggers.early_stopping_trigger`), 535
`elapsed_time` (`pytorch_pfn_extras.training.extension.ExtensionsManagerProtocol` property), 322
`elapsed_time` (`pytorch_pfn_extras.training.extensions.best_value_extensions_manager_protocol` property), 413
`elapsed_time` (`pytorch_pfn_extras.training.extensions.evaluator_extensions_manager_protocol` property), 423
`elapsed_time` (`pytorch_pfn_extras.training.extensions.fail_on_non_number_extensions_manager_protocol` property), 429
`elapsed_time` (`pytorch_pfn_extras.training.extensions.log_report.ExtensionsManagerProtocol` property), 432
`elapsed_time` (`pytorch_pfn_extras.training.extensions.lr_scheduler.ExtensionsManagerProtocol` property), 438
`elapsed_time` (`pytorch_pfn_extras.training.extensions.micro_average.ExtensionsManagerProtocol` property), 444
`elapsed_time` (`pytorch_pfn_extras.training.extensions.parameter_statistics_extensions_manager_protocol` property), 448
`elapsed_time` (`pytorch_pfn_extras.training.extensions.plot_report.ExtensionsManagerProtocol` property), 453
`elapsed_time` (`pytorch_pfn_extras.training.extensions.print_report.ExtensionsManagerProtocol` property), 459
`elapsed_time` (`pytorch_pfn_extras.training.extensions.profile_report.ExtensionsManagerProtocol` property), 467
`elapsed_time` (`pytorch_pfn_extras.training.extensions.progress_bar.ExtensionsManagerProtocol` property), 473
`elapsed_time` (`pytorch_pfn_extras.training.extensions.slack.ExtensionsManagerProtocol` property), 477
`elapsed_time` (`pytorch_pfn_extras.training.extensions.timeline_trace.ExtensionsManagerProtocol` property), 498
`elapsed_time` (`pytorch_pfn_extras.training.extensions.util.ExtensionsManagerProtocol` property), 504
`elapsed_time` (`pytorch_pfn_extras.training.extensions.value_observation.ExtensionsManagerProtocol` property), 510
`elapsed_time` (`pytorch_pfn_extras.training.extensions.variable_statistics_extensions_manager_protocol` property), 513
`elapsed_time` (`pytorch_pfn_extras.training.ExtensionsManagerProtocol` property), 309
`elapsed_time` (`pytorch_pfn_extras.training.triggers.early_stopping_trigger_extensions_manager_protocol` property), 538
`elapsed_time` (`pytorch_pfn_extras.training.triggers.function_trigger.ExtensionsManagerProtocol` property), 540
`elapsed_time` (`pytorch_pfn_extras.training.triggers.interval_trigger.ExtensionsManagerProtocol` property), 543
`elapsed_time` (`pytorch_pfn_extras.training.triggers.manual_schedule_trigger_extensions_manager_protocol` property), 546
`elapsed_time` (`pytorch_pfn_extras.training.triggers.minmax_value_trigger_extensions_manager_protocol` property), 551
`elapsed_time` (`pytorch_pfn_extras.training.triggers.once_trigger.ExtensionsManagerProtocol` property), 554
`elapsed_time` (`pytorch_pfn_extras.training.triggers.time_trigger.ExtensionsManagerProtocol` property), 558
`enable()` (`pytorch_pfn_extras.profiler.ChromeTracer` method), 275
`enable()` (`pytorch_pfn_extras.profiler.Tracer` method), 279
`enable()` (`pytorch_pfn_extras.training.extensions.timeline_trace.Tracer` method), 502
`enable_global_trace()` (in module `pytorch_pfn_extras.profiler`), 271
`enable_thread_trace()` (in module `pytorch_pfn_extras.profiler`), 272
`encoding` (`pytorch_pfn_extras.training.extensions.evaluator.TextIOProtocol` property), 428

encoding (pytorch_pfn_extras.training.extensions.util.TextLogger property), 508
 Ensure (class in pytorch_pfn_extras.nn), 106
 Ensure (class in pytorch_pfn_extras.nn.modules.ensure_shape), 130
 ensure() (in module pytorch_pfn_extras.nn), 105
 ensure() (in module pytorch_pfn_extras.nn.modules.ensure_shape), 129
 epoch (pytorch_pfn_extras.engine.Trainer property), 83
 epoch (pytorch_pfn_extras.training.extension.ExtensionsManager property), 322
 epoch (pytorch_pfn_extras.training.extensions.best_value.ExtensionsManager property), 413
 epoch (pytorch_pfn_extras.training.extensions.evaluator.ExtensionsManager property), 423
 epoch (pytorch_pfn_extras.training.extensions.fail_on_non_monotonicity.ExtensionsManager property), 429
 epoch (pytorch_pfn_extras.training.extensions.log_report.ExtensionsManager property), 432
 epoch (pytorch_pfn_extras.training.extensions.lr_scheduler.ExtensionsManager property), 438
 epoch (pytorch_pfn_extras.training.extensions.micro_average.ExtensionsManager property), 444
 epoch (pytorch_pfn_extras.training.extensions.parameter_statistics.ExtensionsManager property), 448
 epoch (pytorch_pfn_extras.training.extensions.plot_report.ExtensionsManager property), 453
 epoch (pytorch_pfn_extras.training.extensions.print_report.ExtensionsManager property), 459
 epoch (pytorch_pfn_extras.training.extensions.profile_report.ExtensionsManager property), 467
 epoch (pytorch_pfn_extras.training.extensions.progress_bar.ExtensionsManager property), 473
 epoch (pytorch_pfn_extras.training.extensions.slack.ExtensionsManager property), 477
 epoch (pytorch_pfn_extras.training.extensions.timeline_trace.ExtensionsManager property), 498
 epoch (pytorch_pfn_extras.training.extensions.util.ExtensionsManager property), 504
 epoch (pytorch_pfn_extras.training.extensions.value_observation.ExtensionsManager property), 510
 epoch (pytorch_pfn_extras.training.extensions.variable_statistics.ExtensionsManager property), 513
 epoch (pytorch_pfn_extras.training.ExtensionsManagerProtocol property), 309
 epoch (pytorch_pfn_extras.training.Trainer property), 314
 epoch (pytorch_pfn_extras.training.triggers.early_stopping_trigger.ExtensionsManager property), 538
 epoch (pytorch_pfn_extras.training.triggers.function_trigger.ExtensionsManager property), 540
 epoch (pytorch_pfn_extras.training.triggers.interval_trigger.ExtensionsManager property), 543
 epoch (pytorch_pfn_extras.training.triggers.manual_schedule_trigger.ExtensionsManager property), 546
 epoch (pytorch_pfn_extras.training.triggers.minmax_value_trigger.ExtensionsManager property), 551
 epoch (pytorch_pfn_extras.training.triggers.once_trigger.ExtensionsManager property), 554
 epoch (pytorch_pfn_extras.training.triggers.time_trigger.ExtensionsManager property), 558
 epoch_detail (pytorch_pfn_extras.engine.Trainer property), 83
 epoch_detail (pytorch_pfn_extras.training.extension.ExtensionsManager property), 322
 epoch_detail (pytorch_pfn_extras.training.extensions.best_value.ExtensionsManager property), 413
 epoch_detail (pytorch_pfn_extras.training.extensions.evaluator.ExtensionsManager property), 423
 epoch_detail (pytorch_pfn_extras.training.extensions.fail_on_non_monotonicity.ExtensionsManager property), 426
 epoch_detail (pytorch_pfn_extras.training.extensions.log_report.ExtensionsManager property), 429
 epoch_detail (pytorch_pfn_extras.training.extensions.lr_scheduler.ExtensionsManager property), 432
 epoch_detail (pytorch_pfn_extras.training.extensions.micro_average.ExtensionsManager property), 438
 epoch_detail (pytorch_pfn_extras.training.extensions.parameter_statistics.ExtensionsManager property), 444
 epoch_detail (pytorch_pfn_extras.training.extensions.plot_report.ExtensionsManager property), 448
 epoch_detail (pytorch_pfn_extras.training.extensions.print_report.ExtensionsManager property), 453
 epoch_detail (pytorch_pfn_extras.training.extensions.profile_report.ExtensionsManager property), 459
 epoch_detail (pytorch_pfn_extras.training.extensions.progress_bar.ExtensionsManager property), 467
 epoch_detail (pytorch_pfn_extras.training.extensions.slack.ExtensionsManager property), 473
 epoch_detail (pytorch_pfn_extras.training.extensions.timeline_trace.ExtensionsManager property), 477
 epoch_detail (pytorch_pfn_extras.training.extensions.util.ExtensionsManager property), 498
 epoch_detail (pytorch_pfn_extras.training.extensions.value_observation.ExtensionsManager property), 504
 epoch_detail (pytorch_pfn_extras.training.extensions.variable_statistics.ExtensionsManager property), 510
 epoch_detail (pytorch_pfn_extras.training.ExtensionsManagerProtocol property), 309
 epoch_detail (pytorch_pfn_extras.training.Trainer property), 314
 epoch_detail (pytorch_pfn_extras.training.triggers.early_stopping_trigger.ExtensionsManager property), 538
 epoch_detail (pytorch_pfn_extras.training.triggers.function_trigger.ExtensionsManager property), 540

epoch_detail (pytorch_pfn_extras.training.triggers.interval_trigger.ExtensibleManagerProtocol
 property), 543 eval_pre_step() (py-
 epoch_detail (pytorch_pfn_extras.training.triggers.manual_scheduler.trigger_extensions.manager.BaseProtocol
 property), 546 method), 289
 epoch_detail (pytorch_pfn_extras.training.triggers.minimal_eval_pre_step_extensions_manager.Protocol (py-
 property), 551 torch_pfn_extras.runtime.PyTorchRuntime
 epoch_detail (pytorch_pfn_extras.training.triggers.once_trigger.ExtensibleManagerProtocol
 property), 554 eval_setup() (pytorch_pfn_extras.handler.BaseHandler
 epoch_detail (pytorch_pfn_extras.training.triggers.time_trigger.ExtensibleManagerProtocol
 property), 558 eval_setup() (pytorch_pfn_extras.handler.Handler
 eps (pytorch_pfn_extras.nn.modules.lazy_batchnorm.LazyBatchNormMethod), 99
 attribute), 158 eval_step() (pytorch_pfn_extras.handler.BaseHandler
 eps (pytorch_pfn_extras.nn.modules.lazy_batchnorm.LazyBatchNormMethod), 88
 attribute), 161 eval_step() (pytorch_pfn_extras.handler.BaseLogic
 eps (pytorch_pfn_extras.nn.modules.lazy_batchnorm.LazyBatchNormMethod), 91
 attribute), 164 eval_step() (pytorch_pfn_extras.handler.CodeBlockLogic
 erase_console() (py- method), 96
 torch_pfn_extras.training.extensions.util.ProgressEval_step() (pytorch_pfn_extras.handler.Handler
 method), 506 method), 100
 errors (pytorch_pfn_extras.training.extensions.evaluator.TextIOEval_step() (pytorch_pfn_extras.handler.Logic
 property), 428 method), 103
 errors (pytorch_pfn_extras.training.extensions.util.TextIO evaluate() (pytorch_pfn_extras.training.extensions.Evaluator
 property), 508 method), 336
 eval_func (pytorch_pfn_extras.training.extensions.Evaluator evaluate() (pytorch_pfn_extras.training.extensions.evaluator.Evaluator
 attribute), 334 method), 422
 eval_func (pytorch_pfn_extras.training.extensions.evaluator.Evaluator evaluate() (pytorch_pfn_extras.training.extensions.evaluator.IgniteEvaluator
 attribute), 420 method), 425
 eval_func() (pytorch_pfn_extras.training.extensions.Evaluator evaluate() (pytorch_pfn_extras.training.extensions.IgniteEvaluator
 method), 336 method), 340
 eval_func() (pytorch_pfn_extras.training.extensions.evaluator.Evaluator class in pytorch_pfn_extras.engine), 79
 method), 422 Evaluator (class in pytorch_pfn_extras.training), 301
 eval_hook (pytorch_pfn_extras.training.extensions.Evaluator Evaluator (class in py-
 attribute), 334 torch_pfn_extras.training.extensions), 333
 eval_hook (pytorch_pfn_extras.training.extensions.evaluator.Evaluator (class in py-
 attribute), 420 torch_pfn_extras.training.extensions.evaluator),
 eval_loop_begin() (py- 419
 torch_pfn_extras.handler.BaseHandler evaluator (pytorch_pfn_extras.engine.Trainer prop-
 method), 87 erty), 83
 eval_loop_end() (py- evaluator (pytorch_pfn_extras.training.Trainer prop-
 torch_pfn_extras.handler.BaseHandler erty), 315
 method), 87 execute() (pytorch_pfn_extras.runtime.BaseRuntime
 eval_loop_end() (pytorch_pfn_extras.handler.Handler method), 289
 method), 99 execute() (pytorch_pfn_extras.runtime.PyTorchRuntime
 eval_post_step() (py- method), 294
 torch_pfn_extras.handler.BaseHandler exists() (pytorch_pfn_extras.onnx.load.Path method),
 method), 88 261
 eval_post_step() (py- exists() (pytorch_pfn_extras.onnx.unstrip_tensor.Path
 torch_pfn_extras.handler.Handler method),
 99 method), 267
 eval_post_step() (py- expanduser() (pytorch_pfn_extras.onnx.load.Path
 torch_pfn_extras.runtime.BaseRuntime method), 261
 method), 289 expanduser() (pytorch_pfn_extras.onnx.unstrip_tensor.Path
 method), 267
 eval_post_step() (py- export() (in module pytorch_pfn_extras.onnx), 251
 torch_pfn_extras.runtime.PyTorchRuntime export_testcase() (in module py-

`torch_pfn_extras.onnx`), 251
`extend()` (`pytorch_pfn_extras.engine.Trainer` method), 83
`extend()` (`pytorch_pfn_extras.training.Trainer` method), 315
`ExtendedSequential` (class in `pytorch_pfn_extras.nn`), 109
`ExtendedSequential` (class in `pytorch_pfn_extras.nn.modules.extended_sequential`), 133
`Extension` (class in `pytorch_pfn_extras.training`), 302
`Extension` (class in `pytorch_pfn_extras.training.extension`), 318
`ExtensionEntry` (class in `pytorch_pfn_extras.training`), 304
`ExtensionEntry` (class in `pytorch_pfn_extras.training.extension`), 320
`ExtensionsManager` (class in `pytorch_pfn_extras.training`), 305
`ExtensionsManagerProtocol` (class in `pytorch_pfn_extras.training`), 308
`ExtensionsManagerProtocol` (class in `pytorch_pfn_extras.training.extension`), 321
`ExtensionsManagerProtocol` (class in `pytorch_pfn_extras.training.extensions.best_value`), 413
`ExtensionsManagerProtocol` (class in `pytorch_pfn_extras.training.extensions.evaluator`), 423
`ExtensionsManagerProtocol` (class in `pytorch_pfn_extras.training.extensions.fail_on_non_number`), 428
`ExtensionsManagerProtocol` (class in `pytorch_pfn_extras.training.extensions.log_report`), 432
`ExtensionsManagerProtocol` (class in `pytorch_pfn_extras.training.extensions.lr_scheduler`), 437
`ExtensionsManagerProtocol` (class in `pytorch_pfn_extras.training.extensions.micro_average`), 443
`ExtensionsManagerProtocol` (class in `pytorch_pfn_extras.training.extensions.parameter_statistics`), 447
`ExtensionsManagerProtocol` (class in `pytorch_pfn_extras.training.extensions.plot_report`), 452
`ExtensionsManagerProtocol` (class in `pytorch_pfn_extras.training.extensions.print_report`), 458
`ExtensionsManagerProtocol` (class in `pytorch_pfn_extras.training.extensions.profile_report`), 466
`ExtensionsManagerProtocol` (class in `pytorch_pfn_extras.training.extensions.progress_bar`), 472
`ExtensionsManagerProtocol` (class in `pytorch_pfn_extras.training.extensions.slack`), 476
`ExtensionsManagerProtocol` (class in `pytorch_pfn_extras.training.extensions.timeline_trace`), 497
`ExtensionsManagerProtocol` (class in `pytorch_pfn_extras.training.extensions.util`), 504
`ExtensionsManagerProtocol` (class in `pytorch_pfn_extras.training.extensions.value_observation`), 509
`ExtensionsManagerProtocol` (class in `pytorch_pfn_extras.training.extensions.variable_statistics_plot`), 512
`ExtensionsManagerProtocol` (class in `pytorch_pfn_extras.training.triggers.early_stopping_trigger`), 537
`ExtensionsManagerProtocol` (class in `pytorch_pfn_extras.training.triggers.function_trigger`), 539
`ExtensionsManagerProtocol` (class in `pytorch_pfn_extras.training.triggers.interval_trigger`), 542
`ExtensionsManagerProtocol` (class in `pytorch_pfn_extras.training.triggers.manual_schedule_trigger`), 546
`ExtensionsManagerProtocol` (class in `pytorch_pfn_extras.training.triggers.minmax_value_trigger`), 550
`ExtensionsManagerProtocol` (class in `pytorch_pfn_extras.training.triggers.once_trigger`), 554
`ExtensionsManagerProtocol` (class in `pytorch_pfn_extras.training.triggers.time_trigger`), 557
F
`FailOnNonNumber` (class in `pytorch_pfn_extras.training.extensions`), 337
`FailOnNonNumber` (class in `pytorch_pfn_extras.training.extensions.fail_on_non_number`), 430
`fetch()` (`pytorch_pfn_extras.dataset.tabular.tabular_dataset.TabularDataset` method), 71
`fetch()` (`pytorch_pfn_extras.dataset.TabularDataset` method), 59
`fileno()` (`pytorch_pfn_extras.onnx.load.IO` method), 256
`fileno()` (`pytorch_pfn_extras.training.extensions.print_report.IO` method), 461
`filter_and_sort_entries()` (in module `py-`

`torch_pfn_extras.training.extensions.print_report` (method), 457
`filter_state_objects()` (in module `pytorch_pfn_extras.engine`), 78
`filter_state_objects_dict()` (in module `pytorch_pfn_extras.engine`), 78
`finalize()` (`pytorch_pfn_extras.profiler.ChromeTracer` method), 275
`finalize()` (`pytorch_pfn_extras.profiler.TimeSummary` method), 277
`finalize()` (`pytorch_pfn_extras.profiler.Tracer` method), 279
`finalize()` (`pytorch_pfn_extras.training.Extension` method), 303
`finalize()` (`pytorch_pfn_extras.training.extension.Extensio` method), 319
`finalize()` (`pytorch_pfn_extras.training.extensions.log_report` method), 435
`finalize()` (`pytorch_pfn_extras.training.extensions.LogReport` method), 344
`finalize()` (`pytorch_pfn_extras.training.extensions.plot_report` method), 456
`finalize()` (`pytorch_pfn_extras.training.extensions.PlotReport` method), 357
`finalize()` (`pytorch_pfn_extras.training.extensions.profile_report` method), 471
`finalize()` (`pytorch_pfn_extras.training.extensions.ProfileReport` method), 361
`finalize()` (`pytorch_pfn_extras.training.extensions.progress_bar` method), 475
`finalize()` (`pytorch_pfn_extras.training.extensions.ProgressBar` method), 363
`finalize()` (`pytorch_pfn_extras.training.extensions.snapshot_writers` method), 488
`finalize()` (`pytorch_pfn_extras.training.extensions.snapshot_writers` method), 491
`finalize()` (`pytorch_pfn_extras.training.extensions.snapshot_writers` method), 492
`finalize()` (`pytorch_pfn_extras.training.extensions.snapshot_writers` method), 496
`finalize()` (`pytorch_pfn_extras.training.extensions.timeline_trace` method), 501
`finalize()` (`pytorch_pfn_extras.training.extensions.timeline_trace` method), 503
`finalize()` (`pytorch_pfn_extras.training.extensions.TimelineTrace` method), 372
`finalize()` (`pytorch_pfn_extras.training.extensions.variable_statistics` method), 519
`finalize()` (`pytorch_pfn_extras.training.extensions.VariableStatistics` method), 376
`finalize()` (`pytorch_pfn_extras.training.ExtensionsManager` method), 307
`finalize()` (`pytorch_pfn_extras.writing.QueueWriter` method), 571
`finalize()` (`pytorch_pfn_extras.writing.StandardWriter` method), 574
`finalize()` (`pytorch_pfn_extras.writing.TensorBoardWriter` method), 576
`finalize()` (`pytorch_pfn_extras.writing.Writer` method), 579
`find_inplace()` (in module `pytorch_pfn_extras.torchscript`), 297
`finished()` (`pytorch_pfn_extras.training.triggers.once_trigger.OnceTrigger` attribute), 555
`finished()` (`pytorch_pfn_extras.training.triggers.once_trigger.OnceTrigger` property), 556
`finished()` (`pytorch_pfn_extras.training.triggers.OnceTrigger` attribute), 533
`finished()` (`pytorch_pfn_extras.training.triggers.OnceTrigger` property), 533
`flush()` (`pytorch_pfn_extras.onnx.load.IO` method), 257
`flush()` (`pytorch_pfn_extras.profiler.ChromeTracer` method), 275
`flush()` (`pytorch_pfn_extras.profiler.Tracer` method), 279
`flush()` (`pytorch_pfn_extras.training.extensions.print_report.IO` method), 462
`flush()` (`pytorch_pfn_extras.training.extensions.timeline_trace.Tracer` method), 503
`flush()` (`pytorch_pfn_extras.training.extensions.util.ProgressBar` method), 506
`forward()` (in module `pytorch_pfn_extras.handler`), 85
`forward()` (`pytorch_pfn_extras.nn.Ensure` method), 108
`forward()` (`pytorch_pfn_extras.nn.LazyLinear` method), 128
`forward()` (`pytorch_pfn_extras.nn.modules.ensure_shape.Ensure` method), 108
`forward()` (`pytorch_pfn_extras.nn.modules.lazy_linear.LazyLinear` method), 128
`forward()` (`pytorch_pfn_extras.nn.parallel.distributed.DistributedDataParallel` method), 235
`forward()` (`pytorch_pfn_extras.nn.parallel.DistributedDataParallel` method), 235
`from_data()` (in module `pytorch_pfn_extras.dataset.tabular`), 63
`from_ndarray()` (in module `pytorch_pfn_extras`), 39
`from_numpy_dtype()` (in module `pytorch_pfn_extras`), 39
`from_keys()` (`pytorch_pfn_extras.nn.parallel.distributed.OrderedDict` method), 245
`from_keys()` (`pytorch_pfn_extras.nn.parallel.distributed.OrderedDict` method), 245
`from_keys()` (`pytorch_pfn_extras.training.extensions.profile_report.OrderedDict` method), 468
`from_keys()` (`pytorch_pfn_extras.handler.CodeBlock` attribute), 95
`FunctionTrigger` (class in `pytorch_pfn_extras.training.triggers`), 527
`FunctionTrigger` (class in `pytorch_pfn_extras.training.triggers.function_trigger`), 527

| | | | |
|---|--|---|------|
| 541 | | get_extension() | (py- |
| | | torch_pfn_extras.training.extensions.parameter_statistics.ExtensionsManagerProtocol | |
| | | method), 448 | |
| G | | | |
| get_all_iterators() | (py- | get_extension() | (py- |
| torch_pfn_extras.training.extensions.Evaluator | | torch_pfn_extras.training.extensions.plot_report.ExtensionsManagerProtocol | |
| method), 336 | | method), 453 | |
| get_all_iterators() | (py- | get_extension() | (py- |
| torch_pfn_extras.training.extensions.evaluator.Evaluator | | torch_pfn_extras.training.extensions.print_report.ExtensionsManagerProtocol | |
| method), 422 | | method), 459 | |
| get_all_targets() | (py- | get_extension() | (py- |
| torch_pfn_extras.training.extensions.Evaluator | | torch_pfn_extras.training.extensions.profile_report.ExtensionsManagerProtocol | |
| method), 336 | | method), 467 | |
| get_all_targets() | (py- | get_extension() | (py- |
| torch_pfn_extras.training.extensions.evaluator.Evaluator | | torch_pfn_extras.training.extensions.progress_bar.ExtensionsManagerProtocol | |
| method), 422 | | method), 473 | |
| get_current_reporter() | (in module py- | get_extension() | (py- |
| torch_pfn_extras.reporting), 280 | | torch_pfn_extras.training.extensions.slack.ExtensionsManagerProtocol | |
| get_data() | (pytorch_pfn_extras.training.extensions.variable_statistics_plot.Reservoir | get_extension() | (py- |
| method), 515 | | method), 477 | |
| get_default_comparer() | (in module py- | get_extension() | (py- |
| torch_pfn_extras.utils.comparer), 561 | | torch_pfn_extras.training.extensions.timeline_trace.ExtensionsManagerProtocol | |
| method), 498 | | | |
| get_example() | (pytorch_pfn_extras.dataset.tabular.tabular_dataset.TabularDataset | get_extension() | (py- |
| method), 71 | | torch_pfn_extras.training.extensions.util.ExtensionsManagerProtocol | |
| get_example() | (pytorch_pfn_extras.dataset.TabularDataset | method), 504 | |
| method), 59 | | get_extension() | (py- |
| get_examples() | (pytorch_pfn_extras.dataset.tabular.delegate_dataset.DelegateDataset | method), 510 | |
| method), 67 | | torch_pfn_extras.training.extensions.value_observation.ExtensionsManagerProtocol | |
| get_examples() | (pytorch_pfn_extras.dataset.tabular.DelegateDataset | get_extension() | (py- |
| method), 65 | | torch_pfn_extras.training.extensions.variable_statistics_plot.ExtensionsManagerProtocol | |
| get_examples() | (pytorch_pfn_extras.dataset.tabular.tabular_dataset.TabularDataset | method), 512 | |
| method), 71 | | get_extension() | (py- |
| get_examples() | (pytorch_pfn_extras.dataset.TabularDataset | torch_pfn_extras.training.ExtensionsManagerProtocol | |
| method), 59 | | method), 309 | |
| get_extension() | (py- | get_extension() | (py- |
| torch_pfn_extras.training.extension.ExtensionsManagerProtocol | | torch_pfn_extras.training.triggers.early_stopping_trigger.ExtensionsManagerProtocol | |
| method), 322 | | method), 538 | |
| get_extension() | (py- | get_extension() | (py- |
| torch_pfn_extras.training.extensions.best_value.ExtensionsManagerProtocol | | torch_pfn_extras.training.triggers.function_trigger.ExtensionsManagerProtocol | |
| method), 413 | | method), 540 | |
| get_extension() | (py- | get_extension() | (py- |
| torch_pfn_extras.training.extensions.evaluator.ExtensionsManagerProtocol | | torch_pfn_extras.training.triggers.interval_trigger.ExtensionsManagerProtocol | |
| method), 423 | | method), 543 | |
| get_extension() | (py- | get_extension() | (py- |
| torch_pfn_extras.training.extensions.fail_on_non_number.ExtensionsManagerProtocol | | torch_pfn_extras.training.triggers.manual_schedule_trigger.ExtensionsManagerProtocol | |
| method), 429 | | method), 546 | |
| get_extension() | (py- | get_extension() | (py- |
| torch_pfn_extras.training.extensions.log_report.ExtensionsManagerProtocol | | torch_pfn_extras.training.triggers.minmax_value_trigger.ExtensionsManagerProtocol | |
| method), 432 | | method), 551 | |
| get_extension() | (py- | get_extension() | (py- |
| torch_pfn_extras.training.extensions.lr_scheduler.ExtensionsManagerProtocol | | torch_pfn_extras.training.triggers.once_trigger.ExtensionsManagerProtocol | |
| method), 438 | | method), 554 | |
| get_extension() | (py- | get_extension() | (py- |
| torch_pfn_extras.training.extensions.micro_average.ExtensionsManagerProtocol | | torch_pfn_extras.training.triggers.time_trigger.ExtensionsManagerProtocol | |
| method), 444 | | method), 558 | |

[get_foreach_wrapper\(\)](#) (in module `pytorch_pfn_extras.nn.parallel.distributed`), 239
[get_iterator\(\)](#) (`pytorch_pfn_extras.training.extensions.Evaluator` method), 337
[get_iterator\(\)](#) (`pytorch_pfn_extras.training.extensions.evaluator.Evaluator` method), 422
[get_lines\(\)](#) (`pytorch_pfn_extras.training.extensions.util.ProgressBar` method), 506
[get_log_report\(\)](#) (`pytorch_pfn_extras.training.extensions.print_report.PrintReport` method), 464
[get_log_report\(\)](#) (`pytorch_pfn_extras.training.extensions.PrintReport` method), 358
[get_logger\(\)](#) (in module `pytorch_pfn_extras.logging`), 105
[get_optimizer\(\)](#) (`pytorch_pfn_extras.engine.Trainer` method), 83
[get_optimizer\(\)](#) (`pytorch_pfn_extras.training.Trainer` method), 315
[get_target\(\)](#) (`pytorch_pfn_extras.training.extensions.Evaluator` method), 337
[get_target\(\)](#) (`pytorch_pfn_extras.training.extensions.evaluator.Evaluator` method), 422
[get_time_summary\(\)](#) (in module `pytorch_pfn_extras.profiler`), 272
[get_time_summary\(\)](#) (in module `pytorch_pfn_extras.training.extensions.profile_report`), 465
[get_tracer\(\)](#) (in module `pytorch_pfn_extras.profiler`), 272
[get_tracer\(\)](#) (in module `pytorch_pfn_extras.training.extensions.timeline_trace`), 497
[get_training_length\(\)](#) (`pytorch_pfn_extras.training.trigger.IntervalTrigger` method), 522
[get_training_length\(\)](#) (`pytorch_pfn_extras.training.triggers.early_stopping_trigger.EarlyStoppingTrigger` method), 537
[get_training_length\(\)](#) (`pytorch_pfn_extras.training.triggers.EarlyStoppingTrigger` method), 527
[get_training_length\(\)](#) (`pytorch_pfn_extras.training.triggers.interval_trigger.IntervalTrigger` method), 545
[get_training_length\(\)](#) (`pytorch_pfn_extras.training.triggers.IntervalTrigger` method), 529
[get_trigger\(\)](#) (in module `pytorch_pfn_extras.training.trigger`), 521
[get_value\(\)](#) (`pytorch_pfn_extras.dataset.shared_dataset.Cache` method), 61
[get_value\(\)](#) (`pytorch_pfn_extras.dataset.shared_dataset.InfiniteCache` method), 62
[get_worker_info\(\)](#) (in module `pytorch_pfn_extras.dataloaders.data_loader`), 51
[get_worker_info\(\)](#) (in module `pytorch_pfn_extras`), 39
[glob\(\)](#) (`pytorch_pfn_extras.onnx.load.Path` method), 261
[glob\(\)](#) (`pytorch_pfn_extras.onnx.unstrip_tensor.Path` method), 267
[group\(\)](#) (`pytorch_pfn_extras.onnx.load.Path` method), 261
[group\(\)](#) (`pytorch_pfn_extras.onnx.unstrip_tensor.Path` method), 267
[groups](#) (`pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv1d` attribute), 185
[groups](#) (`pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv2d` attribute), 188
[groups](#) (`pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv3d` attribute), 191

H

[Handler](#) (`pytorch_pfn_extras.handler.Handler` class in `pytorch_pfn_extras.handler`), 98
[home\(\)](#) (`pytorch_pfn_extras.onnx.load.Path` class method), 261
[home\(\)](#) (`pytorch_pfn_extras.onnx.unstrip_tensor.Path` class method), 267

I

[IgniteEvaluator](#) (class in `pytorch_pfn_extras.training.extensions`), 338
[IgniteEvaluator](#) (class in `pytorch_pfn_extras.training.extensions.evaluator`), 424
[IgniteExtensionsManager](#) (class in `pytorch_pfn_extras.training`), 310
[in_channels](#) (`pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv1d` attribute), 185
[in_channels](#) (`pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv2d` attribute), 188
[in_channels](#) (`pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv3d` attribute), 191
[in_cooldown](#) (`pytorch_pfn_extras.training.extensions.lr_scheduler.ReduceLrOnPlateau` property), 442
[in_features](#) (`pytorch_pfn_extras.nn.modules.lazy_linear.LazyLinear` attribute), 215
[InfiniteCache](#) (class in `pytorch_pfn_extras.dataset.shared_dataset`), 62
[initialize\(\)](#) (`pytorch_pfn_extras.profiler.TimeSummary` method), 277
[initialize\(\)](#) (`pytorch_pfn_extras.training.Extension` method), 303

`initialize()` (`pytorch_pfn_extras.training.extension.Extension` method), 319
`initialize()` (`pytorch_pfn_extras.training.extensions.print_report.PrintReport` method), 465
`initialize()` (`pytorch_pfn_extras.training.extensions.snapshot_writers.Writer` method), 496
`initialize()` (`pytorch_pfn_extras.training.extensions.timeline_trace.IntervalTrigger` method), 501
`initialize()` (`pytorch_pfn_extras.training.extensions.TimelineTracer` method), 372
`initialize()` (`pytorch_pfn_extras.writing.Writer` method), 579
`initialize_module()` (`pytorch_pfn_extras.runtime.BaseRuntime` method), 289
`initialize_module()` (`pytorch_pfn_extras.runtime.PyTorchRuntime` method), 294
`initialize_ompi_environment()` (in module `pytorch_pfn_extras.distributed`), 74
`initialize_writer()` (`pytorch_pfn_extras.profiler.ChromeTracer` method), 275
`initialize_writer()` (`pytorch_pfn_extras.profiler.Tracer` method), 280
`initialize_writer()` (`pytorch_pfn_extras.training.extensions.timeline_trace.IntervalTrigger` method), 503
`intermediate_value()` (in module `pytorch_pfn_extras.utils.comparer`), 561
`IntervalTrigger` (class in `pytorch_pfn_extras.training.trigger`), 522
`IntervalTrigger` (class in `pytorch_pfn_extras.training.triggers`), 529
`IntervalTrigger` (class in `pytorch_pfn_extras.training.triggers.interval_trigger`), 544
`IO` (class in `pytorch_pfn_extras.onnx.load`), 255
`IO` (class in `pytorch_pfn_extras.training.extensions.print_report.PrintReport`), 460
`is_async` (`pytorch_pfn_extras.training.Extension` attribute), 303
`is_async` (`pytorch_pfn_extras.training.extension.Extension` attribute), 319
`is_before_training` (`pytorch_pfn_extras.engine.Trainer` property), 83
`is_before_training` (`pytorch_pfn_extras.training.extension.ExtensionsManagerProtocol` property), 322
`is_before_training` (`pytorch_pfn_extras.training.Trainer` property), 315
`is_before_training` (`pytorch_pfn_extras.training.extensions.best_value.ExtensionsManagerProtocol` property), 414
`is_before_training` (`pytorch_pfn_extras.training.extensions.evaluator.ExtensionsManagerProtocol` property), 424
`is_before_training` (`pytorch_pfn_extras.training.extensions.fail_on_non_number.ExtensionsManagerProtocol` property), 429
`is_before_training` (`pytorch_pfn_extras.training.extensions.log_report.ExtensionsManagerProtocol` property), 433
`is_before_training` (`pytorch_pfn_extras.training.extensions.lr_scheduler.ExtensionsManagerProtocol` property), 438
`is_before_training` (`pytorch_pfn_extras.training.extensions.micro_average.ExtensionsManagerProtocol` property), 444
`is_before_training` (`pytorch_pfn_extras.training.extensions.parameter_statistics.ExtensionsManagerProtocol` property), 448
`is_before_training` (`pytorch_pfn_extras.training.extensions.plot_report.ExtensionsManagerProtocol` property), 453
`is_before_training` (`pytorch_pfn_extras.training.extensions.print_report.ExtensionsManagerProtocol` property), 459
`is_before_training` (`pytorch_pfn_extras.training.extensions.profile_report.ExtensionsManagerProtocol` property), 467
`is_before_training` (`pytorch_pfn_extras.training.extensions.progress_bar.ExtensionsManagerProtocol` property), 473
`is_before_training` (`pytorch_pfn_extras.training.extensions.slack.ExtensionsManagerProtocol` property), 477
`is_before_training` (`pytorch_pfn_extras.training.extensions.timeline_trace.ExtensionsManagerProtocol` property), 498
`is_before_training` (`pytorch_pfn_extras.training.extensions.util.ExtensionsManagerProtocol` property), 505
`is_before_training` (`pytorch_pfn_extras.training.extensions.value_observation.ExtensionsManagerProtocol` property), 510
`is_before_training` (`pytorch_pfn_extras.training.extensions.variable_statistics_plot.ExtensionsManagerProtocol` property), 513
`is_before_training` (`pytorch_pfn_extras.training.ExtensionsManagerProtocol` property), 309
`is_before_training` (`pytorch_pfn_extras.training.Trainer` property), 315
`is_before_training` (`pytorch_pfn_extras.training.Trainer` property), 315

| | |
|--|--|
| <code>torch_pfn_extras.training.triggers.early_stopping_trigger.ExtensionsManagerProtocol</code> (property), 538 | <code>is_file()</code> (pytorch_pfn_extras.onnx.load.Path method), 261 |
| <code>is_before_training</code> (pytorch_pfn_extras.training.triggers.function_trigger.ExtensionsManagerProtocol (property), 540 | <code>is_file()</code> (pytorch_pfn_extras.onnx.unstrip_tensor.Path method), 267 |
| <code>is_before_training</code> (pytorch_pfn_extras.training.triggers.interval_trigger.ExtensionsManagerProtocol (property), 543 | <code>is_large_tensor()</code> (in module py- torch_pfn_extras.onnx), 253 |
| <code>is_before_training</code> (pytorch_pfn_extras.training.triggers.manual_schedule_trigger.ExtensionsManagerProtocol (property), 547 | <code>is_large_tensor()</code> (in module py- torch_pfn_extras.onnx.strip_large_tensor), 253 |
| <code>is_before_training</code> (pytorch_pfn_extras.training.triggers.minmax_value_trigger.ExtensionsManagerProtocol (property), 551 | <code>is_leaf</code> (pytorch_pfn_extras.nn.modules.lazy.UninitializedParameter property), 154 |
| <code>is_before_training</code> (pytorch_pfn_extras.training.triggers.once_trigger.ExtensionsManagerProtocol (property), 555 | <code>is_leaf</code> (pytorch_pfn_extras.nn.modules.lazy_batchnorm.UninitializedPar property), 182 |
| <code>is_before_training</code> (pytorch_pfn_extras.training.triggers.time_trigger.ExtensionsManagerProtocol (property), 558 | <code>is_leaf</code> (pytorch_pfn_extras.nn.modules.lazy_conv.UninitializedParameter property), 200 |
| <code>is_better()</code> (pytorch_pfn_extras.training.extensions.lr_scheduler.LRScheduler method), 442 | <code>is_leaf</code> (pytorch_pfn_extras.nn.modules.lazy_linear.UninitializedParameter property), 231 |
| <code>is_block_device()</code> (pytorch_pfn_extras.onnx.load.Path method), 261 | <code>is_socket()</code> (pytorch_pfn_extras.onnx.load.Path method), 261 |
| <code>is_block_device()</code> (pytorch_pfn_extras.onnx.unstrip_tensor.Path method), 267 | <code>is_socket()</code> (pytorch_pfn_extras.onnx.unstrip_tensor.Path method), 268 |
| <code>is_cached()</code> (pytorch_pfn_extras.dataset.shared_dataset.GatherDataset method), 62 | <code>is_symlink()</code> (pytorch_pfn_extras.onnx.load.Path method), 261 |
| <code>is_cached()</code> (pytorch_pfn_extras.dataset.shared_dataset.InstantiatedDataset method), 62 | <code>is_symlink()</code> (pytorch_pfn_extras.onnx.unstrip_tensor.Path method), 268 |
| <code>is_cached()</code> (pytorch_pfn_extras.dataset.shared_dataset.SharedDataset method), 63 | <code>is_sync_device()</code> (pytorch_pfn_extras.training.extensions.print_report.IO method), 462 |
| <code>is_cached()</code> (pytorch_pfn_extras.dataset.SharedDataset method), 56 | <code>ItemNotFoundException</code> , 61, 63 |
| <code>is_char_device()</code> (pytorch_pfn_extras.onnx.load.Path method), 261 | <code>items()</code> (pytorch_pfn_extras.nn.parallel.distributed.OrderedDict method), 245 |
| <code>is_char_device()</code> (pytorch_pfn_extras.onnx.unstrip_tensor.Path method), 267 | <code>items()</code> (pytorch_pfn_extras.training.extensions.profile_report.OrderedDict method), 469 |
| <code>is_dir()</code> (pytorch_pfn_extras.onnx.load.Path method), 261 | <code>iteration</code> (pytorch_pfn_extras.engine.Trainer prop- erty), 83 |
| <code>is_dir()</code> (pytorch_pfn_extras.onnx.unstrip_tensor.Path method), 267 | <code>iteration</code> (pytorch_pfn_extras.training.extension.ExtensionsManagerProt property), 323 |
| <code>is_epoch_last_iter()</code> (pytorch_pfn_extras.engine.Trainer method), 83 | <code>iteration</code> (pytorch_pfn_extras.training.extensions.best_value.ExtensionsM property), 414 |
| <code>is_epoch_last_iter()</code> (pytorch_pfn_extras.training.Trainer method), 315 | <code>iteration</code> (pytorch_pfn_extras.training.extensions.evaluator.ExtensionsM property), 424 |
| <code>is_fifo()</code> (pytorch_pfn_extras.onnx.load.Path method), 261 | <code>iteration</code> (pytorch_pfn_extras.training.extensions.fail_on_non_number.E property), 430 |
| <code>is_fifo()</code> (pytorch_pfn_extras.onnx.unstrip_tensor.Path method), 267 | <code>iteration</code> (pytorch_pfn_extras.training.extensions.log_report.ExtensionsM property), 433 |
| | <code>iteration</code> (pytorch_pfn_extras.training.extensions.lr_scheduler.Extension property), 438 |
| | <code>iteration</code> (pytorch_pfn_extras.training.extensions.micro_average.Extension property), 441 |

- property), 444
- iteration(pytorch_pfn_extras.training.extensions.parameter_statistics.ExtensionsManagerProtocol property), 448
- iteration(pytorch_pfn_extras.training.extensions.plot_report.ExtensionsManagerProtocol property), 453
- iteration(pytorch_pfn_extras.training.extensions.print_report.ExtensionsManagerProtocol property), 459
- iteration(pytorch_pfn_extras.training.extensions.profile_report.ExtensionsManagerProtocol property), 467
- iteration(pytorch_pfn_extras.training.extensions.progress_bar.ExtensionsManagerProtocol property), 473
- iteration(pytorch_pfn_extras.training.extensions.slack.ExtensionsManagerProtocol property), 477
- iteration(pytorch_pfn_extras.training.extensions.timeline_trace.ExtensionsManagerProtocol property), 499
- iteration(pytorch_pfn_extras.training.extensions.util.ExtensionsManagerProtocol property), 505
- iteration(pytorch_pfn_extras.training.extensions.value_observation.ExtensionsManagerProtocol property), 510
- iteration(pytorch_pfn_extras.training.extensions.variable_statistics_plot.ExtensionsManagerProtocol property), 513
- ## K
- kernel_size(pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv1d attribute), 185
- kernel_size(pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv2d attribute), 188
- kernel_size(pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv3d attribute), 191
- keys(pytorch_pfn_extras.dataset.tabular.delegate_dataset.DelegateDataset property), 68
- keys(pytorch_pfn_extras.dataset.tabular.DelegateDataset property), 65
- keys(pytorch_pfn_extras.dataset.tabular.tabular_dataset.TabularDataset property), 71
- keys(pytorch_pfn_extras.dataset.TabularDataset property), 59
- keys()(pytorch_pfn_extras.nn.parallel.distributed.OrderedDict method), 246
- keys()(pytorch_pfn_extras.training.extensions.profile_report.OrderedDict method), 469
- ## L
- iteration(pytorch_pfn_extras.training.ExtensionsManagerProtocol property), 309
- iteration(pytorch_pfn_extras.training.Trainer property), 315
- iteration(pytorch_pfn_extras.training.triggers.early_stopping_trigger.ExtensionsManagerProtocol property), 539
- iteration(pytorch_pfn_extras.training.triggers.function_trigger.ExtensionsManagerProtocol property), 540
- iteration(pytorch_pfn_extras.training.triggers.interval_trigger.ExtensionsManagerProtocol property), 544
- iteration(pytorch_pfn_extras.training.triggers.manual_schedule_trigger.ExtensionsManagerProtocol property), 547
- iteration(pytorch_pfn_extras.training.triggers.minmax_value_trigger.ExtensionsManagerProtocol property), 551
- iteration(pytorch_pfn_extras.training.triggers.once_trigger.ExtensionsManagerProtocol property), 555
- iteration(pytorch_pfn_extras.training.triggers.time_trigger.ExtensionsManagerProtocol property), 558
- IterationStatus (class in pytorch_pfn_extras.training.extensions.evaluator), 426
- iterdir()(pytorch_pfn_extras.onnx.load.Path method), 261
- iterdir()(pytorch_pfn_extras.onnx.unstrip_tensor.Path method), 268
- ## J
- join()(pytorch_pfn_extras.dataset.tabular.tabular_dataset.TabularDataset method), 71
- join()(pytorch_pfn_extras.dataset.TabularDataset method), 59
- lazy_buffer_names (pytorch_pfn_extras.nn.modules.lazy.LazyInitializationMixin attribute), 137
- lazy_buffer_names (pytorch_pfn_extras.nn.modules.lazy_batchnorm.LazyInitializationMixin attribute), 165
- lazy_buffer_names (pytorch_pfn_extras.nn.modules.lazy_conv.LazyInitializationMixin attribute), 192
- lazy_buffer_names (pytorch_pfn_extras.nn.modules.lazy_linear.LazyInitializationMixin attribute), 211
- lazy_parameter_names (pytorch_pfn_extras.nn.LazyLinear attribute), 128
- lazy_parameter_names (pytorch_pfn_extras.nn.modules.lazy.LazyInitializationMixin attribute), 138
- lazy_parameter_names (pytorch_pfn_extras.nn.modules.lazy_batchnorm.LazyInitializationMixin attribute), 165
- lazy_parameter_names (pytorch_pfn_extras.nn.modules.lazy_conv.LazyInitializationMixin attribute), 193
- lazy_parameter_names (pytorch_pfn_extras.nn.modules.lazy_linear.LazyInitializationMixin attribute), 212
- lazy_parameters_determined (pytorch_pfn_extras.nn.modules.lazy_linear.LazyLinear attribute), 215
- lazy_parameters_determined (pytorch_pfn_extras.nn.modules.lazy.LazyInitializationMixin property), 138

| | | |
|--|--|--|
| lazy_parameters_determined | (pytorch_pfn_extras.nn.modules.lazy_batchnorm.LazyInitializationMixin property), 165 | line_buffering (pytorch_pfn_extras.training.extensions.util.TextIO property), 508 |
| lazy_parameters_determined | (pytorch_pfn_extras.nn.modules.lazy_conv.LazyInitializationMixin property), 193 | link_to() (pytorch_pfn_extras.onnx.load.Path method), 261 |
| lazy_parameters_determined | (pytorch_pfn_extras.nn.modules.lazy_linear.LazyInitializationMixin property), 212 | link_to() (pytorch_pfn_extras.onnx.unstrip_tensor.Path method), 268 |
| LazyBatchNorm1d (class in pytorch_pfn_extras.nn), 111 | | load_chrome_trace_as_json() (in module pytorch_pfn_extras.profiler), 272 |
| LazyBatchNorm1d (class in pytorch_pfn_extras.nn.modules.lazy_batchnorm), 156 | | load_model() (in module pytorch_pfn_extras.onnx), 253 |
| LazyBatchNorm2d (class in pytorch_pfn_extras.nn), 114 | | load_model() (in module pytorch_pfn_extras.onnx.load), 255 |
| LazyBatchNorm2d (class in pytorch_pfn_extras.nn.modules.lazy_batchnorm), 159 | | load_path() (pytorch_pfn_extras.config.Config class method), 43 |
| LazyBatchNorm3d (class in pytorch_pfn_extras.nn), 116 | | load_path_with_optuna_types() (in module pytorch_pfn_extras.config_types), 43 |
| LazyBatchNorm3d (class in pytorch_pfn_extras.nn.modules.lazy_batchnorm), 161 | | load_state_dict() (pytorch_pfn_extras.engine.StateObjectProtocol method), 81 |
| LazyConv1d (class in pytorch_pfn_extras.nn), 119 | | load_state_dict() (pytorch_pfn_extras.engine.Trainer method), 83 |
| LazyConv1d (class in pytorch_pfn_extras.nn.modules.lazy_conv), 183 | | load_state_dict() (pytorch_pfn_extras.handler.CodeBlock method), 95 |
| LazyConv2d (class in pytorch_pfn_extras.nn), 121 | | load_state_dict() (pytorch_pfn_extras.nn.parallel.distributed.DistributedDataParallel method), 243 |
| LazyConv2d (class in pytorch_pfn_extras.nn.modules.lazy_conv), 186 | | load_state_dict() (pytorch_pfn_extras.nn.parallel.DistributedDataParallel method), 236 |
| LazyConv3d (class in pytorch_pfn_extras.nn), 123 | | load_state_dict() (pytorch_pfn_extras.profiler.ChromeTracer method), 276 |
| LazyConv3d (class in pytorch_pfn_extras.nn.modules.lazy_conv), 189 | | load_state_dict() (pytorch_pfn_extras.profiler.Tracer method), 280 |
| LazyInitializationMixin (class in pytorch_pfn_extras.nn.modules.lazy), 137 | | load_state_dict() (pytorch_pfn_extras.reporting.DictSummary method), 283 |
| LazyInitializationMixin (class in pytorch_pfn_extras.nn.modules.lazy_batchnorm), 164 | | load_state_dict() (pytorch_pfn_extras.reporting.Summary method), 287 |
| LazyInitializationMixin (class in pytorch_pfn_extras.nn.modules.lazy_conv), 192 | | load_state_dict() (pytorch_pfn_extras.training.Extension method), 304 |
| LazyInitializationMixin (class in pytorch_pfn_extras.nn.modules.lazy_linear), 211 | | load_state_dict() (pytorch_pfn_extras.training.extension.Extension method), 320 |
| LazyLinear (class in pytorch_pfn_extras.nn), 126 | | load_state_dict() (pytorch_pfn_extras.training.extension.ExtensionEntry method), 321 |
| LazyLinear (class in pytorch_pfn_extras.nn.modules.lazy_linear), 212 | | load_state_dict() (pytorch_pfn_extras.training.ExtensionEntry method), 321 |
| lchmod() (pytorch_pfn_extras.onnx.load.Path method), 261 | | load_state_dict() (pytorch_pfn_extras.training.ExtensionEntry method), 321 |
| lchmod() (pytorch_pfn_extras.onnx.unstrip_tensor.Path method), 268 | | load_state_dict() (pytorch_pfn_extras.training.ExtensionEntry method), 321 |
| line_buffering (pytorch_pfn_extras.training.extensions.util.TextIO property), 428 | | load_state_dict() (pytorch_pfn_extras.training.ExtensionEntry method), 321 |

| | | |
|---|---|----------|
| method), 305 | method), 435 | |
| load_state_dict() | (py- load_state_dict() | (py- |
| torch_pfn_extras.training.extensions.accumulate.average_accumulate_base | torch_pfn_extras.training.extensions.LogReport | |
| method), 385 | method), 344 | |
| load_state_dict() | (py- load_state_dict() | (py- |
| torch_pfn_extras.training.extensions.accumulate.average_accumulate_base | torch_pfn_extras.training.extensions.Lr_scheduler.LRScheduler | |
| method), 387 | method), 440 | |
| load_state_dict() | (py- load_state_dict() | (py- |
| torch_pfn_extras.training.extensions.accumulate.average_accumulate_base | torch_pfn_extras.training.extensions.Lr_scheduler.ReduceLROnPlateau | |
| method), 388 | method), 442 | |
| load_state_dict() | (py- load_state_dict() | (py- |
| torch_pfn_extras.training.extensions.accumulate.max_accumulate_base | torch_pfn_extras.training.extensions.LRScheduler | |
| method), 390 | method), 342 | |
| load_state_dict() | (py- load_state_dict() | (py- |
| torch_pfn_extras.training.extensions.accumulate.max_accumulate_base | torch_pfn_extras.training.extensions.micro_average.MicroAverage | |
| method), 392 | method), 447 | |
| load_state_dict() | (py- load_state_dict() | (py- |
| torch_pfn_extras.training.extensions.accumulate.max_accumulate_base | torch_pfn_extras.training.extensions.MicroAverage | |
| method), 393 | method), 349 | |
| load_state_dict() | (py- load_state_dict() | (py- |
| torch_pfn_extras.training.extensions.accumulate.min_accumulate_base | torch_pfn_extras.training.extensions.plot_report.PlotReport | |
| method), 396 | method), 456 | |
| load_state_dict() | (py- load_state_dict() | (py- |
| torch_pfn_extras.training.extensions.accumulate.min_accumulate_base | torch_pfn_extras.training.extensions.PlotReport | |
| method), 398 | method), 357 | |
| load_state_dict() | (py- load_state_dict() | (py- |
| torch_pfn_extras.training.extensions.accumulate.min_accumulate_base | torch_pfn_extras.training.extensions.print_report.PrintReport | |
| method), 399 | method), 465 | |
| load_state_dict() | (py- load_state_dict() | (py- |
| torch_pfn_extras.training.extensions.accumulate.standard_deviation_accumulate_base | torch_pfn_extras.training.extensions.PrintReport | |
| method), 401 | method), 359 | |
| load_state_dict() | (py- load_state_dict() | (py- |
| torch_pfn_extras.training.extensions.accumulate.standard_deviation_accumulate_base | torch_pfn_extras.training.extensions.ProfileReport | |
| method), 403 | method), 471 | |
| load_state_dict() | (py- load_state_dict() | (py- |
| torch_pfn_extras.training.extensions.accumulate.standard_deviation_accumulate_base | torch_pfn_extras.training.extensions.ProfileReport | |
| method), 404 | method), 361 | |
| load_state_dict() | (py- load_state_dict() | (py- |
| torch_pfn_extras.training.extensions.accumulate.unbiased_standard_deviation_accumulate_base | torch_pfn_extras.training.extensions.TimelineTracer | |
| method), 407 | method), 501 | |
| load_state_dict() | (py- load_state_dict() | (py- |
| torch_pfn_extras.training.extensions.accumulate.unbiased_standard_deviation_accumulate_base | torch_pfn_extras.training.extensions.Tracer | |
| method), 408 | method), 503 | |
| load_state_dict() | (py- load_state_dict() | (py- |
| torch_pfn_extras.training.extensions.accumulate.unbiased_standard_deviation_accumulate_base | torch_pfn_extras.training.extensions.LombScatterTheoremDeviationSum | |
| method), 410 | method), 372 | |
| load_state_dict() | (py- load_state_dict() | (py- |
| torch_pfn_extras.training.extensions.best_value.BestValue | torch_pfn_extras.training.IgniteExtensionsManager | |
| method), 412 | method), 312 | |
| load_state_dict() | (py- load_state_dict() | (py- |
| torch_pfn_extras.training.extensions.BestValue | torch_pfn_extras.training.StateObjectProtocol | |
| method), 331 | method), 312 | |
| load_state_dict() | (py- load_state_dict() | (py- |
| torch_pfn_extras.training.extensions.log_report.LogReport | torch_pfn_extras.training.Trainer | method), |

[315](#)
[load_state_dict\(\)](#) (pytorch_pfn_extras.training.trigger.Trigger method), [523](#)
[load_state_dict\(\)](#) (pytorch_pfn_extras.training.triggers.BestValueTrigger method), [525](#)
[load_state_dict\(\)](#) (pytorch_pfn_extras.training.triggers.function_trigger.Trigger method), [542](#)
[load_state_dict\(\)](#) (pytorch_pfn_extras.training.triggers.FunctionTrigger method), [528](#)
[load_state_dict\(\)](#) (pytorch_pfn_extras.training.triggers.minmax_value_trigger.Trigger method), [550](#)
[load_state_dict\(\)](#) (pytorch_pfn_extras.training.triggers.once_trigger.OnceTrigger method), [556](#)
[load_state_dict\(\)](#) (pytorch_pfn_extras.training.triggers.OnceTrigger method), [533](#)
[load_state_dict\(\)](#) (pytorch_pfn_extras.training.triggers.time_trigger.TimeTrigger method), [559](#)
[load_state_dict\(\)](#) (pytorch_pfn_extras.training.triggers.TimeTrigger method), [534](#)
[log](#) (pytorch_pfn_extras.training.extensions.log_report.LogReport property), [436](#)
[log](#) (pytorch_pfn_extras.training.extensions.LogReport property), [345](#)
[Logic](#) (class in pytorch_pfn_extras.handler), [102](#)
[LogReport](#) (class in pytorch_pfn_extras.training.extensions), [342](#)
[LogReport](#) (class in pytorch_pfn_extras.training.extensions.log_report), [433](#)
[LogWriterSaveFunc](#) (class in pytorch_pfn_extras.training.extensions.log_report), [436](#)
[LRScheduler](#) (class in pytorch_pfn_extras.training.extensions), [340](#)
[LRScheduler](#) (class in pytorch_pfn_extras.training.extensions.lr_scheduler), [439](#)
[lstat\(\)](#) (pytorch_pfn_extras.onnx.load.Path method), [261](#)
[lstat\(\)](#) (pytorch_pfn_extras.onnx.unstrip_tensor.Path method), [268](#)
M
[make_extension\(\)](#) (in module pytorch_pfn_extras.training), [298](#)
[make_extension\(\)](#) (in module pytorch_pfn_extras.training.extension), [317](#)
[make_statistics\(\)](#) (pytorch_pfn_extras.reporting.DictSummary method), [283](#)
[make_statistics\(\)](#) (pytorch_pfn_extras.reporting.Summary method), [287](#)
[manager](#) (pytorch_pfn_extras.engine.Trainer property), [84](#)
[manager](#) (pytorch_pfn_extras.training.Trainer property), [315](#)
[ManualScheduleTrigger](#) (class in pytorch_pfn_extras.training.triggers), [530](#)
[ManualScheduleTrigger](#) (class in pytorch_pfn_extras.training.triggers.manual_schedule_trigger), [547](#)
[map\(\)](#) (pytorch_pfn_extras.runtime.BaseRuntime method), [290](#)
[map\(\)](#) (pytorch_pfn_extras.runtime.PyTorchRuntime method), [294](#)
[materialize\(\)](#) (pytorch_pfn_extras.nn.modules.lazy.UninitializedParameter method), [155](#)
[materialize\(\)](#) (pytorch_pfn_extras.nn.modules.lazy_batchnorm.UninitializedParameter method), [182](#)
[materialize\(\)](#) (pytorch_pfn_extras.nn.modules.lazy_conv.UninitializedParameter method), [210](#)
[materialize\(\)](#) (pytorch_pfn_extras.nn.modules.lazy_linear.UninitializedParameter method), [232](#)
[matplotlib_savefun\(\)](#) (in module pytorch_pfn_extras.training.extensions.plot_report), [452](#)
[matplotlib_savefun\(\)](#) (in module pytorch_pfn_extras.training.extensions.variable_statistics_plot), [511](#)
[MaxAccumulate](#) (class in pytorch_pfn_extras.training.extensions), [345](#)
[MaxAccumulate](#) (class in pytorch_pfn_extras.training.extensions.accumulate), [379](#)
[MaxAccumulate](#) (class in pytorch_pfn_extras.training.extensions.accumulate.max_accumulate), [390](#)
[MaxSummary](#) (class in pytorch_pfn_extras.training.extensions.accumulate.max_accumulate), [391](#)
[MaxValue](#) (class in pytorch_pfn_extras.training.extensions), [346](#)
[MaxValue](#) (class in pytorch_pfn_extras.training.extensions.best_value), [414](#)
[MaxValueTrigger](#) (class in pytorch_pfn_extras.training.triggers), [531](#)

MaxValueTrigger (class in pytorch_pfn_extras.training.triggers), 532
 torch_pfn_extras.training.triggers.minmax_value_trigger (class in pytorch_pfn_extras.training.triggers.minmax_value_trigger), 552
 may_fire() (pytorch_pfn_extras.training.trigger.IntervalTrigger method), 523
 may_fire() (pytorch_pfn_extras.training.trigger.Trigger method), 523
 may_fire() (pytorch_pfn_extras.training.triggers.BestValueTrigger method), 268
 may_fire() (pytorch_pfn_extras.training.triggers.early_stopping_trigger.EarlyStoppingTrigger method), 537
 may_fire() (pytorch_pfn_extras.training.triggers.EarlyStoppingTrigger property), 65
 may_fire() (pytorch_pfn_extras.training.triggers.function_trigger.FunctionTrigger method), 542
 may_fire() (pytorch_pfn_extras.training.triggers.FunctionTrigger property), 59
 may_fire() (pytorch_pfn_extras.training.triggers.interval_trigger.IntervalTrigger method), 545
 may_fire() (pytorch_pfn_extras.training.triggers.IntervalTrigger method), 530
 may_fire() (pytorch_pfn_extras.training.triggers.manual_trigger.ManualScheduleTrigger method), 548
 may_fire() (pytorch_pfn_extras.training.triggers.ManualScheduleTrigger property), 323
 may_fire() (pytorch_pfn_extras.training.triggers.minmax_value_trigger.BestValueTrigger method), 550
 may_fire() (pytorch_pfn_extras.training.triggers.once_trigger.OnceTrigger method), 556
 may_fire() (pytorch_pfn_extras.training.triggers.OnceTrigger property), 533
 meta (pytorch_pfn_extras.ops.OpDesc attribute), 270
 MicroAverage (class in pytorch_pfn_extras.training.extensions), 347
 MicroAverage (class in pytorch_pfn_extras.training.extensions.micro_average), 445
 MinAccumulate (class in pytorch_pfn_extras.training.extensions), 349
 MinAccumulate (class in pytorch_pfn_extras.training.extensions.accumulate), 380
 MinAccumulate (class in pytorch_pfn_extras.training.extensions.accumulate.min_accumulate), 396
 MinSummary (class in pytorch_pfn_extras.training.extensions.accumulate.min_summary), 397
 MinValue (class in pytorch_pfn_extras.training.extensions), 350
 MinValue (class in pytorch_pfn_extras.training.extensions.best_value), 415
 MinValueTrigger (class in pytorch_pfn_extras.training.triggers), 532
 mkdir() (pytorch_pfn_extras.onnx.load.Path method), 262
 mkdir() (pytorch_pfn_extras.onnx.unstrip_tensor.Path method), 268
 mode (pytorch_pfn_extras.dataset.tabular.delegate_dataset.DelegateDataset property), 58
 mode (pytorch_pfn_extras.dataset.tabular.DelegateDataset property), 65
 mode (pytorch_pfn_extras.dataset.tabular.tabular_dataset.TabularDataset property), 59
 mode (pytorch_pfn_extras.dataset.TabularDataset property), 59
 mode (pytorch_pfn_extras.onnx.load.IO property), 257
 mode (pytorch_pfn_extras.training.extensions.print_report.IO property), 462
 ModelComarar (class in pytorch_pfn_extras.utils.comparer), 564
 models (pytorch_pfn_extras.training.extension.ExtensionsManagerProtocol property), 84
 models (pytorch_pfn_extras.training.extension.ExtensionsManagerProtocol property), 323
 models (pytorch_pfn_extras.training.extensions.best_value.ExtensionsManagerProtocol property), 414
 models (pytorch_pfn_extras.training.extensions.evaluator.ExtensionsManagerProtocol property), 424
 models (pytorch_pfn_extras.training.extensions.fail_on_non_number.ExtensionsManagerProtocol property), 430
 models (pytorch_pfn_extras.training.extensions.log_report.ExtensionsManagerProtocol property), 433
 models (pytorch_pfn_extras.training.extensions.lr_scheduler.ExtensionsManagerProtocol property), 438
 models (pytorch_pfn_extras.training.extensions.micro_average.ExtensionsManagerProtocol property), 444
 models (pytorch_pfn_extras.training.extensions.parameter_statistics.ExtensionsManagerProtocol property), 448
 models (pytorch_pfn_extras.training.extensions.plot_report.ExtensionsManagerProtocol property), 453
 models (pytorch_pfn_extras.training.extensions.print_report.ExtensionsManagerProtocol property), 459
 models (pytorch_pfn_extras.training.extensions.profile_report.ExtensionsManagerProtocol property), 467
 models (pytorch_pfn_extras.training.extensions.progress_bar.ExtensionsManagerProtocol property), 473
 models (pytorch_pfn_extras.training.extensions.slack.ExtensionsManagerProtocol property), 477
 models (pytorch_pfn_extras.training.extensions.timeline_trace.ExtensionsManagerProtocol property), 499
 models (pytorch_pfn_extras.training.extensions.util.ExtensionsManagerProtocol property), 505
 models (pytorch_pfn_extras.training.extensions.value_observation.ExtensionsManagerProtocol property), 510

[models \(pytorch_pfn_extras.training.extensions.variable_statistics_extension.Parallel, 232\)](#)
[property\), 513](#)
[models \(pytorch_pfn_extras.training.ExtensionsManagerProtocol, 238\)](#)
[property\), 309](#)
[models \(pytorch_pfn_extras.training.Trainer, 315\)](#)
[property\), 315](#)
[models \(pytorch_pfn_extras.training.triggers.early_stopping_trigger.ExtensionsManagerProtocol, 238\)](#)
[property\), 539](#)
[models \(pytorch_pfn_extras.training.triggers.function_trigger.ExtensionsManagerProtocol, 238\)](#)
[property\), 540](#)
[models \(pytorch_pfn_extras.training.triggers.interval_trigger.ExtensionsManagerProtocol, 238\)](#)
[property\), 544](#)
[models \(pytorch_pfn_extras.training.triggers.manual_schedule_trigger.ExtensionsManagerProtocol, 238\)](#)
[property\), 547](#)
[models \(pytorch_pfn_extras.training.triggers.minmax_value_trigger.ExtensionsManagerProtocol, 238\)](#)
[property\), 551](#)
[models \(pytorch_pfn_extras.training.triggers.once_trigger.ExtensionsManagerProtocol, 238\)](#)
[property\), 555](#)
[models \(pytorch_pfn_extras.training.triggers.time_trigger.ExtensionsManagerProtocol, 238\)](#)
[property\), 558](#)
[module](#)
[pytorch_pfn_extras, 37](#)
[pytorch_pfn_extras.config, 42](#)
[pytorch_pfn_extras.config_types, 43](#)
[pytorch_pfn_extras.cuda, 44](#)
[pytorch_pfn_extras.dataloaders, 45](#)
[pytorch_pfn_extras.dataloaders.data_loader, 48](#)
[pytorch_pfn_extras.dataloaders.utils, 55](#)
[pytorch_pfn_extras.dataset, 56](#)
[pytorch_pfn_extras.dataset.shared_dataset, 61](#)
[pytorch_pfn_extras.dataset.tabular, 63](#)
[pytorch_pfn_extras.dataset.tabular.delegate_dataset, 66](#)
[pytorch_pfn_extras.dataset.tabular.tabular_dataset, 68](#)
[pytorch_pfn_extras.distributed, 73](#)
[pytorch_pfn_extras.engine, 75](#)
[pytorch_pfn_extras.handler, 85](#)
[pytorch_pfn_extras.logging, 105](#)
[pytorch_pfn_extras.nn, 105](#)
[pytorch_pfn_extras.nn.modules, 129](#)
[pytorch_pfn_extras.nn.modules.ensure_shape, 129](#)
[pytorch_pfn_extras.nn.modules.extended_sequential, 133](#)
[pytorch_pfn_extras.nn.modules.lazy, 136](#)
[pytorch_pfn_extras.nn.modules.lazy_batchnorm, 156](#)
[pytorch_pfn_extras.nn.modules.lazy_conv, 183](#)
[pytorch_pfn_extras.nn.modules.lazy_linear, 210](#)
[pytorch_pfn_extras.nn.parallel, 232](#)
[pytorch_pfn_extras.nn.parallel.distributed, 238](#)
[pytorch_pfn_extras.onnx, 249](#)
[pytorch_pfn_extras.onnx.load, 254](#)
[pytorch_pfn_extras.onnx.pfto_exporter, 260](#)
[pytorch_pfn_extras.onnx.strip_large_tensor, 263](#)
[pytorch_pfn_extras.onnx.symbolic_registry, 264](#)
[pytorch_pfn_extras.onnx.unstrip_tensor, 265](#)
[pytorch_pfn_extras.ops, 269](#)
[pytorch_pfn_extras.profiler, 271](#)
[pytorch_pfn_extras.reporting, 280](#)
[pytorch_pfn_extras.runtime, 287](#)
[pytorch_pfn_extras.testing, 297](#)
[pytorch_pfn_extras.torchscript, 297](#)
[pytorch_pfn_extras.training, 298](#)
[pytorch_pfn_extras.training.extension, 317](#)
[pytorch_pfn_extras.training.extensions, 323](#)
[pytorch_pfn_extras.training.extensions.accumulate, 378](#)
[pytorch_pfn_extras.training.extensions.accumulate.average, 383](#)
[pytorch_pfn_extras.training.extensions.accumulate.max, 388](#)
[pytorch_pfn_extras.training.extensions.accumulate.min, 394](#)
[pytorch_pfn_extras.training.extensions.accumulate.stan, 409](#)
[pytorch_pfn_extras.training.extensions.accumulate.unbi, 405](#)
[pytorch_pfn_extras.training.extensions.best_value, 410](#)
[pytorch_pfn_extras.training.extensions.evaluator, 417](#)
[pytorch_pfn_extras.training.extensions.fail_on_non_num, 428](#)
[pytorch_pfn_extras.training.extensions.log_report, 431](#)
[pytorch_pfn_extras.training.extensions.lr_scheduler, 437](#)
[pytorch_pfn_extras.training.extensions.micro_average, 443](#)
[pytorch_pfn_extras.training.extensions.parameter_stat, 447](#)
[pytorch_pfn_extras.training.extensions.plot_report, 452](#)
[pytorch_pfn_extras.training.extensions.print_report, 457](#)

pytorch_pfn_extras.training.extensions.profile_report, 295
 465 move_to_end() (pytorch_pfn_extras.nn.parallel.distributed.OrderedDict
 pytorch_pfn_extras.training.extensions.progress_bar, method), 246
 472 move_to_end() (pytorch_pfn_extras.training.extensions.profile_report.
 pytorch_pfn_extras.training.extensions.slack, method), 469
 476 multiprocessing_context (py-
 pytorch_pfn_extras.training.extensions.snapshot_writer, pytorch_pfn_extras.dataLoaders.DataLoader
 483 property), 48
 pytorch_pfn_extras.training.extensions.timeline, multiprocessing_context (py-
 497 torch_pfn_extras.dataLoaders.dataLoader.DataLoader
 pytorch_pfn_extras.training.extensions.util, property), 54
 503
 pytorch_pfn_extras.training.extensions.value_observation,
 508 name (pytorch_pfn_extras.onnx.load.IO property), 257
 pytorch_pfn_extras.training.extensions.variable_statistics_plot,
 511 name (pytorch_pfn_extras.training.Extension attribute),
 304
 pytorch_pfn_extras.training.manager, 519 name (pytorch_pfn_extras.training.extension.Extension
 pytorch_pfn_extras.training.metrics, 520 attribute), 320
 pytorch_pfn_extras.training.trigger, 521 name (pytorch_pfn_extras.training.extensions.print_report.IO
 pytorch_pfn_extras.training.triggers, 524 property), 462
 pytorch_pfn_extras.training.triggers.early_stopping_trigger,
 535 name (pytorch_pfn_extras.training.extensions.print_report.IO
 pytorch_pfn_extras.training.triggers.function_trigger,
 539 needs_model_state (py-
 304
 pytorch_pfn_extras.training.triggers.interval_trigger,
 542 needs_model_state (py-
 torch_pfn_extras.training.extension.Extension
 pytorch_pfn_extras.training.triggers.manual_schedule_trigger,
 545 attribute), 320
 needs_model_state (py-
 pytorch_pfn_extras.training.triggers.minmax_value_trigger,
 548 torch_pfn_extras.training.extensions.fail_on_non_number.FailOn
 attribute), 431
 pytorch_pfn_extras.training.triggers.once_trigger,
 553 needs_model_state (py-
 torch_pfn_extras.training.extensions.FailOnNonNumber
 pytorch_pfn_extras.training.triggers.time_trigger,
 557 attribute), 338
 pytorch_pfn_extras.utils, 560 newlines (pytorch_pfn_extras.training.extensions.evaluator.TextIO
 pytorch_pfn_extras.utils.checkpoint, 560 property), 428
 pytorch_pfn_extras.utils.comparer, 560 newlines (pytorch_pfn_extras.training.extensions.util.TextIO
 pytorch_pfn_extras.writing, 566 property), 508
 no_grad() (in module pytorch_pfn_extras.onnx), 253
 momentum (pytorch_pfn_extras.nn.modules.lazy_batchnorm.LazyBatchNorm1d,
 attribute), 158 no_sync() (pytorch_pfn_extras.nn.parallel.distributed.DistributedDataParallel
 method), 243
 momentum (pytorch_pfn_extras.nn.modules.lazy_batchnorm.LazyBatchNorm2d,
 attribute), 161 no_sync() (pytorch_pfn_extras.nn.parallel.DistributedDataParallel
 method), 236
 momentum (pytorch_pfn_extras.nn.modules.lazy_batchnorm.LazyBatchNorm3d,
 attribute), 164 num_batches_tracked (py-
 torch_pfn_extras.nn.modules.lazy_batchnorm.LazyBatchNorm1d
 move_cursor_up() (py-
 attribute), 158
 torch_pfn_extras.training.extensions.util.ProgressBar, num_batches_tracked (py-
 method), 506 torch_pfn_extras.nn.modules.lazy_batchnorm.LazyBatchNorm2d
 move_module() (pytorch_pfn_extras.runtime.BaseRuntime
 attribute), 161
 method), 290 num_batches_tracked (py-
 move_module() (pytorch_pfn_extras.runtime.PyTorchRuntime
 attribute), 164
 method), 295 torch_pfn_extras.nn.modules.lazy_batchnorm.LazyBatchNorm3d
 move_tensor() (pytorch_pfn_extras.runtime.BaseRuntime
 attribute), 158
 method), 290 num_features (pytorch_pfn_extras.nn.modules.lazy_batchnorm.LazyBatch
 move_tensor() (pytorch_pfn_extras.runtime.PyTorchRuntime
 attribute), 158

num_features (pytorch_pfn_extras.nn.modules.lazy_batch_normalization.
 attribute), 161
 num_features (pytorch_pfn_extras.nn.modules.lazy_batch_normalization.
 attribute), 164
 num_workers (pytorch_pfn_extras.dataloaders.DataLoader
 attribute), 48
 num_workers (pytorch_pfn_extras.dataloaders.data_loader.
 attribute), 54
O
 observation (pytorch_pfn_extras.reporting.Reporter
 attribute), 284
 observation (pytorch_pfn_extras.training.extension.ExtensionsManagerProtocol
 property), 323
 observation (pytorch_pfn_extras.training.extensions.best_value.ExtensionsManagerProtocol
 property), 414
 observation (pytorch_pfn_extras.training.extensions.evaluator.ExtensionsManagerProtocol
 property), 424
 observation (pytorch_pfn_extras.training.extensions.fail_on_non_number.ExtensionsManagerProtocol
 property), 430
 observation (pytorch_pfn_extras.training.extensions.log_report.ExtensionsManagerProtocol
 property), 433
 observation (pytorch_pfn_extras.training.extensions.lr_scheduler.ExtensionsManagerProtocol
 property), 438
 observation (pytorch_pfn_extras.training.extensions.micro_average.ExtensionsManagerProtocol
 property), 444
 observation (pytorch_pfn_extras.training.extensions.parameter_statistics.ExtensionsManagerProtocol
 property), 448
 observation (pytorch_pfn_extras.training.extensions.plot_report.ExtensionsManagerProtocol
 property), 453
 observation (pytorch_pfn_extras.training.extensions.print_report.ExtensionsManagerProtocol
 property), 459
 observation (pytorch_pfn_extras.training.extensions.profile_report.ExtensionsManagerProtocol
 property), 467
 observation (pytorch_pfn_extras.training.extensions.progress_bar.ExtensionsManagerProtocol
 property), 473
 observation (pytorch_pfn_extras.training.extensions.slack.ExtensionsManagerProtocol
 property), 477
 observation (pytorch_pfn_extras.training.extensions.timeline_trace.ExtensionsManagerProtocol
 property), 499
 observation (pytorch_pfn_extras.training.extensions.util.ExtensionsManagerProtocol
 property), 505
 observation (pytorch_pfn_extras.training.extensions.value_observation.ExtensionsManagerProtocol
 property), 510
 observation (pytorch_pfn_extras.training.extensions.variable_statistics.ExtensionsManagerProtocol
 property), 513
 observation (pytorch_pfn_extras.training.ExtensionsManagerProtocol
 property), 309
 observation (pytorch_pfn_extras.training.triggers.early_stopping_trigger.ExtensionsManagerProtocol
 property), 539
 observation (pytorch_pfn_extras.training.triggers.function_trigger.ExtensionsManagerProtocol
 property), 540
 observation (pytorch_pfn_extras.training.triggers.interval_trigger.ExtensionsManagerProtocol
 property), 544
 observation (pytorch_pfn_extras.training.triggers.manual_schedule_trigger.ExtensionsManagerProtocol
 property), 547
 observation (pytorch_pfn_extras.training.triggers.minmax_value_trigger.ExtensionsManagerProtocol
 property), 551
 observation (pytorch_pfn_extras.training.triggers.once_trigger.ExtensionsManagerProtocol
 property), 555
 observation (pytorch_pfn_extras.training.triggers.time_trigger.ExtensionsManagerProtocol
 property), 558
 observe_lr() (in module pytorch_pfn_extras.training.extensions), 323
 observe_lr() (in module pytorch_pfn_extras.training.extensions.value_observation), 508
 observe_value() (in module pytorch_pfn_extras.training.extensions), 324
 observe_value() (in module pytorch_pfn_extras.training.extensions.value_observation), 509
 on_error() (pytorch_pfn_extras.training.ExtensionManagerProtocol method), 304
 on_error() (pytorch_pfn_extras.training.ExtensionManagerProtocol method), 320
 OnceTrigger (class in pytorch_pfn_extras.training.triggers), 532
 OnceTrigger (class in pytorch_pfn_extras.training.triggers.once_trigger), 555
 op (pytorch_pfn_extras.ops.OpDesc attribute), 270
 OpDesc (class in pytorch_pfn_extras.ops), 270
 open() (pytorch_pfn_extras.onnx.load.Path method), 262
 open() (pytorch_pfn_extras.onnx.unstrip_tensor.Path method), 268
 optimizers (pytorch_pfn_extras.engine.Trainer property), 84
 optimizers (pytorch_pfn_extras.handler.CodeBlock attribute), 95
 optimizers (pytorch_pfn_extras.training.extension.ExtensionsManagerProtocol property), 323
 optimizers (pytorch_pfn_extras.training.extensions.best_value.ExtensionsManagerProtocol property), 414
 optimizers (pytorch_pfn_extras.training.extensions.evaluator.ExtensionsManagerProtocol property), 424
 optimizers (pytorch_pfn_extras.training.extensions.fail_on_non_number.ExtensionsManagerProtocol property), 430
 optimizers (pytorch_pfn_extras.training.extensions.log_report.ExtensionsManagerProtocol property), 433
 optimizers (pytorch_pfn_extras.training.extensions.lr_scheduler.ExtensionsManagerProtocol property), 438
 optimizers (pytorch_pfn_extras.training.extensions.micro_average.ExtensionsManagerProtocol property), 444
 optimizers (pytorch_pfn_extras.training.extensions.parameter_statistics.ExtensionsManagerProtocol property), 448
 optimizers (pytorch_pfn_extras.training.extensions.plot_report.ExtensionsManagerProtocol property), 453
 optimizers (pytorch_pfn_extras.training.extensions.print_report.ExtensionsManagerProtocol property), 459
 optimizers (pytorch_pfn_extras.training.extensions.profile_report.ExtensionsManagerProtocol property), 467
 optimizers (pytorch_pfn_extras.training.extensions.progress_bar.ExtensionsManagerProtocol property), 473
 optimizers (pytorch_pfn_extras.training.extensions.slack.ExtensionsManagerProtocol property), 477
 optimizers (pytorch_pfn_extras.training.extensions.timeline_trace.ExtensionsManagerProtocol property), 499
 optimizers (pytorch_pfn_extras.training.extensions.util.ExtensionsManagerProtocol property), 505
 optimizers (pytorch_pfn_extras.training.extensions.value_observation.ExtensionsManagerProtocol property), 510
 optimizers (pytorch_pfn_extras.training.extensions.variable_statistics.ExtensionsManagerProtocol property), 513
 optimizers (pytorch_pfn_extras.training.ExtensionsManagerProtocol property), 309
 optimizers (pytorch_pfn_extras.training.triggers.early_stopping_trigger.ExtensionsManagerProtocol property), 539
 optimizers (pytorch_pfn_extras.training.triggers.function_trigger.ExtensionsManagerProtocol property), 540
 optimizers (pytorch_pfn_extras.training.triggers.interval_trigger.ExtensionsManagerProtocol property), 544

property), 453

optimizers (pytorch_pfn_extras.training.extensions.print_report.TrainingExtensionsManagerProtocol), 438

property), 459

optimizers (pytorch_pfn_extras.training.extensions.profile_report.TrainingExtensionsManagerProtocol), 444

property), 467

optimizers (pytorch_pfn_extras.training.extensions.progress_bar.TrainingExtensionsManagerProtocol), 448

property), 473

optimizers (pytorch_pfn_extras.training.extensions.slack.TrainingExtensionsManagerProtocol), 453

property), 477

optimizers (pytorch_pfn_extras.training.extensions.timeline_report.TrainingExtensionsManagerProtocol), 459

property), 499

optimizers (pytorch_pfn_extras.training.extensions.util.ExtensionsManagerProtocol), 505

property), 510

optimizers (pytorch_pfn_extras.training.extensions.value_observation.TrainingExtensionsManagerProtocol), 477

property), 513

optimizers (pytorch_pfn_extras.training.extensions.variable_statistics.TrainingExtensionsManagerProtocol), 499

optimizers (pytorch_pfn_extras.training.ExtensionsManagerProtocol), 309

optimizers (pytorch_pfn_extras.training.Trainer property), 315

optimizers (pytorch_pfn_extras.training.triggers.early_stopping.TrainingExtensionsManagerProtocol), 539

optimizers (pytorch_pfn_extras.training.triggers.function_trigger.TrainingExtensionsManagerProtocol), 309

optimizers (pytorch_pfn_extras.training.triggers.interval_trigger.TrainingExtensionsManagerProtocol), 539

optimizers (pytorch_pfn_extras.training.triggers.manual_schedule.TrainingExtensionsManagerProtocol), 540

optimizers (pytorch_pfn_extras.training.triggers.minmax_value.TrainingExtensionsManagerProtocol), 544

optimizers (pytorch_pfn_extras.training.triggers.once_trigger.TrainingExtensionsManagerProtocol), 547

optimizers (pytorch_pfn_extras.training.triggers.time_trigger.TrainingExtensionsManagerProtocol), 551

optuna_types() (in module pytorch_pfn_extras.config_types), 44

OrderedDict (class in pytorch_pfn_extras.nn.parallel.distributed), 245

OrderedDict (class in pytorch_pfn_extras.training.extensions.profile_report), 468

out (pytorch_pfn_extras.training.extension.ExtensionsManagerProtocol), 323

out (pytorch_pfn_extras.training.extensions.best_value.ExtensionsManagerProtocol), 414

out (pytorch_pfn_extras.training.extensions.evaluator.ExtensionsManagerProtocol), 424

out (pytorch_pfn_extras.training.extensions.fail_on_non_number.ExtensionsManagerProtocol), 430

out (pytorch_pfn_extras.training.extensions.log_report.ExtensionsManagerProtocol), 433

out (pytorch_pfn_extras.training.extensions.lr_scheduler.ExtensionsManagerProtocol), 438

out (pytorch_pfn_extras.training.extensions.micro_average.ExtensionsManagerProtocol), 444

out (pytorch_pfn_extras.training.extensions.parameter_statistics.ExtensionsManagerProtocol), 448

out (pytorch_pfn_extras.training.extensions.plot_report.ExtensionsManagerProtocol), 453

out (pytorch_pfn_extras.training.extensions.profile_report.ExtensionsManagerProtocol), 459

out (pytorch_pfn_extras.training.extensions.progress_bar.ExtensionsManagerProtocol), 467

out (pytorch_pfn_extras.training.extensions.slack.ExtensionsManagerProtocol), 477

out (pytorch_pfn_extras.training.extensions.timeline_trace.ExtensionsManagerProtocol), 499

out (pytorch_pfn_extras.training.extensions.util.ExtensionsManagerProtocol), 505

out (pytorch_pfn_extras.training.extensions.value_observation.ExtensionsManagerProtocol), 510

out (pytorch_pfn_extras.training.extensions.variable_statistics_plot.ExtensionsManagerProtocol), 513

out (pytorch_pfn_extras.training.extensions.value_observation.ExtensionsManagerProtocol), 309

out (pytorch_pfn_extras.training.extensions.value_observation.ExtensionsManagerProtocol), 539

out (pytorch_pfn_extras.training.extensions.value_observation.ExtensionsManagerProtocol), 540

out (pytorch_pfn_extras.training.extensions.value_observation.ExtensionsManagerProtocol), 544

out (pytorch_pfn_extras.training.extensions.value_observation.ExtensionsManagerProtocol), 547

out (pytorch_pfn_extras.training.extensions.value_observation.ExtensionsManagerProtocol), 551

out (pytorch_pfn_extras.training.extensions.value_observation.ExtensionsManagerProtocol), 555

out (pytorch_pfn_extras.training.extensions.value_observation.ExtensionsManagerProtocol), 558

out_channels (pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv1d attribute), 185

out_channels (pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv2d attribute), 188

out_channels (pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv3d attribute), 191

out_channels (pytorch_pfn_extras.nn.modules.lazy_linear.LazyLinear attribute), 215

output_padding (pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv1d attribute), 186

output_padding (pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv2d attribute), 188

output_padding (pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv3d attribute), 191

OutputsCompendium (class in py-

`torch_pfn_extras.utils.comparer)`, 565
`overload()` (in module `pytorch_pfn_extras.reporting`), 281
`owner()` (`pytorch_pfn_extras.onnx.load.Path` method), 262
`owner()` (`pytorch_pfn_extras.onnx.unstrip_tensor.Path` method), 268
P
`padding` (`pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv1d` attribute), 186
`padding` (`pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv2d` attribute), 188
`padding` (`pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv3d` attribute), 191
`padding_mode` (`pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv1d` attribute), 186
`padding_mode` (`pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv2d` attribute), 188
`padding_mode` (`pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv3d` attribute), 191
`ParameterStatistics` (class in `pytorch_pfn_extras.training.extensions`), 352
`ParameterStatistics` (class in `pytorch_pfn_extras.training.extensions.parameter_statistics`), 449
`Path` (class in `pytorch_pfn_extras.onnx.load`), 258
`Path` (class in `pytorch_pfn_extras.onnx.unstrip_tensor`), 265
`percentile()` (in module `pytorch_pfn_extras.training.extensions.variable_statistics_plot`), 511
`pin_memory` (`pytorch_pfn_extras.dataloaders.DataLoader` attribute), 48
`pin_memory` (`pytorch_pfn_extras.dataloaders.dataloader.DataLoader` attribute), 54
`pin_memory_device` (`pytorch_pfn_extras.dataloaders.DataLoader` attribute), 48
`pin_memory_device` (`pytorch_pfn_extras.dataloaders.dataloader.DataLoader` attribute), 54
`PlotReport` (class in `pytorch_pfn_extras.training.extensions`), 354
`PlotReport` (class in `pytorch_pfn_extras.training.extensions.plot_report`), 454
`pop()` (`pytorch_pfn_extras.nn.parallel.distributed.OrderedDict` method), 246
`pop()` (`pytorch_pfn_extras.training.extensions.profile_report.OrderedDict` method), 469
`popitem()` (`pytorch_pfn_extras.nn.parallel.distributed.OrderedDict` method), 246
`popitem()` (`pytorch_pfn_extras.training.extensions.profile_report.OrderedDict` method), 469
`prefetch_factor` (`pytorch_pfn_extras.dataloaders.DataLoader` attribute), 48
`prefetch_factor` (`pytorch_pfn_extras.dataloaders.dataloader.DataLoader` attribute), 54
`PrintReport` (class in `pytorch_pfn_extras.training.extensions`), 357
`PrintReport` (class in `pytorch_pfn_extras.training.extensions.print_report`), 463
`PrintReportCLI` (in module `pytorch_pfn_extras.training.extensions`), 359
`priority` (`pytorch_pfn_extras.training.Extension` attribute), 302, 304
`priority` (`pytorch_pfn_extras.training.extension.Extension` attribute), 318, 320
`priority` (`pytorch_pfn_extras.training.extensions.accumulate.average_accumulate` attribute), 385
`priority` (`pytorch_pfn_extras.training.extensions.accumulate.max_accumulate` attribute), 390
`priority` (`pytorch_pfn_extras.training.extensions.accumulate.min_accumulate` attribute), 396
`priority` (`pytorch_pfn_extras.training.extensions.accumulate.standard_deviation` attribute), 401
`priority` (`pytorch_pfn_extras.training.extensions.accumulate.unbiased_standard_deviation` attribute), 407
`priority` (`pytorch_pfn_extras.training.extensions.Evaluator` attribute), 337
`priority` (`pytorch_pfn_extras.training.extensions.evaluator.Evaluator` attribute), 422
`priority` (`pytorch_pfn_extras.training.extensions.micro_average.MicroAverage` attribute), 447
`priority` (`pytorch_pfn_extras.training.extensions.MicroAverage` attribute), 349
`priority` (`pytorch_pfn_extras.training.extensions.parameter_statistics.ParameterStatistics` attribute), 451
`priority` (`pytorch_pfn_extras.training.extensions.ParameterStatistics` attribute), 354
`ProcessQueueWriter` (class in `pytorch_pfn_extras.training.extensions.snapshot_writers`), 484
`ProcessQueueWriter` (class in `pytorch_pfn_extras.writing`), 567
`ProcessWriter` (class in `pytorch_pfn_extras.training.extensions.snapshot_writers`), 485
`ProcessWriter` (class in `pytorch_pfn_extras.writing`), 568
`ProfileReport` (class in `pytorch_pfn_extras.training.extensions`), 359
`ProfileReport` (class in `pytorch_pfn_extras.training.extensions.print_report`), 463

| | |
|--|---|
| <i>torch_pfn_extras.training.extensions.profile_report</i> , 469 | <i>pytorch_pfn_extras.nn.modules.lazy_batchnorm</i> module, 156 |
| <i>ProgressBar</i> (class in <i>pytorch_pfn_extras.training.extensions</i>), 362 | <i>pytorch_pfn_extras.nn.modules.lazy_conv</i> module, 183 |
| <i>ProgressBar</i> (class in <i>pytorch_pfn_extras.training.extensions.progress_bar</i>), 474 | <i>pytorch_pfn_extras.nn.modules.lazy_linear</i> module, 210 |
| <i>ProgressBar</i> (class in <i>pytorch_pfn_extras.training.extensions.util</i>), 505 | <i>pytorch_pfn_extras.nn.parallel</i> module, 232 |
| <i>ProgressBarCLI</i> (in module <i>pytorch_pfn_extras.training.extensions</i>), 364 | <i>pytorch_pfn_extras.nn.parallel.distributed</i> module, 238 |
| <i>pytorch_pfn_extras</i> module, 37 | <i>pytorch_pfn_extras.onnx</i> module, 249 |
| <i>pytorch_pfn_extras.config</i> module, 42 | <i>pytorch_pfn_extras.onnx.load</i> module, 254 |
| <i>pytorch_pfn_extras.config_types</i> module, 43 | <i>pytorch_pfn_extras.onnx.pfto_exporter</i> module, 263 |
| <i>pytorch_pfn_extras.cuda</i> module, 44 | <i>pytorch_pfn_extras.onnx.strip_large_tensor</i> module, 263 |
| <i>pytorch_pfn_extras.dataloaders</i> module, 45 | <i>pytorch_pfn_extras.onnx.symbolic_registry</i> module, 264 |
| <i>pytorch_pfn_extras.dataloaders.data_loader</i> module, 48 | <i>pytorch_pfn_extras.onnx.unstrip_tensor</i> module, 264 |
| <i>pytorch_pfn_extras.dataloaders.utils</i> module, 55 | <i>pytorch_pfn_extras.ops</i> module, 269 |
| <i>pytorch_pfn_extras.dataset</i> module, 56 | <i>pytorch_pfn_extras.profiler</i> module, 271 |
| <i>pytorch_pfn_extras.dataset.shared_dataset</i> module, 61 | <i>pytorch_pfn_extras.reporting</i> module, 280 |
| <i>pytorch_pfn_extras.dataset.tabular</i> module, 63 | <i>pytorch_pfn_extras.runtime</i> module, 287 |
| <i>pytorch_pfn_extras.dataset.tabular.delegate_data_loader</i> module, 66 | <i>pytorch_pfn_extras.testing</i> module, 297 |
| <i>pytorch_pfn_extras.dataset.tabular.tabular_data_loader</i> module, 68 | <i>pytorch_pfn_extras.torchscript</i> module, 297 |
| <i>pytorch_pfn_extras.distributed</i> module, 73 | <i>pytorch_pfn_extras.training</i> module, 298 |
| <i>pytorch_pfn_extras.engine</i> module, 75 | <i>pytorch_pfn_extras.training.extension</i> module, 317 |
| <i>pytorch_pfn_extras.handler</i> module, 85 | <i>pytorch_pfn_extras.training.extensions</i> module, 323 |
| <i>pytorch_pfn_extras.logging</i> module, 105 | <i>pytorch_pfn_extras.training.extensions.accumulate</i> module, 378 |
| <i>pytorch_pfn_extras.nn</i> module, 105 | <i>pytorch_pfn_extras.training.extensions.accumulate.average</i> module, 383 |
| <i>pytorch_pfn_extras.nn.modules</i> module, 129 | <i>pytorch_pfn_extras.training.extensions.accumulate.max_accuracy</i> module, 388 |
| <i>pytorch_pfn_extras.nn.modules.ensure_shape</i> module, 129 | <i>pytorch_pfn_extras.training.extensions.accumulate.min_accuracy</i> module, 394 |
| <i>pytorch_pfn_extras.nn.modules.extended_sequential</i> module, 133 | <i>pytorch_pfn_extras.training.extensions.accumulate.standard_deviation</i> module, 399 |
| <i>pytorch_pfn_extras.nn.modules.lazy</i> module, 136 | <i>pytorch_pfn_extras.training.extensions.accumulate.unbiased_variance</i> module, 405 |
| | <i>pytorch_pfn_extras.training.extensions.best_value</i> module, 410 |

pytorch_pfn_extras.training.extensions.evaluator module, 417
 pytorch_pfn_extras.training.extensions.fail_on_non_number module, 428
 pytorch_pfn_extras.training.extensions.log_report module, 431
 pytorch_pfn_extras.training.extensions.lr_scheduler module, 437
 pytorch_pfn_extras.training.extensions.micro_average module, 443
 pytorch_pfn_extras.training.extensions.parameter_statistics module, 447
 pytorch_pfn_extras.training.extensions.plot_report module, 452
 pytorch_pfn_extras.training.extensions.print_report module, 457
 pytorch_pfn_extras.training.extensions.profile_report module, 465
 pytorch_pfn_extras.training.extensions.progress_bar module, 472
 pytorch_pfn_extras.training.extensions.slack module, 476
 pytorch_pfn_extras.training.extensions.snapshot_writers module, 483
 pytorch_pfn_extras.training.extensions.timeline_trace module, 497
 pytorch_pfn_extras.training.extensions.util module, 503
 pytorch_pfn_extras.training.extensions.value_observation module, 508
 pytorch_pfn_extras.training.extensions.variable_statistics_plot module, 511
 pytorch_pfn_extras.training.manager module, 519
 pytorch_pfn_extras.training.metrics module, 520
 pytorch_pfn_extras.training.trigger module, 521
 pytorch_pfn_extras.training.triggers module, 524
 pytorch_pfn_extras.training.triggers.early_stopping_trigger module, 535
 pytorch_pfn_extras.training.triggers.function_trigger module, 539
 pytorch_pfn_extras.training.triggers.interval_trigger module, 542
 pytorch_pfn_extras.training.triggers.manual_schedule_trigger module, 545
 pytorch_pfn_extras.training.triggers.minmax_value_trigger module, 548
 pytorch_pfn_extras.training.triggers.once_trigger module, 553
 pytorch_pfn_extras.training.triggers.time_trigger module, 557
 pytorch_pfn_extras.utils module, 560
 pytorch_pfn_extras.utils.checkpoint module, 560
 pytorch_pfn_extras.utils.comparer module, 560
 pytorch_pfn_extras.writing module, 566
 PyTorchRuntime (class in pytorch_pfn_extras.runtime), 292
 QueueWriter (class in pytorch_pfn_extras.training.extensions.snapshot_writers), 486
 QueueWriter (class in pytorch_pfn_extras.writing), 570

Q

R

raw_models (pytorch_pfn_extras.training.extension.ExtensionsManagerProtocol property), 323
 raw_models (pytorch_pfn_extras.training.extensions.best_value.ExtensionsManagerProtocol property), 414
 raw_models (pytorch_pfn_extras.training.extensions.evaluator.ExtensionsManagerProtocol property), 424
 raw_models (pytorch_pfn_extras.training.extensions.fail_on_non_number.ExtensionsManagerProtocol property), 430
 raw_models (pytorch_pfn_extras.training.extensions.log_report.ExtensionsManagerProtocol property), 433
 raw_models (pytorch_pfn_extras.training.extensions.lr_scheduler.ExtensionsManagerProtocol property), 439
 raw_models (pytorch_pfn_extras.training.extensions.micro_average.ExtensionsManagerProtocol property), 445
 raw_models (pytorch_pfn_extras.training.extensions.parameter_statistics.ExtensionsManagerProtocol property), 449
 raw_models (pytorch_pfn_extras.training.extensions.plot_report.ExtensionsManagerProtocol property), 454
 raw_models (pytorch_pfn_extras.training.extensions.print_report.ExtensionsManagerProtocol property), 460
 raw_models (pytorch_pfn_extras.training.extensions.profile_report.ExtensionsManagerProtocol property), 468
 raw_models (pytorch_pfn_extras.training.extensions.progress_bar.ExtensionsManagerProtocol property), 474
 raw_models (pytorch_pfn_extras.training.extensions.slack.ExtensionsManagerProtocol property), 478
 raw_models (pytorch_pfn_extras.training.extensions.timeline_trace.ExtensionsManagerProtocol property), 499
 raw_models (pytorch_pfn_extras.training.extensions.util.ExtensionsManagerProtocol property), 505
 raw_models (pytorch_pfn_extras.training.extensions.value_observation.ExtensionsManagerProtocol property), 511
 raw_models (pytorch_pfn_extras.training.extensions.variable_statistics_plot.ExtensionsManagerProtocol property), 514
 raw_models (pytorch_pfn_extras.training.ExtensionsManagerProtocol property), 310

`raw_models` (pytorch_pfn_extras.training.triggers.early_stopping_trigger.ExtendedMinimizerProtocol.ops), 269
 property), 539
`raw_models` (pytorch_pfn_extras.training.triggers.function_trigger.ExtensionsManagerProtocol.parallel.distributed.DistributedDataParallel
 property), 541
 method), 244
`raw_models` (pytorch_pfn_extras.training.triggers.interval_trigger.ExtensionsManagerProtocol (py-
 property), 544
 torch_pfn_extras.nn.parallel.DistributedDataParallel
`raw_models` (pytorch_pfn_extras.training.triggers.manual_schedule_trigger.ExtensionsManagerProtocol
 property), 547
 register_statistics() (py-
`raw_models` (pytorch_pfn_extras.training.triggers.minmax_value_trigger.ExtensionsManagerProtocol.parameter_statistics.Parameter
 property), 552
 method), 451
`raw_models` (pytorch_pfn_extras.training.triggers.once_trigger.ExtensionsManagerProtocol (py-
 property), 555
 torch_pfn_extras.training.extensions.ParameterStatistics
`raw_models` (pytorch_pfn_extras.training.triggers.time_trigger.ExtensionsManagerProtocol
 property), 559
 rename() (pytorch_pfn_extras.onnx.load.Path method),
`read`() (pytorch_pfn_extras.onnx.load.IO method), 257
 262
`read`() (pytorch_pfn_extras.training.extensions.print_report_io.rename() (pytorch_pfn_extras.onnx.unstrip_tensor.Path
 method), 462
 method), 268
`read_bytes`() (pytorch_pfn_extras.onnx.load.Path repeat() (pytorch_pfn_extras.nn.ExtendedSequential
 method), 262
 method), 111
`read_bytes`() (pytorch_pfn_extras.onnx.unstrip_tensor.Path repeat() (pytorch_pfn_extras.nn.modules.extended_sequential.ExtendedS
 method), 268
 method), 135
`read_text`() (pytorch_pfn_extras.onnx.load.Path replace() (pytorch_pfn_extras.onnx.load.Path method),
 method), 262
 262
`read_text`() (pytorch_pfn_extras.onnx.unstrip_tensor.Path replace() (pytorch_pfn_extras.onnx.unstrip_tensor.Path
 method), 268
 method), 268
`readable`() (pytorch_pfn_extras.onnx.load.IO method), report() (in module pytorch_pfn_extras.reporting), 281
 257
 report() (pytorch_pfn_extras.profiler.TimeSummary
`readable`() (pytorch_pfn_extras.training.extensions.print_report_io.IO method), 277
 method), 462
 report() (pytorch_pfn_extras.reporting.Reporter
`readline`() (pytorch_pfn_extras.onnx.load.IO method),
 257
 method), 285
`readline`() (pytorch_pfn_extras.training.extensions.print_report_io.IO torch_pfn_extras.training.extensions.parameter_statistics.Parameter
 method), 462
 attribute), 451
`readlines`() (pytorch_pfn_extras.onnx.load.IO report_key_template (py-
 method), 257
 torch_pfn_extras.training.extensions.ParameterStatistics
`readlines`() (pytorch_pfn_extras.training.extensions.print_report_io.attribute), 354
 method), 462
`record`() (in module pytorch_pfn_extras.reporting), 282
 torch_pfn_extras.nn.parallel.distributed),
 Reporter (class in pytorch_pfn_extras.reporting), 284
 239
`record`() (in module pytorch_pfn_extras.profiler), 273
`record_function` (class in pytorch_pfn_extras.nn.parallel.distributed),
 reporter (pytorch_pfn_extras.training.extension.ExtensionsManagerProto
 247
 property), 323
`record_function`() (in module pytorch_pfn_extras.profiler), 273
 reporter (pytorch_pfn_extras.training.extensions.best_value.ExtensionsM
 property), 414
`record_iterable`() (in module pytorch_pfn_extras.profiler), 273
 reporter (pytorch_pfn_extras.training.extensions.evaluator.ExtensionsMa
 property), 424
`reduce`() (in module pytorch_pfn_extras.onnx.strip_large_tensor),
 reporter (pytorch_pfn_extras.training.extensions.fail_on_non_number.Ex
 property), 430
 reporter (pytorch_pfn_extras.training.extensions.log_report.ExtensionsM
 property), 433
 reporter (pytorch_pfn_extras.training.extensions.lr_scheduler.ExtensionsS
 property), 439
`ReduceLROnPlateau` (class in pytorch_pfn_extras.training.extensions.lr_scheduler),
 reporter (pytorch_pfn_extras.training.extensions.micro_average.Extension
 property), 445
 441
 reporter (pytorch_pfn_extras.training.extensions.parameter_statistics.Ext

[illegible]

| | |
|---|--|
| method), 519 | torch_pfn_extras.training.extensions.snapshot_writers), |
| save_plot_using_module() (pytorch_pfn_extras.training.extensions.VariableStatisticsSimpleWriter (class in pytorch_pfn_extras.writing), method), 376 | 572 |
| scope() (pytorch_pfn_extras.reporting.Reporter method), 286 | Slack (class in pytorch_pfn_extras.training.extensions), 364 |
| scoped_anchor() (in module pytorch_pfn_extras.onnx), 253 | Slack (class in pytorch_pfn_extras.training.extensions.slack), 478 |
| seek() (pytorch_pfn_extras.onnx.load.IO method), 257 | SlackWebhook (class in pytorch_pfn_extras.training.extensions), 367 |
| seek() (pytorch_pfn_extras.training.extensions.print_report.IO method), 462 | SlackWebhook (class in pytorch_pfn_extras.training.extensions.slack), 481 |
| seekable() (pytorch_pfn_extras.onnx.load.IO method), 257 | slice (pytorch_pfn_extras.dataset.tabular.tabular_dataset.TabularDataset property), 72 |
| seekable() (pytorch_pfn_extras.training.extensions.print_report.IO method), 462 | slice (pytorch_pfn_extras.dataset.TabularDataset property), 59 |
| set_evaluator_handlers() (pytorch_pfn_extras.training.extensions.evaluator.IgniteEvaluator method), 425 | snapshot() (in module pytorch_pfn_extras.training.extensions), 324 |
| set_evaluator_handlers() (pytorch_pfn_extras.training.extensions.IgniteEvaluator method), 340 | snapshot_object() (in module pytorch_pfn_extras.training.extensions), 326 |
| set_ignite_handlers() (pytorch_pfn_extras.training.IgniteExtensionsManager method), 312 | SnapshotMode (class in pytorch_pfn_extras.training.extensions), 369 |
| set_optimizer() (pytorch_pfn_extras.engine.Trainer method), 84 | StandardDeviationAccumulate (class in pytorch_pfn_extras.training.extensions), 369 |
| set_optimizer() (pytorch_pfn_extras.training.Trainer method), 316 | StandardDeviationAccumulate (class in pytorch_pfn_extras.training.extensions.accumulate), 381 |
| setdefault() (pytorch_pfn_extras.nn.parallel.distributed.StandardDeviationAccumulate method), 246 | StandardDeviationAccumulate (class in pytorch_pfn_extras.training.extensions.accumulate.standard_deviation), 402 |
| setdefault() (pytorch_pfn_extras.training.extensions.profile_report.OrderedDict method), 469 | StandardDeviationSummary (class in pytorch_pfn_extras.training.extensions.accumulate.standard_deviation), 403 |
| SHARDED (pytorch_pfn_extras.training.extensions.SnapshotMode attribute), 369 | StandardWriter (class in pytorch_pfn_extras.training.extensions.snapshot_writers), 490 |
| share_memory_() (pytorch_pfn_extras.nn.modules.lazy.UninitializedParameter method), 155 | StandardWriter (class in pytorch_pfn_extras.writing), 572 |
| share_memory_() (pytorch_pfn_extras.nn.modules.lazy_batchnorm.UninitializedParameter method), 183 | stat() (pytorch_pfn_extras.onnx.load.Path method), 262 |
| share_memory_() (pytorch_pfn_extras.nn.modules.lazy_conv.UninitializedParameter method), 210 | stat() (pytorch_pfn_extras.onnx.unstrip_tensor.Path method), 269 |
| share_memory_() (pytorch_pfn_extras.nn.modules.lazy_linear.UninitializedParameter method), 232 | state (pytorch_pfn_extras.handler.CodeBlock attribute), 138 |
| SharedDataset (class in pytorch_pfn_extras.dataset), 56 | state_dict() (pytorch_pfn_extras.engine.StateObjectProtocol method), 81 |
| SharedDataset (class in pytorch_pfn_extras.dataset.shared_dataset), 62 | state_dict() (pytorch_pfn_extras.engine.Trainer method), 84 |
| signature (pytorch_pfn_extras.ops.OpDesc attribute), 270 | state_dict() (pytorch_pfn_extras.handler.CodeBlock method), 95 |
| SimpleWriter (class in pytorch_pfn_extras.writing), 572 | state_dict() (pytorch_pfn_extras.nn.modules.lazy.LazyInitializationMixin method), 138 |
| SimpleWriter (class in pytorch_pfn_extras.writing), 572 | state_dict() (pytorch_pfn_extras.nn.modules.lazy_batchnorm.LazyInitializationMixin method), 138 |

method), 556
 state_dict() (pytorch_pfn_extras.training.triggers.OnceTrigger property), 511
 method), 534
 state_dict() (pytorch_pfn_extras.training.triggers.time_trigger.TimeTrigger property), 514
 method), 559
 state_dict() (pytorch_pfn_extras.training.triggers.TimeTrigger property), 310
 method), 535
 StateObjectProtocol (class in pytorch_pfn_extras.engine), 81
 StateObjectProtocol (class in pytorch_pfn_extras.training), 312
 Statistician (class in pytorch_pfn_extras.training.extensions.variable_statistics), 515
 step() (pytorch_pfn_extras.training.extensions.lr_scheduler.ReduceLROnPlateau method), 442
 step_by_value() (pytorch_pfn_extras.training.extensions.lr_scheduler.LRScheduler static method), 440
 step_by_value() (pytorch_pfn_extras.training.extensions.LRScheduler static method), 342
 stop_trigger (pytorch_pfn_extras.engine.Trainer property), 84
 stop_trigger (pytorch_pfn_extras.training.extension.ExtensionsManagerProtocol property), 323
 stop_trigger (pytorch_pfn_extras.training.extensions.best_value.ExtensionsManagerProtocol property), 414
 stop_trigger (pytorch_pfn_extras.training.extensions.evaluator.ExtensionsManagerProtocol property), 424
 stop_trigger (pytorch_pfn_extras.training.extensions.fail_summary.ExtensionsManagerProtocol property), 430
 stop_trigger (pytorch_pfn_extras.training.extensions.log_summary.ExtensionsManagerProtocol property), 433
 stop_trigger (pytorch_pfn_extras.training.extensions.lr_scheduler.ExtensionsManagerProtocol property), 439
 stop_trigger (pytorch_pfn_extras.training.extensions.micro_average.ExtensionsManagerProtocol property), 445
 stop_trigger (pytorch_pfn_extras.training.extensions.params_extensions.ExtensionsManagerProtocol property), 449
 stop_trigger (pytorch_pfn_extras.training.extensions.plot_report.ExtensionsManagerProtocol property), 454
 stop_trigger (pytorch_pfn_extras.training.extensions.print_report.ExtensionsManagerProtocol property), 460
 stop_trigger (pytorch_pfn_extras.training.extensions.progress_bar.ExtensionsManagerProtocol property), 474
 stop_trigger (pytorch_pfn_extras.training.extensions.slack.ExtensionsManagerProtocol property), 478
 stop_trigger (pytorch_pfn_extras.training.extensions.timeline_trace_extensions.ExtensionsManagerProtocol property), 499
 stop_trigger (pytorch_pfn_extras.training.extensions.util.ExtensionsManagerProtocol property), 505
 stop_trigger (pytorch_pfn_extras.training.extensions.value_observation.ExtensionsManagerProtocol property), 511
 stop_trigger (pytorch_pfn_extras.training.extensions.variable_statistics.ExtensionsManagerProtocol property), 514
 stop_trigger (pytorch_pfn_extras.training.ExtensionsManagerProtocol property), 310
 stop_trigger (pytorch_pfn_extras.training.Trainer property), 316
 stop_trigger (pytorch_pfn_extras.training.triggers.early_stopping_trigger.ExtensionsManagerProtocol property), 539
 stop_trigger (pytorch_pfn_extras.training.triggers.function_trigger.ExtensionsManagerProtocol property), 541
 stop_trigger (pytorch_pfn_extras.training.triggers.interval_trigger.ExtensionsManagerProtocol property), 544
 stop_trigger (pytorch_pfn_extras.training.triggers.manual_schedule_trigger.ExtensionsManagerProtocol property), 547
 stop_trigger (pytorch_pfn_extras.training.triggers.minmax_value_trigger.ExtensionsManagerProtocol property), 552
 stop_trigger (pytorch_pfn_extras.training.triggers.once_trigger.ExtensionsManagerProtocol property), 555
 stop_trigger (pytorch_pfn_extras.training.triggers.time_trigger.ExtensionsManagerProtocol property), 559
 stream() (in module pytorch_pfn_extras.cuda), 44
 stride (pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv1d attribute), 86
 stride (pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv2d attribute), 88
 stride (pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv3d attribute), 89
 Summary (class in pytorch_pfn_extras.reporting), 286
 summary (pytorch_pfn_extras.reporting.SummaryBase method), 277
 SummaryBase (class in pytorch_pfn_extras.training.extensions.accumulate.average_accumulate_extensions), 387
 SummaryBase (class in pytorch_pfn_extras.training.extensions.accumulate.max_accumulate_extensions), 393
 SummaryBase (class in pytorch_pfn_extras.training.extensions.accumulate.min_accumulate_extensions), 400
 SummaryBase (class in pytorch_pfn_extras.training.extensions.accumulate.standard_deviation_extensions), 404
 SummaryBase (class in pytorch_pfn_extras.training.extensions.accumulate.unbiased_standard_deviation_extensions), 406
 symlink_to() (pytorch_pfn_extras.onnx.load.Path method), 261
 symlink_to() (pytorch_pfn_extras.onnx.unstrip_tensor.Path method), 269
 synchronize() (pytorch_pfn_extras.profiler.TimeSummary method), 277

T

- `T_destination` (`pytorch_pfn_extras.nn.parallel.distributed.DistributedDataParallel` attribute), 242
- `T_destination` (`pytorch_pfn_extras.nn.parallel.DistributedDataParallel` attribute), 235
- `TabularDataset` (class in `pytorch_pfn_extras.dataset`), 56
- `TabularDataset` (class in `pytorch_pfn_extras.dataset.tabular.tabular_dataset`), 69
- `tell()` (`pytorch_pfn_extras.onnx.load.IO` method), 257
- `tell()` (`pytorch_pfn_extras.training.extensions.print_report` method), 462
- `TensorBoardWriter` (class in `pytorch_pfn_extras.training.extensions.snapshot_writers`), 491
- `TensorBoardWriter` (class in `pytorch_pfn_extras.writing`), 575
- `Text` (in module `pytorch_pfn_extras.onnx.load`), 263
- `TextIO` (class in `pytorch_pfn_extras.training.extensions.evaluator`), 426
- `TextIO` (class in `pytorch_pfn_extras.training.extensions.util`), 507
- `ThreadQueueWriter` (class in `pytorch_pfn_extras.training.extensions.snapshot_writers`), 493
- `ThreadQueueWriter` (class in `pytorch_pfn_extras.writing`), 576
- `ThreadWriter` (class in `pytorch_pfn_extras.training.extensions.snapshot_writers`), 494
- `ThreadWriter` (class in `pytorch_pfn_extras.writing`), 577
- `TimelineTrace` (class in `pytorch_pfn_extras.training.extensions`), 370
- `TimelineTrace` (class in `pytorch_pfn_extras.training.extensions.timeline_trace`), 499
- `timeout` (`pytorch_pfn_extras.dataloaders.DataLoader` attribute), 48
- `timeout` (`pytorch_pfn_extras.dataloaders.dataloader.DataLoader` attribute), 54
- `TimeSummary` (class in `pytorch_pfn_extras.profiler`), 276
- `TimeTrigger` (class in `pytorch_pfn_extras.training.triggers`), 534
- `TimeTrigger` (class in `pytorch_pfn_extras.training.triggers.time_trigger`), 559
- `to()` (in module `pytorch_pfn_extras`), 40
- `to_dataframe()` (`pytorch_pfn_extras.training.extensions.log_report.LogReport` method), 436
- `to_dataframe()` (`pytorch_pfn_extras.training.extensions.LogReport` method), 345
- `torch_autocast()` (in module `pytorch_pfn_extras.handler`), 85
- `touch()` (`pytorch_pfn_extras.onnx.load.Path` method), 262
- `touch()` (`pytorch_pfn_extras.onnx.unstrip_tensor.Path` method), 269
- `trace()` (`pytorch_pfn_extras.runtime.BaseRuntime` class method), 290
- `trace()` (`pytorch_pfn_extras.runtime.PyTorchRuntime` class method), 295
- `TraceableDataset` (class in `pytorch_pfn_extras.profiler`), 278
- `Tracer` (class in `pytorch_pfn_extras.profiler`), 278
- `Tracer` (class in `pytorch_pfn_extras.training.extensions.timeline_trace`), 502
- `track_running_stats` (`pytorch_pfn_extras.nn.modules.lazy_batchnorm.LazyBatchNorm1d` attribute), 158
- `track_running_stats` (`pytorch_pfn_extras.nn.modules.lazy_batchnorm.LazyBatchNorm2d` attribute), 161
- `track_running_stats` (`pytorch_pfn_extras.nn.modules.lazy_batchnorm.LazyBatchNorm3d` attribute), 164
- `train_cleanup()` (`pytorch_pfn_extras.handler.BaseHandler` method), 88
- `train_cleanup()` (`pytorch_pfn_extras.handler.Handler` method), 100
- `train_cleanup()` (`pytorch_pfn_extras.runtime.BaseRuntime` method), 291
- `train_cleanup()` (`pytorch_pfn_extras.runtime.PyTorchRuntime` method), 295
- `train_epoch_begin()` (`pytorch_pfn_extras.handler.BaseHandler` method), 88
- `train_epoch_begin()` (`pytorch_pfn_extras.handler.BaseLogic` method), 91
- `train_epoch_begin()` (`pytorch_pfn_extras.handler.CodeBlockLogic` method), 97
- `train_epoch_begin()` (`pytorch_pfn_extras.handler.Handler` method), 100
- `train_epoch_begin()` (`pytorch_pfn_extras.handler.Logic` method), 103
- `train_epoch_begin()` (`pytorch_pfn_extras.runtime.BaseRuntime` method), 291
- `train_epoch_begin()` (`pytorch_pfn_extras.runtime.PyTorchRuntime` method), 295

`method`), 295
`train_epoch_end()` (`pytorch_pfn_extras.handler.BaseHandler` `method`), 89
`train_epoch_end()` (`pytorch_pfn_extras.handler.BaseLogic` `method`), 91
`train_epoch_end()` (`pytorch_pfn_extras.handler.CodeBlockLogic` `method`), 97
`train_epoch_end()` (`pytorch_pfn_extras.handler.Handler` `method`), 100
`train_epoch_end()` (`pytorch_pfn_extras.handler.Logic` `method`), 103
`train_epoch_end()` (`pytorch_pfn_extras.runtime.BaseRuntime` `method`), 291
`train_epoch_end()` (`pytorch_pfn_extras.runtime.PyTorchRuntime` `method`), 296
`train_post_step()` (`pytorch_pfn_extras.handler.BaseHandler` `method`), 89
`train_post_step()` (`pytorch_pfn_extras.handler.Handler` `method`), 100
`train_post_step()` (`pytorch_pfn_extras.runtime.BaseRuntime` `method`), 291
`train_post_step()` (`pytorch_pfn_extras.runtime.PyTorchRuntime` `method`), 296
`train_pre_step()` (`pytorch_pfn_extras.runtime.BaseRuntime` `method`), 291
`train_pre_step()` (`pytorch_pfn_extras.runtime.PyTorchRuntime` `method`), 296
`train_setup()` (`pytorch_pfn_extras.handler.BaseHandler` `method`), 89
`train_setup()` (`pytorch_pfn_extras.handler.Handler` `method`), 101
`train_step()` (`pytorch_pfn_extras.handler.BaseHandler` `method`), 89
`train_step()` (`pytorch_pfn_extras.handler.BaseLogic` `method`), 91
`train_step()` (`pytorch_pfn_extras.handler.ClousureLogic` `method`), 93
`train_step()` (`pytorch_pfn_extras.handler.CodeBlockLogic` `method`), 97
`train_step()` (`pytorch_pfn_extras.handler.Handler` `method`), 101
`train_step()` (`pytorch_pfn_extras.handler.Logic` `method`), 103
`train_step_optimizers()` (`pytorch_pfn_extras.handler.BaseLogic` `method`), 92
`train_step_optimizers()` (`pytorch_pfn_extras.handler.ClousureLogic` `method`), 93
`train_step_optimizers()` (`pytorch_pfn_extras.handler.Logic` `method`), 104
`train_validation_begin()` (`pytorch_pfn_extras.handler.BaseHandler` `method`), 89
`train_validation_begin()` (`pytorch_pfn_extras.handler.BaseLogic` `method`), 92
`train_validation_begin()` (`pytorch_pfn_extras.handler.CodeBlockLogic` `method`), 97
`train_validation_begin()` (`pytorch_pfn_extras.handler.Handler` `method`), 101
`train_validation_begin()` (`pytorch_pfn_extras.handler.Logic` `method`), 104
`train_validation_begin()` (`pytorch_pfn_extras.runtime.BaseRuntime` `method`), 292
`train_validation_begin()` (`pytorch_pfn_extras.runtime.PyTorchRuntime` `method`), 296
`train_validation_end()` (`pytorch_pfn_extras.handler.BaseHandler` `method`), 90
`train_validation_end()` (`pytorch_pfn_extras.handler.BaseLogic` `method`), 92
`train_validation_end()` (`pytorch_pfn_extras.handler.CodeBlockLogic` `method`), 98
`train_validation_end()` (`pytorch_pfn_extras.handler.Handler` `method`), 101
`train_validation_end()` (`pytorch_pfn_extras.handler.Logic` `method`), 104
`train_validation_end()` (`pytorch_pfn_extras.runtime.BaseRuntime` `method`), 292
`train_validation_end()` (`pytorch_pfn_extras.runtime.PyTorchRuntime` `method`), 297
Trainer (class in `pytorch_pfn_extras.engine`), 81
Trainer (class in `pytorch_pfn_extras.training`), 313

[training \(pytorch_pfn_extras.nn.Ensure attribute\), 108](#)
[training \(pytorch_pfn_extras.nn.modules.ensure_shape.EnsureShape attribute\), 132](#)
[training \(pytorch_pfn_extras.nn.modules.extended_sequential.ExtendedSequential class in pytorch_pfn_extras.nn.modules.extended_sequential\), 135](#)
[training \(pytorch_pfn_extras.nn.modules.lazy_batchnorm.LazyBatchNorm1d attribute\), 158](#)
[training \(pytorch_pfn_extras.nn.modules.lazy_batchnorm.LazyBatchNorm2d attribute\), 161](#)
[training \(pytorch_pfn_extras.nn.modules.lazy_batchnorm.LazyBatchNorm3d attribute\), 164](#)
[training \(pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv1d attribute\), 186](#)
[training \(pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv2d attribute\), 188](#)
[training \(pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv3d attribute\), 191](#)
[training \(pytorch_pfn_extras.nn.modules.lazy_linear.LazyLinear attribute\), 215](#)
[training \(pytorch_pfn_extras.nn.parallel.distributed.DistributedDataParallel attribute\), 245](#)
[training \(pytorch_pfn_extras.nn.parallel.DistributedDataParallel attribute\), 237](#)
[transform\(\) \(pytorch_pfn_extras.dataset.tabular.tabular_dataset.TabularDataset method\), 72](#)
[transform\(\) \(pytorch_pfn_extras.dataset.TabularDataset method\), 59](#)
[transform_batch\(\) \(pytorch_pfn_extras.dataset.tabular.tabular_dataset.TabularDataset method\), 72](#)
[transform_batch\(\) \(pytorch_pfn_extras.dataset.TabularDataset method\), 60](#)
[transposed \(pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv1d attribute\), 186](#)
[transposed \(pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv2d attribute\), 188](#)
[transposed \(pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv3d attribute\), 191](#)
[Trigger \(class in pytorch_pfn_extras.training.trigger\), 523](#)
[trigger \(pytorch_pfn_extras.training.Extension attribute\), 302, 304](#)
[trigger \(pytorch_pfn_extras.training.extension.Extension attribute\), 318, 320](#)
[trigger \(pytorch_pfn_extras.training.extensions.Evaluator attribute\), 337](#)
[trigger \(pytorch_pfn_extras.training.extensions.evaluator.Evaluator attribute\), 422](#)
[trigger \(pytorch_pfn_extras.training.extensions.Slack attribute\), 366](#)
[trigger \(pytorch_pfn_extras.training.extensions.slack.Slack attribute\), 481](#)
[truncate\(\) \(pytorch_pfn_extras.onnx.load.IO method\), 258](#)
[truncate\(\) \(pytorch_pfn_extras.training.extensions.print_report.IO method\), 463](#)
[TypeVar \(class in pytorch_pfn_extras.nn.parallel.distributed\), 246](#)
[UnbiasedStandardDeviationAccumulate \(class in pytorch_pfn_extras.training.extensions\), 373](#)
[UnbiasedStandardDeviationAccumulate \(class in pytorch_pfn_extras.training.extensions.accumulate\), 382](#)
[UnbiasedStandardDeviationAccumulate \(class in pytorch_pfn_extras.training.extensions.accumulate.unbiased_standard_deviation\), 408](#)
[UnbiasedStandardDeviationSummary \(class in pytorch_pfn_extras.training.extensions.accumulate.unbiased_standard_deviation_summary\), 408](#)
[UninitializedParameter \(class in pytorch_pfn_extras.nn.modules.lazy\), 138](#)
[UninitializedParameter \(class in pytorch_pfn_extras.nn.modules.lazy_batchnorm\), 166](#)
[UninitializedParameter \(class in pytorch_pfn_extras.nn.modules.lazy_conv\), 193](#)
[UninitializedParameter \(class in pytorch_pfn_extras.nn.modules.lazy_linear\), 215](#)
[unlink\(\) \(pytorch_pfn_extras.onnx.load.Path method\), 262](#)
[unlink\(\) \(pytorch_pfn_extras.onnx.unstrip_tensor.Path method\), 269](#)
[unstrip\(\) \(in module pytorch_pfn_extras.onnx.unstrip_tensor\), 264](#)
[update\(\) \(pytorch_pfn_extras.nn.parallel.distributed.OrderedDict method\), 246](#)
[update\(\) \(pytorch_pfn_extras.training.extensions.profile_report.OrderedDict method\), 469](#)
[update\(\) \(pytorch_pfn_extras.training.extensions.util.ProgressBar method\), 506](#)
[update_parameters\(\) \(in module pytorch_pfn_extras.handler\), 85](#)
[update_speed\(\) \(pytorch_pfn_extras.training.extensions.util.ProgressBar method\), 506](#)
[update_via_args\(\) \(pytorch_pfn_extras.config.Config method\), 43](#)

[use_default_mempool_in_cupy\(\)](#) (in module `pytorch_pfn_extras.cuda`), 44
[use_torch_mempool_in_cupy\(\)](#) (in module `pytorch_pfn_extras.cuda`), 45

V

[values\(\)](#) (`pytorch_pfn_extras.nn.parallel.distributed.OrderedDict` method), 246
[values\(\)](#) (`pytorch_pfn_extras.training.extensions.profile_report.OrderedDict` method), 469
[Variable](#) (class in `pytorch_pfn_extras.nn.parallel.distributed`), 247
[VariableStatisticsPlot](#) (class in `pytorch_pfn_extras.training.extensions`), 374
[VariableStatisticsPlot](#) (class in `pytorch_pfn_extras.training.extensions.variable_statistics_plot`), 516

W

[weight](#) (`pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv1d` attribute), 186
[weight](#) (`pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv2d` attribute), 188
[weight](#) (`pytorch_pfn_extras.nn.modules.lazy_conv.LazyConv3d` attribute), 191
[weight](#) (`pytorch_pfn_extras.nn.modules.lazy_linear.LazyLinear` attribute), 215
[with_converter\(\)](#) (`pytorch_pfn_extras.dataset.tabular.tabular_dataset.TabularDataset` method), 73
[with_converter\(\)](#) (`pytorch_pfn_extras.dataset.TabularDataset` method), 60
[writable\(\)](#) (`pytorch_pfn_extras.onnx.load.IO` method), 258
[writable\(\)](#) (`pytorch_pfn_extras.training.extensions.print_report.IO` method), 463
[write\(\)](#) (`pytorch_pfn_extras.onnx.load.IO` method), 258
[write\(\)](#) (`pytorch_pfn_extras.training.extensions.print_report.IO` method), 463
[write_bytes\(\)](#) (`pytorch_pfn_extras.onnx.load.Path` method), 263
[write_bytes\(\)](#) (`pytorch_pfn_extras.onnx.unstrip_tensor.Path` method), 269
[write_text\(\)](#) (`pytorch_pfn_extras.onnx.load.Path` method), 263
[write_text\(\)](#) (`pytorch_pfn_extras.onnx.unstrip_tensor.Path` method), 269
[writelines\(\)](#) (`pytorch_pfn_extras.onnx.load.IO` method), 258
[writelines\(\)](#) (`pytorch_pfn_extras.training.extensions.print_report.IO` method), 463
[Writer](#) (class in `pytorch_pfn_extras.training.extensions.snapshot_writers`), 495
[Writer](#) (class in `pytorch_pfn_extras.writing`), 578
[writer](#) (`pytorch_pfn_extras.training.extension.ExtensionsManagerProtocol` property), 323
[writer](#) (`pytorch_pfn_extras.training.extensions.best_value.ExtensionsManagerProtocol` property), 414
[writer](#) (`pytorch_pfn_extras.training.extensions.evaluator.ExtensionsManagerProtocol` property), 424
[writer](#) (`pytorch_pfn_extras.training.extensions.fail_on_non_number.ExtensionsManagerProtocol` property), 430
[writer](#) (`pytorch_pfn_extras.training.extensions.log_report.ExtensionsManagerProtocol` property), 433
[writer](#) (`pytorch_pfn_extras.training.extensions.lr_scheduler.ExtensionsManagerProtocol` property), 439
[writer](#) (`pytorch_pfn_extras.training.extensions.micro_average.ExtensionsManagerProtocol` property), 445
[writer](#) (`pytorch_pfn_extras.training.extensions.parameter_statistics.ExtensionsManagerProtocol` property), 449
[writer](#) (`pytorch_pfn_extras.training.extensions.plot_report.ExtensionsManagerProtocol` property), 454
[writer](#) (`pytorch_pfn_extras.training.extensions.print_report.ExtensionsManagerProtocol` property), 460
[writer](#) (`pytorch_pfn_extras.training.extensions.profile_report.ExtensionsManagerProtocol` property), 468
[writer](#) (`pytorch_pfn_extras.training.extensions.progress_bar.ExtensionsManagerProtocol` property), 474
[writer](#) (`pytorch_pfn_extras.training.extensions.slack.ExtensionsManagerProtocol` property), 478
[writer](#) (`pytorch_pfn_extras.training.extensions.timeline_trace.ExtensionsManagerProtocol` property), 499
[writer](#) (`pytorch_pfn_extras.training.extensions.util.ExtensionsManagerProtocol` property), 505
[writer](#) (`pytorch_pfn_extras.training.extensions.value_observation.ExtensionsManagerProtocol` property), 511
[writer](#) (`pytorch_pfn_extras.training.extensions.variable_statistics_plot.ExtensionsManagerProtocol` property), 514
[writer](#) (`pytorch_pfn_extras.training.ExtensionsManagerProtocol` property), 310
[writer](#) (`pytorch_pfn_extras.training.triggers.early_stopping_trigger.ExtensionsManagerProtocol` property), 539
[writer](#) (`pytorch_pfn_extras.training.triggers.function_trigger.ExtensionsManagerProtocol` property), 541
[writer](#) (`pytorch_pfn_extras.training.triggers.interval_trigger.ExtensionsManagerProtocol` property), 544
[writer](#) (`pytorch_pfn_extras.training.triggers.manual_schedule_trigger.ExtensionsManagerProtocol` property), 547
[writer](#) (`pytorch_pfn_extras.training.triggers.minmax_value_trigger.ExtensionsManagerProtocol` property), 552
[writer](#) (`pytorch_pfn_extras.training.triggers.once_trigger.ExtensionsManagerProtocol` property), 555
[writer](#) (`pytorch_pfn_extras.training.triggers.time_trigger.ExtensionsManagerProtocol` property), 559