PyTorch-Lightning-Bolts Documentation

Release 0.1.0

PyTorchLightning et al.

Sep 09, 2020

START HERE

1 Introduction Guide 1

2 Model quality control 11

3 Build a Callback 15

4 Info Callbacks 17

5 Variational Callbacks 19

6 Lightning DataModule 21

7 Sklearn Datamodule 27

8 Vision DataModules 31

9 Losses 41

10 Reinforcement Learning 43

11 Bolts Loggers 45

12 How to use models 47

13 Autoencoders 55

14 Classic ML Models 59

15 Convolutional Architectures 63

16 GANs 67

17 Reinforcement Learning 71

18 Self-supervised Learning 101

19 Self-supervised learning Transforms 109

20 Self-supervised learning 119

21 Semi-supervised learning 121

22 Self-supervised Learning Contrastive tasks 123

23 Indices and tables 127

Python Module Index 233

Index 235

CHAPTER

ONE

INTRODUCTION GUIDE

Welcome to PyTorch Lightning Bolts!

Bolts is a Deep learning research and production toolbox of:

- SOTA pretrained models.
- Model components.
- Callbacks.
- Losses.
- Datasets.

The Main goal of bolts is to enable trying new ideas as fast as possible!

All models are tested (daily), benchmarked, documented and work on CPUs, TPUs, GPUs and 16-bit precision.

some examples!

```
from pl_bolts.models import VAE, GPT2, ImageGPT, PixelCNN
from pl_bolts.models.self_supervised import AMDIM, CPCV2, SimCLR, MocoV2
from pl_bolts.models import LinearRegression, LogisticRegression
from pl_bolts.models.gans import GAN
from pl_bolts.callbacks import PrintTableMetricsCallback
from pl_bolts.datamodules import FashionMNISTDataModule, CIFAR10DataModule,␣
↪ImagenetDataModule
```

Bolts are built for rapid idea iteration - subclass, override and train!

```
from pl_bolts.models import ImageGPT
from pl_bolts.self_supervised import SimCLR

class VideoGPT(ImageGPT):

    def training_step(self, batch, batch_idx):
        x, y = batch
        x = _shape_input(x)

        logits = self.gpt(x)
        simclr_features = self.simclr(x)

        # ------------------
        # do something new with GPT logits + simclr_features
        # ------------------

        loss = self.criterion(logits.view(-1, logits.size(-1)), x.view(-1).long())
```

(continues on next page)

(continued from previous page)

```
        logs = {"loss": loss}
        return {"loss": loss, "log": logs}
```

Mix and match data, modules and components as you please!

```
model = GAN(datamodule=ImagenetDataModule(PATH))
model = GAN(datamodule=FashionMNISTDataModule(PATH))
model = ImageGPT(datamodule=FashionMNISTDataModule(PATH))
```

And train on any hardware accelerator

```
import pytorch_lightning as pl

model = ImageGPT(datamodule=FashionMNISTDataModule(PATH))

# cpus
pl.Trainer.fit(model)

# gpus
pl.Trainer(gpus=8).fit(model)

# tpus
pl.Trainer(tpu_cores=8).fit(model)
```

Or pass in any dataset of your choice

```
model = ImageGPT()
Trainer().fit(
    model,
    train_dataloader=DataLoader(...),
    val_dataloader=DataLoader(...)
)
```

1.1 Community Built

Bolts are built-by the Lightning community and contributed to bolts. The lightning team guarantees that contributions are:

1. Rigorously Tested (CPUs, GPUs, TPUs).
2. Rigorously Documented.
3. Standardized via PyTorch Lightning.
4. Optimized for speed.
5. Checked for correctness.

1.1.1 How to contribute

We accept contributions directly to Bolts or via your own repository.

Note: We encourage you to have your own repository so we can link to it via our docs!

To contribute:

1. Submit a pull request to Bolts (we will help you finish it!).
2. We'll help you add tests.
3. We'll help you refactor models to work on (GPU, TPU, CPU)..
4. We'll help you remove bottlenecks in your model.
5. We'll help you write up documentation.
6. We'll help you pretrain expensive models and host weights for you.
7. We'll create proper attribution for you and link to your repo.
8. Once all of this is ready, we will merge into bolts.

After your model or other contribution is in bolts, our team will make sure it maintains compatibility with the other components of the library!

1.1.2 Contribution ideas

Don't have something to contribute? Ping us on Slack or look at our Github issues!

We'll help and guide you through the implementation / conversion

---

1.2 When to use Bolts

1.2.1 For pretrained models

Most bolts have pretrained weights trained on various datasets or algorithms. This is useful when you don't have enough data, time or money to do your own training.

For example, you could use a pretrained VAE to generate features for an image dataset.

```
from pl_bolts.models.autoencoders import VAE
from pl_bolts.models.self_supervised import CPCV2

model1 = VAE(pretrained='imagenet2012')
encoder = model1.encoder
encoder.freeze()

# bolts are pretrained on different datasets
model2 = CPCV2(encoder='resnet18', pretrained='imagenet128').freeze()
model3 = CPCV2(encoder='resnet18', pretrained='stl10').freeze()
```

(continues on next page)

(continued from previous page)

```
for (x, y) in own_data
    features = encoder(x)
    feat2 = model2(x)
    feat3 = model3(x)

# which is better?
```

1.2.2 To finetune on your data

If you have your own data, finetuning can often increase the performance. Since this is pure PyTorch you can use any finetuning protocol you prefer.

Example 1: Unfrozen finetune

```
# unfrozen finetune
model = CPCV2(encoder='resnet18', pretrained='imagenet128')
resnet18 = model.encoder
# don't call .freeze()

classifier = LogisticRegression()

for (x, y) in own_data:
    feats = resnet18(x)
    y_hat = classifier(feats)
```

Example 2: Freeze then unfreeze

```
# FREEZE!
model = CPCV2(encoder='resnet18', pretrained='imagenet128')
resnet18 = model.encoder
resnet18.freeze()

classifier = LogisticRegression()

for epoch in epochs:
    for (x, y) in own_data:
        feats = resnet18(x)
        y_hat = classifier(feats)
        loss = cross_entropy_with_logits(y_hat, y)

    # UNFREEZE after 10 epochs
    if epoch == 10:
        resnet18.unfreeze()
```

1.2.3 For research

Here is where bolts is very different than other libraries with models. It's not just designed for production, but each module is written to be easily extended for research.

```
from pl_bolts.models import ImageGPT
from pl_bolts.self_supervised import SimCLR

class VideoGPT(ImageGPT):

    def training_step(self, batch, batch_idx):
        x, y = batch
        x = _shape_input(x)

        logits = self.gpt(x)
        simclr_features = self.simclr(x)

        # ------------------
        # do something new with GPT logits + simclr_features
        # ------------------

        loss = self.criterion(logits.view(-1, logits.size(-1)), x.view(-1).long())

        logs = {"loss": loss}
        return {"loss": loss, "log": logs}
```

Or perhaps your research is in self_supervised_learning and you want to do a new SimCLR. In this case, the only thing you want to change is the loss.

By subclassing you can focus on changing a single piece of a system without worrying that the other parts work (because if they are in Bolts, then they do and we've tested it).

```
# subclass SimCLR and change ONLY what you want to try
class ComplexCLR(SimCLR):

    def init_loss(self):
        return self.new_xent_loss

    def new_xent_loss(self):
        out = torch.cat([out_1, out_2], dim=0) n_samples = len(out)

        # Full similarity matrix
        cov = torch.mm(out, out.t().contiguous())
        sim = torch.exp(cov / temperature)

        # Negative similarity
        mask = ~torch.eye(n_samples, device=sim.device).bool()
        neg = sim.masked_select(mask).view(n_samples, -1).sum(dim=-1)

        # ------------------
        # some new thing we want to do
        # ------------------

        # Positive similarity :
        pos = torch.exp(torch.sum(out_1 * out_2, dim=-1) / temperature)
        pos = torch.cat([pos, pos], dim=0)
        loss = -torch.log(pos / neg).mean()
```

(continues on next page)

(continued from previous page)

```
        return loss
```

1.3 Callbacks

Callbacks are arbitrary programs which can run at any points in time within a training loop in Lightning.

Bolts houses a collection of callbacks that are community contributed and can work in any Lightning Module!

```
from pl_bolts.callbacks import PrintTableMetricsCallback
import pytorch_lightning as pl

trainer = pl.Trainer(callbacks=[PrintTableMetricsCallback()])
```

1.4 DataModules

In PyTorch, working with data has these major elements.

1. Downloading, saving and preparing the dataset.
2. Splitting into train, val and test.
3. For each split, applying different transforms

A DataModule groups together those actions into a single reproducible *DataModule* that can be shared around to guarantee:

1. Consistent data preprocessing (download, splits, etc. . .)
2. The same exact splits
3. The same exact transforms

```
from pl_bolts.datamodules import ImagenetDataModule

dm = ImagenetDataModule(data_dir=PATH)

# standard PyTorch!
train_loader = dm.train_dataloader()
val_loader = dm.val_dataloader()
test_loader = dm.test_dataloader()

Trainer().fit(
    model,
    train_loader,
    val_loader
)
```

But when paired with PyTorch LightningModules (all bolts models), you can plug and play full dataset definitions with the same splits, transforms, etc. . .

```
imagenet = ImagenetDataModule(PATH)
model = VAE(datamodule=imagenet)
model = ImageGPT(datamodule=imagenet)
model = GAN(datamodule=imagenet)
```

We even have prebuilt modules to bridge the gap between Numpy, Sklearn and PyTorch

```
from sklearn.datasets import load_boston
from pl_bolts.datamodules import SklearnDataModule

X, y = load_boston(return_X_y=True)
datamodule = SklearnDataModule(X, y)

model = LitModel(datamodule)
```

1.5 Regression Heroes

In case your job or research doesn't need a "hammer", we offer implementations of Classic ML models which benefit from lightning's multi-GPU and TPU support.

So, now you can run huge workloads scalably, without needing to do any engineering. For instance, here we can run Logistic Regression on Imagenet (each epoch takes about 3 minutes)!

```
from pl_bolts.models.regression import LogisticRegression

imagenet = ImagenetDataModule(PATH)

# 224 x 224 x 3
pixels_per_image = 150_528
model = LogisticRegression(input_dim=pixels_per_image, num_classes=1000)
model.prepare_data = imagenet.prepare_data

trainer = Trainer(gpus=2)
trainer.fit(
    model,
    imagenet.train_dataloader(batch_size=256),
    imagenet.val_dataloader(batch_size=256)
)
```

1.5.1 Linear Regression

Here's an example for Linear regression

```
import pytorch_lightning as pl
from pl_bolts.datamodules import SklearnDataModule
from sklearn.datasets import load_boston

# link the numpy dataset to PyTorch
X, y = load_boston(return_X_y=True)
loaders = SklearnDataModule(X, y)

# training runs training batches while validating against a validation set
```

(continues on next page)

(continued from previous page)

```
model = LinearRegression()
trainer = pl.Trainer(num_gpus=8)
trainer.fit(model, loaders.train_dataloader(), loaders.val_dataloader())
```

Once you're done, you can run the test set if needed.

```
trainer.test(test_dataloaders=loaders.test_dataloader())
```

But more importantly, you can scale up to many GPUs, TPUs or even CPUs

```
# 8 GPUs
trainer = pl.Trainer(num_gpus=8)

# 8 TPU cores
trainer = pl.Trainer(tpu_cores=8)

# 32 GPUs
trainer = pl.Trainer(num_gpus=8, num_nodes=4)

# 128 CPUs
trainer = pl.Trainer(num_processes=128)
```

1.5.2 Logistic Regression

Here's an example for Logistic regression

```
from sklearn.datasets import load_iris
from pl_bolts.models.regression import LogisticRegression
from pl_bolts.datamodules import SklearnDataModule
import pytorch_lightning as pl

# use any numpy or sklearn dataset
X, y = load_iris(return_X_y=True)
dm = SklearnDataModule(X, y)

# build model
model = LogisticRegression(input_dim=4, num_classes=3)

# fit
trainer = pl.Trainer(tpu_cores=8, precision=16)
trainer.fit(model, dm.train_dataloader(), dm.val_dataloader())

trainer.test(test_dataloaders=dm.test_dataloader(batch_size=12))
```

Any input will be flattened across all dimensions except the firs one (batch). This means images, sound, etc... work out of the box.

```
# create dataset
dm = MNISTDataModule(num_workers=0, data_dir=tmpdir)

model = LogisticRegression(input_dim=28 * 28, num_classes=10, learning_rate=0.001)
model.prepare_data = dm.prepare_data
model.train_dataloader = dm.train_dataloader
model.val_dataloader = dm.val_dataloader
model.test_dataloader = dm.test_dataloader
```

(continues on next page)

(continued from previous page)

```
trainer = pl.Trainer(max_epochs=2)
trainer.fit(model)
trainer.test(model)
# {test_acc: 0.92}
```

But more importantly, you can scale up to many GPUs, TPUs or even CPUs

```
# 8 GPUs
trainer = pl.Trainer(num_gpus=8)

# 8 TPUs
trainer = pl.Trainer(tpu_cores=8)

# 32 GPUs
trainer = pl.Trainer(num_gpus=8, num_nodes=4)

# 128 CPUs
trainer = pl.Trainer(num_processes=128)
```

1.6 Regular PyTorch

Everything in bolts also works with regular PyTorch since they are all just nn.Modules!

However, if you train using Lightning you don't have to deal with engineering code :)

1.7 Command line support

Any bolt module can also be trained from the command line

```
cd pl_bolts/models/autoencoders/basic_vae
python basic_vae_pl_module.py
```

Each script accepts Argparse arguments for both the lightning trainer and the model

```
python basic_vae_pl_module.py -latent_dim 32 --batch_size 32 --gpus 4 --max_epochs 12
```

CHAPTER

TWO

MODEL QUALITY CONTROL

For bolts to be added to the library we have a **rigorous** quality control checklist

2.1 Bolts vs my own repo

We hope you keep your own repo still! We want to link to it to let people know. However, by adding your contribution to bolts you get these **additional** benefits!

1. More visibility! (more people all over the world use your code)
2. We test your code on every PR (CPUs, GPUs, TPUs).
3. We host the docs (and test on every PR).
4. We help you build thorough, beautiful documentation.
5. We help you build robust tests.
6. We'll pretrain expensive models for you and host weights.
7. We will improve the speed of your models!
8. Eligible for invited talks to discuss your implementation.
9. Lightning Swag + involvement in the broader contributor community :)

Note: You still get to keep your attribution and be recognized for your work!

Note: Bolts is a community library built by incredible people like you!

2.2 Contribution requirements

2.2.1 Benchmarked

Models have known performance results on common baseline datasets.

2.2.2 Device agnostic

Models must work on CPUs, GPUs and TPUs without changing code. We help authors with this.

```
# bad
encoder.to(device)
```

2.2.3 Fast

We inspect models for computational inefficiencies and help authors meet the bar. Granted, sometimes the approaches are slow for mathematical reasons. But anything related to engineering we help overcome.

```
# bad
mtx = ...
for xi in rows:
    for yi in cols
        mxt[xi, yi] = ...

# good
x = x.item().numpy()
x = np.some_fx(x)
x = torch.tensor(x)
```

2.2.4 Tested

Models are tested on every PR (on CPUs, GPUs and soon TPUs).

- Live build
- Tests

2.2.5 Modular

Models are modularized to be extended and reused easily.

```
# GOOD!
class LitVAE(pl.LightningModule):

    def init_prior(self, ...):
        # enable users to override interesting parts of each model

    def init_posterior(self, ...):
        # enable users to override interesting parts of each model

# BAD
class LitVAE(pl.LightningModule):

    def __init__(self):
        self.prior = ...
        self.posterior = ...
```

2.2.6 Attribution

Any models and weights that are contributed are attributed to you as the author(s).

We request that each contribution have:

- The original paper link
- The list of paper authors
- The link to the original paper code (if available)
- The link to your repo
- Your name and your team's name as the implementation authors.
- Your team's affiliation
- Any generated examples, or result plots.
- Hyperparameters configurations for the results.

Thank you for all your amazing contributions!

2.3 The bar seems high

If your model doesn't yet meet this bar, no worries! Please open the PR and our team of core contributors will help you get there!

2.4 Do you have contribution ideas?

Yes! Check the Github issues for requests from the Lightning team and the community! We'll even work with you to finish your implementation! Then we'll help you pretrain it and cover the compute costs when possible.

CHAPTER

THREE

BUILD A CALLBACK

This module houses a collection of callbacks that can be passed into the trainer

```
from pl_bolts.callbacks import PrintTableMetricsCallback
import pytorch_lightning as pl

trainer = pl.Trainer(callbacks=[PrintTableMetricsCallback()])

# loss|train_loss|val_loss|epoch
# ──────────────────────────────────
# 2.2541470527648926|2.2541470527648926|2.2158432006835938|0
```

3.1 What is a Callback

A callback is a self-contained program that can be intertwined into a training pipeline without polluting the main research logic.

3.2 Create a Callback

Creating a callback is simple:

```
from pytorch_lightning.callbacks import Callback

class MyCallback(Callback)
    def on_epoch_end(self, trainer, pl_module):
        # do something
```

Please refer to Callback docs for a full list of the 20+ hooks available.

CHAPTER

FOUR

INFO CALLBACKS

These callbacks give all sorts of useful information during training.

4.1 Print Table Metrics

This callbacks prints training metrics to a table. It's very bare-bones for speed purposes.

class pl_bolts.callbacks.printing.**PrintTableMetricsCallback**

Bases: pytorch_lightning.callbacks.Callback

Prints a table with the metrics in columns on every epoch end

Example:

```
from pl_bolts.callbacks import PrintTableMetricsCallback

callback = PrintTableMetricsCallback()
```

pass into trainer like so:

```
trainer = pl.Trainer(callbacks=[callback])
trainer.fit(...)

# ------------------------------
# at the end of every epoch it will print
# ------------------------------

# loss|train_loss|val_loss|epoch
# ——————————————————————————
# 2.2541470527648926|2.2541470527648926|2.2158432006835938|0
```

CHAPTER

FIVE

VARIATIONAL CALLBACKS

Useful callbacks for GANs, variational-autoencoders or anything with latent spaces.

5.1 Latent Dim Interpolator

Interpolates latent dims.

class pl_bolts.callbacks.variational.**LatentDimInterpolator**(*interpolate_epoch_interval=20*, *range_start=-5*, *range_end=5*)

Bases: pytorch_lightning.callbacks.Callback

Interpolates the latent space for a model by setting all dims to zero and stepping through the first two dims increasing one unit at a time.

Default interpolates between [-5, 5] (-5, -4, -3, . . . , 3, 4, 5)

Example:

```
from pl_bolts.callbacks import LatentDimInterpolator

Trainer(callbacks=[LatentDimInterpolator()])
```

Parameters

- **interpolate_epoch_interval** –
- **range_start** – default -5
- **range_end** – default 5

CHAPTER

SIX

LIGHTNING DATAMODULE

Datasets in PyTorch, Lightning and general Deep learning research have 4 main parts:

1. A train split + dataloader
2. A val split + dataloader
3. A test split + dataloader
4. A step to download, split, etc...

Step 4, also needs special care to make sure that it's only done on 1 GPU in a multi-GPU set-up. In addition, there are other challenges such as models that are built using information from the dataset such as needing to know image dimensions or number of classes.

A datamodule simplifies all of these parts and integrates seamlessly into Lightning.

```
class LitModel(pl.LightningModule):

    def __init__(self, datamodule):
        c, w, h = datamodule.size()
        self.l1 = nn.Linear(128, datamodule.num_classes)
        self.datamodule = datamodule

    def prepare_data(self):
        self.datamodule.prepare_data()

    def train_dataloader(self)
        return self.datamodule.train_dataloader()

    def val_dataloader(self)
        return self.datamodule.val_dataloader()

    def test_dataloader(self)
        return self.datamodule.test_dataloader()
```

DataModules can also be used with plain PyTorch

```
from pl_bolts.datamodules import MNISTDataModule, CIFAR10DataModule

datamodule = CIFAR10DataModule(PATH)
train_loader = datamodule.train_dataloader()
val_loader = datamodule.train_dataloader()
test_loader = datamodule.train_dataloader()
```

An advantage is that you can parametrize the data of your LightningModule

```
model = LitModel(datamodule = CIFAR10DataModule(PATH))
model = LitModel(datamodule = ImagenetDataModule(PATH))
```

Or even bridge between SKLearn or numpy datasets

```
from sklearn.datasets import load_boston
from pl_bolts.datamodules import SklearnDataModule

X, y = load_boston(return_X_y=True)
datamodule = SklearnDataModule(X, y)

model = LitModel(datamodule)
```

6.1 DataModule Advantages

Datamodules have two advantages:

1. You can guarantee that the exact same train, val and test splits can be used across models.
2. You can parameterize your model to be dataset agnostic.

Example:

```
from pl_bolts.datamodules import STL10DataModule, CIFAR10DataModule

# use the same dataset on different models (with exactly the same splits)
stl10_model = LitModel(STL10DataModule(PATH))
stl10_model = CoolModel(STL10DataModule(PATH))

# or make your model dataset agnostic
cifar10_model = LitModel(CIFAR10DataModule(PATH))
```

6.2 Build a DataModule

Use this to build your own consistent train, validation, test splits.

Example:

```
from pl_bolts.datamodules import LightningDataModule

class MyDataModule(LightningDataModule):

    def __init__(self,...):

    def prepare_data(self):
        # download and do something to your data

    def train_dataloader(self, batch_size):
        return DataLoader(...)

    def val_dataloader(self, batch_size):
        return DataLoader(...)
```

(continues on next page)

(continued from previous page)

```
    def test_dataloader(self, batch_size):
        return DataLoader(...)
```

Then use this in any model you want.

Example:

```
class LitModel(pl.LightningModule):

    def __init__(self, data_module=MyDataModule(PATH)):
        super().__init()
        self.dm = data_module

    def prepare_data(self):
        self.dm.prepare_data()

    def train_dataloader(self):
        return self.dm.train_dataloader()

    def val_dataloader(self):
        return self.dm.val_dataloader()

    def test_dataloader(self):
        return self.dm.test_dataloader()
```

6.2.1 DataModule class

class pl_bolts.datamodules.lightning_datamodule.**LightningDataModule**(*train_transforms=None*, *val_transforms=None*, *test_transforms=None*)

Bases: object

A DataModule standardizes the training, val, test splits, data preparation and transforms. The main advantage is consistent data splits and transforms across models.

Example:

```
class MyDataModule(LightningDataModule):

    def __init__(self):
        super().__init__()

    def prepare_data(self):
        # download, split, etc...

    def train_dataloader(self):
        train_split = Dataset(...)
        return DataLoader(train_split)

    def val_dataloader(self):
        val_split = Dataset(...)
        return DataLoader(val_split)

    def test_dataloader(self):
        test_split = Dataset(...)
        return DataLoader(test_split)
```

A DataModule implements 4 key methods

1. **prepare_data** (things to do on 1 GPU not on every GPU in distributed mode)
2. **train_dataloader** the training dataloader.
3. **val_dataloader** the val dataloader.
4. **test_dataloader** the test dataloader.

This allows you to share a full dataset without explaining what the splits, transforms or download process is.

classmethod add_argparse_args(*parent_parser*)
Extends existing argparse by default *LightningDataModule* attributes.

Return type ArgumentParser

classmethod from_argparse_args(*args*, ***kwargs*)
Create an instance from CLI arguments.

Parameters

- **args** (Union[Namespace, ArgumentParser]) – The parser or namespace to take arguments from. Only known arguments will be parsed and passed to the *LightningDataModule*.
- ****kwargs** – Additional keyword arguments that may override ones in the parser or namespace. These must be valid Trainer arguments.

Example:

```
parser = ArgumentParser(add_help=False)
parser = LightningDataModule.add_argparse_args(parser)
module = LightningDataModule.from_argparse_args(args)
```

classmethod get_init_arguments_and_types()
Scans the Trainer signature and returns argument names, types and default values.

Returns (argument name, set with argument types, argument default value).

Return type List with tuples of 3 values

abstract prepare_data(**args*, ***kwargs*)
Use this to download and prepare data. In distributed (GPU, TPU), this will only be called once. This is called before requesting the dataloaders:

Warning: Do not assign anything to the model in this step since this will only be called on 1 GPU.

Pseudocode:

```
model.prepare_data()
model.train_dataloader()
model.val_dataloader()
model.test_dataloader()
```

Example:

```
def prepare_data(self):
    download_imagenet()
    clean_imagenet()
    cache_imagenet()
```

size(*dim=None*)
Return the dimension of each input Either as a tuple or list of tuples

Return type Union[Tuple, int]

abstract test_dataloader(**args*, ***kwargs*)
Implement a PyTorch DataLoader for training.

Return type Union[DataLoader, List[DataLoader]]

Returns Single PyTorch DataLoader.

Note: Lightning adds the correct sampler for distributed and arbitrary hardware. There is no need to set it yourself.

Note: You can also return a list of DataLoaders

Example:

```
def test_dataloader(self):
    dataset = MNIST(root=PATH, train=False, transform=transforms.ToTensor(),␣
↪download=False)
    loader = torch.utils.data.DataLoader(dataset=dataset, shuffle=False)
    return loader
```

abstract train_dataloader(**args*, ***kwargs*)
Implement a PyTorch DataLoader for training.

Return type DataLoader

Returns Single PyTorch DataLoader.

Note: Lightning adds the correct sampler for distributed and arbitrary hardware. There is no need to set it yourself.

Example:

```
def train_dataloader(self):
    dataset = MNIST(root=PATH, train=True, transform=transforms.ToTensor(),␣
↪download=False)
    loader = torch.utils.data.DataLoader(dataset=dataset)
    return loader
```

abstract val_dataloader(**args*, ***kwargs*)
Implement a PyTorch DataLoader for training.

Return type Union[DataLoader, List[DataLoader]]

Returns Single PyTorch DataLoader.

Note: Lightning adds the correct sampler for distributed and arbitrary hardware. There is no need to set it yourself.

Note: You can also return a list of DataLoaders

Example:

```
def val_dataloader(self):
    dataset = MNIST(root=PATH, train=False, transform=transforms.ToTensor(),␣
↪download=False)
    loader = torch.utils.data.DataLoader(dataset=dataset, shuffle=False)
    return loader
```

CHAPTER

SEVEN

SKLEARN DATAMODULE

Utilities to map sklearn or numpy datasets to PyTorch Dataloaders with automatic data splits and GPU/TPU support.

```
from sklearn.datasets import load_boston
from pl_bolts.datamodules import SklearnDataModule

X, y = load_boston(return_X_y=True)
loaders = SklearnDataModule(X, y)

train_loader = loaders.train_dataloader(batch_size=32)
val_loader = loaders.val_dataloader(batch_size=32)
test_loader = loaders.test_dataloader(batch_size=32)
```

Or build your own torch datasets

```
from sklearn.datasets import load_boston
from pl_bolts.datamodules import SklearnDataset

X, y = load_boston(return_X_y=True)
dataset = SklearnDataset(X, y)
loader = DataLoader(dataset)
```

7.1 Sklearn Dataset Class

Transforms a sklearn or numpy dataset to a PyTorch Dataset.

class pl_bolts.datamodules.sklearn_datamodule.**SklearnDataset**(*X*, *y*, *X_transform=None*, *y_transform=None*)

Bases: torch.utils.data.Dataset

Mapping between numpy (or sklearn) datasets to PyTorch datasets.

Parameters

- **X** (ndarray) – Numpy ndarray
- **y** (ndarray) – Numpy ndarray
- **X_transform** (Optional[Any]) – Any transform that works with Numpy arrays
- **y_transform** (Optional[Any]) – Any transform that works with Numpy arrays

Example

```
>>> from sklearn.datasets import load_boston
>>> from pl_bolts.datamodules import SklearnDataset
...
>>> X, y = load_boston(return_X_y=True)
>>> dataset = SklearnDataset(X, y)
>>> len(dataset)
506
```

7.2 Sklearn DataModule Class

Automatically generates the train, validation and test splits for a Numpy dataset. They are set up as dataloaders for convenience. Optionally, you can pass in your own validation and test splits.

class pl_bolts.datamodules.sklearn_datamodule.**SklearnDataModule**(*X*, *y*, *x_val=None*, *y_val=None*, *x_test=None*, *y_test=None*, *val_split=0.2*, *test_split=0.1*, *num_workers=2*, *random_state=1234*, *shuffle=True*, **args*, ***kwargs*)

Bases: *pl_bolts.datamodules.lightning_datamodule.LightningDataModule*

Automatically generates the train, validation and test splits for a Numpy dataset. They are set up as dataloaders for convenience. Optionally, you can pass in your own validation and test splits.

Example

```
>>> from sklearn.datasets import load_boston
>>> from pl_bolts.datamodules import SklearnDataModule
...
>>> X, y = load_boston(return_X_y=True)
>>> loaders = SklearnDataModule(X, y)
...
>>> # train set
>>> train_loader = loaders.train_dataloader(batch_size=32)
>>> len(train_loader.dataset)
355
>>> len(train_loader)
11
>>> # validation set
>>> val_loader = loaders.val_dataloader(batch_size=32)
>>> len(val_loader.dataset)
100
```

(continues on next page)

(continued from previous page)

```
>>> len(val_loader)
3
>>> # test set
>>> test_loader = loaders.test_dataloader(batch_size=32)
>>> len(test_loader.dataset)
51
>>> len(test_loader)
1
```

test_dataloader(*batch_size=16*)

Implement a PyTorch DataLoader for training.

Returns Single PyTorch `DataLoader`.

Note: Lightning adds the correct sampler for distributed and arbitrary hardware. There is no need to set it yourself.

Note: You can also return a list of DataLoaders

Example:

```
def test_dataloader(self):
    dataset = MNIST(root=PATH, train=False, transform=transforms.ToTensor(),␣
↪download=False)
    loader = torch.utils.data.DataLoader(dataset=dataset, shuffle=False)
    return loader
```

train_dataloader(*batch_size=16*)

Implement a PyTorch DataLoader for training.

Returns Single PyTorch `DataLoader`.

Note: Lightning adds the correct sampler for distributed and arbitrary hardware. There is no need to set it yourself.

Example:

```
def train_dataloader(self):
    dataset = MNIST(root=PATH, train=True, transform=transforms.ToTensor(),␣
↪download=False)
    loader = torch.utils.data.DataLoader(dataset=dataset)
    return loader
```

val_dataloader(*batch_size=16*)

Implement a PyTorch DataLoader for training.

Returns Single PyTorch `DataLoader`.

Note: Lightning adds the correct sampler for distributed and arbitrary hardware. There is no need to set it yourself.

Note: You can also return a list of DataLoaders

Example:

```
def val_dataloader(self):
    dataset = MNIST(root=PATH, train=False, transform=transforms.ToTensor(),␣
↪download=False)
    loader = torch.utils.data.DataLoader(dataset=dataset, shuffle=False)
    return loader
```

CHAPTER

EIGHT

VISION DATAMODULES

The following are pre-built datamodules for computer-vision.

8.1 Supervised learning

These are standard vision datasets with the train, test, val splits pre-generated in DataLoaders with the standard transforms (and Normalization) values

8.1.1 MNIST

class pl_bolts.datamodules.mnist_datamodule.**MNISTDataModule**(*data_dir*, *val_split=5000*, *num_workers=16*, *normalize=False*, **args*, ***kwargs*)

Bases: *pl_bolts.datamodules.lightning_datamodule.LightningDataModule*

Standard MNIST, train, val, test splits and transforms

Transforms:

```
mnist_transforms = transform_lib.Compose([
    transform_lib.ToTensor()
])
```

Example:

```
from pl_bolts.datamodules import MNISTDataModule

dm = MNISTDataModule('.')
model = LitModel(datamodule=dm)
```

Parameters

- **data_dir** (str) – where to save/load the data
- **val_split** (int) – how many of the training images to use for the validation split
- **num_workers** (int) – how many workers to use for loading data
- **normalize** (bool) – If true applies image normalize

prepare_data()
Saves MNIST files to data_dir

test_dataloader(*batch_size=32*, *transforms=None*)
MNIST test set uses the test split

Parameters

- **batch_size** – size of batch
- **transforms** – custom transforms

train_dataloader(*batch_size=32*, *transforms=None*)
MNIST train set removes a subset to use for validation

Parameters

- **batch_size** – size of batch
- **transforms** – custom transforms

val_dataloader(*batch_size=32*, *transforms=None*)
MNIST val set uses a subset of the training set for validation

Parameters

- **batch_size** – size of batch
- **transforms** – custom transforms

property num_classes
Return: 10

8.1.2 FashionMNIST

class pl_bolts.datamodules.fashion_mnist_datamodule.**FashionMNISTDataModule**(*data_dir*, *val_split=5000*, *num_workers=16*, **args*, ***kwargs*)

Bases: *pl_bolts.datamodules.lightning_datamodule.LightningDataModule*

Standard FashionMNIST, train, val, test splits and transforms

Transforms:

```
mnist_transforms = transform_lib.Compose([
    transform_lib.ToTensor()
])
```

Example:

```
from pl_bolts.datamodules import FashionMNISTDataModule

dm = FashionMNISTDataModule('.')
model = LitModel(datamodule=dm)
```

Parameters

- **data_dir** (str) – where to save/load the data
- **val_split** (int) – how many of the training images to use for the validation split

- **num_workers** (int) – how many workers to use for loading data

prepare_data()
Saves FashionMNIST files to data_dir

test_dataloader(*batch_size=32*, *transforms=None*)
FashionMNIST test set uses the test split

Parameters

- **batch_size** – size of batch
- **transforms** – custom transforms

train_dataloader(*batch_size=32*, *transforms=None*)
FashionMNIST train set removes a subset to use for validation

Parameters

- **batch_size** – size of batch
- **transforms** – custom transforms

val_dataloader(*batch_size=32*, *transforms=None*)
FashionMNIST val set uses a subset of the training set for validation

Parameters

- **batch_size** – size of batch
- **transforms** – custom transforms

property num_classes
Return: 10

8.1.3 CIFAR-10

class pl_bolts.datamodules.cifar10_datamodule.**CIFAR10DataModule**(*data_dir*, *val_split=5000*, *num_workers=16*, **args*, ***kwargs*)

Bases: *pl_bolts.datamodules.lightning_datamodule.LightningDataModule*

Standard CIFAR10, train, val, test splits and transforms

Transforms:

```
mnist_transforms = transform_lib.Compose([
    transform_lib.ToTensor(),
    transforms.Normalize(
        mean=[x / 255.0 for x in [125.3, 123.0, 113.9]],
        std=[x / 255.0 for x in [63.0, 62.1, 66.7]]
    )
])
```

Example:

```
from pl_bolts.datamodules import CIFAR10DataModule

dm = CIFAR10DataModule(PATH)
model = LitModel(datamodule=dm)
```

Or you can set your own transforms

Example:

```
dm.train_transforms = ...
dm.test_transforms = ...
dm.val_transforms  = ...
```

Parameters

- **data_dir** – where to save/load the data
- **val_split** – how many of the training images to use for the validation split
- **num_workers** – how many workers to use for loading data

prepare_data()
Saves CIFAR10 files to data_dir

test_dataloader(*batch_size*)
CIFAR10 test set uses the test split

Parameters

- **batch_size** – size of batch
- **transforms** – custom transforms

train_dataloader(*batch_size*)
CIFAR train set removes a subset to use for validation

Parameters **batch_size** – size of batch

val_dataloader(*batch_size*)
CIFAR10 val set uses a subset of the training set for validation

Parameters **batch_size** – size of batch

property num_classes
Return: 10

8.1.4 Imagenet

class pl_bolts.datamodules.imagenet_datamodule.**ImagenetDataModule**(*data_dir*, *meta_dir=None*, *num_imgs_per_val_class=50*, *image_size=224*, *num_workers=16*, **args*, ***kwargs*)

Bases: *pl_bolts.datamodules.lightning_datamodule.LightningDataModule*

Imagenet train, val and test dataloaders.

The train set is the imagenet train.

The val set is taken from the train set with *num_imgs_per_val_class* images per class. For example if *num_imgs_per_val_class=2* then there will be 2,000 images in the validation set.

The test set is the official imagenet validation set.

Example:

```
from pl_bolts.datamodules import ImagenetDataModule

datamodule = ImagenetDataModule(IMAGENET_PATH)
```

Parameters

- **data_dir** (str) – path to the imagenet dataset file
- **meta_dir** (Optional[str]) – path to meta.bin file
- **num_imgs_per_val_class** (int) – how many images per class for the validation set
- **image_size** (int) – final image size
- **num_workers** (int) – how many data workers

prepare_data()

This method already assumes you have imagenet2012 downloaded. It validates the data using the meta.bin.

> **Warning:** Please download imagenet on your own first.

test_dataloader(*batch_size*, *num_images_per_class=-1*, *transforms=None*)

Uses the validation split of imagenet2012 for testing

Parameters

- **batch_size** – the batch size
- **num_images_per_class** – how many images per class to test on
- **transforms** – the transforms

train_dataloader(*batch_size*)

Uses the train split of imagenet2012 and puts away a portion of it for the validation split

Parameters

- **batch_size** – the batch size
- **transforms** – the transforms

train_transform()

The standard imagenet transforms

```
transform_lib.Compose([
    transform_lib.RandomResizedCrop(self.image_size),
    transform_lib.RandomHorizontalFlip(),
    transform_lib.ToTensor(),
    transform_lib.Normalize(
        mean=[0.485, 0.456, 0.406],
        std=[0.229, 0.224, 0.225]
    ),
])
```

val_dataloader(*batch_size*, *transforms=None*)

Uses the part of the train split of imagenet2012 that was not used for training via *num_imgs_per_val_class*

Parameters

- **batch_size** – the batch size
- **transforms** – the transforms

val_transform()
The standard imagenet transforms for validation

```
transform_lib.Compose([
    transform_lib.Resize(self.image_size + 32),
    transform_lib.CenterCrop(self.image_size),
    transform_lib.ToTensor(),
    transform_lib.Normalize(
        mean=[0.485, 0.456, 0.406],
        std=[0.229, 0.224, 0.225]
    ),
])
```

property num_classes
Return:

1000

8.2 Semi-supervised learning

The following datasets have support for unlabeled training and semi-supervised learning where only a few examples are labeled.

8.2.1 Imagenet (ssl)

class pl_bolts.datamodules.ssl_imagenet_datamodule.**SSLImagenetDataModule**(*data_dir*, *meta_dir=None*, *num_workers=16*, **args*, ***kwargs*)

Bases: *pl_bolts.datamodules.lightning_datamodule.LightningDataModule*

prepare_data()
Use this to download and prepare data. In distributed (GPU, TPU), this will only be called once. This is called before requesting the dataloaders:

> **Warning:** Do not assign anything to the model in this step since this will only be called on 1 GPU.

Pseudocode:

```
model.prepare_data()
model.train_dataloader()
model.val_dataloader()
model.test_dataloader()
```

Example:

```
def prepare_data(self):
    download_imagenet()
    clean_imagenet()
    cache_imagenet()
```

test_dataloader(*batch_size*, *num_images_per_class*, *add_normalize=False*)
Implement a PyTorch DataLoader for training.

Returns Single PyTorch `DataLoader`.

Note: Lightning adds the correct sampler for distributed and arbitrary hardware. There is no need to set it yourself.

Note: You can also return a list of DataLoaders

Example:

```
def test_dataloader(self):
    dataset = MNIST(root=PATH, train=False, transform=transforms.ToTensor(),␣
↪download=False)
    loader = torch.utils.data.DataLoader(dataset=dataset, shuffle=False)
    return loader
```

train_dataloader(*batch_size*, *num_images_per_class=-1*, *add_normalize=False*)
Implement a PyTorch DataLoader for training.

Returns Single PyTorch `DataLoader`.

Note: Lightning adds the correct sampler for distributed and arbitrary hardware. There is no need to set it yourself.

Example:

```
def train_dataloader(self):
    dataset = MNIST(root=PATH, train=True, transform=transforms.ToTensor(),␣
↪download=False)
    loader = torch.utils.data.DataLoader(dataset=dataset)
    return loader
```

val_dataloader(*batch_size*, *num_images_per_class=50*, *add_normalize=False*)
Implement a PyTorch DataLoader for training.

Returns Single PyTorch `DataLoader`.

Note: Lightning adds the correct sampler for distributed and arbitrary hardware. There is no need to set it yourself.

Note: You can also return a list of DataLoaders

Example:

```
def val_dataloader(self):
    dataset = MNIST(root=PATH, train=False, transform=transforms.ToTensor(),␣
↪download=False)
    loader = torch.utils.data.DataLoader(dataset=dataset, shuffle=False)
    return loader
```

8.2.2 STL-10

class pl_bolts.datamodules.stl10_datamodule.**STL10DataModule**(*data_dir*, *unlabeled_val_split=5000*, *train_val_split=500*, *num_workers=16*, **args*, ***kwargs*)

Bases: *pl_bolts.datamodules.lightning_datamodule.LightningDataModule*

Standard STL-10, train, val, test splits and transforms. STL-10 has support for doing validation splits on the labeled or unlabeled splits

Transforms:

```
mnist_transforms = transform_lib.Compose([
    transform_lib.ToTensor(),
    transforms.Normalize(
        mean=(0.43, 0.42, 0.39),
        std=(0.27, 0.26, 0.27)
    )
])
```

Example:

```
from pl_bolts.datamodules import STL10DataModule

dm = STL10DataModule(PATH)
model = LitModel(datamodule=dm)
```

Parameters

- **data_dir** (str) – where to save/load the data
- **unlabeled_val_split** (int) – how many images from the unlabeled training split to use for validation
- **train_val_split** (int) – how many images from the labeled training split to use for validation
- **num_workers** (int) – how many workers to use for loading data

prepare_data()
Downloads the unlabeled, train and test split

test_dataloader(*batch_size*)
Loads the test split of STL10

Parameters

- **batch_size** – the batch size
- **transforms** – the transforms

train_dataloader(*batch_size*)
Loads the 'unlabeled' split minus a portion set aside for validation via *unlabeled_val_split*.

Parameters **batch_size** – the batch size

train_dataloader_mixed(*batch_size*)
Loads a portion of the 'unlabeled' training data and 'train' (labeled) data. both portions have a subset removed for validation via *unlabeled_val_split* and *train_val_split*

Parameters

- **batch_size** – the batch size
- **transforms** – a sequence of transforms

val_dataloader(*batch_size*)
Loads a portion of the 'unlabeled' training data set aside for validation The val dataset = (unlabeled - train_val_split)

Parameters

- **batch_size** – the batch size
- **transforms** – a sequence of transforms

val_dataloader_mixed(*batch_size*)
Loads a portion of the 'unlabeled' training data set aside for validation along with the portion of the 'train' dataset to be used for validation

unlabeled_val = (unlabeled - train_val_split)

labeled_val = (train- train_val_split)

full_val = unlabeled_val + labeled_val

Parameters

- **batch_size** – the batch size
- **transforms** – a sequence of transforms

CHAPTER

NINE

LOSSES

This package lists common losses across research domains (This is a work in progress. If you have any losses you want to contribute, please submit a PR!)

Note: this module is a work in progress

9.1 Your Loss

We're cleaning up many of our losses, but in the meantime, submit a PR to add your loss here!

CHAPTER

TEN

REINFORCEMENT LEARNING

These are common losses used in RL.

10.1 DQN Loss

pl_bolts.losses.rl.**dqn_loss**(*batch*, *net*, *target_net*, *gamma=0.99*)

Calculates the mse loss using a mini batch from the replay buffer

Parameters

- **batch** (Tuple[Tensor, Tensor]) – current mini batch of replay data
- **net** (Module) – main training network
- **target_net** (Module) – target network of the main training network
- **gamma** (float) – discount factor

Return type Tensor

Returns loss

10.2 Double DQN Loss

pl_bolts.losses.rl.**double_dqn_loss**(*batch*, *net*, *target_net*, *gamma=0.99*)

Calculates the mse loss using a mini batch from the replay buffer. This uses an improvement to the original DQN loss by using the double dqn. This is shown by using the actions of the train network to pick the value from the target network. This code is heavily commented in order to explain the process clearly

Parameters

- **batch** (Tuple[Tensor, Tensor]) – current mini batch of replay data
- **net** (Module) – main training network
- **target_net** (Module) – target network of the main training network
- **gamma** (float) – discount factor

Return type Tensor

Returns loss

10.3 Per DQN Loss

pl_bolts.losses.rl.**per_dqn_loss**(*batch*, *batch_weights*, *net*, *target_net*, *gamma=0.99*)

Calculates the mse loss with the priority weights of the batch from the PER buffer

Parameters

- **batch** (Tuple[Tensor, Tensor]) – current mini batch of replay data
- **batch_weights** (List) – how each of these samples are weighted in terms of priority
- **net** (Module) – main training network
- **target_net** (Module) – target network of the main training network
- **gamma** (float) – discount factor

Return type Tuple[Tensor, ndarray]

Returns loss and batch_weights

CHAPTER

ELEVEN

BOLTS LOGGERS

The loggers in this package are being considered to be added to the main PyTorch Lightning repository. These loggers may be more unstable, in development, or not fully tested yet.

Note: This module is a work in progress

11.1 allegro.ai TRAINS

allegro.ai is a third-party logger. To use *TrainsLogger* as your logger do the following. First, install the package:

```
pip install trains
```

Then configure the logger and pass it to the `Trainer`:

```
from pl_bolts.loggers import TrainsLogger
trains_logger = TrainsLogger(
    project_name='examples',
    task_name='pytorch lightning test',
)
trainer = Trainer(logger=trains_logger)
```

```
class MyModule(LightningModule):
    def __init__(self):
        some_img = fake_image()
        self.logger.experiment.log_image('debug', 'generated_image_0', some_img, 0)
```

See also:

TrainsLogger docs.

11.1.1 Your Logger

Add your loggers here!

CHAPTER

TWELVE

HOW TO USE MODELS

Models are meant to be "bolted" onto your research or production cases.

Bolts are meant to be used in the following ways

12.1 Predicting on your data

Most bolts have pretrained weights trained on various datasets or algorithms. This is useful when you don't have enough data, time or money to do your own training.

For example, you could use a pretrained VAE to generate features for an image dataset.

```
from pl_bolts.models.autoencoders import VAE

model = VAE(pretrained='imagenet2012')
encoder = model.encoder
encoder.freeze()

for (x, y) in own_data
    features = encoder(x)
```

The advantage of bolts is that each system can be decomposed and used in interesting ways. For instance, this resnet18 was trained using self-supervised learning (no labels) on Imagenet, and thus might perform better than the same resnet18 trained with labels

```
# trained without labels
from pl_bolts.models.self_supervised import CPCV2

model = CPCV2(encoder='resnet18', pretrained='imagenet128')
resnet18_unsupervised = model.encoder.freeze()

# trained with labels
from torchvision.models import resnet18
resnet18_supervised = resnet18(pretrained=True)

# perhaps the features when trained without labels are much better for classification␣
↪or other tasks
x = image_sample()
unsup_feats = resnet18_unsupervised(x)
sup_feats = resnet18_supervised(x)

# which one will be better?
```

Bolts are often trained on more than just one dataset.

```
model = CPCV2(encoder='resnet18', pretrained='stl10')
```

12.2 Finetuning on your data

If you have a little bit of data and can pay for a bit of training, it's often better to finetune on your own data.

To finetune you have two options unfrozen finetuning or unfrozen later.

12.2.1 Unfrozen Finetuning

In this approach, we load the pretrained model and unfreeze from the beginning

```
model = CPCV2(encoder='resnet18', pretrained='imagenet128')
resnet18 = model.encoder
# don't call .freeze()

classifier = LogisticRegression()

for (x, y) in own_data:
    feats = resnet18(x)
    y_hat = classifier(feats)
    ...
```

Or as a LightningModule

```
class FineTuner(pl.LightningModule):

    def __init__(self, encoder):
        self.encoder = encoder
        self.classifier = LogisticRegression()

    def training_step(self, batch, batch_idx):
        (x, y) = batch
        feats = self.encoder(x)
        y_hat = self.classifier(feats)
        loss = cross_entropy_with_logits(y_hat, y)
        return loss

trainer = Trainer(gpus=2)
model = FineTuner(resnet18)
trainer.fit(model)
```

Sometimes this works well, but more often it's better to keep the encoder frozen for a while

12.2.2 Freeze then unfreeze

The approach that works best most often is to freeze first then unfreeze later

```
# freeze!
model = CPCV2(encoder='resnet18', pretrained='imagenet128')
resnet18 = model.encoder
resnet18.freeze()

classifier = LogisticRegression()

for epoch in epochs:
    for (x, y) in own_data:
        feats = resnet18(x)
        y_hat = classifier(feats)
        loss = cross_entropy_with_logits(y_hat, y)

    # unfreeze after 10 epochs
    if epoch == 10:
        resnet18.unfreeze()
```

Note: In practice, unfreezing later works MUCH better.

Or in Lightning as a Callback so you don't pollute your research code.

```
class UnFreezeCallback(Callback):

    def on_epoch_end(self, trainer, pl_module):
        if trainer.current_epoch == 10.
            encoder.unfreeze()

trainer = Trainer(gpus=2, callbacks=[UnFreezeCallback()])
model = FineTuner(resnet18)
trainer.fit(model)
```

Unless you still need to mix it into your research code.

```
class FineTuner(pl.LightningModule):

    def __init__(self, encoder):
        self.encoder = encoder
        self.classifier = LogisticRegression()

    def training_step(self, batch, batch_idx):

        # option 1 - (not recommended because it's messy)
        if self.trainer.current_epoch == 10:
            self.encoder.unfreeze()

        (x, y) = batch
        feats = self.encoder(x)
        y_hat = self.classifier(feats)
        loss = cross_entropy_with_logits(y_hat, y)
        return loss

    def on_epoch_end(self, trainer, pl_module):
```

(continues on next page)

(continued from previous page)

```
        # a hook is cleaner (but a callback is much better)
        if self.trainer.current_epoch == 10:
            self.encoder.unfreeze()
```

12.2.3 Hyperparameter search

For finetuning to work well, you should try many versions of the model hyperparameters. Otherwise you're unlikely to get the most value out of your data.

```
learning_rates = [0.01, 0.001, 0.0001]
hidden_dim = [128, 256, 512]

for lr in learning_rates:
    for hd in hidden_dim:
        vae = VAE(hidden_dim=hd, learning_rate=lr)
        trainer = Trainer()
        trainer.fit(vae)
```

12.3 Train from scratch

If you do have enough data and compute resources, then you could try training from scratch.

```
# get data
train_data = DataLoader(YourDataset)
val_data = DataLoader(YourDataset)

# use any bolts model without pretraining
model = VAE()

# fit!
trainer = Trainer(gpus=2)
trainer.fit(model, train_data, val_data)
```

Note: For this to work well, make sure you have enough data and time to train these models!

12.4 For research

What separates bolts from all the other libraries out there is that bolts is built by and used by AI researchers. This means every single bolt is modularized so that it can be easily extended or mixed with arbitrary parts of the rest of the code-base.

12.4.1 Extending work

Perhaps a research project requires modifying a part of a know approach. In this case, you're better off only changing that part of a system that is already know to perform well. Otherwise, you risk not implementing the work correctly.

Example 1: Changing the prior or approx posterior of a VAE

```
from pl_bolts.models.autoencoders import VAE

class MyVAEFlavor(VAE):

    def init_prior(self, z_mu, z_std):
        P = MyPriorDistribution

        # default is standard normal
        # P = distributions.normal.Normal(loc=torch.zeros_like(z_mu), scale=torch.
↪ones_like(z_std))
        return P

    def init_posterior(self, z_mu, z_std):
        Q = MyPosteriorDistribution
        # default is normal(z_mu, z_sigma)
        # Q = distributions.normal.Normal(loc=z_mu, scale=z_std)
        return Q
```

And of course train it with lightning.

```
model = MyVAEFlavor()
trainer = Trainer()
trainer.fit(model)
```

In just a few lines of code you changed something fundamental about a VAE... This means you can iterate through ideas much faster knowing that the bolt implementation and the training loop are CORRECT and TESTED.

If your model doesn't work with the new P, Q, then you can discard that research idea much faster than trying to figure out if your VAE implementation was correct, or if your training loop was correct.

Example 2: Changing the generator step of a GAN

```
from pl_bolts.models.gans import GAN

class FancyGAN(GAN):

    def generator_step(self, x):
        # sample noise
        z = torch.randn(x.shape[0], self.hparams.latent_dim)
        z = z.type_as(x)

        # generate images
        self.generated_imgs = self(z)

        # ground truth result (ie: all real)
        real = torch.ones(x.size(0), 1)
        real = real.type_as(x)
        g_loss = self.generator_loss(real)

        tqdm_dict = {'g_loss': g_loss}
        output = OrderedDict({
            'loss': g_loss,
```

(continues on next page)

(continued from previous page)

```
            'progress_bar': tqdm_dict,
            'log': tqdm_dict
        })
        return output
```

Example 3: Changing the way the loss is calculated in a contrastive self-supervised learning approach

```
from pl_bolts.models.self_supervised import AMDIM

class MyDIM(AMDIM):

    def validation_step(self, batch, batch_nb):
        [img_1, img_2], labels = batch

        # generate features
        r1_x1, r5_x1, r7_x1, r1_x2, r5_x2, r7_x2 = self.forward(img_1, img_2)

        # Contrastive task
        loss, lgt_reg = self.contrastive_task((r1_x1, r5_x1, r7_x1), (r1_x2, r5_x2,␣
↪r7_x2))
        unsupervised_loss = loss.sum() + lgt_reg

        result = {
            'val_nce': unsupervised_loss
        }
        return result
```

12.4.2 Importing parts

All the bolts are modular. This means you can also arbitrarily mix and match fundamental blocks from across approaches.

Example 1: Use the VAE encoder for a GAN as a generator

```
from pl_bolts.models.gans import GAN
from pl_bolts.models.autoencoders.basic_vae import Encoder

class FancyGAN(GAN):

    def init_generator(self, img_dim):
        generator = Encoder(...)
        return generator

trainer = Trainer(...)
trainer.fit(FancyGAN())
```

Example 2: Use the contrastive task of AMDIM in CPC

```
from pl_bolts.models.self_supervised import AMDIM, CPCV2

default_amdim_task = AMDIM().contrastive_task
model = CPCV2(contrastive_task=default_amdim_task, encoder='cpc_default')
# you might need to modify the cpc encoder depending on what you use
```

12.4.3 Compose new ideas

You may also be interested in creating completely new approaches that mix and match all sorts of different pieces together

```
# this model is for illustration purposes, it makes no research sense but it's␣
↪intended to show
# that you can be as creative and expressive as you want.
class MyNewContrastiveApproach(pl.LightningModule):

    def __init__(self):
        suoer().__init_()

        self.gan = GAN()
        self.vae = VAE()
        self.amdim = AMDIM()
        self.cpc = CPCV2

    def training_step(self, batch, batch_idx):
        (x, y) = batch

        feat_a = self.gan.generator(x)
        feat_b = self.vae.encoder(x)

        unsup_loss = self.amdim(feat_a) + self.cpc(feat_b)

        vae_loss = self.vae._step(batch)
        gan_loss = self.gan.generator_loss(x)

        return unsup_loss + vae_loss + gan_loss
```

CHAPTER

THIRTEEN

AUTOENCODERS

This section houses autoencoders and variational autoencoders.

13.1 Basic AE

This is the simplest autoencoder. You can use it like so

```
from pl_bolts.models.autoencoders import AE

model = AE()
trainer = Trainer()
trainer.fit(model)
```

You can override any part of this AE to build your own variation.

```
from pl_bolts.models.autoencoders import AE

class MyAEFlavor(AE):

    def init_encoder(self, hidden_dim, latent_dim, input_width, input_height):
        encoder = YourSuperFancyEncoder(...)
        return encoder
```

class pl_bolts.models.autoencoders.**AE**(*datamodule=None*, *input_channels=1*, *input_height=28*, *input_width=28*, *latent_dim=32*, *batch_size=32*, *hidden_dim=128*, *learning_rate=0.001*, *num_workers=8*, *data_dir='.'*, ***kwargs*)

Bases: pytorch_lightning.LightningModule

Arg:

datamodule: the datamodule (train, val, test splits) input_channels: num of image channels input_height: image height input_width: image width latent_dim: emb dim for encoder batch_size: the batch size hidden_dim: the encoder dim learning_rate: the learning rate num_workers: num dataloader workers data_dir: where to store data

---

13.1.1 Variational Autoencoders

13.2 Basic VAE

Use the VAE like so.

```
from pl_bolts.models.autoencoders import VAE

model = VAE()
trainer = Trainer()
trainer.fit(model)
```

You can override any part of this VAE to build your own variation.

```
from pl_bolts.models.autoencoders import VAE

class MyVAEFlavor(VAE):

    def get_posterior(self, mu, std):
        # do something other than the default
        # P = self.get_distribution(self.prior, loc=torch.zeros_like(mu), scale=torch.
↪ones_like(std))

        return P
```

class pl_bolts.models.autoencoders.**VAE**(*hidden_dim=128*, *latent_dim=32*, *input_channels=3*, *input_width=224*, *input_height=224*, *batch_size=32*, *learning_rate=0.001*, *data_dir='.'*, *datamodule=None*, *pretrained=None*, ***kwargs*)

Bases: pytorch_lightning.LightningModule

Standard VAE with Gaussian Prior and approx posterior.

Model is available pretrained on different datasets:

Example:

```
# not pretrained
vae = VAE()

# pretrained on imagenet
vae = VAE(pretrained='imagenet')

# pretrained on cifar10
vae = VAE(pretrained='cifar10')
```

Parameters

- **hidden_dim** (int) – encoder and decoder hidden dims
- **latent_dim** (int) – latenet code dim
- **input_channels** (int) – num of channels of the input image.
- **input_width** (int) – image input width
- **input_height** (int) – image input height
- **batch_size** (int) – the batch size

- **the learning rate** (*learning_rate"*) –
- **data_dir** (str) – the directory to store data
- **datamodule** (Optional[*LightningDataModule*]) – The Lightning DataModule
- **pretrained** (Optional[str]) – Load weights pretrained on a dataset

CHAPTER

FOURTEEN

CLASSIC ML MODELS

This module implements classic machine learning models in PyTorch Lightning, including linear regression and logistic regression. Unlike other libraries that implement these models, here we use PyTorch to enable multi-GPU, multi-TPU and half-precision training.

14.1 Linear Regression

Linear regression fits a linear model between a real-valued target variable (y) and one or more features (X). We estimate the regression coefficients $beta$ that minimizes the mean squared error between the predicted and true target values.

```
from pl_bolts.models.regression import LinearRegression
import pytorch_lightning as pl
from pl_bolts.datamodules import SklearnDataModule
from sklearn.datasets import load_boston

X, y = load_boston(return_X_y=True)
loaders = SklearnDataModule(X, y)

model = LinearRegression(input_dim=13)
trainer = pl.Trainer()
trainer.fit(model, loaders.train_dataloader(), loaders.val_dataloader())
trainer.test(test_dataloaders=loaders.test_dataloader())
```

class pl_bolts.models.regression.linear_regression.**LinearRegression**(*input_dim*, *bias=True*, *learning_rate=0.0001*, *optimizer=torch.optim.Adam*, *l1_strength=None*, *l2_strength=None*, ***kwargs*)

Bases: pytorch_lightning.LightningModule

Linear regression model implementing - with optional L1/L2 regularization $$min_{W} ||(Wx + b) - y ||_2^2 $$

Parameters

- **input_dim** (int) – number of dimensions of the input (1+)

- **bias** (bool) – If false, will not use $+b$
- **learning_rate** (float) – learning_rate for the optimizer
- **optimizer** (Optimizer) – the optimizer to use (default='Adam')
- **l1_strength** (Optional[float]) – L1 regularization strength (default=None)
- **l2_strength** (Optional[float]) – L2 regularization strength (default=None)

14.2 Logistic Regression

Logistic regression is a non-linear model used for classification, i.e. when we have a categorical target variable. This implementation supports both binary and multi-class classification.

To leverage autograd we think of logistic regression as a one-layer neural network with a sigmoid activation. This allows us to support training on GPUs and TPUs.

```
from sklearn.datasets import load_iris
from pl_bolts.models.regression import LogisticRegression
from pl_bolts.datamodules import SklearnDataModule
import pytorch_lightning as pl

# use any numpy or sklearn dataset
X, y = load_iris(return_X_y=True)
dm = SklearnDataModule(X, y)

# build model
model = LogisticRegression(input_dim=4, num_classes=3)

# fit
trainer = pl.Trainer(tpu_cores=8, precision=16)
trainer.fit(model, dm.train_dataloader(), dm.val_dataloader())

trainer.test(test_dataloaders=dm.test_dataloader(batch_size=12))
```

Any input will be flattened across all dimensions except the firs one (batch). This means images, sound, etc... work out of the box.

```
# create dataset
dm = MNISTDataModule(num_workers=0, data_dir=tmpdir)

model = LogisticRegression(input_dim=28 * 28, num_classes=10, learning_rate=0.001)
model.prepare_data = dm.prepare_data
model.train_dataloader = dm.train_dataloader
model.val_dataloader = dm.val_dataloader
model.test_dataloader = dm.test_dataloader

trainer = pl.Trainer(max_epochs=2)
trainer.fit(model)
trainer.test(model)
# {test_acc: 0.92}
```

class pl_bolts.models.regression.logistic_regression.**LogisticRegression**(*input_dim*, *num_classes*, *bias=True*, *learning_rate=0.0001*, *optimizer=torch.optim.Adam*, *l1_strength=0.0*, *l2_strength=0.0*, ***kwargs*)

Bases: pytorch_lightning.LightningModule

Logistic regression model

Parameters

- **input_dim** (int) – number of dimensions of the input (at least 1)
- **num_classes** (int) – number of class labels (binary: 2, multi-class: >2)
- **bias** (bool) – specifies if a constant or intercept should be fitted (equivalent to fit_intercept in sklearn)
- **learning_rate** (float) – learning_rate for the optimizer
- **optimizer** (Optimizer) – the optimizer to use (default='Adam')
- **l1_strength** (float) – L1 regularization strength (default=None)
- **l2_strength** (float) – L2 regularization strength (default=None)

CHAPTER

FIFTEEN

CONVOLUTIONAL ARCHITECTURES

This package lists contributed convolutional architectures.

15.1 GPT-2

class pl_bolts.models.vision.image_gpt.gpt2.**GPT2**(*embed_dim*, *heads*, *layers*, *num_positions*, *vocab_size*, *num_classes*)

Bases: pytorch_lightning.LightningModule

GPT-2 from language Models are Unsupervised Multitask Learners

Paper by: Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever

Implementation contributed by:

- Teddy Koker

Example:

```
from pl_bolts.models import GPT2

seq_len = 17
batch_size = 32
vocab_size = 16
x = torch.randint(0, vocab_size, (seq_len, batch_size))
model = GPT2(embed_dim=32, heads=2, layers=2, num_positions=2, vocab_size=vocab_
↪size, num_classes=4)
results = model(x)
```

forward(*x*, *classify=False*)

Expect input as shape [sequence len, batch] If classify, return classification logits

15.2 Image GPT

class pl_bolts.models.vision.image_gpt.igpt_module.**ImageGPT**(*datamodule=None*, *embed_dim=16*, *heads=2*, *layers=2*, *pixels=28*, *vocab_size=16*, *num_classes=10*, *classify=False*, *batch_size=64*, *learning_rate=0.01*, *steps=25000*, *data_dir='.'*, *num_workers=8*, ***kwargs*)

Bases: pytorch_lightning.LightningModule

Paper: Generative Pretraining from Pixels [original paper code].

Paper by: Mark Che, Alec Radford, Rewon Child, Jeff Wu, Heewoo Jun, Prafulla Dhariwal, David Luan, Ilya Sutskever

Implementation contributed by:

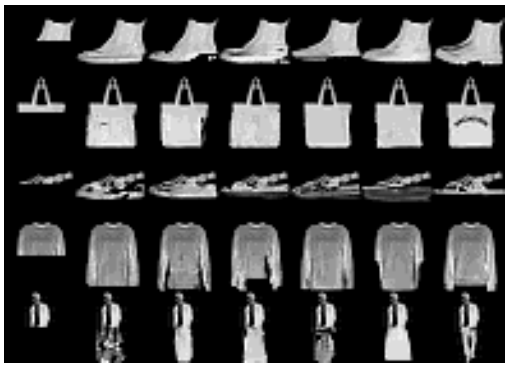
- Teddy Koker

Original repo with results and more implementation details:

- https://github.com/teddykoker/image-gpt

Example Results (Photo credits: Teddy Koker):





Default arguments:

Table 1: Argument Defaults

Argument	Default	iGPT-S (Chen et al.)
–embed_dim	16	512
–heads	2	8
–layers	8	24
–pixels	28	32
–vocab_size	16	512
–num_classes	10	10
–batch_size	64	128
–learning_rate	0.01	0.01
–steps	25000	1000000

Example:

```
import pytorch_lightning as pl
from pl_bolts.models.vision import ImageGPT

dm = MNISTDataModule('.')
model = ImageGPT(dm)

pl.Trainer(gpu=4).fit(model)
```

As script:

```
cd pl_bolts/models/vision/image_gpt
python igpt_module.py --learning_rate 1e-2 --batch_size 32 --gpus 4
```

Parameters

- **datamodule** (Optional[*LightningDataModule*]) – LightningDataModule
- **embed_dim** (int) – the embedding dim
- **heads** (int) – number of attention heads
- **layers** (int) – number of layers
- **pixels** (int) – number of input pixels
- **vocab_size** (int) – vocab size
- **num_classes** (int) – number of classes in the input
- **classify** (bool) – true if should classify
- **batch_size** (int) – the batch size
- **learning_rate** (float) – learning rate
- **steps** (int) – number of steps for cosine annealing
- **data_dir** (str) – where to store data
- **num_workers** (int) – num_data workers

15.3 Pixel CNN

class pl_bolts.models.vision.pixel_cnn.**PixelCNN**(*input_channels*, *hidden_channels=256*, *num_blocks=5*)

Bases: torch.nn.Module

Implementation of Pixel CNN.

Paper authors: Aaron van den Oord, Nal Kalchbrenner, Oriol Vinyals, Lasse Espeholt, Alex Graves, Koray Kavukcuoglu

Implemented by:

- William Falcon

Example:

```
>>> from pl_bolts.models.vision import PixelCNN
>>> import torch
...
>>> model = PixelCNN(input_channels=3)
>>> x = torch.rand(5, 3, 64, 64)
>>> out = model(x)
...
>>> out.shape
torch.Size([5, 3, 64, 64])
```

CHAPTER

SIXTEEN

GANS

Collection of Generative Adversarial Networks

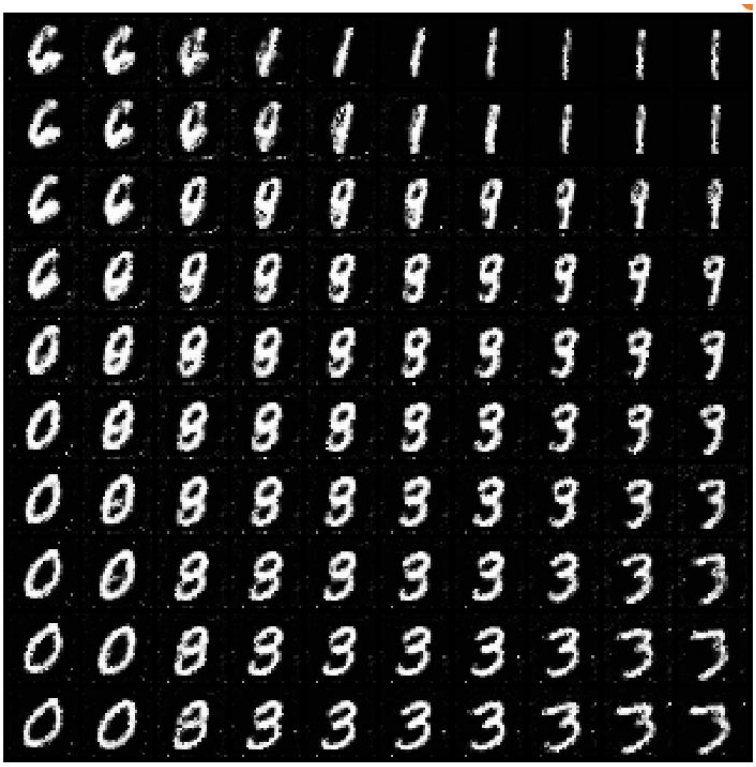
16.1 Basic GAN

This is a vanilla GAN. This model can work on any dataset size but results are shown for MNIST. Replace the encoder, decoder or any part of the training loop to build a new method, or simply finetune on your data.

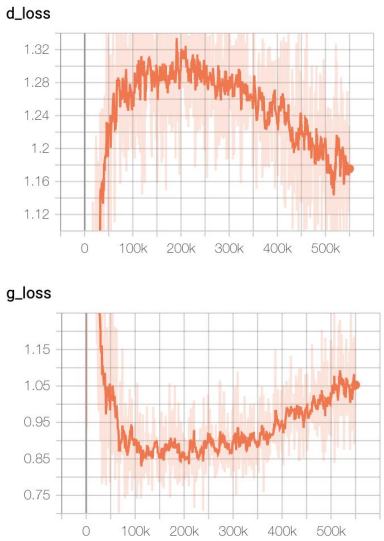
Implemented by:

- William Falcon

Example outputs:



Loss curves:




```
from pl_bolts.models.gans import GAN
...
gan = GAN()
trainer = Trainer()
trainer.fit(gan)
```

class pl_bolts.models.gans.**GAN**(*datamodule=None*, *latent_dim=32*, *batch_size=100*, *learning_rate=0.0002*, *data_dir=''*, *num_workers=8*, ***kwargs*)

Bases: pytorch_lightning.LightningModule

Vanilla GAN implementation.

Example:

```
from pl_bolts.models.gan import GAN

m = GAN()
Trainer(gpus=2).fit(m)
```

Example CLI:

```
# mnist
python  basic_gan_module.py --gpus 1

# imagenet
python  basic_gan_module.py --gpus 1 --dataset 'imagenet2012'
--data_dir /path/to/imagenet/folder/ --meta_dir ~/path/to/meta/bin/folder
--batch_size 256 --learning_rate 0.0001
```

Parameters

- **datamodule** (Optional[*LightningDataModule*]) – the datamodule (train, val, test splits)
- **latent_dim** (int) – emb dim for encoder
- **batch_size** (int) – the batch size
- **learning_rate** (float) – the learning rate
- **data_dir** (str) – where to store data

- **num_workers** (int) – data workers

forward(*z*)

Generates an image given input noise z

Example:

```
z = torch.rand(batch_size, latent_dim)
gan = GAN.load_from_checkpoint(PATH)
img = gan(z)
```

CHAPTER

SEVENTEEN

REINFORCEMENT LEARNING

This module is a collection of common RL approaches implemented in Lightning.

17.1 Module authors

Contributions by: Donal Byrne

- DQN
- Double DQN
- Dueling DQN
- Noisy DQN
- NStep DQN
- Prioritized Experience Replay DQN
- Reinforce
- Vanilla Policy Gradient

Note: RL models currently only support CPU and single GPU training with *distributed_backend=dp*. Full GPU support will be added in later updates.

17.2 DQN Models

The following models are based on DQN

17.2.1 Deep-Q-Network (DQN)

DQN model introduced in Playing Atari with Deep Reinforcement Learning. Paper authors: Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, Martin Riedmiller.

Original implementation by: Donal Byrne

The DQN was introduced in Playing Atari with Deep Reinforcement Learning by researchers at DeepMind. This took the concept of tabular Q learning and scaled it to much larger problems by apporximating the Q function using a deep neural network.

The goal behind DQN was to take the simple control method of Q learning and scale it up in order to solve complicated tasks. As well as this, the method needed to be stable. The DQN solves these issues with the following additions.

Approximated Q Function

Storing Q values in a table works well in theory, but is completely unscalable. Instead, the authors apporximate the Q function using a deep neural network. This allows the DQN to be used for much more complicated tasks

Replay Buffer

Similar to supervised learning, the DQN learns on randomly sampled batches of previous data stored in an Experience Replay Buffer. The 'target' is calculated using the Bellman equation

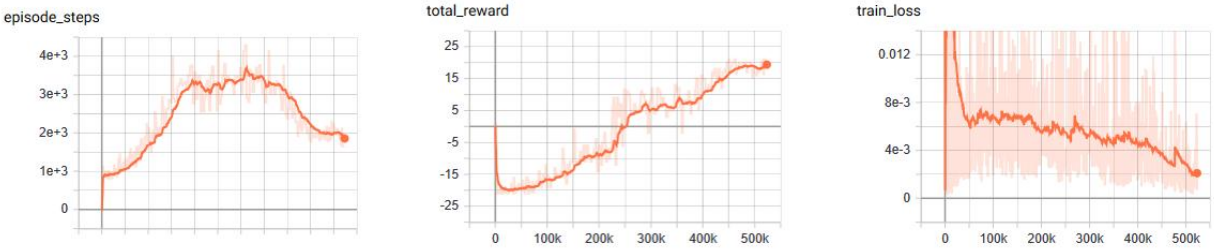
$$Q(s,a) < -(r + \gamma \max_{a' \in A} Q(s',a'))^2$$

and then we optimize using SGD just like a standard supervised learning problem.

$$L = (Q(s,a) - (r + \gamma \max_{a' \in A} Q(s',a'))^2$$

DQN Results

DQN: Pong



Example:

```
from pl_bolts.models.rl import DQN
dqn = DQN("PongNoFrameskip-v4")
trainer = Trainer()
trainer.fit(dqn)
```

class pl_bolts.models.rl.dqn_model.**DQN**(*env*, *gpus=0*, *eps_start=1.0*, *eps_end=0.02*, *eps_last_frame=150000*, *sync_rate=1000*, *gamma=0.99*, *learning_rate=0.0001*, *batch_size=32*, *replay_size=100000*, *warm_start_size=10000*, *num_samples=500*, ***kwargs*)

Bases: pytorch_lightning.LightningModule

Basic DQN Model

PyTorch Lightning implementation of DQN

Paper authors: Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, Martin Riedmiller.

Model implemented by:

- *Donal Byrne <https://github.com/djbyrne>*

Example

```
>>> from pl_bolts.models.rl.dqn_model import DQN
...
>>> model = DQN("PongNoFrameskip-v4")
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

Parameters

- **env** (str) – gym environment tag
- **gpus** (int) – number of gpus being used
- **eps_start** (float) – starting value of epsilon for the epsilon-greedy exploration
- **eps_end** (float) – final value of epsilon for the epsilon-greedy exploration
- **eps_last_frame** (int) – the final frame in for the decrease of epsilon. At this frame espilon = eps_end
- **sync_rate** (int) – the number of iterations between syncing up the target network with the train network
- **gamma** (float) – discount factor
- **learning_rate** (float) – learning rate
- **batch_size** (int) – size of minibatch pulled from the DataLoader
- **replay_size** (int) – total capacity of the replay buffer
- **warm_start_size** (int) – how many random steps through the environment to be carried out at the start of training to fill the buffer with a starting point
- **num_samples** (int) – the number of samples to pull from the dataset iterator and feed to the DataLoader

---

Note:

This example is based on: https://github.com/PacktPublishing/Deep-Reinforcement-Learning-Hands-On-Second-Edition/blob/master/Chapter06/02_dqn_pong.py

---

Note: Currently only supports CPU and single GPU training with *distributed_backend=dp*

---

static add_model_specific_args(*arg_parser*)
Adds arguments for DQN model

Note: these params are fine tuned for Pong env

Parameters **arg_parser** (ArgumentParser) – parent parser

Return type ArgumentParser

build_networks()
Initializes the DQN train and target networks

Return type None

configure_optimizers()
Initialize Adam optimizer

Return type List[Optimizer]

forward(*x*)
Passes in a state x through the network and gets the q_values of each action as an output

Parameters **x** (Tensor) – environment state

Return type Tensor

Returns q values

populate(*warm_start*)
Populates the buffer with initial experience

Return type None

prepare_data()
Initialize the Replay Buffer dataset used for retrieving experiences

Return type None

test_dataloader()
Get test loader

Return type DataLoader

test_epoch_end(*outputs*)
Log the avg of the test results

Return type Dict[str, Tensor]

test_step(**args*, ***kwargs*)
Evaluate the agent for 10 episodes

Return type Dict[str, Tensor]

train_dataloader()
Get train loader

Return type DataLoader

training_step(*batch*, *_*)
Carries out a single step through the environment to update the replay buffer. Then calculates loss based on the minibatch recieved

Parameters

- **batch** (Tuple[Tensor, Tensor]) – current mini batch of replay data
- **_** – batch number, not used

Return type `OrderedDict`

Returns Training loss and log metrics

17.2.2 Double DQN

Double DQN model introduced in Deep Reinforcement Learning with Double Q-learning Paper authors: Hado van Hasselt, Arthur Guez, David Silver

Original implementation by: Donal Byrne

The original DQN tends to overestimate Q values during the Bellman update, leading to instability and is harmful to training. This is due to the max operation in the Bellman equation.

We are constantly taking the max of our agents estimates during our update. This may seem reasonable, if we could trust these estimates. However during the early stages of training, the estimates for these values will be off center and can lead to instability in training until our estimates become more reliable

The Double DQN fixes this overestimation by choosing actions for the next state using the main trained network but uses the values of these actions from the more stable target network. So we are still going to take the greedy action, but the value will be less "optimisitc" because it is chosen by the target network.

DQN expected return

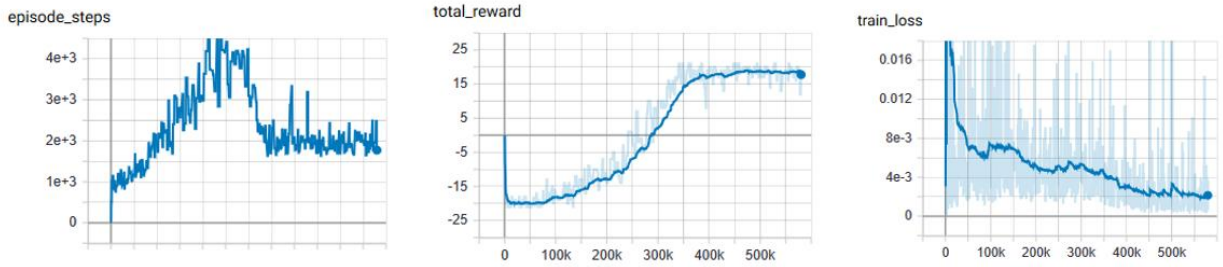
$$Q(s_t, a_t) = r_t + \gamma * \max_{Q'}(S_{t+1}, a)$$

Double DQN expected return

$$Q(s_t, a_t) = r_t + \gamma * \max Q'(S_{t+1}, \arg\max_{Q}(S_{t+1}, a))$$

Double DQN Results

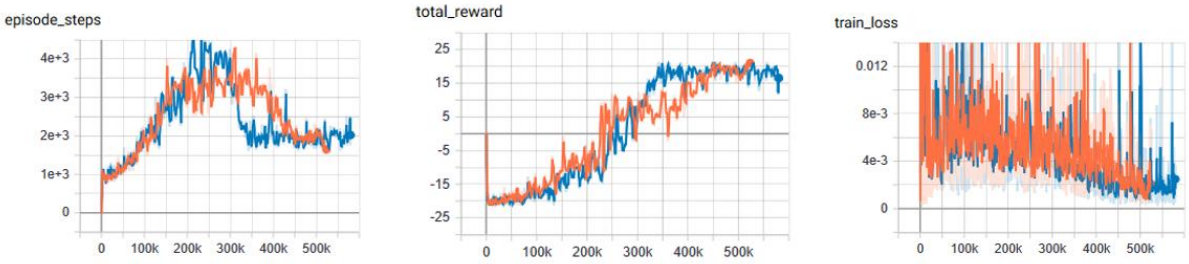
Double DQN: Pong



DQN vs Double DQN: Pong

orange: DQN

blue: Double DQN



Example:

```
from pl_bolts.models.rl import DoubleDQN
ddqn = DoubleDQN("PongNoFrameskip-v4")
trainer = Trainer()
trainer.fit(ddqn)
```

class pl_bolts.models.rl.double_dqn_model.**DoubleDQN**(*env*, *gpus=0*, *eps_start=1.0*, *eps_end=0.02*, *eps_last_frame=150000*, *sync_rate=1000*, *gamma=0.99*, *learning_rate=0.0001*, *batch_size=32*, *replay_size=100000*, *warm_start_size=10000*, *num_samples=500*, ***kwargs*)

Bases: *pl_bolts.models.rl.dqn_model.DQN*

Double Deep Q-network (DDQN) PyTorch Lightning implementation of Double DQN

Paper authors: Hado van Hasselt, Arthur Guez, David Silver

Model implemented by:

- *Donal Byrne <https://github.com/djbyrne>*

Example

```
>>> from pl_bolts.models.rl.double_dqn_model import DoubleDQN
...
>>> model = DoubleDQN("PongNoFrameskip-v4")
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

Parameters

- **env** (str) – gym environment tag
- **gpus** (int) – number of gpus being used
- **eps_start** (float) – starting value of epsilon for the epsilon-greedy exploration
- **eps_end** (float) – final value of epsilon for the epsilon-greedy exploration

- **eps_last_frame** (int) – the final frame in for the decrease of epsilon. At this frame espilon = eps_end
- **sync_rate** (int) – the number of iterations between syncing up the target network with the train network
- **gamma** (float) – discount factor
- **lr** – learning rate
- **batch_size** (int) – size of minibatch pulled from the DataLoader
- **replay_size** (int) – total capacity of the replay buffer
- **warm_start_size** (int) – how many random steps through the environment to be carried out at the start of training to fill the buffer with a starting point
- **sample_len** – the number of samples to pull from the dataset iterator and feed to the DataLoader

---

Note:

This example is based on: https://github.com/PacktPublishing/Deep-Reinforcement-Learning-Hands-On-Second-Edition/blob/master/Chapter08/03_dqn_double.py

---

Note: Currently only supports CPU and single GPU training with *distributed_backend=dp*

PyTorch Lightning implementation of DQN

Paper authors: Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, Martin Riedmiller.

Model implemented by:

- *Donal Byrne <https://github.com/djbyrne>*

Example

```
>>> from pl_bolts.models.rl.dqn_model import DQN
...
>>> model = DQN("PongNoFrameskip-v4")
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

Parameters

- **env** (str) – gym environment tag
- **gpus** (int) – number of gpus being used
- **eps_start** (float) – starting value of epsilon for the epsilon-greedy exploration
- **eps_end** (float) – final value of epsilon for the epsilon-greedy exploration

- **eps_last_frame** (int) – the final frame in for the decrease of epsilon. At this frame espilon = eps_end
- **sync_rate** (int) – the number of iterations between syncing up the target network with the train network
- **gamma** (float) – discount factor
- **learning_rate** (float) – learning rate
- **batch_size** (int) – size of minibatch pulled from the DataLoader
- **replay_size** (int) – total capacity of the replay buffer
- **warm_start_size** (int) – how many random steps through the environment to be carried out at the start of training to fill the buffer with a starting point
- **num_samples** (int) – the number of samples to pull from the dataset iterator and feed to the DataLoader

---

Note:

This example is based on: https://github.com/PacktPublishing/Deep-Reinforcement-Learning-Hands-On-Second-Edition/blob/master/Chapter06/02_dqn_pong.py

---

Note: Currently only supports CPU and single GPU training with *distributed_backend=dp*

---

training_step(*batch*, _)

Carries out a single step through the environment to update the replay buffer. Then calculates loss based on the minibatch recieved

Parameters

- **batch** (Tuple[Tensor, Tensor]) – current mini batch of replay data
- _ – batch number, not used

Return type OrderedDict

Returns Training loss and log metrics

17.2.3 Dueling DQN

Dueling DQN model introduced in Dueling Network Architectures for Deep Reinforcement Learning Paper authors: Ziyu Wang, Tom Schaul, Matteo Hessel, Hado van Hasselt, Marc Lanctot, Nando de Freitas

Original implementation by: Donal Byrne

The Q value that we are trying to approximate can be divided into two parts, the value state V(s) and the 'advantage' of actions in that state A(s, a). Instead of having one full network estimate the entire Q value, Dueling DQN uses two estimator heads in order to separate the estimation of the two parts.

The value is the same as in value iteration. It is the discounted expected reward achieved from state s. Think of the value as the 'base reward' from being in state s.

The advantage tells us how much 'extra' reward we get from taking action a while in state s. The advantage bridges the gap between Q(s, a) and V(s) as Q(s, a) = V(s) + A(s, a).

In the paper [Dueling Network Architectures for Deep Reinforcement Learning](https://arxiv.org/abs/1511.06581) the network uses two heads, one outputs the value state and the other outputs the advantage. This leads to better training stability, faster convergence and overall better results. The V head outputs a single scalar (the state value), while the advantage head outputs a tensor equal to the size of the action space, containing an advantage value for each action in state s.

Changing the network architecture is not enough, we also need to ensure that the advantage mean is 0. This is done by subtracting the mean advantage from the Q value. This essentially pulls the mean advantage to 0.

$$Q(s,a) = V(s) + A(s,a) - 1/N * \sum_{k}(A(s,k)$$

Dueling DQN Benefits

- **Ability to efficiently learn the state value function. In the dueling network, every Q update also updates the Value** stream, where as in DQN only the value of the chosen action is updated. This provides a better approximation of the values
- **The differences between total Q values for a given state are quite small in relation to the magnitude of Q. The** difference in the Q values between the best action and the second best action can be very small, while the average state value can be much larger. The differences in scale can introduce noise, which may lead to the greedy policy switching the priority of these actions. The seperate estimators for state value and advantage makes the Dueling DQN robust to this type of scenario

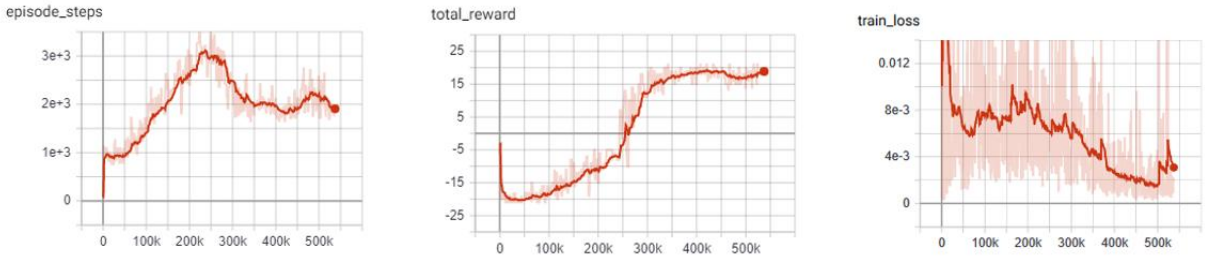
Dueling DQN Results

The results below a noticeable improvement from the original DQN network.

Dueling DQN baseline: Pong

Similar to the results of the DQN baseline, the agent has a period where the number of steps per episodes increase as it begins to hold its own against the heuristic oppoent, but then the steps per episode quickly begins to drop as it gets better and starts to beat its opponent faster and faster. There is a noticable point at step ~250k where the agent goes from losing to winning.

As you can see by the total rewards, the dueling network's training progression is very stable and continues to trend upward until it finally plateus.



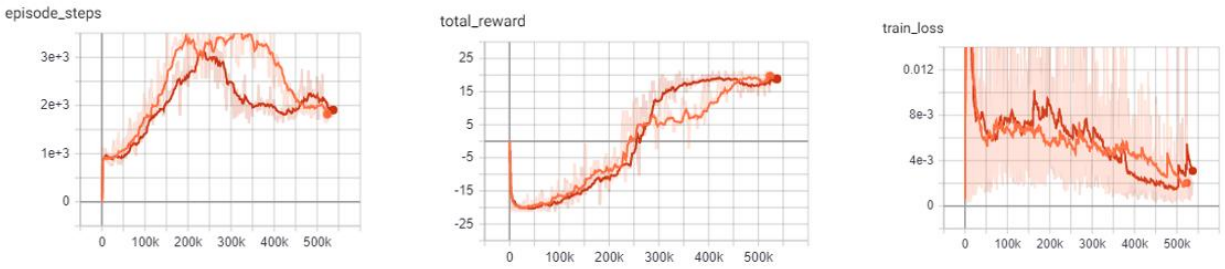
DQN vs Dueling DQN: Pong

In comparison to the base DQN, we see that the Dueling network's training is much more stable and is able to reach a score in the high teens faster than the DQN agent. Even though the Dueling network is more stable and out performs DQN early in training, by the end of training the two networks end up at the same point.

This could very well be due to the simplicity of the Pong environment.

- Orange: DQN

- Red: Dueling DQN



Example:

```
from pl_bolts.models.rl import DuelingDQN
dueling_dqn = DuelingDQN("PongNoFrameskip-v4")
trainer = Trainer()
trainer.fit(dueling_dqn)
```

class pl_bolts.models.rl.dueling_dqn_model.**DuelingDQN**(*env*, *gpus=0*, *eps_start=1.0*, *eps_end=0.02*, *eps_last_frame=150000*, *sync_rate=1000*, *gamma=0.99*, *learning_rate=0.0001*, *batch_size=32*, *replay_size=100000*, *warm_start_size=10000*, *num_samples=500*, ***kwargs*)

Bases: *pl_bolts.models.rl.dqn_model.DQN*

PyTorch Lightning implementation of Dueling DQN

Paper authors: Ziyu Wang, Tom Schaul, Matteo Hessel, Hado van Hasselt, Marc Lanctot, Nando de Freitas

Model implemented by:

- *Donal Byrne <https://github.com/djbyrne>*

Example

```
>>> from pl_bolts.models.rl.dueling_dqn_model import DuelingDQN
...
>>> model = DuelingDQN("PongNoFrameskip-v4")
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

Parameters

- **env** (str) – gym environment tag
- **gpus** (int) – number of gpus being used

- **eps_start** (float) – starting value of epsilon for the epsilon-greedy exploration
- **eps_end** (float) – final value of epsilon for the epsilon-greedy exploration
- **eps_last_frame** (int) – the final frame in for the decrease of epsilon. At this frame espilon = eps_end
- **sync_rate** (int) – the number of iterations between syncing up the target network with the train network
- **gamma** (float) – discount factor
- **lr** – learning rate
- **batch_size** (int) – size of minibatch pulled from the DataLoader
- **replay_size** (int) – total capacity of the replay buffer
- **warm_start_size** (int) – how many random steps through the environment to be carried out at the start of training to fill the buffer with a starting point
- **sample_len** – the number of samples to pull from the dataset iterator and feed to the DataLoader

---

Note: Currently only supports CPU and single GPU training with *distributed_backend=dp*

PyTorch Lightning implementation of DQN

Paper authors: Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, Martin Riedmiller.

Model implemented by:

- *Donal Byrne <https://github.com/djbyrne>*

Example

```
>>> from pl_bolts.models.rl.dqn_model import DQN
...
>>> model = DQN("PongNoFrameskip-v4")
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

Parameters

- **env** (str) – gym environment tag
- **gpus** (int) – number of gpus being used
- **eps_start** (float) – starting value of epsilon for the epsilon-greedy exploration
- **eps_end** (float) – final value of epsilon for the epsilon-greedy exploration
- **eps_last_frame** (int) – the final frame in for the decrease of epsilon. At this frame espilon = eps_end
- **sync_rate** (int) – the number of iterations between syncing up the target network with the train network

- **gamma** (float) – discount factor
- **learning_rate** (float) – learning rate
- **batch_size** (int) – size of minibatch pulled from the DataLoader
- **replay_size** (int) – total capacity of the replay buffer
- **warm_start_size** (int) – how many random steps through the environment to be carried out at the start of training to fill the buffer with a starting point
- **num_samples** (int) – the number of samples to pull from the dataset iterator and feed to the DataLoader

---

Note:

This example is based on: https://github.com/PacktPublishing/Deep-Reinforcement-Learning-Hands-On-Second-Edition/blob/master/Chapter06/02_dqn_pong.py

---

Note: Currently only supports CPU and single GPU training with *distributed_backend=dp*

---

build_networks()
Initializes the Dueling DQN train and target networks

Return type None

17.2.4 Noisy DQN

Noisy DQN model introduced in Noisy Networks for Exploration Paper authors: Meire Fortunato, Mohammad Gheshlaghi Azar, Bilal Piot, Jacob Menick, Ian Osband, Alex Graves, Vlad Mnih, Remi Munos, Demis Hassabis, Olivier Pietquin, Charles Blundell, Shane Legg

Original implementation by: Donal Byrne

Up until now the DQN agent uses a seperate exploration policy, generally epsilon-greedy where start and end values are set for its exploration. [Noisy Networks For Exploration](https://arxiv.org/abs/1706.10295) introduces a new exploration strategy by adding noise parameters to the weightsof the fully connect layers which get updated during backpropagation of the network. The noise parameters drive the exploration of the network instead of simply taking random actions more frequently at the start of training and less frequently towards the end.The of authors of propose two ways of doing this.

During the optimization step a new set of noisy parameters are sampled. During training the agent acts according to the fixed set of parameters. At the next optimization step, the parameters are updated with a new sample. This ensures the agent always acts based on the parameters that are drawn from the current noise distribution.

The authors propose two methods of injecting noise to the network.

1) **Independent Gaussian Noise: This injects noise per weight. For each weight a random value is taken from** the distribution. Noise parameters are stored inside the layer and are updated during backpropagation. The output of the layer is calculated as normal.
2) **Factorized Gaussian Noise: This injects nosier per input/ouput. In order to minimize the number of random values** this method stores two random vectors, one with the size of the input and the other with the size of the output. Using these two vectors, a random matrix is generated for the layer by calculating the outer products of the vector

Noisy DQN Benefits

- **Improved exploration function. Instead of just performing completely random actions, we add decreasing amount of noise**
  and uncertainty to our policy allowing to explore while still utilising its policy
- **The fact that this method is automatically tuned means that we do not have to tune hyper parameters for**
  epsilon-greedy!

Note: for now I have just implemented the Independant Gaussian as it has been reported there isn't much difference in results for these benchmark environments.

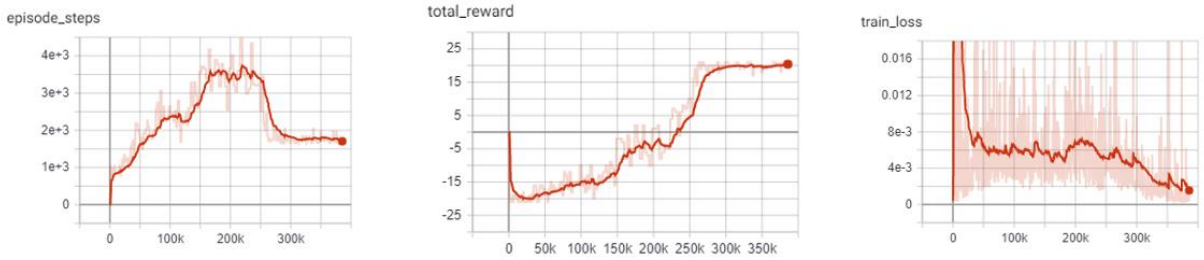
In order to update the basic DQN to a Noisy DQN we need to do the following

Noisy DQN Results

The results below improved stability and faster performance growth.

Noisy DQN baseline: Pong

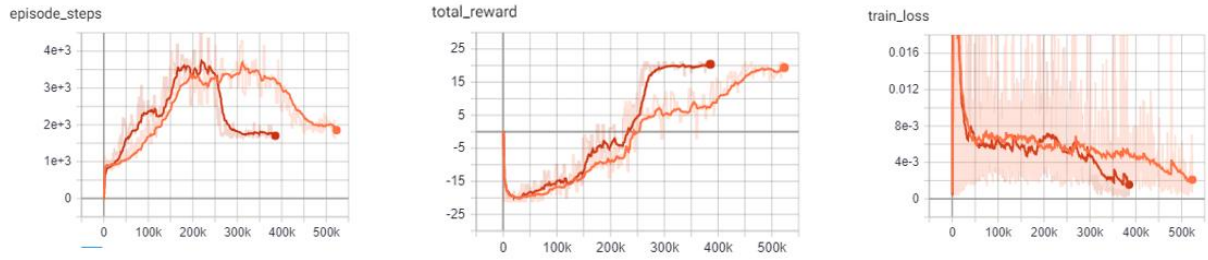
Similar to the other improvements, the average score of the agent reaches positive numbers around the 250k mark and steadily increases till convergence.



DQN vs Dueling DQN: Pong

In comparison to the base DQN, the Noisy DQN is more stable and is able to converge on an optimal policy much faster than the original. It seems that the replacement of the epsilon-greedy strategy with network noise provides a better form of exploration.

- Orange: DQN
- Red: Noisy DQN



Example:

```
from pl_bolts.models.rl import NoisyDQN
noisy_dqn = NoisyDQN("PongNoFrameskip-v4")
trainer = Trainer()
trainer.fit(noisy_dqn)
```

class pl_bolts.models.rl.noisy_dqn_model.**NoisyDQN**(*env*, *gpus=0*, *eps_start=1.0*, *eps_end=0.02*, *eps_last_frame=150000*, *sync_rate=1000*, *gamma=0.99*, *learning_rate=0.0001*, *batch_size=32*, *replay_size=100000*, *warm_start_size=10000*, *num_samples=500*, ***kwargs*)

Bases: *pl_bolts.models.rl.dqn_model.DQN*

PyTorch Lightning implementation of Noisy DQN

Paper authors: Meire Fortunato, Mohammad Gheshlaghi Azar, Bilal Piot, Jacob Menick, Ian Osband, Alex Graves, Vlad Mnih, Remi Munos, Demis Hassabis, Olivier Pietquin, Charles Blundell, Shane Legg

Model implemented by:

- *Donal Byrne <https://github.com/djbyrne>*

Example

```
>>> from pl_bolts.models.rl.n_step_dqn_model import NStepDQN
...
>>> model = NStepDQN("PongNoFrameskip-v4")
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

Parameters

- **env** (str) – gym environment tag
- **gpus** (int) – number of gpus being used
- **eps_start** (float) – starting value of epsilon for the epsilon-greedy exploration
- **eps_end** (float) – final value of epsilon for the epsilon-greedy exploration
- **eps_last_frame** (int) – the final frame in for the decrease of epsilon. At this frame espilon = eps_end
- **sync_rate** (int) – the number of iterations between syncing up the target network with the train network
- **gamma** (float) – discount factor
- **lr** – learning rate
- **batch_size** (int) – size of minibatch pulled from the DataLoader
- **replay_size** (int) – total capacity of the replay buffer

- **warm_start_size** (int) – how many random steps through the environment to be carried out at the start of
- **to fill the buffer with a starting point** (*training*) –
- **sample_len** – the number of samples to pull from the dataset iterator and feed to the DataLoader

---

Note: Currently only supports CPU and single GPU training with *distributed_backend=dp*

PyTorch Lightning implementation of DQN

Paper authors: Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, Martin Riedmiller.

Model implemented by:

- *Donal Byrne <https://github.com/djbyrne>*

Example

```
>>> from pl_bolts.models.rl.dqn_model import DQN
...
>>> model = DQN("PongNoFrameskip-v4")
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

Parameters

- **env** (str) – gym environment tag
- **gpus** (int) – number of gpus being used
- **eps_start** (float) – starting value of epsilon for the epsilon-greedy exploration
- **eps_end** (float) – final value of epsilon for the epsilon-greedy exploration
- **eps_last_frame** (int) – the final frame in for the decrease of epsilon. At this frame espilon = eps_end
- **sync_rate** (int) – the number of iterations between syncing up the target network with the train network
- **gamma** (float) – discount factor
- **learning_rate** (float) – learning rate
- **batch_size** (int) – size of minibatch pulled from the DataLoader
- **replay_size** (int) – total capacity of the replay buffer
- **warm_start_size** (int) – how many random steps through the environment to be carried out at the start of training to fill the buffer with a starting point
- **num_samples** (int) – the number of samples to pull from the dataset iterator and feed to the DataLoader

Note:

This example is based on: https://github.com/PacktPublishing/Deep-Reinforcement-Learning-Hands-On-Second-Edition/blob/master/Chapter06/02_dqn_pong.py

Note: Currently only supports CPU and single GPU training with *distributed_backend=dp*

build_networks()
Initializes the Noisy DQN train and target networks

Return type None

on_train_start()
Set the agents epsilon to 0 as the exploration comes from the network

Return type None

training_step(*batch*, _)
Carries out a single step through the environment to update the replay buffer. Then calculates loss based on the minibatch recieved

Parameters

- **batch** (Tuple[Tensor, Tensor]) – current mini batch of replay data
- _ – batch number, not used

Return type OrderedDict

Returns Training loss and log metrics

17.2.5 N-Step DQN

N-Step DQN model introduced in Learning to Predict by the Methods of Temporal Differences Paper authors: Richard S. Sutton

Original implementation by: Donal Byrne

N Step DQN was introduced in Learning to Predict by the Methods of Temporal Differences. This method improves upon the original DQN by updating our Q values with the expected reward from multiple steps in the future as opposed to the expected reward from the immediate next state. When getting the Q values for a state action pair using a single step which looks like this

$$Q(s_t, a_t) = r_t + \gamma \max_a Q(s_{t+1}, a_{t+1})$$

but because the Q function is recursive we can continue to roll this out into multiple steps, looking at the expected return for each step into the future.

$$Q(s_t, a_t) = r_t + \gamma r_{t+1} + \gamma^2 \max_{a'} Q(s_{t+2}, a')$$

The above example shows a 2-Step look ahead, but this could be rolled out to the end of the episode, which is just Monte Carlo learning. Although we could just do a monte carlo update and look forward to the end of the episode, it wouldn't be a good idea. Every time we take another step into the future, we are basing our approximation off our current policy. For a large portion of training, our policy is going to be less than optimal. For example, at the start of training, our policy will be in a state of high exploration, and will be little better than random.

Note: For each rollout step you must scale the discount factor accordingly by the number of steps. As you can see from the equation above, the second gamma value is to the power of 2. If we rolled this out one step further, we would use gamma to the power of 3 and so.

So if we are aproximating future rewards off a bad policy, chances are those approximations are going to be pretty bad and every time we unroll our update equation, the worse it will get. The fact that we are using an off policy method like DQN with a large replay buffer will make this even worse, as there is a high chance that we will be training on experiences using an old policy that was worse than our current policy.

So we need to strike a balance between looking far enough ahead to improve the convergence of our agent, but not so far that are updates become unstable. In general, small values of 2-4 work best.

N-Step Benefits

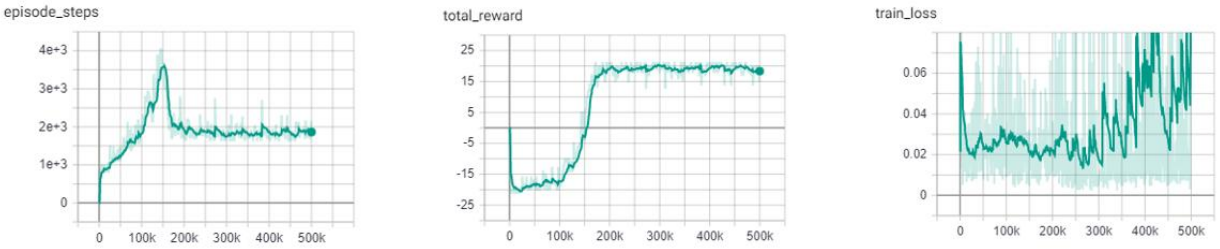
- Multi-Step learning is capable of learning faster than typical 1 step learning methods.
- **Note that this method introduces a new hyperparameter n. Although n=4 is generally a good starting point and provides** good results across the board.

N-Step Results

As expected, the N-Step DQN converges much faster than the standard DQN, however it also adds more instability to the loss of the agent. This can be seen in the following experiments.

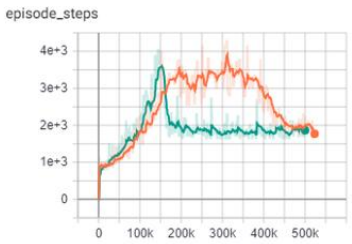
N-Step DQN: Pong

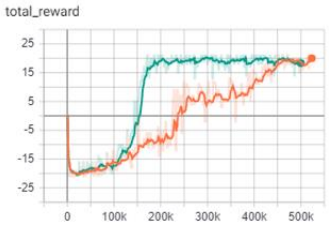
The N-Step DQN shows the greatest increase in performance with respect to the other DQN variations. After less than 150k steps the agent begins to consistently win games and achieves the top score after ~170K steps. This is reflected in the sharp peak of the total episode steps and of course, the total episode rewards.

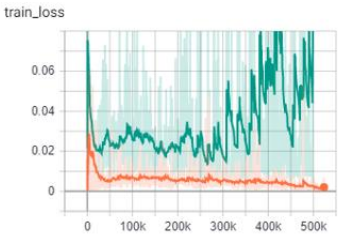


DQN vs N-Step DQN: Pong

This improvement is shown in stark contrast to the base DQN, which only begins to win games after 250k steps and requires over twice as many steps (450k) as the N-Step agent to achieve the high score of 21. One important thing to notice is the large increase in the loss of the N-Step agent. This is expected as the agent is building its expected reward off approximations of the future states. The large the size of N, the greater the instability. Previous literature, listed below, shows the best results for the Pong environment with an N step between 3-5. For these experiments I opted with an N step of 4.







Example:

```
from pl_bolts.models.rl import NStepDQN
n_step_dqn = NStepDQN("PongNoFrameskip-v4")
trainer = Trainer()
trainer.fit(n_step_dqn)
```

class pl_bolts.models.rl.n_step_dqn_model.**NStepDQN**(*env*, *gpus=0*, *eps_start=1.0*, *eps_end=0.02*, *eps_last_frame=150000*, *sync_rate=1000*, *gamma=0.99*, *learning_rate=0.0001*, *batch_size=32*, *replay_size=100000*, *warm_start_size=10000*, *num_samples=500*, *n_steps=4*, ***kwargs*)

Bases: *pl_bolts.models.rl.dqn_model.DQN*

NStep DQN Model

PyTorch Lightning implementation of N-Step DQN

Paper authors: Richard Sutton

Model implemented by:

- *Donal Byrne <https://github.com/djbyrne>*

Example

```
>>> from pl_bolts.models.rl.dqn_model import DQN
...
>>> model = NStepDQN("PongNoFrameskip-v4")
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

Parameters

- **env** (str) – gym environment tag
- **gpus** (int) – number of gpus being used
- **eps_start** (float) – starting value of epsilon for the epsilon-greedy exploration
- **eps_end** (float) – final value of epsilon for the epsilon-greedy exploration

- **eps_last_frame** (int) – the final frame in for the decrease of epsilon. At this frame espilon = eps_end
- **sync_rate** (int) – the number of iterations between syncing up the target network with the train network
- **gamma** (float) – discount factor
- **learning_rate** (float) – learning rate
- **batch_size** (int) – size of minibatch pulled from the DataLoader
- **replay_size** (int) – total capacity of the replay buffer
- **warm_start_size** (int) – how many random steps through the environment to be carried out at the start of training to fill the buffer with a starting point
- **num_samples** (int) – the number of samples to pull from the dataset iterator and feed to the DataLoader
- **n_steps** – number of steps to approximate and use in the bellman update

Note: Currently only supports CPU and single GPU training with *distributed_backend=dp*

17.2.6 Prioritized Experience Replay DQN

Double DQN model introduced in Prioritized Experience Replay Paper authors: Tom Schaul, John Quan, Ioannis Antonoglou, David Silver

Original implementation by: Donal Byrne

The standard DQN uses a buffer to break up the correlation between experiences and uniform random samples for each batch. Instead of just randomly sampling from the buffer prioritized experience replay (PER) prioritizes these samples based on training loss. This concept was introduced in the paper Prioritized Experience Replay

Essentially we want to train more on the samples that sunrise the agent.

The priority of each sample is defined below where

$$P(i) = P_i^{\alpha} / \sum_k P_k^{\alpha}$$

where pi is the priority of the ith sample in the buffer and is the number that shows how much emphasis we give to the priority. If = 0 , our sampling will become uniform as in the classic DQN method. Larger values for put more stress on samples with higher priority

Its important that new samples are set to the highest priority so that they are sampled soon. This however introduces bias to new samples in our dataset. In order to compensate for this bias, the value of the weight is defined as

$$w_i = (N.P(i))^{-\beta}$$

Where beta is a hyper parameter between 0-1. When beta is 1 the bias is fully compensated. However authors noted that in practice it is better to start beta with a small value near 0 and slowly increase it to 1.

PER Benefits

- **The benefits of this technique are that the agent sees more samples that it struggled with and gets more** chances to improve upon it.

Memory Buffer

First step is to replace the standard experience replay buffer with the prioritized experience replay buffer. This is pretty large (100+ lines) so I wont go through it here. There are two buffers implemented. The first is a naive list based buffer found in memory.PERBuffer and the second is more efficient buffer using a Sum Tree datastructure.

The list based version is simpler, but has a sample complexity of O(N). The Sum Tree in comparison has a complexity of O(1) for sampling and O(logN) for updating priorities.

Update loss function

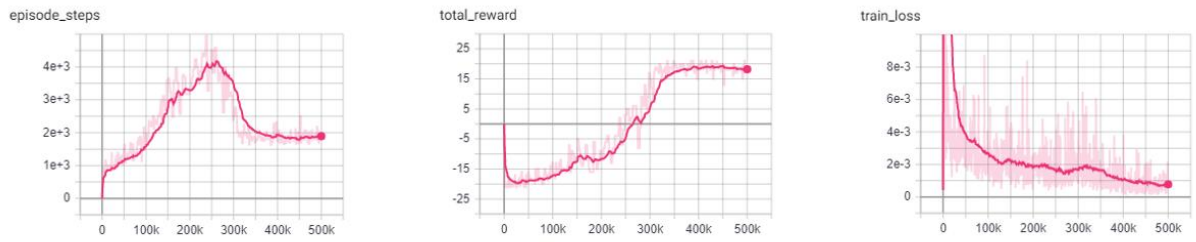
The next thing we do is to use the sample weights that we get from PER. Add the following code to the end of the loss function. This applies the weights of our sample to the batch loss. Then we return the mean loss and weighted loss for each datum, with the addition of a small epsilon value.

PER Results

The results below show improved stability and faster performance growth.

PER DQN: Pong

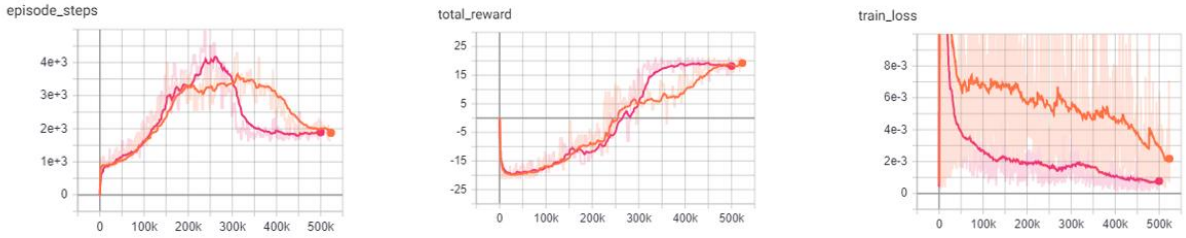
Similar to the other improvements, we see that PER improves the stability of the agents training and shows to converged on an optimal policy faster.



DQN vs PER DQN: Pong

In comparison to the base DQN, the PER DQN does show improved stability and performance. As expected, the loss of the PER DQN is siginificantly lower. This is the main objective of PER by focusing on experiences with high loss.

It is important to note that loss is not the only metric we should be looking at. Although the agent may have very low loss during training, it may still perform poorly due to lack of exploration.



- Orange: DQN

- Pink: PER DQN

Example:

```
from pl_bolts.models.rl import PERDQN
per_dqn = PERDQN("PongNoFrameskip-v4")
trainer = Trainer()
trainer.fit(per_dqn)
```

class pl_bolts.models.rl.per_dqn_model.**PERDQN**(*env*, *gpus=0*, *eps_start=1.0*, *eps_end=0.02*, *eps_last_frame=150000*, *sync_rate=1000*, *gamma=0.99*, *learning_rate=0.0001*, *batch_size=32*, *replay_size=100000*, *warm_start_size=10000*, *num_samples=500*, ***kwargs*)

Bases: *pl_bolts.models.rl.dqn_model.DQN*

PyTorch Lightning implementation of DQN With Prioritized Experience Replay

Paper authors: Tom Schaul, John Quan, Ioannis Antonoglou, David Silver

Model implemented by:

- *Donal Byrne <https://github.com/djbyrne>*

Example

```
>>> from pl_bolts.models.rl.per_dqn_model import PERDQN
...
>>> model = PERDQN("PongNoFrameskip-v4")
```

Train:

```
trainer = Trainer()
trainer.fit(model)

Args:
    env: gym environment tag
    gpus: number of gpus being used
    eps_start: starting value of epsilon for the epsilon-greedy exploration
    eps_end: final value of epsilon for the epsilon-greedy exploration
    eps_last_frame: the final frame in for the decrease of epsilon. At this frame␣
↪espilon = eps_end
    sync_rate: the number of iterations between syncing up the target network␣
↪with the train network
    gamma: discount factor
    learning_rate: learning rate
    batch_size: size of minibatch pulled from the DataLoader
    replay_size: total capacity of the replay buffer
    warm_start_size: how many random steps through the environment to be carried␣
↪out at the start of
        training to fill the buffer with a starting point
    num_samples: the number of samples to pull from the dataset iterator and feed␣
↪to the DataLoader

.. note::
    This example is based on:
```

(continues on next page)

(continued from previous page)

```
    https://github.com/PacktPublishing/Deep-Reinforcement-Learning-Hands-On-
↪Second-Edition             /blob/master/Chapter08/05_dqn_prio_replay.py

.. note:: Currently only supports CPU and single GPU training with `distributed_
↪backend=dp`
```

PyTorch Lightning implementation of DQN

Paper authors: Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, Martin Riedmiller.

Model implemented by:

- *Donal Byrne <https://github.com/djbyrne>*

Example

```
>>> from pl_bolts.models.rl.dqn_model import DQN
...
>>> model = DQN("PongNoFrameskip-v4")
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

Parameters

- **env** (str) – gym environment tag
- **gpus** (int) – number of gpus being used
- **eps_start** (float) – starting value of epsilon for the epsilon-greedy exploration
- **eps_end** (float) – final value of epsilon for the epsilon-greedy exploration
- **eps_last_frame** (int) – the final frame in for the decrease of epsilon. At this frame espilon = eps_end
- **sync_rate** (int) – the number of iterations between syncing up the target network with the train network
- **gamma** (float) – discount factor
- **learning_rate** (float) – learning rate
- **batch_size** (int) – size of minibatch pulled from the DataLoader
- **replay_size** (int) – total capacity of the replay buffer
- **warm_start_size** (int) – how many random steps through the environment to be carried out at the start of training to fill the buffer with a starting point
- **num_samples** (int) – the number of samples to pull from the dataset iterator and feed to the DataLoader

---

Note:

This example is based on: https://github.com/PacktPublishing/Deep-Reinforcement-Learning-Hands-On-Second-Edition/blob/master/Chapter06/02_dqn_pong.py

---

Note: Currently only supports CPU and single GPU training with *distributed_backend=dp*

---

prepare_data()
Initialize the Replay Buffer dataset used for retrieving experiences

Return type None

training_step(*batch*, *_*)
Carries out a single step through the environment to update the replay buffer. Then calculates loss based on the minibatch recieved

Parameters

- **batch** – current mini batch of replay data
- _ – batch number, not used

Return type OrderedDict

Returns Training loss and log metrics

17.3 Policy Gradient Models

The following models are based on Policy gradient

17.3.1 REINFORCE

REINFORCE model introduced in Policy Gradient Methods For Reinforcement Learning With Function Approximation Paper authors: Richard S. Sutton, David McAllester, Satinder Singh, Yishay Mansour

Original implementation by: Donal Byrne

Example:

```
from pl_bolts.models.rl import Reinforce
reinforce = Reinforce("CartPole-v0")
trainer = Trainer()
trainer.fit(reinforce)
```

class pl_bolts.models.rl.reinforce_model.**Reinforce**(*env*, *gamma=0.99*, *lr=0.0001*, *batch_size=32*, *batch_episodes=4*, ***kwargs*)

Bases: pytorch_lightning.LightningModule

Basic REINFORCE Policy Model

PyTorch Lightning implementation of REINFORCE

Paper authors: Richard S. Sutton, David McAllester, Satinder Singh, Yishay Mansour

Model implemented by:

- *Donal Byrne <https://github.com/djbyrne>*

Example

```
>>> from pl_bolts.models.rl.reinforce_model import Reinforce
...
>>> model = Reinforce("PongNoFrameskip-v4")
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

Parameters

- **env** (str) – gym environment tag
- **gamma** (float) – discount factor
- **lr** (float) – learning rate
- **batch_size** (int) – size of minibatch pulled from the DataLoader
- **batch_episodes** (int) – how many episodes to rollout for each batch of training

---

Note:

This example is based on: https://github.com/PacktPublishing/Deep-Reinforcement-Learning-Hands-On-Second-Edition/blob/master/Chapter11/02_cartpole_reinforce.py

---

Note: Currently only supports CPU and single GPU training with *distributed_backend=dp*

_dataloader()
Initialize the Replay Buffer dataset used for retrieving experiences

Return type DataLoader

static add_model_specific_args(*arg_parser*)
Adds arguments for DQN model

Note: these params are fine tuned for Pong env

Parameters **arg_parser** – the current argument parser to add to

Return type ArgumentParser

Returns arg_parser with model specific cargs added

build_networks()
Initializes the DQN train and target networks

Return type None

calc_qvals(*rewards*)
Takes in the rewards for each batched episode and returns list of qvals for each batched episode

Parameters **rewards** (List[List]) – list of rewards for each episodes in the batch

Return type List[List]

Returns List of qvals for each episodes

configure_optimizers()
Initialize Adam optimizer

Return type List[Optimizer]

static flatten_batch(*batch_actions*, *batch_qvals*, *batch_rewards*, *batch_states*)
Takes in the outputs of the processed batch and flattens the several episodes into a single tensor for each batched output

Parameters

- **batch_actions** (List[List[Tensor]]) – actions taken in each batch episodes
- **batch_qvals** (List[List[Tensor]]) – Q vals for each batch episode
- **batch_rewards** (List[List[Tensor]]) – reward for each batch episode
- **batch_states** (List[List[Tuple[Tensor, Tensor]]]) – states for each batch episodes

Return type Tuple[Tensor, Tensor, Tensor, Tensor]

Returns The input batched results flattend into a single tensor

forward(*x*)
Passes in a state x through the network and gets the q_values of each action as an output

Parameters **x** (Tensor) – environment state

Return type Tensor

Returns q values

get_device(*batch*)
Retrieve device currently being used by minibatch

Return type str

loss(*batch_qvals*, *batch_states*, *batch_actions*)
Calculates the mse loss using a batch of states, actions and Q values from several episodes. These have all been flattend into a single tensor.

Parameters

- **batch_qvals** (List[Tensor]) – current mini batch of q values
- **batch_actions** (List[Tensor]) – current batch of actions
- **batch_states** (List[Tensor]) – current batch of states

Return type Tensor

Returns loss

process_batch(*batch*)
Takes in a batch of episodes and retrieves the q vals, the states and the actions for the batch

Parameters **batch** (List[List[*Experience*]]) – list of episodes, each containing a list of Experiences

Return type Tuple[List[Tensor], List[Tensor], List[Tensor]]

Returns q_vals, states and actions used for calculating the loss

train_dataloader()
Get train loader

Return type DataLoader

training_step(*batch*, *_*)

Carries out a single step through the environment to update the replay buffer. Then calculates loss based on the minibatch recieved

Parameters

- **batch** (Tuple[Tensor, Tensor]) – current mini batch of replay data
- *_* – batch number, not used

Return type OrderedDict

Returns Training loss and log metrics

17.3.2 Vanilla Policy Gradient

Vanilla Policy Gradient model introduced in Policy Gradient Methods For Reinforcement Learning With Function Approximation Paper authors: Richard S. Sutton, David McAllester, Satinder Singh, Yishay Mansour

Original implementation by: Donal Byrne

Example:

```
from pl_bolts.models.rl import PolicyGradient
vpg = PolicyGradient("CartPole-v0")
trainer = Trainer()
trainer.fit(vpg)
```

class pl_bolts.models.rl.vanilla_policy_gradient_model.**PolicyGradient**(*env*, *gamma=0.99*, *lr=0.0001*, *batch_size=32*, *entropy_beta=0.01*, *batch_episodes=4*, **args*, ***kwargs*)

Bases: pytorch_lightning.LightningModule

Vanilla Policy Gradient Model

PyTorch Lightning implementation of Vanilla Policy Gradient

Paper authors: Richard S. Sutton, David McAllester, Satinder Singh, Yishay Mansour

Model implemented by:

- *Donal Byrne <https://github.com/djbyrne>*

Example

```
>>> from pl_bolts.models.rl.vanilla_policy_gradient_model import PolicyGradient
...
>>> model = PolicyGradient("PongNoFrameskip-v4")
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

Parameters

- **env** (str) – gym environment tag
- **gamma** (float) – discount factor
- **lr** (float) – learning rate
- **batch_size** (int) – size of minibatch pulled from the DataLoader
- **batch_episodes** (int) – how many episodes to rollout for each batch of training
- **entropy_beta** (float) – dictates the level of entropy per batch

---

Note:

This example is based on: https://github.com/PacktPublishing/Deep-Reinforcement-Learning-Hands-On-Second-Edition/blob/master/Chapter11/04_cartpole_pg.py

---

Note: Currently only supports CPU and single GPU training with *distributed_backend=dp*

_dataloader()
Initialize the Replay Buffer dataset used for retrieving experiences

Return type DataLoader

static add_model_specific_args(*arg_parser*)
Adds arguments for DQN model

Note: these params are fine tuned for Pong env

Parameters **parent** –

Return type ArgumentParser

build_networks()
Initializes the DQN train and target networks

Return type None

calc_entropy_loss(*log_prob*, *logits*)
Calculates the entropy to be added to the loss function :type _sphinx_paramlinks_pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradient.calc_entropy_loss.log_prob: Tensor :param _sphinx_paramlinks_pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradient.calc_entropy_loss.l log probabilities for each action :type _sphinx_paramlinks_pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradient Tensor :param _sphinx_paramlinks_pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradient.calc_entropy_loss.l the raw outputs of the network

Return type Tensor

Returns entropy penalty for each state

static calc_policy_loss(*batch_actions*, *batch_qvals*, *batch_states*, *logits*)
Calculate the policy loss give the batch outputs and logits :type _sphinx_paramlinks_pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradient.calc_policy_loss.batch_actions: Tensor :param _sphinx_paramlinks_pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradient.calc_policy_loss.ba actions from batched episodes :type _sphinx_paramlinks_pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradient.c Tensor :param _sphinx_paramlinks_pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradient.calc_policy_loss.ba Q values from batched episodes :type _sphinx_paramlinks_pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradien Tensor :param _sphinx_paramlinks_pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradient.calc_policy_loss.ba states from batched episodes :type _sphinx_paramlinks_pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradient.ca Tensor :param _sphinx_paramlinks_pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradient.calc_policy_loss.lo raw output of the network given the batch_states

Return type Tuple[List, Tensor]

Returns policy loss

calc_qvals(*rewards*)
Takes in the rewards for each batched episode and returns list of qvals for each batched episode

Parameters **rewards** (List[Tensor]) – list of rewards for each episodes in the batch

Return type List[Tensor]

Returns List of qvals for each episodes

configure_optimizers()
Initialize Adam optimizer

Return type List[Optimizer]

static flatten_batch(*batch_actions*, *batch_qvals*, *batch_rewards*, *batch_states*)
Takes in the outputs of the processed batch and flattens the several episodes into a single tensor for each batched output

Parameters

- **batch_actions** (List[List[Tensor]]) – actions taken in each batch episodes
- **batch_qvals** (List[List[Tensor]]) – Q vals for each batch episode
- **batch_rewards** (List[List[Tensor]]) – reward for each batch episode
- **batch_states** (List[List[Tuple[Tensor, Tensor]]]) – states for each batch episodes

Return type Tuple[Tensor, Tensor, Tensor, Tensor]

Returns The input batched results flattend into a single tensor

forward(*x*)
Passes in a state x through the network and gets the q_values of each action as an output

Parameters **x** (Tensor) – environment state

Return type Tensor

Returns q values

get_device(*batch*)
Retrieve device currently being used by minibatch

Return type str

loss(*batch_qvals*, *batch_states*, *batch_actions*)
Calculates the mse loss using a batch of states, actions and Q values from several episodes. These have all been flattend into a single tensor.

Parameters

- **batch_qvals** (List[Tensor]) – current mini batch of q values
- **batch_actions** (List[Tensor]) – current batch of actions
- **batch_states** (List[Tensor]) – current batch of states

Return type Tensor

Returns loss

process_batch(*batch*)
Takes in a batch of episodes and retrieves the q vals, the states and the actions for the batch

Parameters **batch** (List[List[*Experience*]]) – list of episodes, each containing a list of Experiences

Return type Tuple[List[Tensor], List[Tensor], List[Tensor]]

Returns q_vals, states and actions used for calculating the loss

train_dataloader()
Get train loader

Return type DataLoader

training_step(*batch*, _)
Carries out a single step through the environment to update the replay buffer. Then calculates loss based on the minibatch recieved

Parameters

- **batch** (Tuple[Tensor, Tensor]) – current mini batch of replay data
- _ – batch number, not used

Return type OrderedDict

Returns Training loss and log metrics

CHAPTER

EIGHTEEN

SELF-SUPERVISED LEARNING

This bolts module houses a collection of all self-supervised learning models.

Self-supervised learning extracts representations of an input by solving a pretext task. In this package, we implement many of the current state-of-the-art self-supervised algorithms.

Self-supervised models are trained with unlabeled datasets

18.1 Use cases

Here are some use cases for the self-supervised package.

18.1.1 Extracting image features

The models in this module are trained unsupervised and thus can capture better image representations (features).

In this example, we'll load a resnet 18 which was pretrained on imagenet using CPC as the pretext task.

Example:

```
from pl_bolts.models.self_supervised import CPCV2

# load resnet18 pretrained using CPC on imagenet
model = CPCV2(pretrained='resnet18')
cpc_resnet18 = model.encoder
cpc_resnet18.freeze()

# it supports any torchvision resnet
model = CPCV2(pretrained='resnet50')
```

This means you can now extract image representations that were pretrained via unsupervised learning.

Example:

```
my_dataset = SomeDataset()
for batch in my_dataset:
    x, y = batch
    out = cpc_resnet18(x)
```

18.1.2 Train with unlabeled data

These models are perfect for training from scratch when you have a huge set of unlabeled images

```
from pl_bolts.models.self_supervised import SimCLR
from pl_bolts.models.self_supervised.simclr import SimCLREvalDataTransform,␣
↪SimCLRTrainDataTransform


train_dataset = MyDataset(transforms=SimCLRTrainDataTransform())
val_dataset = MyDataset(transforms=SimCLREvalDataTransform())

# simclr needs a lot of compute!
model = SimCLR()
trainer = Trainer(tpu_cores=128)
trainer.fit(
    model,
    DataLoader(train_dataset),
    DataLoader(val_dataset),
)
```

18.1.3 Research

Mix and match any part, or subclass to create your own new method

```
from pl_bolts.models.self_supervised import CPCV2
from pl_bolts.losses.self_supervised_learning import FeatureMapContrastiveTask

amdim_task = FeatureMapContrastiveTask(comparisons='01, 11, 02', bidirectional=True)
model = CPCV2(contrastive_task=amdim_task)
```

18.2 Contrastive Learning Models

Contrastive self-supervised learning (CSL) is a self-supervised learning approach where we generate representations of instances such that similar instances are near each other and far from dissimilar ones. This is often done by comparing triplets of positive, anchor and negative representations.

In this section, we list Lightning implementations of popular contrastive learning approaches.

18.2.1 AMDIM

class pl_bolts.models.self_supervised.**AMDIM**(*datamodule='cifar10'*, *encoder='amdim_encoder'*, *contrastive_task=torch.nn.Module*, *image_channels=3*, *image_height=32*, *encoder_feature_dim=320*, *embedding_fx_dim=1280*, *conv_block_depth=10*, *use_bn=False*, *tclip=20.0*, *learning_rate=0.0002*, *data_dir=''*, *num_classes=10*, *batch_size=200*, ***kwargs*)

Bases: pytorch_lightning.LightningModule

PyTorch Lightning implementation of Augmented Multiscale Deep InfoMax (AMDIM)

Paper authors: Philip Bachman, R Devon Hjelm, William Buchwalter.

Model implemented by: William Falcon

This code is adapted to Lightning using the original author repo (the original repo).

Example

```
>>> from pl_bolts.models.self_supervised import AMDIM
...
>>> model = AMDIM(encoder='resnet18')
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

Parameters

- **datamodule** (Union[str, *LightningDataModule*]) – A LightningDatamodule
- **encoder** (Union[str, Module, LightningModule]) – an encoder string or model
- **image_channels** (int) – 3
- **image_height** (int) – pixels
- **encoder_feature_dim** (int) – Called *ndf* in the paper, this is the representation size for the encoder.
- **embedding_fx_dim** (int) – Output dim of the embedding function (*nrkhs* in the paper) (Reproducing Kernel Hilbert Spaces).
- **conv_block_depth** (int) – Depth of each encoder block,
- **use_bn** (bool) – If true will use batchnorm.
- **tclip** (int) – soft clipping non-linearity to the scores after computing the regularization term and before computing the log-softmax. This is the 'second trick' used in the paper
- **learning_rate** (int) – The learning rate
- **data_dir** (str) – Where to store data
- **num_classes** (int) – How many classes in the dataset
- **batch_size** (int) – The batch size

18.2.2 CPC (V2)

class pl_bolts.models.self_supervised.**CPCV2**(*datamodule=None*, *encoder='cpc_encoder'*, *patch_size=8*, *patch_overlap=4*, *online_ft=True*, *task='cpc'*, *num_workers=4*, *learning_rate=0.0001*, *data_dir=''*, *batch_size=32*, *pretrained=None*, ***kwargs*)

Bases: pytorch_lightning.LightningModule

PyTorch Lightning implementation of Data-Efficient Image Recognition with Contrastive Predictive Coding

Paper authors: (Olivier J. Hénaff, Aravind Srinivas, Jeffrey De Fauw, Ali Razavi, Carl Doersch, S. M. Ali Eslami, Aaron van den Oord).

Model implemented by:

- William Falcon
- Tullie Murrell

Example

```
>>> from pl_bolts.models.self_supervised import CPCV2
...
>>> model = CPCV2()
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

CLI command:

```
# cifar10
python cpc_module.py --gpus 1

# imagenet
python cpc_module.py
    --gpus 8
    --dataset imagenet2012
    --data_dir /path/to/imagenet/
    --meta_dir /path/to/folder/with/meta.bin/
    --batch_size 32
```

Some uses:

```
# load resnet18 pretrained using CPC on imagenet
model = CPCV2(encoder='resnet18', pretrained=True)
resnet18 = model.encoder
renset18.freeze()

# it supportes any torchvision resnet
model = CPCV2(encoder='resnet50', pretrained=True)

# use it as a feature extractor
x = torch.rand(2, 3, 224, 224)
out = model(x)
```

Parameters

- **datamodule** (Optional[*LightningDataModule*]) – A Datamodule (optional). Otherwise set the dataloaders directly
- **encoder** (Union[str, Module, LightningModule]) – A string for any of the resnets in torchvision, or the original CPC encoder, or a custon nn.Module encoder
- **patch_size** (int) – How big to make the image patches
- **patch_overlap** (int) – How much overlap should each patch have.

- **online_ft** (int) – Enable a 1024-unit MLP to fine-tune online
- **task** (str) – Which self-supervised task to use ('cpc', 'amdim', etc...)
- **num_workers** (int) – num dataloader worksers
- **learning_rate** (int) – what learning rate to use
- **data_dir** (str) – where to store data
- **batch_size** (int) – batch size
- **pretrained** (Optional[str]) – If true, will use the weights pretrained (using CPC) on Imagenet

18.2.3 Moco (V2)

class pl_bolts.models.self_supervised.**MocoV2**(*base_encoder='resnet18'*, *emb_dim=128*, *num_negatives=65536*, *encoder_momentum=0.999*, *softmax_temperature=0.07*, *learning_rate=0.03*, *momentum=0.9*, *weight_decay=0.0001*, *datamodule=None*, *data_dir='./'*, *batch_size=256*, *use_mlp=False*, *num_workers=8*, **args*, ***kwargs*)

Bases: pytorch_lightning.LightningModule

PyTorch Lightning implementation of Moco

Paper authors: Xinlei Chen, Haoqi Fan, Ross Girshick, Kaiming He.

Code adapted from facebookresearch/moco to Lightning by:

- William Falcon

Example

```
>>> from pl_bolts.models.self_supervised import MocoV2
...
>>> model = MocoV2()
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

CLI command:

```
# cifar10
python moco2_module.py --gpus 1

# imagenet
python moco2_module.py
    --gpus 8
    --dataset imagenet2012
    --data_dir /path/to/imagenet/
    --meta_dir /path/to/folder/with/meta.bin/
    --batch_size 32
```

Parameters

- **base_encoder** (Union[str, Module]) – torchvision model name or torch.nn.Module
- **emb_dim** (int) – feature dimension (default: 128)
- **num_negatives** (int) – queue size; number of negative keys (default: 65536)
- **encoder_momentum** (float) – moco momentum of updating key encoder (default: 0.999)
- **softmax_temperature** (float) – softmax temperature (default: 0.07)
- **learning_rate** (float) – the learning rate
- **momentum** (float) – optimizer momentum
- **weight_decay** (float) – optimizer weight decay
- **datamodule** (Optional[*LightningDataModule*]) – the DataModule (train, val, test dataloaders)
- **data_dir** (str) – the directory to store data
- **batch_size** (int) – batch size
- **use_mlp** (bool) – add an mlp to the encoders
- **num_workers** (int) – workers for the loaders

_batch_shuffle_ddp(*x*)
Batch shuffle, for making use of BatchNorm. *** Only support DistributedDataParallel (DDP) model. ***

_batch_unshuffle_ddp(*x*, *idx_unshuffle*)
Undo batch shuffle. *** Only support DistributedDataParallel (DDP) model. ***

_momentum_update_key_encoder()
Momentum update of the key encoder

forward(*img_q*, *img_k*)

Input: im_q: a batch of query images im_k: a batch of key images

Output: logits, targets

init_encoders(*base_encoder*)
Override to add your own encoders

18.2.4 SimCLR

class pl_bolts.models.self_supervised.**SimCLR**(*datamodule=None*, *data_dir=''*, *learning_rate=6e-05*, *weight_decay=0.0005*, *input_height=32*, *batch_size=128*, *online_ft=False*, *num_workers=4*, *optimizer='lars'*, *lr_sched_step=30.0*, *lr_sched_gamma=0.5*, *lars_momentum=0.9*, *lars_eta=0.001*, *loss_temperature=0.5*, ***kwargs*)

Bases: pytorch_lightning.LightningModule

PyTorch Lightning implementation of SIMCLR

Paper authors: Ting Chen, Simon Kornblith, Mohammad Norouzi, Geoffrey Hinton.

Model implemented by:

- William Falcon
- Tullie Murrell

Example

```
>>> from pl_bolts.models.self_supervised import SimCLR
...
>>> model = SimCLR()
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

CLI command:

```
# cifar10
python simclr_module.py --gpus 1

# imagenet
python simclr_module.py
    --gpus 8
    --dataset imagenet2012
    --data_dir /path/to/imagenet/
    --meta_dir /path/to/folder/with/meta.bin/
    --batch_size 32
```

Parameters

- **datamodule** (Optional[*LightningDataModule*]) – The datamodule
- **data_dir** (str) – directory to store data
- **learning_rate** (float) – the learning rate
- **weight_decay** (float) – optimizer weight decay
- **input_height** (int) – image input height
- **batch_size** (int) – the batch size
- **online_ft** (bool) – whether to tune online or not
- **num_workers** (int) – number of workers
- **optimizer** (str) – optimizer name
- **lr_sched_step** (float) – step for learning rate scheduler
- **lr_sched_gamma** (float) – gamma for learning rate scheduler
- **lars_momentum** (float) – the mom param for lars optimizer
- **lars_eta** (float) – for lars optimizer
- **loss_temperature** (float) – float = 0.

CHAPTER

NINETEEN

SELF-SUPERVISED LEARNING TRANSFORMS

These transforms are used in various self-supervised learning approaches.

19.1 CPC transforms

Transforms used for CPC

19.1.1 CIFAR-10 Train (c)

class pl_bolts.models.self_supervised.cpc.transforms.**CPCTrainTransformsCIFAR10**(*patch_size=8*, *overlap=4*)

Bases: object

Transforms used for CPC:

Parameters

- **patch_size** – size of patches when cutting up the image into overlapping patches
- **overlap** – how much to overlap patches

Transforms:

```
random_flip
img_jitter
col_jitter
rnd_gray
transforms.ToTensor()
normalize
Patchify(patch_size=patch_size, overlap_size=patch_size // 2)
```

Example:

```
# in a regular dataset
CIFAR10(..., transforms=CPCTrainTransformsCIFAR10())

# in a DataModule
module = CIFAR10DataModule(PATH)
train_loader = module.train_dataloader(batch_size=32,␣
↪transforms=CPCTrainTransformsCIFAR10())
```

__call__(*inp*)
Call self as a function.

19.1.2 CIFAR-10 Eval (c)

class pl_bolts.models.self_supervised.cpc.transforms.**CPCEvalTransformsCIFAR10**(*patch_size=8*, *overlap=4*)

Bases: object

Transforms used for CPC:

Parameters

- **patch_size** – size of patches when cutting up the image into overlapping patches
- **overlap** – how much to overlap patches

Transforms:

```
random_flip
transforms.ToTensor()
normalize
Patchify(patch_size=patch_size, overlap_size=overlap)
```

Example:

```
# in a regular dataset
CIFAR10(..., transforms=CPCEvalTransformsCIFAR10())

# in a DataModule
module = CIFAR10DataModule(PATH)
train_loader = module.train_dataloader(batch_size=32,␣
↪transforms=CPCEvalTransformsCIFAR10())
```

__call__(*inp*)
Call self as a function.

19.1.3 Imagenet Train (c)

class pl_bolts.models.self_supervised.cpc.transforms.**CPCTrainTransformsImageNet128**(*patch_size=*
overlap=16)

Bases: object

Transforms used for CPC:

Parameters

- **patch_size** – size of patches when cutting up the image into overlapping patches
- **overlap** – how much to overlap patches

Transforms:

```
random_flip
transforms.ToTensor()
normalize
Patchify(patch_size=patch_size, overlap_size=patch_size // 2)
```

Example:

```
# in a regular dataset
Imagenet(..., transforms=CPCTrainTransformsImageNet128())

# in a DataModule
module = ImagenetDataModule(PATH)
train_loader = module.train_dataloader(batch_size=32,␣
↪transforms=CPCTrainTransformsImageNet128())
```

__call__(*inp*)
Call self as a function.

19.1.4 Imagenet Eval (c)

class pl_bolts.models.self_supervised.cpc.transforms.**CPCEvalTransformsImageNet128**(*patch_size=*, *overlap=16*)

Bases: object

Transforms used for CPC:

Parameters

- **patch_size** – size of patches when cutting up the image into overlapping patches
- **overlap** – how much to overlap patches

Transforms:

```
random_flip
transforms.ToTensor()
normalize
Patchify(patch_size=patch_size, overlap_size=patch_size // 2)
```

Example:

```
# in a regular dataset
Imagenet(..., transforms=CPCEvalTransformsImageNet128())

# in a DataModule
module = ImagenetDataModule(PATH)
train_loader = module.train_dataloader(batch_size=32,␣
↪transforms=CPCEvalTransformsImageNet128())
```

__call__(*inp*)
Call self as a function.

19.1.5 STL-10 Train (c)

class pl_bolts.models.self_supervised.cpc.transforms.**CPCTrainTransformsSTL10**(*patch_size=16*, *overlap=8*)

Bases: object

Transforms used for CPC:

Parameters

- **patch_size** – size of patches when cutting up the image into overlapping patches
- **overlap** – how much to overlap patches

Transforms:

```
random_flip
img_jitter
col_jitter
rnd_gray
transforms.ToTensor()
normalize
Patchify(patch_size=patch_size, overlap_size=patch_size // 2)
```

Example:

```
# in a regular dataset
STL10(..., transforms=CPCTrainTransformsSTL10())

# in a DataModule
module = STL10DataModule(PATH)
train_loader = module.train_dataloader(batch_size=32,␣
↪transforms=CPCTrainTransformsSTL10())
```

__call__(*inp*)

Call self as a function.

19.1.6 STL-10 Eval (c)

class pl_bolts.models.self_supervised.cpc.transforms.**CPCEvalTransformsSTL10**(*patch_size=16*, *overlap=8*)

Bases: object

Transforms used for CPC:

Parameters

- **patch_size** – size of patches when cutting up the image into overlapping patches
- **overlap** – how much to overlap patches

Transforms:

```
random_flip
transforms.ToTensor()
normalize
Patchify(patch_size=patch_size, overlap_size=patch_size // 2)
```

Example:

```
# in a regular dataset
STL10(..., transforms=CPCEvalTransformsSTL10())

# in a DataModule
module = STL10DataModule(PATH)
train_loader = module.train_dataloader(batch_size=32,␣
↪transforms=CPCEvalTransformsSTL10())
```

__call__(*inp*)
Call self as a function.

19.2 AMDIM transforms

Transforms used for AMDIM

19.2.1 CIFAR-10 Train (a)

class pl_bolts.models.self_supervised.amdim.transforms.**AMDIMTrainTransformsCIFAR10**
Bases: object

Transforms applied to AMDIM

Transforms:

```
img_jitter,
col_jitter,
rnd_gray,
transforms.ToTensor(),
normalize
```

Example:

```
x = torch.rand(5, 3, 32, 32)

transform = AMDIMTrainTransformsCIFAR10()
(view1, view2) = transform(x)
```

__call__(*inp*)
Call self as a function.

19.2.2 CIFAR-10 Eval (a)

class pl_bolts.models.self_supervised.amdim.transforms.**AMDIMEvalTransformsCIFAR10**
Bases: object

Transforms applied to AMDIM

Transforms:

```
transforms.ToTensor(),
normalize
```

Example:

```
x = torch.rand(5, 3, 32, 32)

transform = AMDIMEvalTransformsCIFAR10()
(view1, view2) = transform(x)
```

__call__(*inp*)
Call self as a function.

19.2.3 Imagenet Train (a)

class pl_bolts.models.self_supervised.amdim.transforms.**AMDIMTrainTransformsImageNet128**(*heigh*
Bases: object

Transforms applied to AMDIM

Transforms:

```
img_jitter,
col_jitter,
rnd_gray,
transforms.ToTensor(),
normalize
```

Example:

```
x = torch.rand(5, 3, 128, 128)

transform = AMDIMTrainTransformsSTL10()
(view1, view2) = transform(x)
```

__call__(*inp*)
Call self as a function.

19.2.4 Imagenet Eval (a)

class pl_bolts.models.self_supervised.amdim.transforms.**AMDIMEvalTransformsImageNet128**(*height*
Bases: object

Transforms applied to AMDIM

Transforms:

```
transforms.Resize(height + 6, interpolation=3),
transforms.CenterCrop(height),
transforms.ToTensor(),
normalize
```

Example:

```
x = torch.rand(5, 3, 128, 128)

transform = AMDIMEvalTransformsImageNet128()
view1 = transform(x)
```

__call__(*inp*)
Call self as a function.

19.2.5 STL-10 Train (a)

class pl_bolts.models.self_supervised.amdim.transforms.**AMDIMTrainTransformsSTL10**(*height=64*)
Bases: object

Transforms applied to AMDIM

Transforms:

```
img_jitter,
col_jitter,
rnd_gray,
transforms.ToTensor(),
normalize
```

Example:

```
x = torch.rand(5, 3, 64, 64)

transform = AMDIMTrainTransformsSTL10()
(view1, view2) = transform(x)
```

__call__(*inp*)
Call self as a function.

19.2.6 STL-10 Eval (a)

class pl_bolts.models.self_supervised.amdim.transforms.**AMDIMEvalTransformsSTL10**(*height=64*)
Bases: object

Transforms applied to AMDIM

Transforms:

```
transforms.Resize(height + 6, interpolation=3),
transforms.CenterCrop(height),
transforms.ToTensor(),
normalize
```

Example:

```
x = torch.rand(5, 3, 64, 64)

transform = AMDIMTrainTransformsSTL10()
view1 = transform(x)
```

__call__(*inp*)
Call self as a function.

19.3 MOCO V2 transforms

Transforms used for MOCO V2

19.3.1 CIFAR-10 Train (m2)

class pl_bolts.models.self_supervised.moco.transforms.**Moco2TrainCIFAR10Transforms**(*height=32*)
Bases: object

Moco 2 augmentation: https://arxiv.org/pdf/2003.04297.pdf

__call__(*inp*)
Call self as a function.

19.3.2 CIFAR-10 Eval (m2)

class pl_bolts.models.self_supervised.moco.transforms.**Moco2EvalCIFAR10Transforms**(*height=32*)
Bases: object

Moco 2 augmentation: https://arxiv.org/pdf/2003.04297.pdf

__call__(*inp*)
Call self as a function.

19.3.3 Imagenet Train (m2)

class pl_bolts.models.self_supervised.moco.transforms.**Moco2TrainSTL10Transforms**(*height=64*)
Bases: object

Moco 2 augmentation: https://arxiv.org/pdf/2003.04297.pdf

__call__(*inp*)
Call self as a function.

19.3.4 Imagenet Eval (m2)

class pl_bolts.models.self_supervised.moco.transforms.**Moco2EvalSTL10Transforms**(*height=64*)
Bases: object

Moco 2 augmentation: https://arxiv.org/pdf/2003.04297.pdf

__call__(*inp*)
Call self as a function.

19.3.5 STL-10 Train (m2)

class pl_bolts.models.self_supervised.moco.transforms.**Moco2TrainImagenetTransforms**(*height=128*

Bases: object

Moco 2 augmentation: https://arxiv.org/pdf/2003.04297.pdf

__call__(*inp*)

Call self as a function.

19.3.6 STL-10 Eval (m2)

class pl_bolts.models.self_supervised.moco.transforms.**Moco2EvalImagenetTransforms**(*height=128*)

Bases: object

Moco 2 augmentation: https://arxiv.org/pdf/2003.04297.pdf

__call__(*inp*)

Call self as a function.

19.4 SimCLR transforms

Transforms used for SimCLR

19.4.1 Train (sc)

class pl_bolts.models.self_supervised.simclr.simclr_transforms.**SimCLRTrainDataTransform**(*inp*
s=1

Bases: object

Transforms for SimCLR

Transform:

```
RandomResizedCrop(size=self.input_height)
RandomHorizontalFlip()
RandomApply([color_jitter], p=0.8)
RandomGrayscale(p=0.2)
GaussianBlur(kernel_size=int(0.1 * self.input_height))
transforms.ToTensor()
```

Example:

```
from pl_bolts.models.self_supervised.simclr.transforms import␣
↪SimCLRTrainDataTransform

transform = SimCLRTrainDataTransform(input_height=32)
x = sample()
(xi, xj) = transform(x)
```

__call__(*sample*)

Call self as a function.

19.4.2 Eval (sc)

class pl_bolts.models.self_supervised.simclr.simclr_transforms.**SimCLREvalDataTransform**(*input_height*, *s=1*)

Bases: object

Transforms for SimCLR

Transform:

```
Resize(input_height + 10, interpolation=3)
transforms.CenterCrop(input_height),
transforms.ToTensor()
```

Example:

```
from pl_bolts.models.self_supervised.simclr.transforms import␣
↪SimCLREvalDataTransform

transform = SimCLREvalDataTransform(input_height=32)
x = sample()
(xi, xj) = transform(x)
```

__call__(*sample*)

Call self as a function.

CHAPTER

TWENTY

SELF-SUPERVISED LEARNING

Collection of useful functions for self-supervised learning

20.1 Identity class

Example:

```
from pl_bolts.utils import Identity
```

class pl_bolts.utils.self_supervised.**Identity**

Bases: torch.nn.Module

An identity class to replace arbitrary layers in pretrained models

Example:

```
from pl_bolts.utils import Identity

model = resnet18()
model.fc = Identity()
```

20.2 SSL-ready resnets

Torchvision resnets with the fc layers removed and with the ability to return all feature maps instead of just the last one.

Example:

```
from pl_bolts.utils.self_supervised import torchvision_ssl_encoder

resnet = torchvision_ssl_encoder('resnet18', pretrained=False, return_all_feature_
↪maps=True)
x = torch.rand(3, 3, 32, 32)

feat_maps = resnet(x)
```

pl_bolts.utils.self_supervised.**torchvision_ssl_encoder**(*name*, *pretrained=False*, *return_all_feature_maps=False*)

CHAPTER

TWENTYONE

SEMI-SUPERVISED LEARNING

Collection of utilities for semi-supervised learning where some part of the data is labeled and the other part is not.

21.1 Balanced classes

Example:

```
from pl_bolts.utils.semi_supervised import balance_classes
```

pl_bolts.utils.semi_supervised.**balance_classes**(*X*, *Y*, *batch_size*)

Makes sure each batch has an equal amount of data from each class. Perfect balance

Parameters

- **X** (ndarray) – input features
- **Y** (list) – mixed labels (ints)
- **batch_size** (int) – the ultimate batch size

21.2 half labeled batches

Example:

```
from pl_bolts.utils.semi_supervised import balance_classes
```

pl_bolts.utils.semi_supervised.**generate_half_labeled_batches**(*smaller_set_X*, *smaller_set_Y*, *larger_set_X*, *larger_set_Y*, *batch_size*)

Given a labeled dataset and an unlabeled dataset, this function generates a joint pair where half the batches are labeled and the other half is not

CHAPTER

TWENTYTWO

SELF-SUPERVISED LEARNING CONTRASTIVE TASKS

This section implements popular contrastive learning tasks used in self-supervised learning.

22.1 FeatureMapContrastiveTask

This task compares sets of feature maps.

In general the feature map comparison pretext task uses triplets of features. Here are the abstract steps of comparison.

Generate multiple views of the same image

```
x1_view_1 = data_augmentation(x1)
x1_view_2 = data_augmentation(x1)
```

Use a different example to generate additional views (usually within the same batch or a pool of candidates)

```
x2_view_1 = data_augmentation(x2)
x2_view_2 = data_augmentation(x2)
```

Pick 3 views to compare, these are the anchor, positive and negative features

```
anchor = x1_view_1
positive = x1_view_2
negative = x2_view_1
```

Generate feature maps for each view

```
(a0, a1, a2) = encoder(anchor)
(p0, p1, p2) = encoder(positive)
```

Make a comparison for a set of feature maps

```
phi = some_score_function()

# the '01' comparison
score = phi(a0, p1)

# and can be bidirectional
score = phi(p0, a1)
```

In practice the contrastive task creates a BxB matrix where B is the batch size. The diagonals for set 1 of feature maps are the anchors, the diagonals of set 2 of the feature maps are the positives, the non-diagonals of set 1 are the negatives.

class pl_bolts.losses.self_supervised_learning.**FeatureMapContrastiveTask**(*comparisons='00, 11'*, *tclip=10.0*, *bidirectional=True*)

Bases: torch.nn.Module

Performs an anchor, positive negative pair comparison for each each tuple of feature maps passed.

```
# extract feature maps
pos_0, pos_1, pos_2 = encoder(x_pos)
anc_0, anc_1, anc_2 = encoder(x_anchor)

# compare only the 0th feature maps
task = FeatureMapContrastiveTask('00')
loss, regularizer = task((pos_0), (anc_0))

# compare (pos_0 to anc_1) and (pos_0, anc_2)
task = FeatureMapContrastiveTask('01, 02')
losses, regularizer = task((pos_0, pos_1, pos_2), (anc_0, anc_1, anc_2))
loss = losses.sum()

# compare (pos_1 vs a anc_random)
task = FeatureMapContrastiveTask('0r')
loss, regularizer = task((pos_0, pos_1, pos_2), (anc_0, anc_1, anc_2))
```

Parameters

- **comparisons** (str) – groupings of feature map indices to compare (zero indexed, 'r' means random) ex: '00, 1r'
- **tclip** (float) – stability clipping value
- **bidirectional** (bool) – if true, does the comparison both ways

```
# with bidirectional the comparisons are done both ways
task = FeatureMapContrastiveTask('01, 02')

# will compare the following:
# 01: (pos_0, anc_1), (anc_0, pos_1)
# 02: (pos_0, anc_2), (anc_0, pos_2)
```

forward(*anchor_maps*, *positive_maps*)

Takes in a set of tuples, each tuple has two feature maps with all matching dimensions

Example

```
>>> import torch
>>> from pytorch_lightning import seed_everything
>>> seed_everything(0)
0
>>> a1 = torch.rand(3, 5, 2, 2)
>>> a2 = torch.rand(3, 5, 2, 2)
>>> b1 = torch.rand(3, 5, 2, 2)
>>> b2 = torch.rand(3, 5, 2, 2)
```

(continues on next page)

(continued from previous page)

```
...
>>> task = FeatureMapContrastiveTask('01, 11')
...
>>> losses, regularizer = task((a1, a2), (b1, b2))
>>> losses
tensor([2.2351, 2.1902])
>>> regularizer
tensor(0.0324)
```

static parse_map_indexes(*comparisons*)

Example:

```
>>> FeatureMapContrastiveTask.parse_map_indexes('11')
[(1, 1)]
>>> FeatureMapContrastiveTask.parse_map_indexes('11,59')
[(1, 1), (5, 9)]
>>> FeatureMapContrastiveTask.parse_map_indexes('11,59, 2r')
[(1, 1), (5, 9), (2, -1)]
```

22.2 Context prediction tasks

The following tasks aim to predict a target using a context representation.

22.2.1 CPCContrastiveTask

This is the predictive task from CPC (v2).

```
task = CPCTask(num_input_channels=32)

# (batch, channels, rows, cols)
# this should be thought of as 49 feature vectors, each with 32 dims
Z = torch.random.rand(3, 32, 7, 7)

loss = task(Z)
```

class pl_bolts.losses.self_supervised_learning.**CPCTask**(*num_input_channels*, *target_dim=64*, *embed_scale=0.1*)

Bases: torch.nn.Module

Loss used in CPC

CHAPTER

TWENTYTHREE

INDICES AND TABLES

- genindex
- modindex
- search

Logo

23.1 PyTorch Lightning Bolts

Pretrained SOTA Deep Learning models, callbacks and more for research and production with PyTorch Lightning and PyTorch

---

23.1.1 Trending contributors

23.1.2 Continuous Integration

| System / PyTorch ver. | 1.4 (min. req.) | 1.5 (latest) | | :—: | :—: | :—: | | Linux py3.6 / py3.7 / py3.8 | CI testing | CI testing | | OSX py3.6 / py3.7 / py3.8 | CI testing | CI testing | | Windows py3.6 / py3.7 / py3.8 | wip | wip |

23.1.3 Install

```
pip install pytorch-lightning-bolts
```

23.1.4 Docs

- master
- stable
- 0.1.0

23.1.5 What is Bolts

Bolts is a Deep learning research and production toolbox of:

- SOTA pretrained models.
- Model components.
- Callbacks.
- Losses.
- Datasets.

23.1.6 Main Goals of Bolts

The main goal of Bolts is to enable rapid model idea iteration.

Example 1: Finetuning on data

```
from pl_bolts.models.self_supervised import SimCLR
from pl_bolts.models.self_supervised.simclr.transforms import␣
↪SimCLRTrainDataTransform, SimCLREvalDataTransform
import pytorch_lightning as pl

# data
train_data = DataLoader(MyDataset(transforms=SimCLRTrainDataTransform(input_
↪height=32)))
val_data = DataLoader(MyDataset(transforms=SimCLREvalDataTransform(input_height=32)))

# model
model = SimCLR(pretrained='imagenet2012')

# train!
trainer = pl.Trainer(gpus=8)
trainer.fit(model, train_data, val_data)
```

Example 2: Subclass and ideate

```
from pl_bolts.models import ImageGPT
from pl_bolts.self_supervised import SimCLR

class VideoGPT(ImageGPT):

    def training_step(self, batch, batch_idx):
        x, y = batch
        x = _shape_input(x)

        logits = self.gpt(x)
        simclr_features = self.simclr(x)

        # ------------------
        # do something new with GPT logits + simclr_features
        # ------------------
```

(continues on next page)

(continued from previous page)

```
        loss = self.criterion(logits.view(-1, logits.size(-1)), x.view(-1).long())

        logs = {"loss": loss}
        return {"loss": loss, "log": logs}
```

23.1.7 Who is Bolts for?

- Corporate production teams
- Professional researchers
- Ph.D. students
- Linear + Logistic regression heroes

23.1.8 I don't need deep learning

Great! We have LinearRegression and LogisticRegression implementations with numpy and sklearn bridges for datasets! But our implementations work on multiple GPUs, TPUs and scale dramatically...

Check out our Linear Regression on TPU demo

```
from pl_bolts.models.regression import LinearRegression
from pl_bolts.datamodules import SklearnDataModule

# sklearn dataset
X, y = load_boston(return_X_y=True)
loaders = SklearnDataModule(X, y)

model = LinearRegression(input_dim=13)
trainer = pl.Trainer(num_tpu_cores=1)
trainer.fit(model, loaders.train_dataloader(), loaders.val_dataloader())
trainer.test(test_dataloaders=loaders.test_dataloader())
```

23.1.9 Is this another model zoo?

No!

Bolts is unique because models are implemented using PyTorch Lightning and structured so that they can be easily subclassed and iterated on.

For example, you can override the elbo loss of a VAE, or the generator_step of a GAN to quickly try out a new idea. The best part is that all the models are benchmarked so you won't waste time trying to "reproduce" or find the bugs with your implementation.

23.1.10 Team

Bolts is supported by the PyTorch Lightning team and the PyTorch Lightning community!

23.2 pl_bolts.callbacks package

Collection of PyTorchLightning callbacks

23.2.1 Submodules

pl_bolts.callbacks.printing module

class pl_bolts.callbacks.printing.**PrintTableMetricsCallback**
Bases: pytorch_lightning.callbacks.Callback

Prints a table with the metrics in columns on every epoch end

Example:

```
from pl_bolts.callbacks import PrintTableMetricsCallback

callback = PrintTableMetricsCallback()
```

pass into trainer like so:

```
trainer = pl.Trainer(callbacks=[callback])
trainer.fit(...)

# ------------------------------
# at the end of every epoch it will print
# ------------------------------

# loss|train_loss|val_loss|epoch
# ──────────────────────────────
# 2.2541470527648926|2.2541470527648926|2.2158432006835938|0
```

on_epoch_end(*trainer*, *pl_module*)

pl_bolts.callbacks.printing.**dicts_to_table**(*dicts*, *keys=None*, *pads=None*, *fcodes=None*, *convert_headers=None*, *header_names=None*, *skip_none_lines=False*, *replace_values=None*)

Generate ascii table from dictionary Taken from (https://stackoverflow.com/questions/40056747/print-a-list-of-dictionaries-in-table-form)

Parameters

- **dicts** (List[Dict]) – input dictionary list; empty lists make keys OR header_names mandatory
- **keys** (Optional[List[str]]) – order list of keys to generate columns for; no key/dict-key should suffix with '____' else adjust code-suffix
- **pads** (Optional[List[str]]) – indicate padding direction and size, eg <10 to right pad alias left-align
- **fcodes** (Optional[List[str]]) – formating codes for respective column type, eg .3f

- **convert_headers** (Optional[Dict[str, Callable]]) – apply converters(dict) on column keys k, eg timestamps
- **header_names** (Optional[List[str]]) – supply for custom column headers instead of keys
- **skip_none_lines** (bool) – skip line if contains None
- **replace_values** (Optional[Dict[str, Any]]) – specify per column keys k a map from seen value to new value; new value must comply with the columns fcode; CAUTION: modifies input (due speed)

Example

```
>>> a = {'a': 1, 'b': 2}
>>> b = {'a': 3, 'b': 4}
>>> print(dicts_to_table([a, b]))
a|b
—
1|2
3|4
```

pl_bolts.callbacks.variational module

class pl_bolts.callbacks.variational.**LatentDimInterpolator**(*interpolate_epoch_interval=20*, *range_start=-5*, *range_end=5*)

Bases: pytorch_lightning.callbacks.Callback

Interpolates the latent space for a model by setting all dims to zero and stepping through the first two dims increasing one unit at a time.

Default interpolates between [-5, 5] (-5, -4, -3, . . . , 3, 4, 5)

Example:

```
from pl_bolts.callbacks import LatentDimInterpolator

Trainer(callbacks=[LatentDimInterpolator()])
```

Parameters

- **interpolate_epoch_interval** –
- **range_start** – default -5
- **range_end** – default 5

interpolate_latent_space(*model*, *latent_dim*)

on_epoch_end(*trainer*, *pl_module*)

23.3 pl_bolts.datamodules package

23.3.1 Submodules

pl_bolts.datamodules.base_dataset module

class pl_bolts.datamodules.base_dataset.**LightDataset**(**args*, ***kwargs*)

Bases: abc.ABC, torch.utils.data.Dataset

_download_from_url(*base_url*, *data_folder*, *file_name*)

static _prepare_subset(*full_data*, *full_targets*, *num_samples*, *labels*)

Prepare a subset of a common dataset.

Return type Tuple[Tensor, Tensor]

DATASET_NAME = 'light'

cache_folder_name: str = None

property cached_folder_path

Return type str

data: torch.Tensor = None

dir_path: str = None

normalize: tuple = None

targets: torch.Tensor = None

pl_bolts.datamodules.cifar10_datamodule module

class pl_bolts.datamodules.cifar10_datamodule.**CIFAR10DataModule**(*data_dir*, *val_split=5000*, *num_workers=16*, **args*, ***kwargs*)

Bases: *pl_bolts.datamodules.lightning_datamodule.LightningDataModule*

Standard CIFAR10, train, val, test splits and transforms

Transforms:

```
mnist_transforms = transform_lib.Compose([
    transform_lib.ToTensor(),
    transforms.Normalize(
        mean=[x / 255.0 for x in [125.3, 123.0, 113.9]],
        std=[x / 255.0 for x in [63.0, 62.1, 66.7]]
    )
])
```

Example:

```
from pl_bolts.datamodules import CIFAR10DataModule

dm = CIFAR10DataModule(PATH)
model = LitModel(datamodule=dm)
```

Or you can set your own transforms

Example:

```
dm.train_transforms = ...
dm.test_transforms = ...
dm.val_transforms  = ...
```

Parameters

- **data_dir** – where to save/load the data
- **val_split** – how many of the training images to use for the validation split
- **num_workers** – how many workers to use for loading data

default_transforms()

prepare_data()
Saves CIFAR10 files to data_dir

test_dataloader(*batch_size*)
CIFAR10 test set uses the test split

Parameters

- **batch_size** – size of batch
- **transforms** – custom transforms

train_dataloader(*batch_size*)
CIFAR train set removes a subset to use for validation

Parameters **batch_size** – size of batch

val_dataloader(*batch_size*)
CIFAR10 val set uses a subset of the training set for validation

Parameters **batch_size** – size of batch

extra_args = {}

name: str = 'cifar10'

property num_classes
Return: 10

class pl_bolts.datamodules.cifar10_datamodule.**TinyCIFAR10DataModule**(*data_dir*, *val_split=50*, *num_workers=16*, *num_samples=100*, *labels=(1*, *5*, *8)*, **args*, ***kwargs*)

Bases: *pl_bolts.datamodules.cifar10_datamodule.CIFAR10DataModule*

Standard CIFAR10, train, val, test splits and transforms

Transforms:

```
mnist_transforms = transform_lib.Compose([
    transform_lib.ToTensor(),
    transforms.Normalize(mean=[x / 255.0 for x in [125.3, 123.0, 113.9]],
                         std=[x / 255.0 for x in [63.0, 62.1, 66.7]])
])
```

Example:

```
from pl_bolts.datamodules import CIFAR10DataModule

dm = CIFAR10DataModule(PATH)
model = LitModel(datamodule=dm)
```

Parameters

- **data_dir** (str) – where to save/load the data
- **val_split** (int) – how many of the training images to use for the validation split
- **num_workers** (int) – how many workers to use for loading data
- **num_samples** (int) – number of examples per selected class/label
- **labels** (Optional[Sequence]) – list selected CIFAR10 classes/labels

property num_classes

Return number of classes.

Return type int

pl_bolts.datamodules.cifar10_dataset module

class pl_bolts.datamodules.cifar10_dataset.**CIFAR10**(*data_dir='.'*, *train=True*, *transform=None*, *download=True*)

Bases: *pl_bolts.datamodules.base_dataset.LightDataset*

Customized CIFAR10 dataset for testing Pytorch Lightning without the torchvision dependency.

Part of the code was copied from https://github.com/pytorch/vision/blob/build/v0.5.0/torchvision/datasets/

Parameters

- **data_dir** (str) – Root directory of dataset where CIFAR10/processed/training.pt and CIFAR10/processed/test.pt exist.
- **train** (bool) – If True, creates dataset from training.pt, otherwise from test.pt.
- **download** (bool) – If true, downloads the dataset from the internet and puts it in root directory. If dataset is already downloaded, it is not downloaded again.

Examples

```
>>> from torchvision import transforms
>>> from pl_bolts.transforms.dataset_normalizations import cifar10_normalization
>>> cf10_transforms = transforms.Compose([transforms.ToTensor(), cifar10_
↪normalization()])
>>> dataset = CIFAR10(download=True, transform=cf10_transforms)
>>> len(dataset)
50000
>>> torch.bincount(dataset.targets)
tensor([5000, 5000, 5000, 5000, 5000, 5000, 5000, 5000, 5000, 5000])
>>> data, label = dataset[0]
>>> data.shape
torch.Size([3, 32, 32])
>>> label
6
```

Labels:

```
airplane: 0
automobile: 1
bird: 2
cat: 3
deer: 4
dog: 5
frog: 6
horse: 7
ship: 8
truck: 9
```

classmethod _check_exists(*data_folder*, *file_names*)

Return type bool

_extract_archive_save_torch(*download_path*)

_unpickle(*path_folder*, *file_name*)

Return type Tuple[Tensor, Tensor]

download(*data_folder*)

Download the data if it doesn't exist in cached_folder_path already.

Return type None

prepare_data(*download*)

BASE_URL = 'https://www.cs.toronto.edu/~kriz/'

DATASET_NAME = 'CIFAR10'

FILE_NAME = 'cifar-10-python.tar.gz'

TEST_FILE_NAME = 'test.pt'

TRAIN_FILE_NAME = 'training.pt'

cache_folder_name: str = 'complete'

data = None

dir_path = None

labels = {0, 1, 2, 3, 4, 5, 6, 7, 8, 9}

normalize = None

relabel = False

targets = None

class pl_bolts.datamodules.cifar10_dataset.**TrialCIFAR10**(*data_dir='.'*, *train=True*, *transform=None*, *download=False*, *num_samples=100*, *labels=(1, 5, 8)*, *relabel=True*)

Bases: *pl_bolts.datamodules.cifar10_dataset.CIFAR10*

Customized CIFAR10 dataset for testing Pytorch Lightning without the torchvision dependency.

Parameters

- **data_dir** (str) – Root directory of dataset where `CIFAR10/processed/training.pt` and `CIFAR10/processed/test.pt` exist.
- **train** (bool) – If `True`, creates dataset from `training.pt`, otherwise from `test.pt`.
- **download** (bool) – If true, downloads the dataset from the internet and puts it in root directory. If dataset is already downloaded, it is not downloaded again.
- **num_samples** (int) – number of examples per selected class/digit
- **labels** (Optional[Sequence]) – list selected CIFAR10 digits/classes

Examples

```
>>> dataset = TrialCIFAR10(download=True, num_samples=150, labels=(1, 5, 8))
>>> len(dataset)
450
>>> sorted(set([d.item() for d in dataset.targets]))
[1, 5, 8]
>>> torch.bincount(dataset.targets)
tensor([  0, 150,   0,   0,   0, 150,   0,   0, 150])
>>> data, label = dataset[0]
>>> data.shape
torch.Size([3, 32, 32])
```

prepare_data(*download*)

Return type None

data = None

dir_path = None

normalize = None

targets = None

pl_bolts.datamodules.concat_dataset module

class pl_bolts.datamodules.concat_dataset.**ConcatDataset**(**datasets*)
Bases: torch.utils.data.Dataset

pl_bolts.datamodules.fashion_mnist_datamodule module

class pl_bolts.datamodules.fashion_mnist_datamodule.**FashionMNISTDataModule**(*data_dir*, *val_split=5000*, *num_workers=16*, **args*, ***kwargs*)
Bases: *pl_bolts.datamodules.lightning_datamodule.LightningDataModule*

Standard FashionMNIST, train, val, test splits and transforms

Transforms:

```
mnist_transforms = transform_lib.Compose([
    transform_lib.ToTensor()
])
```

Example:

```
from pl_bolts.datamodules import FashionMNISTDataModule

dm = FashionMNISTDataModule('.')
model = LitModel(datamodule=dm)
```

Parameters

- **data_dir** (str) – where to save/load the data
- **val_split** (int) – how many of the training images to use for the validation split
- **num_workers** (int) – how many workers to use for loading data

_default_transforms()

prepare_data()
Saves FashionMNIST files to data_dir

test_dataloader(*batch_size=32*, *transforms=None*)
FashionMNIST test set uses the test split

Parameters

- **batch_size** – size of batch
- **transforms** – custom transforms

train_dataloader(*batch_size=32*, *transforms=None*)
FashionMNIST train set removes a subset to use for validation

Parameters

- **batch_size** – size of batch
- **transforms** – custom transforms

val_dataloader(*batch_size=32*, *transforms=None*)
FashionMNIST val set uses a subset of the training set for validation

Parameters

- **batch_size** – size of batch
- **transforms** – custom transforms

name: str = 'fashion_mnist'

property num_classes
Return: 10

pl_bolts.datamodules.imagenet_datamodule module

class pl_bolts.datamodules.imagenet_datamodule.**ImagenetDataModule**(*data_dir*, *meta_dir=None*, *num_imgs_per_val_class=50*, *image_size=224*, *num_workers=16*, **args*, ***kwargs*)

Bases: *pl_bolts.datamodules.lightning_datamodule.LightningDataModule*

Imagenet train, val and test dataloaders.

The train set is the imagenet train.

The val set is taken from the train set with *num_imgs_per_val_class* images per class. For example if *num_imgs_per_val_class=2* then there will be 2,000 images in the validation set.

The test set is the official imagenet validation set.

Example:

```
from pl_bolts.datamodules import ImagenetDataModule

datamodule = ImagenetDataModule(IMAGENET_PATH)
```

Parameters

- **data_dir** (str) – path to the imagenet dataset file
- **meta_dir** (Optional[str]) – path to meta.bin file
- **num_imgs_per_val_class** (int) – how many images per class for the validation set
- **image_size** (int) – final image size
- **num_workers** (int) – how many data workers

_verify_splits(*data_dir*, *split*)

prepare_data()
This method already assumes you have imagenet2012 downloaded. It validates the data using the meta.bin.

Warning: Please download imagenet on your own first.

test_dataloader(*batch_size*, *num_images_per_class=-1*, *transforms=None*)
Uses the validation split of imagenet2012 for testing

Parameters

- **batch_size** – the batch size
- **num_images_per_class** – how many images per class to test on
- **transforms** – the transforms

train_dataloader(*batch_size*)
Uses the train split of imagenet2012 and puts away a portion of it for the validation split

Parameters

- **batch_size** – the batch size
- **transforms** – the transforms

train_transform()
The standard imagenet transforms

```
transform_lib.Compose([
    transform_lib.RandomResizedCrop(self.image_size),
    transform_lib.RandomHorizontalFlip(),
    transform_lib.ToTensor(),
    transform_lib.Normalize(
        mean=[0.485, 0.456, 0.406],
        std=[0.229, 0.224, 0.225]
    ),
])
```

val_dataloader(*batch_size*, *transforms=None*)
Uses the part of the train split of imagenet2012 that was not used for training via *num_imgs_per_val_class*

Parameters

- **batch_size** – the batch size
- **transforms** – the transforms

val_transform()
The standard imagenet transforms for validation

```
transform_lib.Compose([
    transform_lib.Resize(self.image_size + 32),
    transform_lib.CenterCrop(self.image_size),
    transform_lib.ToTensor(),
    transform_lib.Normalize(
        mean=[0.485, 0.456, 0.406],
        std=[0.229, 0.224, 0.225]
    ),
])
```

name: str = 'imagenet'

property num_classes
Return:

1000

pl_bolts.datamodules.imagenet_dataset module

class pl_bolts.datamodules.imagenet_dataset.**UnlabeledImagenet**(*root*, *split='train'*, *num_classes=-1*, *num_imgs_per_class=-1*, *num_imgs_per_class_val_split=50*, *meta_dir=None*, ***kwargs*)

Bases: torchvision.datasets.ImageNet

Official train set gets split into train, val. (using nb_imgs_per_val_class for each class). Official validation becomes test set

Within each class, we further allow limiting the number of samples per class (for semi-sup lng)

Parameters

- **root** – path of dataset
- **split** (str) –
- **num_classes** (int) – Sets the limit of classes
- **num_imgs_per_class** (int) – Limits the number of images per class
- **num_imgs_per_class_val_split** (int) – How many images per class to generate the val split
- **download** –
- **kwargs** –

classmethod generate_meta_bins(*devkit_dir*)

partition_train_set(*imgs*, *nb_imgs_in_val*)

pl_bolts.datamodules.imagenet_dataset.**_calculate_md5**(*fpath*, *chunk_size=1048576*)

pl_bolts.datamodules.imagenet_dataset.**_check_integrity**(*fpath*, *md5=None*)

pl_bolts.datamodules.imagenet_dataset.**_check_md5**(*fpath*, *md5*, ***kwargs*)

pl_bolts.datamodules.imagenet_dataset.**_is_gzip**(*filename*)

pl_bolts.datamodules.imagenet_dataset.**_is_tar**(*filename*)

pl_bolts.datamodules.imagenet_dataset.**_is_targz**(*filename*)

pl_bolts.datamodules.imagenet_dataset.**_is_tarxz**(*filename*)

pl_bolts.datamodules.imagenet_dataset.**_is_zip**(*filename*)

pl_bolts.datamodules.imagenet_dataset.**_verify_archive**(*root*, *file*, *md5*)

pl_bolts.datamodules.imagenet_dataset.**extract_archive**(*from_path*, *to_path=None*, *remove_finished=False*)

pl_bolts.datamodules.imagenet_dataset.**parse_devkit_archive**(*root*, *file=None*)

Parse the devkit archive of the ImageNet2012 classification dataset and save the meta information in a binary file.

Parameters

- **root** (*str*) – Root directory containing the devkit archive

- **file** (*str, optional*) – Name of devkit archive. Defaults to 'ILSVRC2012_devkit_t12.tar.gz'

pl_bolts.datamodules.lightning_datamodule module

class pl_bolts.datamodules.lightning_datamodule.**LightningDataModule**(*train_transforms=None*, *val_transforms=None*, *test_transforms=None*)

Bases: object

A DataModule standardizes the training, val, test splits, data preparation and transforms. The main advantage is consistent data splits and transforms across models.

Example:

```
class MyDataModule(LightningDataModule):

    def __init__(self):
        super().__init__()

    def prepare_data(self):
        # download, split, etc...

    def train_dataloader(self):
        train_split = Dataset(...)
        return DataLoader(train_split)

    def val_dataloader(self):
        val_split = Dataset(...)
        return DataLoader(val_split)

    def test_dataloader(self):
        test_split = Dataset(...)
        return DataLoader(test_split)
```

A DataModule implements 4 key methods

1. **prepare_data** (things to do on 1 GPU not on every GPU in distributed mode)
2. **train_dataloader** the training dataloader.
3. **val_dataloader** the val dataloader.
4. **test_dataloader** the test dataloader.

This allows you to share a full dataset without explaining what the splits, transforms or download process is.

classmethod add_argparse_args(*parent_parser*)
Extends existing argparse by default *LightningDataModule* attributes.

Return type ArgumentParser

classmethod from_argparse_args(*args*, ***kwargs*)
Create an instance from CLI arguments.

Parameters

- **args** (Union[Namespace, ArgumentParser]) – The parser or namespace to take arguments from. Only known arguments will be parsed and passed to the *LightningDataModule*.

- ****kwargs** – Additional keyword arguments that may override ones in the parser or namespace. These must be valid Trainer arguments.

Example:

```
parser = ArgumentParser(add_help=False)
parser = LightningDataModule.add_argparse_args(parser)
module = LightningDataModule.from_argparse_args(args)
```

classmethod get_init_arguments_and_types()

Scans the Trainer signature and returns argument names, types and default values.

Returns (argument name, set with argument types, argument default value).

Return type List with tuples of 3 values

abstract prepare_data(**args*, ***kwargs*)

Use this to download and prepare data. In distributed (GPU, TPU), this will only be called once. This is called before requesting the dataloaders:

Warning: Do not assign anything to the model in this step since this will only be called on 1 GPU.

Pseudocode:

```
model.prepare_data()
model.train_dataloader()
model.val_dataloader()
model.test_dataloader()
```

Example:

```
def prepare_data(self):
    download_imagenet()
    clean_imagenet()
    cache_imagenet()
```

size(*dim=None*)

Return the dimension of each input Either as a tuple or list of tuples

Return type Union[Tuple, int]

abstract test_dataloader(**args*, ***kwargs*)

Implement a PyTorch DataLoader for training.

Return type Union[DataLoader, List[DataLoader]]

Returns Single PyTorch DataLoader.

Note: Lightning adds the correct sampler for distributed and arbitrary hardware. There is no need to set it yourself.

Note: You can also return a list of DataLoaders

Example:

```
def test_dataloader(self):
    dataset = MNIST(root=PATH, train=False, transform=transforms.ToTensor(),␣
↪download=False)
    loader = torch.utils.data.DataLoader(dataset=dataset, shuffle=False)
    return loader
```

abstract train_dataloader(**args*, ***kwargs*)

Implement a PyTorch DataLoader for training.

Return type DataLoader

Returns Single PyTorch DataLoader.

Note: Lightning adds the correct sampler for distributed and arbitrary hardware. There is no need to set it yourself.

Example:

```
def train_dataloader(self):
    dataset = MNIST(root=PATH, train=True, transform=transforms.ToTensor(),␣
↪download=False)
    loader = torch.utils.data.DataLoader(dataset=dataset)
    return loader
```

abstract val_dataloader(**args*, ***kwargs*)

Implement a PyTorch DataLoader for training.

Return type Union[DataLoader, List[DataLoader]]

Returns Single PyTorch DataLoader.

Note: Lightning adds the correct sampler for distributed and arbitrary hardware. There is no need to set it yourself.

Note: You can also return a list of DataLoaders

Example:

```
def val_dataloader(self):
    dataset = MNIST(root=PATH, train=False, transform=transforms.ToTensor(),␣
↪download=False)
    loader = torch.utils.data.DataLoader(dataset=dataset, shuffle=False)
    return loader
```

name: str = Ellipsis

property test_transforms

property train_transforms

property val_transforms

pl_bolts.datamodules.mnist_datamodule module

class pl_bolts.datamodules.mnist_datamodule.**MNISTDataModule**(*data_dir*, *val_split=5000*, *num_workers=16*, *normalize=False*, **args*, ***kwargs*)

Bases: *pl_bolts.datamodules.lightning_datamodule.LightningDataModule*

Standard MNIST, train, val, test splits and transforms

Transforms:

```
mnist_transforms = transform_lib.Compose([
    transform_lib.ToTensor()
])
```

Example:

```
from pl_bolts.datamodules import MNISTDataModule

dm = MNISTDataModule('.')
model = LitModel(datamodule=dm)
```

Parameters

- **data_dir** (str) – where to save/load the data
- **val_split** (int) – how many of the training images to use for the validation split
- **num_workers** (int) – how many workers to use for loading data
- **normalize** (bool) – If true applies image normalize

_default_transforms()

prepare_data()
Saves MNIST files to data_dir

test_dataloader(*batch_size=32*, *transforms=None*)
MNIST test set uses the test split

Parameters

- **batch_size** – size of batch
- **transforms** – custom transforms

train_dataloader(*batch_size=32*, *transforms=None*)
MNIST train set removes a subset to use for validation

Parameters

- **batch_size** – size of batch
- **transforms** – custom transforms

val_dataloader(*batch_size=32*, *transforms=None*)
MNIST val set uses a subset of the training set for validation

Parameters

- **batch_size** – size of batch

- **transforms** – custom transforms

name: str = 'mnist'

property num_classes
Return: 10

pl_bolts.datamodules.sklearn_datamodule module

class pl_bolts.datamodules.sklearn_datamodule.**SklearnDataModule**(*X*, *y*, *x_val=None*, *y_val=None*, *x_test=None*, *y_test=None*, *val_split=0.2*, *test_split=0.1*, *num_workers=2*, *random_state=1234*, *shuffle=True*, **args*, ***kwargs*)

Bases: *pl_bolts.datamodules.lightning_datamodule.LightningDataModule*

Automatically generates the train, validation and test splits for a Numpy dataset. They are set up as dataloaders for convenience. Optionally, you can pass in your own validation and test splits.

Example

```
>>> from sklearn.datasets import load_boston
>>> from pl_bolts.datamodules import SklearnDataModule
...
>>> X, y = load_boston(return_X_y=True)
>>> loaders = SklearnDataModule(X, y)
...
>>> # train set
>>> train_loader = loaders.train_dataloader(batch_size=32)
>>> len(train_loader.dataset)
355
>>> len(train_loader)
11
>>> # validation set
>>> val_loader = loaders.val_dataloader(batch_size=32)
>>> len(val_loader.dataset)
100
>>> len(val_loader)
3
>>> # test set
>>> test_loader = loaders.test_dataloader(batch_size=32)
>>> len(test_loader.dataset)
51
>>> len(test_loader)
1
```

_init_datasets(*X*, *y*, *x_val*, *y_val*, *x_test*, *y_test*)

test_dataloader(*batch_size=16*)
Implement a PyTorch DataLoader for training.

Returns Single PyTorch `DataLoader`.

Note: Lightning adds the correct sampler for distributed and arbitrary hardware. There is no need to set it yourself.

Note: You can also return a list of DataLoaders

Example:

```
def test_dataloader(self):
    dataset = MNIST(root=PATH, train=False, transform=transforms.ToTensor(),␣
↪download=False)
    loader = torch.utils.data.DataLoader(dataset=dataset, shuffle=False)
    return loader
```

train_dataloader(*batch_size=16*)
Implement a PyTorch DataLoader for training.

Returns Single PyTorch `DataLoader`.

Note: Lightning adds the correct sampler for distributed and arbitrary hardware. There is no need to set it yourself.

Example:

```
def train_dataloader(self):
    dataset = MNIST(root=PATH, train=True, transform=transforms.ToTensor(),␣
↪download=False)
    loader = torch.utils.data.DataLoader(dataset=dataset)
    return loader
```

val_dataloader(*batch_size=16*)
Implement a PyTorch DataLoader for training.

Returns Single PyTorch `DataLoader`.

Note: Lightning adds the correct sampler for distributed and arbitrary hardware. There is no need to set it yourself.

Note: You can also return a list of DataLoaders

Example:

```
def val_dataloader(self):
    dataset = MNIST(root=PATH, train=False, transform=transforms.ToTensor(),␣
↪download=False)
    loader = torch.utils.data.DataLoader(dataset=dataset, shuffle=False)
    return loader
```

name: str = 'sklearn'

class pl_bolts.datamodules.sklearn_datamodule.**SklearnDataset**(*X*, *y*, *X_transform=None*, *y_transform=None*)

Bases: torch.utils.data.Dataset

Mapping between numpy (or sklearn) datasets to PyTorch datasets.

Parameters

- **X** (ndarray) – Numpy ndarray
- **y** (ndarray) – Numpy ndarray
- **X_transform** (Optional[Any]) – Any transform that works with Numpy arrays
- **y_transform** (Optional[Any]) – Any transform that works with Numpy arrays

Example

```
>>> from sklearn.datasets import load_boston
>>> from pl_bolts.datamodules import SklearnDataset
...
>>> X, y = load_boston(return_X_y=True)
>>> dataset = SklearnDataset(X, y)
>>> len(dataset)
506
```

class pl_bolts.datamodules.sklearn_datamodule.**TensorDataModule**(*X*, *y*, *x_val=None*, *y_val=None*, *x_test=None*, *y_test=None*, *val_split=0.2*, *test_split=0.1*, *num_workers=2*, *random_state=1234*, *shuffle=True*, **args*, ***kwargs*)

Bases: *pl_bolts.datamodules.sklearn_datamodule.SklearnDataModule*

Automatically generates the train, validation and test splits for a PyTorch tensor dataset. They are set up as dataloaders for convenience. Optionally, you can pass in your own validation and test splits.

Example

```
>>> from pl_bolts.datamodules import TensorDataModule
>>> import torch
...
>>> # create dataset
>>> X = torch.rand(100, 3)
>>> y = torch.rand(100)
>>> loaders = TensorDataModule(X, y)
...
```

(continues on next page)

(continued from previous page)

```
>>> # train set
>>> train_loader = loaders.train_dataloader(batch_size=10)
>>> len(train_loader.dataset)
70
>>> len(train_loader)
7
>>> # validation set
>>> val_loader = loaders.val_dataloader(batch_size=10)
>>> len(val_loader.dataset)
20
>>> len(val_loader)
2
>>> # test set
>>> test_loader = loaders.test_dataloader(batch_size=10)
>>> len(test_loader.dataset)
10
>>> len(test_loader)
1
```

Automatically generates the train, validation and test splits for a Numpy dataset. They are set up as dataloaders for convenience. Optionally, you can pass in your own validation and test splits.

Example

```
>>> from sklearn.datasets import load_boston
>>> from pl_bolts.datamodules import SklearnDataModule
...
>>> X, y = load_boston(return_X_y=True)
>>> loaders = SklearnDataModule(X, y)
...
>>> # train set
>>> train_loader = loaders.train_dataloader(batch_size=32)
>>> len(train_loader.dataset)
355
>>> len(train_loader)
11
>>> # validation set
>>> val_loader = loaders.val_dataloader(batch_size=32)
>>> len(val_loader.dataset)
100
>>> len(val_loader)
3
>>> # test set
>>> test_loader = loaders.test_dataloader(batch_size=32)
>>> len(test_loader.dataset)
51
>>> len(test_loader)
1
```

class pl_bolts.datamodules.sklearn_datamodule.**TensorDataset**(*X*, *y*, *X_transform=None*, *y_transform=None*)

Bases: torch.utils.data.Dataset

Prepare PyTorch tensor dataset for data loaders.

Parameters

- **X** (Tensor) – PyTorch tensor
- **y** (Tensor) – PyTorch tensor
- **X_transform** (Optional[Any]) – Any transform that works with PyTorch tensors
- **y_transform** (Optional[Any]) – Any transform that works with PyTorch tensors

Example

```
>>> from pl_bolts.datamodules import TensorDataset
...
>>> X = torch.rand(10, 3)
>>> y = torch.rand(10)
>>> dataset = TensorDataset(X, y)
>>> len(dataset)
10
```

pl_bolts.datamodules.ssl_imagenet_datamodule module

class pl_bolts.datamodules.ssl_imagenet_datamodule.**SSLImagenetDataModule**(*data_dir*, *meta_dir=None*, *num_workers=16*, **args*, ***kwargs*)

Bases: *pl_bolts.datamodules.lightning_datamodule.LightningDataModule*

_default_transforms()

_verify_splits(*data_dir*, *split*)

prepare_data()

Use this to download and prepare data. In distributed (GPU, TPU), this will only be called once. This is called before requesting the dataloaders:

> **Warning:** Do not assign anything to the model in this step since this will only be called on 1 GPU.

Pseudocode:

```
model.prepare_data()
model.train_dataloader()
model.val_dataloader()
model.test_dataloader()
```

Example:

```
def prepare_data(self):
    download_imagenet()
    clean_imagenet()
    cache_imagenet()
```

test_dataloader(*batch_size*, *num_images_per_class*, *add_normalize=False*)

Implement a PyTorch DataLoader for training.

Returns Single PyTorch `DataLoader`.

Note: Lightning adds the correct sampler for distributed and arbitrary hardware. There is no need to set it yourself.

Note: You can also return a list of DataLoaders

Example:

```
def test_dataloader(self):
    dataset = MNIST(root=PATH, train=False, transform=transforms.ToTensor(),␣
↪download=False)
    loader = torch.utils.data.DataLoader(dataset=dataset, shuffle=False)
    return loader
```

train_dataloader(*batch_size*, *num_images_per_class=-1*, *add_normalize=False*)
Implement a PyTorch DataLoader for training.

Returns Single PyTorch `DataLoader`.

Note: Lightning adds the correct sampler for distributed and arbitrary hardware. There is no need to set it yourself.

Example:

```
def train_dataloader(self):
    dataset = MNIST(root=PATH, train=True, transform=transforms.ToTensor(),␣
↪download=False)
    loader = torch.utils.data.DataLoader(dataset=dataset)
    return loader
```

val_dataloader(*batch_size*, *num_images_per_class=50*, *add_normalize=False*)
Implement a PyTorch DataLoader for training.

Returns Single PyTorch `DataLoader`.

Note: Lightning adds the correct sampler for distributed and arbitrary hardware. There is no need to set it yourself.

Note: You can also return a list of DataLoaders

Example:

```
def val_dataloader(self):
    dataset = MNIST(root=PATH, train=False, transform=transforms.ToTensor(),␣
↪download=False)
    loader = torch.utils.data.DataLoader(dataset=dataset, shuffle=False)
    return loader
```

name: str = 'imagenet'

property num_classes

pl_bolts.datamodules.stl10_datamodule module

class pl_bolts.datamodules.stl10_datamodule.**STL10DataModule**(*data_dir*, *unlabeled_val_split=5000*, *train_val_split=500*, *num_workers=16*, **args*, ***kwargs*)

Bases: *pl_bolts.datamodules.lightning_datamodule.LightningDataModule*

Standard STL-10, train, val, test splits and transforms. STL-10 has support for doing validation splits on the labeled or unlabeled splits

Transforms:

```
mnist_transforms = transform_lib.Compose([
    transform_lib.ToTensor(),
    transforms.Normalize(
        mean=(0.43, 0.42, 0.39),
        std=(0.27, 0.26, 0.27)
    )
])
```

Example:

```
from pl_bolts.datamodules import STL10DataModule

dm = STL10DataModule(PATH)
model = LitModel(datamodule=dm)
```

Parameters

- **data_dir** (str) – where to save/load the data
- **unlabeled_val_split** (int) – how many images from the unlabeled training split to use for validation
- **train_val_split** (int) – how many images from the labeled training split to use for validation
- **num_workers** (int) – how many workers to use for loading data

default_transforms()

prepare_data()
Downloads the unlabeled, train and test split

test_dataloader(*batch_size*)
Loads the test split of STL10

Parameters

- **batch_size** – the batch size
- **transforms** – the transforms

train_dataloader(*batch_size*)
Loads the 'unlabeled' split minus a portion set aside for validation via *unlabeled_val_split*.

Parameters **batch_size** – the batch size

train_dataloader_mixed(*batch_size*)
Loads a portion of the 'unlabeled' training data and 'train' (labeled) data. both portions have a subset removed for validation via *unlabeled_val_split* and *train_val_split*

Parameters

- **batch_size** – the batch size
- **transforms** – a sequence of transforms

val_dataloader(*batch_size*)
Loads a portion of the 'unlabeled' training data set aside for validation The val dataset = (unlabeled - train_val_split)

Parameters

- **batch_size** – the batch size
- **transforms** – a sequence of transforms

val_dataloader_mixed(*batch_size*)
Loads a portion of the 'unlabeled' training data set aside for validation along with the portion of the 'train' dataset to be used for validation

unlabeled_val = (unlabeled - train_val_split)

labeled_val = (train- train_val_split)

full_val = unlabeled_val + labeled_val

Parameters

- **batch_size** – the batch size
- **transforms** – a sequence of transforms

name: str = 'stl10'

property num_classes

23.4 pl_bolts.metrics package

23.4.1 Submodules

pl_bolts.metrics.aggregation module

pl_bolts.metrics.aggregation.**accuracy**(*preds*, *labels*)

pl_bolts.metrics.aggregation.**mean**(*res*, *key*)

pl_bolts.metrics.aggregation.**precision_at_k**(*output*, *target*, *top_k=(1,)*)
Computes the accuracy over the k top predictions for the specified values of k

23.5 pl_bolts.models package

Collection of PyTorchLightning models

23.5.1 Subpackages

pl_bolts.models.autoencoders package

Here are a VAE and GAN

Subpackages

pl_bolts.models.autoencoders.basic_ae package

AE Template

This is a basic template for implementing an Autoencoder in PyTorch Lightning.

A default encoder and decoder have been provided but can easily be replaced by custom models.

This template uses the MNIST dataset but image data of any dimension can be fed in as long as the image width and image height are even values. For other types of data, such as sound, it will be necessary to change the Encoder and Decoder.

The default encoder and decoder are both convolutional with a 128-dimensional hidden layer and a 32-dimensional latent space. The model accepts arguments for these dimensions (see example below) if you want to use the default encoder + decoder but with different hidden layer and latent layer dimensions. The model also assumes a Gaussian prior and a Gaussian approximate posterior distribution.

```
from pl_bolts.models.autoencoders import AE

model = AE()
trainer = pl.Trainer()
trainer.fit(model)
```

Submodules

pl_bolts.models.autoencoders.basic_ae.basic_ae_module module

class pl_bolts.models.autoencoders.basic_ae.basic_ae_module.**AE**(*datamodule=None*, *input_channels=1*, *input_height=28*, *input_width=28*, *latent_dim=32*, *batch_size=32*, *hidden_dim=128*, *learning_rate=0.001*, *num_workers=8*, *data_dir='.'*, ***kwargs*)

Bases: pytorch_lightning.LightningModule

Arg:

datamodule: the datamodule (train, val, test splits) input_channels: num of image channels input_height: image height input_width: image width latent_dim: emb dim for encoder batch_size: the batch size hidden_dim: the encoder dim learning_rate: the learning rate num_workers: num dataloader workers data_dir: where to store data

_run_step(*batch*)

static add_model_specific_args(*parent_parser*)

configure_optimizers()

forward(*z*)

init_decoder(*hidden_dim*, *latent_dim*)

init_encoder(*hidden_dim*, *latent_dim*, *input_width*, *input_height*)

prepare_data()

test_dataloader()

test_epoch_end(*outputs*)

test_step(*batch*, *batch_idx*)

train_dataloader()

training_step(*batch*, *batch_idx*)

val_dataloader()

validation_epoch_end(*outputs*)

validation_step(*batch*, *batch_idx*)

pl_bolts.models.autoencoders.basic_ae.components module

class pl_bolts.models.autoencoders.basic_ae.components.**AEEncoder**(*hidden_dim*, *latent_dim*, *input_width*, *input_height*)

Bases: torch.nn.Module

Takes as input an image, uses a CNN to extract features which get split into a mu and sigma vector

_calculate_output_dim(*input_width*, *input_height*)

forward(*x*)

class pl_bolts.models.autoencoders.basic_ae.components.**DenseBlock**(*in_dim*, *out_dim*, *drop_p=0.2*)

Bases: torch.nn.Module

forward(*x*)

pl_bolts.models.autoencoders.basic_vae package

VAE Template

This is a basic template for implementing a Variational Autoencoder in PyTorch Lightning.

A default encoder and decoder have been provided but can easily be replaced by custom models.

This template uses the MNIST dataset but image data of any dimension can be fed in as long as the image width and image height are even values. For other types of data, such as sound, it will be necessary to change the Encoder and Decoder.

The default encoder and decoder are both convolutional with a 128-dimensional hidden layer and a 32-dimensional latent space. The model accepts arguments for these dimensions (see example below) if you want to use the default encoder + decoder but with different hidden layer and latent layer dimensions. The model also assumes a Gaussian prior and a Gaussian approximate posterior distribution.

Submodules

pl_bolts.models.autoencoders.basic_vae.basic_vae_module module

class pl_bolts.models.autoencoders.basic_vae.basic_vae_module.**VAE**(*hidden_dim=128*, *latent_dim=32*, *input_channels=3*, *input_width=224*, *input_height=224*, *batch_size=32*, *learning_rate=0.001*, *data_dir='.'*, *datamodule=None*, *pretrained=None*, ***kwargs*)

Bases: pytorch_lightning.LightningModule

Standard VAE with Gaussian Prior and approx posterior.

Model is available pretrained on different datasets:

Example:

```
# not pretrained
vae = VAE()

# pretrained on imagenet
vae = VAE(pretrained='imagenet')

# pretrained on cifar10
vae = VAE(pretrained='cifar10')
```

Parameters

- **hidden_dim** (int) – encoder and decoder hidden dims
- **latent_dim** (int) – latenet code dim
- **input_channels** (int) – num of channels of the input image.
- **input_width** (int) – image input width
- **input_height** (int) – image input height
- **batch_size** (int) – the batch size
- **the learning rate** (*learning_rate"*) –
- **data_dir** (str) – the directory to store data
- **datamodule** (Optional[*LightningDataModule*]) – The Lightning DataModule
- **pretrained** (Optional[str]) – Load weights pretrained on a dataset

_VAE__init_system()

_VAE__set_default_datamodule(*data_dir*)

_VAE__set_pretrained_dims(*pretrained*)

_run_step(*batch*)

static add_model_specific_args(*parent_parser*)

configure_optimizers()

elbo_loss(*x*, *P*, *Q*)

forward(*z*)

get_approx_posterior(*z_mu*, *z_std*)

get_prior(*z_mu*, *z_std*)

init_decoder()

init_encoder()

load_pretrained(*pretrained*)

prepare_data()

test_dataloader()

test_epoch_end(*outputs*)

test_step(*batch*, *batch_idx*)

train_dataloader()

training_step(*batch*, *batch_idx*)

val_dataloader()

validation_epoch_end(*outputs*)

validation_step(*batch*, *batch_idx*)

pl_bolts.models.autoencoders.basic_vae.components module

class pl_bolts.models.autoencoders.basic_vae.components.**Decoder**(*hidden_dim*, *latent_dim*, *input_width*, *input_height*, *input_channels*)

Bases: torch.nn.Module

Takes in latent vars and reconstructs an image

_calculate_output_size(*input_width*, *input_height*)

forward(*z*)

class pl_bolts.models.autoencoders.basic_vae.components.**DenseBlock**(*in_dim*, *out_dim*, *drop_p=0.2*)

Bases: torch.nn.Module

forward(*x*)

class pl_bolts.models.autoencoders.basic_vae.components.**Encoder**(*hidden_dim*, *latent_dim*, *input_channels*, *input_width*, *input_height*)

Bases: torch.nn.Module

Takes as input an image, uses a CNN to extract features which get split into a mu and sigma vector

_calculate_output_dim(*input_width*, *input_height*)

forward(*x*)

pl_bolts.models.gans package

Subpackages

pl_bolts.models.gans.basic package

Submodules

pl_bolts.models.gans.basic.basic_gan_module module

class pl_bolts.models.gans.basic.basic_gan_module.**GAN**(*datamodule=None*, *latent_dim=32*, *batch_size=100*, *learning_rate=0.0002*, *data_dir=''*, *num_workers=8*, ***kwargs*)

Bases: pytorch_lightning.LightningModule

Vanilla GAN implementation.

Example:

```
from pl_bolts.models.gan import GAN

m = GAN()
Trainer(gpus=2).fit(m)
```

Example CLI:

```
# mnist
python  basic_gan_module.py --gpus 1

# imagenet
python  basic_gan_module.py --gpus 1 --dataset 'imagenet2012'
--data_dir /path/to/imagenet/folder/ --meta_dir ~/path/to/meta/bin/folder
--batch_size 256 --learning_rate 0.0001
```

Parameters

- **datamodule** (Optional[*LightningDataModule*]) – the datamodule (train, val, test splits)
- **latent_dim** (int) – emb dim for encoder
- **batch_size** (int) – the batch size

- **learning_rate** (float) – the learning rate
- **data_dir** (str) – where to store data
- **num_workers** (int) – data workers

static add_model_specific_args(*parent_parser*)

configure_optimizers()

discriminator_loss(*x*)

discriminator_step(*x*)

forward(*z*)

Generates an image given input noise z

Example:

```
z = torch.rand(batch_size, latent_dim)
gan = GAN.load_from_checkpoint(PATH)
img = gan(z)
```

generator_loss(*x*)

generator_step(*x*)

init_discriminator(*img_dim*)

init_generator(*img_dim*)

prepare_data()

train_dataloader()

training_epoch_end(*outputs*)

training_step(*batch*, *batch_idx*, *optimizer_idx*)

class pl_bolts.models.gans.basic.basic_gan_module.**ImageGenerator**(**args*, ***kwargs*)

Bases: pytorch_lightning.Callback

on_epoch_end(*trainer*, *pl_module*)

pl_bolts.models.gans.basic.components module

class pl_bolts.models.gans.basic.components.**Discriminator**(*img_shape*, *hidden_dim=1024*)

Bases: torch.nn.Module

forward(*img*)

class pl_bolts.models.gans.basic.components.**Generator**(*latent_dim*, *img_shape*, *hidden_dim=256*)

Bases: torch.nn.Module

forward(*z*)

pl_bolts.models.regression package

Submodules

pl_bolts.models.regression.linear_regression module

class pl_bolts.models.regression.linear_regression.**LinearRegression**(*input_dim*, *bias=True*, *learning_rate=0.0001*, *optimizer=torch.optim.Adam*, *l1_strength=None*, *l2_strength=None*, ***kwargs*)

Bases: pytorch_lightning.LightningModule

Linear regression model implementing - with optional L1/L2 regularization $$min_{W} ||(Wx + b) - y ||_2^2 $$

Parameters

- **input_dim** (int) – number of dimensions of the input (1+)
- **bias** (bool) – If false, will not use $+b$
- **learning_rate** (float) – learning_rate for the optimizer
- **optimizer** (Optimizer) – the optimizer to use (default='Adam')
- **l1_strength** (Optional[float]) – L1 regularization strength (default=None)
- **l2_strength** (Optional[float]) – L2 regularization strength (default=None)

static add_model_specific_args(*parent_parser*)

configure_optimizers()

forward(*x*)

test_epoch_end(*outputs*)

test_step(*batch*, *batch_idx*)

training_step(*batch*, *batch_idx*)

validation_epoch_end(*outputs*)

validation_step(*batch*, *batch_idx*)

pl_bolts.models.regression.logistic_regression module

class pl_bolts.models.regression.logistic_regression.**LogisticRegression**(*input_dim*, *num_classes*, *bias=True*, *learning_rate=0.0001*, *optimizer=torch.optim.Adam*, *l1_strength=0.0*, *l2_strength=0.0*, ***kwargs*)

Bases: pytorch_lightning.LightningModule

Logistic regression model

Parameters

- **input_dim** (int) – number of dimensions of the input (at least 1)
- **num_classes** (int) – number of class labels (binary: 2, multi-class: >2)
- **bias** (bool) – specifies if a constant or intercept should be fitted (equivalent to fit_intercept in sklearn)
- **learning_rate** (float) – learning_rate for the optimizer
- **optimizer** (Optimizer) – the optimizer to use (default='Adam')
- **l1_strength** (float) – L1 regularization strength (default=None)
- **l2_strength** (float) – L2 regularization strength (default=None)

static add_model_specific_args(*parent_parser*)

configure_optimizers()

forward(*x*)

test_epoch_end(*outputs*)

test_step(*batch*, *batch_idx*)

training_step(*batch*, *batch_idx*)

validation_epoch_end(*outputs*)

validation_step(*batch*, *batch_idx*)

pl_bolts.models.rl package

Subpackages

pl_bolts.models.rl.common package

Submodules

pl_bolts.models.rl.common.agents module

Agent module containing classes for Agent logic

Based on the implementations found here: https://github.com/Shmuma/ptan/blob/master/ptan/agent.py

class pl_bolts.models.rl.common.agents.**Agent**(*net*)
Bases: object

Basic agent that always returns 0

__call__(*state*, *device*)
Using the given network, decide what action to carry

Parameters

- **state** (Tensor) – current state of the environment
- **device** (str) – device used for current batch

Return type int

Returns action

class pl_bolts.models.rl.common.agents.**PolicyAgent**(*net*)
Bases: *pl_bolts.models.rl.common.agents.Agent*

Policy based agent that returns an action based on the networks policy

__call__(*state*, *device*)
Takes in the current state and returns the action based on the agents policy

Parameters

- **state** (Tensor) – current state of the environment
- **device** (str) – the device used for the current batch

Return type int

Returns action defined by policy

class pl_bolts.models.rl.common.agents.**ValueAgent**(*net*, *action_space*, *eps_start=1.0*, *eps_end=0.2*, *eps_frames=1000*)
Bases: *pl_bolts.models.rl.common.agents.Agent*

Value based agent that returns an action based on the Q values from the network

__call__(*state*, *device*)
Takes in the current state and returns the action based on the agents policy

Parameters

- **state** (Tensor) – current state of the environment
- **device** (str) – the device used for the current batch

Return type int

Returns action defined by policy

get_action(*state*, *device*)
Returns the best action based on the Q values of the network :type _sphinx_paramlinks_pl_bolts.models.rl.common.agents.ValueAgent.get_action.state: Tensor :param _sphinx_paramlinks_pl_bolts.models.rl.common.agents.ValueAgent.get_action.state: current state of the environment :type _sphinx_paramlinks_pl_bolts.models.rl.common.agents.ValueAgent.get_action.device:

`device` :param _sphinx_paramlinks_pl_bolts.models.rl.common.agents.ValueAgent.get_action.device: the device used for the current batch

Returns action defined by Q values

get_random_action()
returns a random action

Return type `int`

update_epsilon(*step*)
Updates the epsilon value based on the current step

Parameters **step** (`int`) – current global step

Return type `None`

pl_bolts.models.rl.common.cli module

Contains generic arguments used for all models

pl_bolts.models.rl.common.cli.**add_base_args**(*parent*)
Adds arguments for DQN model

Note: these params are fine tuned for Pong env

Parameters **parent** –

Return type `ArgumentParser`

pl_bolts.models.rl.common.experience module

Experience sources to be used as datasets for Ligthning DataLoaders

Based on implementations found here: https://github.com/Shmuma/ptan/blob/master/ptan/experience.py

class pl_bolts.models.rl.common.experience.**EpisodicExperienceStream**(*env*, *agent*, *device*, *episodes=1*)
Bases: *pl_bolts.models.rl.common.experience.ExperienceSource*, `torch.utils.data.IterableDataset`

Basic experience stream that iteratively yield the current experience of the agent in the env

Parameters

- **env** (`Env`) – Environmen that is being used
- **agent** (*Agent*) – Agent being used to make decisions

step()
Carries out a single step in the environment

Return type *Experience*

class pl_bolts.models.rl.common.experience.**ExperienceSource**(*env*, *agent*, *device*)
Bases: `object`

Basic single step experience source

Parameters

- **env** (Env) – Environment that is being used
- **agent** (*Agent*) – Agent being used to make decisions

_reset()
resets the env and state

Return type None

run_episode()
Carries out a single episode and returns the total reward. This is used for testing

Return type float

step()
Takes a single step through the environment

Return type Tuple[*Experience*, float, bool]

class pl_bolts.models.rl.common.experience.**NStepExperienceSource**(*env*, *agent*, *device*, *n_steps=1*)
Bases: *pl_bolts.models.rl.common.experience.ExperienceSource*

Expands upon the basic ExperienceSource by collecting experience across N steps

get_transition_info(*gamma=0.9*)
get the accumulated transition info for the n_step_buffer :param _sphinx_paramlinks_pl_bolts.models.rl.common.experience.NStepExperienceSource.get_transition_info.gamma: discount factor

Return type Tuple[float, array, int]

Returns multi step reward, final observation and done

single_step()
Takes a single step in the environment and appends it to the n-step buffer

Return type *Experience*

Returns Experience

step()
Takes an n-step in the environment

Return type Tuple[*Experience*, float, bool]

Returns Experience

class pl_bolts.models.rl.common.experience.**PrioRLDataset**(*buffer*, *sample_size=1*)
Bases: *pl_bolts.models.rl.common.experience.RLDataset*

Iterable Dataset containing the ExperienceBuffer which will be updated with new experiences during training

Parameters

- **buffer** (*Buffer*) – replay buffer
- **sample_size** (int) – number of experiences to sample at a time

class pl_bolts.models.rl.common.experience.**RLDataset**(*buffer*, *sample_size=1*)
Bases: torch.utils.data.IterableDataset

Iterable Dataset containing the ExperienceBuffer which will be updated with new experiences during training

Parameters

- **buffer** (*Buffer*) – replay buffer

- **sample_size** (int) – number of experiences to sample at a time

pl_bolts.models.rl.common.memory module

Series of memory buffers sued

class pl_bolts.models.rl.common.memory.**Buffer**(*capacity*)
Bases: object

Basic Buffer for storing a single experience at a time

Parameters **capacity** (int) – size of the buffer

append(*experience*)
Add experience to the buffer

Parameters **experience** (*Experience*) – tuple (state, action, reward, done, new_state)

Return type None

sample(**args*)
returns everything in the buffer so far it is then reset

Return type Union[Tuple, List[Tuple]]

Returns a batch of tuple np arrays of state, action, reward, done, next_state

class pl_bolts.models.rl.common.memory.**Experience**(*state*, *action*, *reward*, *done*, *new_state*)
Bases: tuple

Create new instance of Experience(state, action, reward, done, new_state)

_asdict()
Return a new OrderedDict which maps field names to their values.

classmethod **_make**(*iterable*)
Make a new Experience object from a sequence or iterable

_replace(***kwds*)
Return a new Experience object replacing specified fields with new values

_fields = ('state', 'action', 'reward', 'done', 'new_state')

_fields_defaults = {}

property action
Alias for field number 1

property done
Alias for field number 3

property new_state
Alias for field number 4

property reward
Alias for field number 2

property state
Alias for field number 0

class pl_bolts.models.rl.common.memory.**MeanBuffer**(*capacity*)
Bases: object

Stores a deque of items and calculates the mean

add(*val*)
Add to the buffer

Return type None

mean()
Retrieve the mean

Return type float

class pl_bolts.models.rl.common.memory.**MultiStepBuffer**(*buffer_size*, *n_step=2*)
Bases: object

N Step Replay Buffer

Deprecated: use the NStepExperienceSource with the standard ReplayBuffer

append(*experience*)
add an experience to the buffer by collecting n steps of experiences :param _sphinx_paramlinks_pl_bolts.models.rl.common.memory.MultiStepBuffer.append.experience: tuple (state, action, reward, done, next_state)

Return type None

get_transition_info(*gamma=0.9*)
get the accumulated transition info for the n_step_buffer :param _sphinx_paramlinks_pl_bolts.models.rl.common.memory.MultiStepBuffer.get_transition_info.gamma: discount factor

Return type Tuple[float, array, int]

Returns multi step reward, final observation and done

sample(*batch_size*)
Takes a sample of the buffer :type _sphinx_paramlinks_pl_bolts.models.rl.common.memory.MultiStepBuffer.sample.batch_s
int :param _sphinx_paramlinks_pl_bolts.models.rl.common.memory.MultiStepBuffer.sample.batch_size: current batch_size

Return type Tuple

Returns a batch of tuple np arrays of Experiences

class pl_bolts.models.rl.common.memory.**PERBuffer**(*buffer_size*, *prob_alpha=0.6*, *beta_start=0.4*, *beta_frames=100000*)
Bases: *pl_bolts.models.rl.common.memory.ReplayBuffer*

simple list based Prioritized Experience Replay Buffer Based on implementation found here: https://github.com/Shmuma/ptan/blob/master/ptan/experience.py#L371

append(*exp*)
Adds experiences from exp_source to the PER buffer

Parameters **exp** – experience tuple being added to the buffer

Return type None

sample(*batch_size=32*)
Takes a prioritized sample from the buffer

Parameters **batch_size** – size of sample

Return type Tuple

Returns sample of experiences chosen with ranked probability

update_beta(*step*)
Update the beta value which accounts for the bias in the PER

Parameters **step** – current global step

Return type float

Returns beta value for this indexed experience

update_priorities(*batch_indices*, *batch_priorities*)
Update the priorities from the last batch, this should be called after the loss for this batch has been calculated.

Parameters

- **batch_indices** (List) – index of each datum in the batch
- **batch_priorities** (List) – priority of each datum in the batch

Return type None

class pl_bolts.models.rl.common.memory.**ReplayBuffer**(*capacity*)
Bases: *pl_bolts.models.rl.common.memory.Buffer*

Replay Buffer for storing past experiences allowing the agent to learn from them

Parameters **capacity** (int) – size of the buffer

sample(*batch_size*)
Takes a sample of the buffer :type _sphinx_paramlinks_pl_bolts.models.rl.common.memory.ReplayBuffer.sample.batch_size int :param _sphinx_paramlinks_pl_bolts.models.rl.common.memory.ReplayBuffer.sample.batch_size: current batch_size

Return type Tuple

Returns a batch of tuple np arrays of state, action, reward, done, next_state

pl_bolts.models.rl.common.networks module

Series of networks used Based on implementations found here:

class pl_bolts.models.rl.common.networks.**CNN**(*input_shape*, *n_actions*)
Bases: torch.nn.Module

Simple MLP network

Parameters

- **input_shape** – observation shape of the environment
- **n_actions** – number of discrete actions available in the environment

_get_conv_out(*shape*)
Calculates the output size of the last conv layer

Parameters **shape** – input dimensions

Return type int

Returns size of the conv output

forward(*input_x*)
Forward pass through network

Parameters **x** – input to network

Return type Tensor

Returns output of network

class pl_bolts.models.rl.common.networks.**DuelingCNN**(*input_shape*, *n_actions*, *_=128*)
Bases: torch.nn.Module

CNN network with duel heads for val and advantage

Parameters

- **input_shape** (Tuple) – observation shape of the environment
- **n_actions** (int) – number of discrete actions available in the environment
- **hidden_size** – size of hidden layers

_get_conv_out(*shape*)
Calculates the output size of the last conv layer

Parameters **shape** – input dimensions

Return type int

Returns size of the conv output

adv_val(*input_x*)
Gets the advantage and value by passing out of the base network through the value and advantage heads

Parameters **input_x** – input to network

Returns advantage, value

forward(*input_x*)
Forward pass through network. Calculates the Q using the value and advantage

Parameters **input_x** – input to network

Returns Q value

class pl_bolts.models.rl.common.networks.**DuelingMLP**(*input_shape*, *n_actions*, *hidden_size=128*)
Bases: torch.nn.Module

MLP network with duel heads for val and advantage

Parameters

- **input_shape** (Tuple) – observation shape of the environment
- **n_actions** (int) – number of discrete actions available in the environment
- **hidden_size** (int) – size of hidden layers

adv_val(*input_x*)
Gets the advantage and value by passing out of the base network through the value and advantage heads

Parameters **input_x** – input to network

Return type Tuple[Tensor, Tensor]

Returns advantage, value

forward(*input_x*)
Forward pass through network. Calculates the Q using the value and advantage

Parameters **x** – input to network

Returns Q value

class pl_bolts.models.rl.common.networks.**MLP**(*input_shape*, *n_actions*, *hidden_size=128*)
Bases: torch.nn.Module

Simple MLP network

Parameters

- **input_shape** (Tuple) – observation shape of the environment
- **n_actions** (int) – number of discrete actions available in the environment
- **hidden_size** (int) – size of hidden layers

forward(*input_x*)
Forward pass through network

Parameters **x** – input to network

Returns output of network

class pl_bolts.models.rl.common.networks.**NoisyCNN**(*input_shape*, *n_actions*)
Bases: torch.nn.Module

CNN with Noisy Linear layers for exploration

Parameters

- **input_shape** – observation shape of the environment
- **n_actions** – number of discrete actions available in the environment

_get_conv_out(*shape*)
Calculates the output size of the last conv layer

Parameters **shape** – input dimensions

Return type int

Returns size of the conv output

forward(*input_x*)
Forward pass through network

Parameters **x** – input to network

Return type Tensor

Returns output of network

class pl_bolts.models.rl.common.networks.**NoisyLinear**(*in_features*, *out_features*, *sigma_init=0.017*, *bias=True*)
Bases: torch.nn.Linear

Noisy Layer using Independent Gaussian Noise.

based on https://github.com/PacktPublishing/Deep-Reinforcement-Learning-Hands-On-Second-Edition/blob/master/ Chapter08/lib/dqn_extra.py#L19

Parameters

- **in_features** – number of inputs
- **out_features** – number of outputs
- **sigma_init** – initial fill value of noisy weights
- **bias** – flag to include bias to linear layer

forward(*input_x*)
Forward pass of the layer

Parameters **input_x** (Tensor) – input tensor

Return type Tensor

Returns output of the layer

reset_parameters()
initializes or resets the paramseter of the layer

Return type None

pl_bolts.models.rl.common.wrappers module

Set of wrapper functions for gym environments taken from https://github.com/Shmuma/ptan/blob/master/ptan/common/wrappers.py

class pl_bolts.models.rl.common.wrappers.**BufferWrapper**(*env*, *n_steps*, *dtype=numpy.float32*)
Bases: gym.core.ObservationWrapper

"Wrapper for image stacking

observation(*observation*)
convert observation

reset()
reset env

class pl_bolts.models.rl.common.wrappers.**DataAugmentation**(*env=None*)
Bases: gym.core.ObservationWrapper

Carries out basic data augmentation on the env observations

- ToTensor
- GrayScale
- RandomCrop

observation(*obs*)
preprocess the obs

class pl_bolts.models.rl.common.wrappers.**FireResetEnv**(*env=None*)
Bases: gym.core.Wrapper

For environments where the user need to press FIRE for the game to start.

reset()
reset the env

step(*action*)
Take 1 step

class pl_bolts.models.rl.common.wrappers.**ImageToPyTorch**(*env*)
Bases: gym.core.ObservationWrapper

converts image to pytorch format

static observation(*observation*)
convert observation

class pl_bolts.models.rl.common.wrappers.**MaxAndSkipEnv**(*env=None*, *skip=4*)
Bases: gym.core.Wrapper

Return only every *skip*-th frame

reset()
Clear past frame buffer and init. to first obs. from inner env.

step(*action*)
take 1 step

class pl_bolts.models.rl.common.wrappers.**ProcessFrame84**(*env=None*)
Bases: gym.core.ObservationWrapper

preprocessing images from env

observation(*obs*)
preprocess the obs

static process(*frame*)
image preprocessing, formats to 84x84

class pl_bolts.models.rl.common.wrappers.**ScaledFloatFrame**(*env*)
Bases: gym.core.ObservationWrapper

scales the pixels

static observation(*obs*)

class pl_bolts.models.rl.common.wrappers.**ToTensor**(*env=None*)
Bases: gym.core.Wrapper

For environments where the user need to press FIRE for the game to start.

reset()
reset the env and cast to tensor

step(*action*)
Take 1 step and cast to tensor

pl_bolts.models.rl.common.wrappers.**make_env**(*env_name*)
Convert environment with wrappers

Submodules

pl_bolts.models.rl.double_dqn_model module

Double DQN

class pl_bolts.models.rl.double_dqn_model.**DoubleDQN**(*env*, *gpus=0*, *eps_start=1.0*, *eps_end=0.02*, *eps_last_frame=150000*, *sync_rate=1000*, *gamma=0.99*, *learning_rate=0.0001*, *batch_size=32*, *replay_size=100000*, *warm_start_size=10000*, *num_samples=500*, ***kwargs*)
Bases: *pl_bolts.models.rl.dqn_model.DQN*

Double Deep Q-network (DDQN) PyTorch Lightning implementation of Double DQN

Paper authors: Hado van Hasselt, Arthur Guez, David Silver

Model implemented by:

- *Donal Byrne <https://github.com/djbyrne>*

Example

```
>>> from pl_bolts.models.rl.double_dqn_model import DoubleDQN
...
>>> model = DoubleDQN("PongNoFrameskip-v4")
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

Parameters

- **env** (str) – gym environment tag
- **gpus** (int) – number of gpus being used
- **eps_start** (float) – starting value of epsilon for the epsilon-greedy exploration
- **eps_end** (float) – final value of epsilon for the epsilon-greedy exploration
- **eps_last_frame** (int) – the final frame in for the decrease of epsilon. At this frame espilon = eps_end
- **sync_rate** (int) – the number of iterations between syncing up the target network with the train network
- **gamma** (float) – discount factor
- **lr** – learning rate
- **batch_size** (int) – size of minibatch pulled from the DataLoader
- **replay_size** (int) – total capacity of the replay buffer
- **warm_start_size** (int) – how many random steps through the environment to be carried out at the start of training to fill the buffer with a starting point
- **sample_len** – the number of samples to pull from the dataset iterator and feed to the DataLoader

---

Note:

This example is based on: https://github.com/PacktPublishing/Deep-Reinforcement-Learning-Hands-On-Second-Edition/blob/master/Chapter08/03_dqn_double.py

---

Note: Currently only supports CPU and single GPU training with *distributed_backend=dp*

PyTorch Lightning implementation of DQN

Paper authors: Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, Martin Riedmiller.

Model implemented by:

- *Donal Byrne <https://github.com/djbyrne>*

Example

```
>>> from pl_bolts.models.rl.dqn_model import DQN
...
>>> model = DQN("PongNoFrameskip-v4")
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

Parameters

- **env** (str) – gym environment tag
- **gpus** (int) – number of gpus being used
- **eps_start** (float) – starting value of epsilon for the epsilon-greedy exploration
- **eps_end** (float) – final value of epsilon for the epsilon-greedy exploration
- **eps_last_frame** (int) – the final frame in for the decrease of epsilon. At this frame espilon = eps_end
- **sync_rate** (int) – the number of iterations between syncing up the target network with the train network
- **gamma** (float) – discount factor
- **learning_rate** (float) – learning rate
- **batch_size** (int) – size of minibatch pulled from the DataLoader
- **replay_size** (int) – total capacity of the replay buffer
- **warm_start_size** (int) – how many random steps through the environment to be carried out at the start of training to fill the buffer with a starting point
- **num_samples** (int) – the number of samples to pull from the dataset iterator and feed to the DataLoader

---

Note:

This example is based on: https://github.com/PacktPublishing/Deep-Reinforcement-Learning-Hands-On-Second-Edition/blob/master/Chapter06/02_dqn_pong.py

---

Note: Currently only supports CPU and single GPU training with *distributed_backend=dp*

---

training_step(*batch*, *_*)

Carries out a single step through the environment to update the replay buffer. Then calculates loss based on the minibatch recieved

Parameters

- **batch** (Tuple[Tensor, Tensor]) – current mini batch of replay data
- _ – batch number, not used

Return type OrderedDict

Returns Training loss and log metrics

pl_bolts.models.rl.dqn_model module

Deep Q Network

class pl_bolts.models.rl.dqn_model.**DQN**(*env*, *gpus=0*, *eps_start=1.0*, *eps_end=0.02*, *eps_last_frame=150000*, *sync_rate=1000*, *gamma=0.99*, *learning_rate=0.0001*, *batch_size=32*, *replay_size=100000*, *warm_start_size=10000*, *num_samples=500*, ***kwargs*)

Bases: pytorch_lightning.LightningModule

Basic DQN Model

PyTorch Lightning implementation of DQN

Paper authors: Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, Martin Riedmiller.

Model implemented by:

- *Donal Byrne <https://github.com/djbyrne>*

Example

```
>>> from pl_bolts.models.rl.dqn_model import DQN
...
>>> model = DQN("PongNoFrameskip-v4")
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

Parameters

- **env** (str) – gym environment tag
- **gpus** (int) – number of gpus being used
- **eps_start** (float) – starting value of epsilon for the epsilon-greedy exploration
- **eps_end** (float) – final value of epsilon for the epsilon-greedy exploration
- **eps_last_frame** (int) – the final frame in for the decrease of epsilon. At this frame espilon = eps_end
- **sync_rate** (int) – the number of iterations between syncing up the target network with the train network
- **gamma** (float) – discount factor
- **learning_rate** (float) – learning rate

- **batch_size** (int) – size of minibatch pulled from the DataLoader
- **replay_size** (int) – total capacity of the replay buffer
- **warm_start_size** (int) – how many random steps through the environment to be carried out at the start of training to fill the buffer with a starting point
- **num_samples** (int) – the number of samples to pull from the dataset iterator and feed to the DataLoader

---

Note:

This example is based on: https://github.com/PacktPublishing/Deep-Reinforcement-Learning-Hands-On-Second-Edition/blob/master/Chapter06/02_dqn_pong.py

---

Note: Currently only supports CPU and single GPU training with *distributed_backend=dp*

---

static add_model_specific_args(*arg_parser*)
Adds arguments for DQN model

Note: these params are fine tuned for Pong env

Parameters **arg_parser** (ArgumentParser) – parent parser

Return type ArgumentParser

build_networks()
Initializes the DQN train and target networks

Return type None

configure_optimizers()
Initialize Adam optimizer

Return type List[Optimizer]

forward(*x*)
Passes in a state x through the network and gets the q_values of each action as an output

Parameters **x** (Tensor) – environment state

Return type Tensor

Returns q values

populate(*warm_start*)
Populates the buffer with initial experience

Return type None

prepare_data()
Initialize the Replay Buffer dataset used for retrieving experiences

Return type None

test_dataloader()
Get test loader

Return type DataLoader

test_epoch_end(*outputs*)
Log the avg of the test results

Return type Dict[str, Tensor]

test_step(**args*, ***kwargs*)
Evaluate the agent for 10 episodes

Return type Dict[str, Tensor]

train_dataloader()
Get train loader

Return type DataLoader

training_step(*batch*, *_*)
Carries out a single step through the environment to update the replay buffer. Then calculates loss based on the minibatch recieved

Parameters

- **batch** (Tuple[Tensor, Tensor]) – current mini batch of replay data
- **_** – batch number, not used

Return type OrderedDict

Returns Training loss and log metrics

pl_bolts.models.rl.dueling_dqn_model module

Dueling DQN

class pl_bolts.models.rl.dueling_dqn_model.**DuelingDQN**(*env*, *gpus=0*, *eps_start=1.0*, *eps_end=0.02*, *eps_last_frame=150000*, *sync_rate=1000*, *gamma=0.99*, *learning_rate=0.0001*, *batch_size=32*, *replay_size=100000*, *warm_start_size=10000*, *num_samples=500*, ***kwargs*)

Bases: *pl_bolts.models.rl.dqn_model.DQN*

PyTorch Lightning implementation of Dueling DQN

Paper authors: Ziyu Wang, Tom Schaul, Matteo Hessel, Hado van Hasselt, Marc Lanctot, Nando de Freitas

Model implemented by:

- *Donal Byrne <https://github.com/djbyrne>*

Example

```
>>> from pl_bolts.models.rl.dueling_dqn_model import DuelingDQN
...
>>> model = DuelingDQN("PongNoFrameskip-v4")
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

Parameters

- **env** (str) – gym environment tag
- **gpus** (int) – number of gpus being used
- **eps_start** (float) – starting value of epsilon for the epsilon-greedy exploration
- **eps_end** (float) – final value of epsilon for the epsilon-greedy exploration
- **eps_last_frame** (int) – the final frame in for the decrease of epsilon. At this frame espilon = eps_end
- **sync_rate** (int) – the number of iterations between syncing up the target network with the train network
- **gamma** (float) – discount factor
- **lr** – learning rate
- **batch_size** (int) – size of minibatch pulled from the DataLoader
- **replay_size** (int) – total capacity of the replay buffer
- **warm_start_size** (int) – how many random steps through the environment to be carried out at the start of training to fill the buffer with a starting point
- **sample_len** – the number of samples to pull from the dataset iterator and feed to the DataLoader

---

Note: Currently only supports CPU and single GPU training with *distributed_backend=dp*

PyTorch Lightning implementation of DQN

Paper authors: Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, Martin Riedmiller.

Model implemented by:

- *Donal Byrne <https://github.com/djbyrne>*

Example

```
>>> from pl_bolts.models.rl.dqn_model import DQN
...
>>> model = DQN("PongNoFrameskip-v4")
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

Parameters

- **env** (str) – gym environment tag
- **gpus** (int) – number of gpus being used
- **eps_start** (float) – starting value of epsilon for the epsilon-greedy exploration
- **eps_end** (float) – final value of epsilon for the epsilon-greedy exploration
- **eps_last_frame** (int) – the final frame in for the decrease of epsilon. At this frame espilon = eps_end
- **sync_rate** (int) – the number of iterations between syncing up the target network with the train network
- **gamma** (float) – discount factor
- **learning_rate** (float) – learning rate
- **batch_size** (int) – size of minibatch pulled from the DataLoader
- **replay_size** (int) – total capacity of the replay buffer
- **warm_start_size** (int) – how many random steps through the environment to be carried out at the start of training to fill the buffer with a starting point
- **num_samples** (int) – the number of samples to pull from the dataset iterator and feed to the DataLoader

---

Note:

This example is based on: https://github.com/PacktPublishing/Deep-Reinforcement-Learning-Hands-On-Second-Edition/blob/master/Chapter06/02_dqn_pong.py

---

Note: Currently only supports CPU and single GPU training with *distributed_backend=dp*

---

build_networks()

Initializes the Dueling DQN train and target networks

Return type None

pl_bolts.models.rl.n_step_dqn_model module

N Step DQN

class pl_bolts.models.rl.n_step_dqn_model.**NStepDQN**(*env*, *gpus=0*, *eps_start=1.0*, *eps_end=0.02*, *eps_last_frame=150000*, *sync_rate=1000*, *gamma=0.99*, *learning_rate=0.0001*, *batch_size=32*, *replay_size=100000*, *warm_start_size=10000*, *num_samples=500*, *n_steps=4*, ***kwargs*)

Bases: *pl_bolts.models.rl.dqn_model.DQN*

NStep DQN Model

PyTorch Lightning implementation of N-Step DQN

Paper authors: Richard Sutton

Model implemented by:

- *Donal Byrne <https://github.com/djbyrne>*

Example

```
>>> from pl_bolts.models.rl.dqn_model import DQN
...
>>> model = NStepDQN("PongNoFrameskip-v4")
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

Parameters

- **env** (str) – gym environment tag
- **gpus** (int) – number of gpus being used
- **eps_start** (float) – starting value of epsilon for the epsilon-greedy exploration
- **eps_end** (float) – final value of epsilon for the epsilon-greedy exploration
- **eps_last_frame** (int) – the final frame in for the decrease of epsilon. At this frame espilon = eps_end
- **sync_rate** (int) – the number of iterations between syncing up the target network with the train network
- **gamma** (float) – discount factor
- **learning_rate** (float) – learning rate
- **batch_size** (int) – size of minibatch pulled from the DataLoader
- **replay_size** (int) – total capacity of the replay buffer

- **warm_start_size** (int) – how many random steps through the environment to be carried out at the start of training to fill the buffer with a starting point
- **num_samples** (int) – the number of samples to pull from the dataset iterator and feed to the DataLoader
- **n_steps** – number of steps to approximate and use in the bellman update

Note: Currently only supports CPU and single GPU training with *distributed_backend=dp*

pl_bolts.models.rl.noisy_dqn_model module

Noisy DQN

class pl_bolts.models.rl.noisy_dqn_model.**NoisyDQN**(*env*, *gpus=0*, *eps_start=1.0*, *eps_end=0.02*, *eps_last_frame=150000*, *sync_rate=1000*, *gamma=0.99*, *learning_rate=0.0001*, *batch_size=32*, *replay_size=100000*, *warm_start_size=10000*, *num_samples=500*, ***kwargs*)

Bases: *pl_bolts.models.rl.dqn_model.DQN*

PyTorch Lightning implementation of Noisy DQN

Paper authors: Meire Fortunato, Mohammad Gheshlaghi Azar, Bilal Piot, Jacob Menick, Ian Osband, Alex Graves, Vlad Mnih, Remi Munos, Demis Hassabis, Olivier Pietquin, Charles Blundell, Shane Legg

Model implemented by:

- *Donal Byrne <https://github.com/djbyrne>*

Example

```
>>> from pl_bolts.models.rl.n_step_dqn_model import NStepDQN
...
>>> model = NStepDQN("PongNoFrameskip-v4")
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

Parameters

- **env** (str) – gym environment tag
- **gpus** (int) – number of gpus being used
- **eps_start** (float) – starting value of epsilon for the epsilon-greedy exploration
- **eps_end** (float) – final value of epsilon for the epsilon-greedy exploration

- **eps_last_frame** (int) – the final frame in for the decrease of epsilon. At this frame espilon = eps_end
- **sync_rate** (int) – the number of iterations between syncing up the target network with the train network
- **gamma** (float) – discount factor
- **lr** – learning rate
- **batch_size** (int) – size of minibatch pulled from the DataLoader
- **replay_size** (int) – total capacity of the replay buffer
- **warm_start_size** (int) – how many random steps through the environment to be carried out at the start of
- **to fill the buffer with a starting point** (*training*) –
- **sample_len** – the number of samples to pull from the dataset iterator and feed to the DataLoader

---

Note: Currently only supports CPU and single GPU training with *distributed_backend=dp*

PyTorch Lightning implementation of DQN

Paper authors: Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, Martin Riedmiller.

Model implemented by:

- *Donal Byrne <https://github.com/djbyrne>*

Example

```
>>> from pl_bolts.models.rl.dqn_model import DQN
...
>>> model = DQN("PongNoFrameskip-v4")
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

Parameters

- **env** (str) – gym environment tag
- **gpus** (int) – number of gpus being used
- **eps_start** (float) – starting value of epsilon for the epsilon-greedy exploration
- **eps_end** (float) – final value of epsilon for the epsilon-greedy exploration
- **eps_last_frame** (int) – the final frame in for the decrease of epsilon. At this frame espilon = eps_end
- **sync_rate** (int) – the number of iterations between syncing up the target network with the train network
- **gamma** (float) – discount factor

- **learning_rate** (float) – learning rate
- **batch_size** (int) – size of minibatch pulled from the DataLoader
- **replay_size** (int) – total capacity of the replay buffer
- **warm_start_size** (int) – how many random steps through the environment to be carried out at the start of training to fill the buffer with a starting point
- **num_samples** (int) – the number of samples to pull from the dataset iterator and feed to the DataLoader

---

Note:

This example is based on: https://github.com/PacktPublishing/Deep-Reinforcement-Learning-Hands-On-Second-Edition/blob/master/Chapter06/02_dqn_pong.py

---

Note: Currently only supports CPU and single GPU training with *distributed_backend=dp*

---

build_networks()
Initializes the Noisy DQN train and target networks

Return type None

on_train_start()
Set the agents epsilon to 0 as the exploration comes from the network

Return type None

training_step(*batch*, *_*)
Carries out a single step through the environment to update the replay buffer. Then calculates loss based on the minibatch recieved

Parameters

- **batch** (Tuple[Tensor, Tensor]) – current mini batch of replay data
- **_** – batch number, not used

Return type OrderedDict

Returns Training loss and log metrics

pl_bolts.models.rl.per_dqn_model module

Prioritized Experience Replay DQN

class pl_bolts.models.rl.per_dqn_model.**PERDQN**(*env*, *gpus=0*, *eps_start=1.0*, *eps_end=0.02*, *eps_last_frame=150000*, *sync_rate=1000*, *gamma=0.99*, *learning_rate=0.0001*, *batch_size=32*, *replay_size=100000*, *warm_start_size=10000*, *num_samples=500*, ***kwargs*)

Bases: *pl_bolts.models.rl.dqn_model.DQN*

PyTorch Lightning implementation of DQN With Prioritized Experience Replay

Paper authors: Tom Schaul, John Quan, Ioannis Antonoglou, David Silver

Model implemented by:

- *Donal Byrne <https://github.com/djbyrne>*

Example

```
>>> from pl_bolts.models.rl.per_dqn_model import PERDQN
...
>>> model = PERDQN("PongNoFrameskip-v4")
```

Train:

```
trainer = Trainer()
trainer.fit(model)

Args:
    env: gym environment tag
    gpus: number of gpus being used
    eps_start: starting value of epsilon for the epsilon-greedy exploration
    eps_end: final value of epsilon for the epsilon-greedy exploration
    eps_last_frame: the final frame in for the decrease of epsilon. At this frame␣
↪espilon = eps_end
    sync_rate: the number of iterations between syncing up the target network␣
↪with the train network
    gamma: discount factor
    learning_rate: learning rate
    batch_size: size of minibatch pulled from the DataLoader
    replay_size: total capacity of the replay buffer
    warm_start_size: how many random steps through the environment to be carried␣
↪out at the start of
        training to fill the buffer with a starting point
    num_samples: the number of samples to pull from the dataset iterator and feed␣
↪to the DataLoader

.. note::
    This example is based on:
     https://github.com/PacktPublishing/Deep-Reinforcement-Learning-Hands-On-
↪Second-Edition             /blob/master/Chapter08/05_dqn_prio_replay.py

.. note:: Currently only supports CPU and single GPU training with `distributed_
↪backend=dp`
```

PyTorch Lightning implementation of DQN

Paper authors: Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, Martin Riedmiller.

Model implemented by:

- *Donal Byrne <https://github.com/djbyrne>*

Example

```
>>> from pl_bolts.models.rl.dqn_model import DQN
...
>>> model = DQN("PongNoFrameskip-v4")
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

Parameters

- **env** (str) – gym environment tag
- **gpus** (int) – number of gpus being used
- **eps_start** (float) – starting value of epsilon for the epsilon-greedy exploration
- **eps_end** (float) – final value of epsilon for the epsilon-greedy exploration
- **eps_last_frame** (int) – the final frame in for the decrease of epsilon. At this frame espilon = eps_end
- **sync_rate** (int) – the number of iterations between syncing up the target network with the train network
- **gamma** (float) – discount factor
- **learning_rate** (float) – learning rate
- **batch_size** (int) – size of minibatch pulled from the DataLoader
- **replay_size** (int) – total capacity of the replay buffer
- **warm_start_size** (int) – how many random steps through the environment to be carried out at the start of training to fill the buffer with a starting point
- **num_samples** (int) – the number of samples to pull from the dataset iterator and feed to the DataLoader

---

Note:

This example is based on: https://github.com/PacktPublishing/Deep-Reinforcement-Learning-Hands-On-Second-Edition/blob/master/Chapter06/02_dqn_pong.py

---

Note: Currently only supports CPU and single GPU training with *distributed_backend=dp*

---

prepare_data()

Initialize the Replay Buffer dataset used for retrieving experiences

Return type None

training_step(*batch*, _)

Carries out a single step through the environment to update the replay buffer. Then calculates loss based on the minibatch recieved

Parameters

- **batch** – current mini batch of replay data
- _ – batch number, not used

Return type OrderedDict

Returns Training loss and log metrics

pl_bolts.models.rl.reinforce_model module

class pl_bolts.models.rl.reinforce_model.**Reinforce**(*env*, *gamma=0.99*, *lr=0.0001*, *batch_size=32*, *batch_episodes=4*, ***kwargs*)

Bases: pytorch_lightning.LightningModule

Basic REINFORCE Policy Model

PyTorch Lightning implementation of REINFORCE

Paper authors: Richard S. Sutton, David McAllester, Satinder Singh, Yishay Mansour

Model implemented by:

- *Donal Byrne <https://github.com/djbyrne>*

Example

```
>>> from pl_bolts.models.rl.reinforce_model import Reinforce
...
>>> model = Reinforce("PongNoFrameskip-v4")
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

Parameters

- **env** (str) – gym environment tag
- **gamma** (float) – discount factor
- **lr** (float) – learning rate
- **batch_size** (int) – size of minibatch pulled from the DataLoader
- **batch_episodes** (int) – how many episodes to rollout for each batch of training

---

Note:

This example is based on: https://github.com/PacktPublishing/Deep-Reinforcement-Learning-Hands-On-Second-Edition/blob/master/Chapter11/02_cartpole_reinforce.py

---

Note: Currently only supports CPU and single GPU training with *distributed_backend=dp*

_dataloader()
Initialize the Replay Buffer dataset used for retrieving experiences

Return type DataLoader

static add_model_specific_args(*arg_parser*)
Adds arguments for DQN model

Note: these params are fine tuned for Pong env

Parameters **arg_parser** – the current argument parser to add to

Return type ArgumentParser

Returns arg_parser with model specific cargs added

build_networks()
Initializes the DQN train and target networks

Return type None

calc_qvals(*rewards*)
Takes in the rewards for each batched episode and returns list of qvals for each batched episode

Parameters **rewards** (List[List]) – list of rewards for each episodes in the batch

Return type List[List]

Returns List of qvals for each episodes

configure_optimizers()
Initialize Adam optimizer

Return type List[Optimizer]

static flatten_batch(*batch_actions*, *batch_qvals*, *batch_rewards*, *batch_states*)
Takes in the outputs of the processed batch and flattens the several episodes into a single tensor for each batched output

Parameters

- **batch_actions** (List[List[Tensor]]) – actions taken in each batch episodes
- **batch_qvals** (List[List[Tensor]]) – Q vals for each batch episode
- **batch_rewards** (List[List[Tensor]]) – reward for each batch episode
- **batch_states** (List[List[Tuple[Tensor, Tensor]]]) – states for each batch episodes

Return type Tuple[Tensor, Tensor, Tensor, Tensor]

Returns The input batched results flattend into a single tensor

forward(*x*)
Passes in a state x through the network and gets the q_values of each action as an output

Parameters **x** (Tensor) – environment state

Return type Tensor

Returns q values

get_device(*batch*)
Retrieve device currently being used by minibatch

Return type str

loss(*batch_qvals*, *batch_states*, *batch_actions*)
Calculates the mse loss using a batch of states, actions and Q values from several episodes. These have all been flattend into a single tensor.

Parameters

- **batch_qvals** (List[Tensor]) – current mini batch of q values
- **batch_actions** (List[Tensor]) – current batch of actions
- **batch_states** (List[Tensor]) – current batch of states

Return type Tensor

Returns loss

process_batch(*batch*)
Takes in a batch of episodes and retrieves the q vals, the states and the actions for the batch

Parameters **batch** (List[List[*Experience*]]) – list of episodes, each containing a list of Experiences

Return type Tuple[List[Tensor], List[Tensor], List[Tensor]]

Returns q_vals, states and actions used for calculating the loss

train_dataloader()
Get train loader

Return type DataLoader

training_step(*batch*, *_*)
Carries out a single step through the environment to update the replay buffer. Then calculates loss based on the minibatch recieved

Parameters

- **batch** (Tuple[Tensor, Tensor]) – current mini batch of replay data
- **_** – batch number, not used

Return type OrderedDict

Returns Training loss and log metrics

pl_bolts.models.rl.vanilla_policy_gradient_model module

Vanilla Policy Gradient

class pl_bolts.models.rl.vanilla_policy_gradient_model.**PolicyGradient**(*env*, *gamma=0.99*, *lr=0.0001*, *batch_size=32*, *entropy_beta=0.01*, *batch_episodes=4*, **args*, ***kwargs*)

Bases: pytorch_lightning.LightningModule

Vanilla Policy Gradient Model

PyTorch Lightning implementation of Vanilla Policy Gradient

Paper authors: Richard S. Sutton, David McAllester, Satinder Singh, Yishay Mansour

Model implemented by:

- *Donal Byrne <https://github.com/djbyrne>*

Example

```
>>> from pl_bolts.models.rl.vanilla_policy_gradient_model import PolicyGradient
...
>>> model = PolicyGradient("PongNoFrameskip-v4")
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

Parameters

- **env** (str) – gym environment tag
- **gamma** (float) – discount factor
- **lr** (float) – learning rate
- **batch_size** (int) – size of minibatch pulled from the DataLoader
- **batch_episodes** (int) – how many episodes to rollout for each batch of training
- **entropy_beta** (float) – dictates the level of entropy per batch

---

Note:

This example is based on: https://github.com/PacktPublishing/Deep-Reinforcement-Learning-Hands-On-Second-Edition/blob/master/Chapter11/04_cartpole_pg.py

---

Note: Currently only supports CPU and single GPU training with *distributed_backend=dp*

_dataloader()
Initialize the Replay Buffer dataset used for retrieving experiences

Return type DataLoader

static add_model_specific_args(*arg_parser*)
Adds arguments for DQN model

Note: these params are fine tuned for Pong env

Parameters **parent** –

Return type ArgumentParser

build_networks()
Initializes the DQN train and target networks

Return type None

calc_entropy_loss(*log_prob*, *logits*)
Calculates the entropy to be added to the loss function :type _sphinx_paramlinks_pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradient.calc_entropy_loss.log_prob: Tensor :param _sphinx_paramlinks_pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradient.calc_entropy_loss.l log probabilities for each action :type _sphinx_paramlinks_pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradient Tensor :param _sphinx_paramlinks_pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradient.calc_entropy_loss.l the raw outputs of the network

Return type Tensor

Returns entropy penalty for each state

static calc_policy_loss(*batch_actions*, *batch_qvals*, *batch_states*, *logits*)
Calculate the policy loss give the batch outputs and logits :type _sphinx_paramlinks_pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradient.calc_policy_loss.batch_actions: Tensor :param _sphinx_paramlinks_pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradient.calc_policy_loss.ba actions from batched episodes :type _sphinx_paramlinks_pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradient.c Tensor :param _sphinx_paramlinks_pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradient.calc_policy_loss.ba Q values from batched episodes :type _sphinx_paramlinks_pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradien Tensor :param _sphinx_paramlinks_pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradient.calc_policy_loss.ba states from batched episodes :type _sphinx_paramlinks_pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradient.ca Tensor :param _sphinx_paramlinks_pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradient.calc_policy_loss.lo raw output of the network given the batch_states

Return type Tuple[List, Tensor]

Returns policy loss

calc_qvals(*rewards*)
Takes in the rewards for each batched episode and returns list of qvals for each batched episode

Parameters **rewards** (List[Tensor]) – list of rewards for each episodes in the batch

Return type List[Tensor]

Returns List of qvals for each episodes

configure_optimizers()
Initialize Adam optimizer

Return type List[Optimizer]

static flatten_batch(*batch_actions*, *batch_qvals*, *batch_rewards*, *batch_states*)
Takes in the outputs of the processed batch and flattens the several episodes into a single tensor for each batched output

Parameters

- **batch_actions** (List[List[Tensor]]) – actions taken in each batch episodes
- **batch_qvals** (List[List[Tensor]]) – Q vals for each batch episode
- **batch_rewards** (List[List[Tensor]]) – reward for each batch episode
- **batch_states** (List[List[Tuple[Tensor, Tensor]]]) – states for each batch episodes

Return type Tuple[Tensor, Tensor, Tensor, Tensor]

Returns The input batched results flattend into a single tensor

forward(*x*)
Passes in a state x through the network and gets the q_values of each action as an output

Parameters **x** (Tensor) – environment state

Return type Tensor

Returns q values

get_device(*batch*)
Retrieve device currently being used by minibatch

Return type str

loss(*batch_qvals*, *batch_states*, *batch_actions*)
Calculates the mse loss using a batch of states, actions and Q values from several episodes. These have all been flattend into a single tensor.

Parameters

- **batch_qvals** (List[Tensor]) – current mini batch of q values
- **batch_actions** (List[Tensor]) – current batch of actions
- **batch_states** (List[Tensor]) – current batch of states

Return type Tensor

Returns loss

process_batch(*batch*)
Takes in a batch of episodes and retrieves the q vals, the states and the actions for the batch

Parameters **batch** (List[List[*Experience*]]) – list of episodes, each containing a list of Experiences

Return type Tuple[List[Tensor], List[Tensor], List[Tensor]]

Returns q_vals, states and actions used for calculating the loss

train_dataloader()
Get train loader

Return type DataLoader

training_step(*batch*, _)
Carries out a single step through the environment to update the replay buffer. Then calculates loss based on the minibatch recieved

Parameters

- **batch** (Tuple[Tensor, Tensor]) – current mini batch of replay data
- _ – batch number, not used

Return type OrderedDict

Returns Training loss and log metrics

pl_bolts.models.self_supervised package

These models have been pre-trained using self-supervised learning. The models can also be used without pre-training and overwritten for your own research.

Here's an example for using these as pretrained models.

```
from pl_bolts.models.self_supervised import CPCV2

images = get_imagenet_batch()

# extract unsupervised representations
pretrained = CPCV2(pretrained=True)
representations = pretrained(images)

# use these in classification or any downstream task
classifications = classifier(representations)
```

Subpackages

pl_bolts.models.self_supervised.amdim package

Submodules

pl_bolts.models.self_supervised.amdim.amdim_module module

class pl_bolts.models.self_supervised.amdim.amdim_module.**AMDIM**(*datamodule='cifar10'*, *encoder='amdim_encoder'*, *contrastive_task=torch.nn.Module*, *image_channels=3*, *image_height=32*, *encoder_feature_dim=320*, *embedding_fx_dim=1280*, *conv_block_depth=10*, *use_bn=False*, *tclip=20.0*, *learning_rate=0.0002*, *data_dir=''*, *num_classes=10*, *batch_size=200*, ***kwargs*)

Bases: pytorch_lightning.LightningModule

PyTorch Lightning implementation of Augmented Multiscale Deep InfoMax (AMDIM)

Paper authors: Philip Bachman, R Devon Hjelm, William Buchwalter.

Model implemented by: William Falcon

This code is adapted to Lightning using the original author repo (the original repo).

Example

```
>>> from pl_bolts.models.self_supervised import AMDIM
...
>>> model = AMDIM(encoder='resnet18')
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

Parameters

- **datamodule** (Union[str, *LightningDataModule*]) – A LightningDatamodule
- **encoder** (Union[str, Module, LightningModule]) – an encoder string or model
- **image_channels** (int) – 3
- **image_height** (int) – pixels
- **encoder_feature_dim** (int) – Called *ndf* in the paper, this is the representation size for the encoder.
- **embedding_fx_dim** (int) – Output dim of the embedding function (*nrkhs* in the paper) (Reproducing Kernel Hilbert Spaces).
- **conv_block_depth** (int) – Depth of each encoder block,
- **use_bn** (bool) – If true will use batchnorm.
- **tclip** (int) – soft clipping non-linearity to the scores after computing the regularization term and before computing the log-softmax. This is the 'second trick' used in the paper
- **learning_rate** (int) – The learning rate
- **data_dir** (str) – Where to store data
- **num_classes** (int) – How many classes in the dataset
- **batch_size** (int) – The batch size

static add_model_specific_args(*parent_parser*)

configure_optimizers()

forward(*img_1*, *img_2*)

init_encoder()

train_dataloader()

training_step(*batch*, *batch_nb*)

training_step_end(*outputs*)

val_dataloader()

validation_epoch_end(*outputs*)

validation_step(*batch*, *batch_nb*)

pl_bolts.models.self_supervised.amdim.datasets module

class pl_bolts.models.self_supervised.amdim.datasets.**AMDIMPatchesPretraining**
Bases: object

" For pretraining we use the train transform for both train and val.

static cifar10(*dataset_root*, *patch_size*, *patch_overlap*, *split='train'*)

static imagenet(*dataset_root*, *nb_classes*, *patch_size*, *patch_overlap*, *split='train'*)

static stl(*dataset_root*, *patch_size*, *patch_overlap*, *split=None*)

class pl_bolts.models.self_supervised.amdim.datasets.**AMDIMPretraining**
Bases: object

" For pretraining we use the train transform for both train and val.

static cifar10(*dataset_root*, *split='train'*)

static cifar10_tiny(*dataset_root*, *split='train'*)

static get_dataset(*datamodule*, *data_dir*, *split='train'*, ***kwargs*)

static imagenet(*dataset_root*, *nb_classes*, *split='train'*)

static stl(*dataset_root*, *split=None*)

pl_bolts.models.self_supervised.amdim.networks module

class pl_bolts.models.self_supervised.amdim.networks.**AMDIMEncoder**(*dummy_batch*, *num_channels=3*, *encoder_feature_dim=64*, *embedding_fx_dim=512*, *conv_block_depth=3*, *encoder_size=32*, *use_bn=False*)
Bases: torch.nn.Module

_config_modules(*x*, *output_widths*, *n_rkhs*, *use_bn*)
Configure the modules for extracting fake rkhs embeddings for infomax.

_forward_acts(*x*)
Return activations from all layers.

forward(*x*)

init_weights(*init_scale=1.0*)
Run custom weight init for modules...

class pl_bolts.models.self_supervised.amdim.networks.**Conv3x3**(*n_in*, *n_out*, *n_kern*, *n_stride*, *n_pad*, *use_bn=True*, *pad_mode='constant'*)
Bases: torch.nn.Module

forward(*x*)

class pl_bolts.models.self_supervised.amdim.networks.**ConvResBlock**(*n_in*, *n_out*, *width*, *stride*, *pad*, *depth*, *use_bn*)

Bases: torch.nn.Module

forward(*x*)

init_weights(*init_scale=1.0*)

Do a fixup-ish init for each ConvResNxN in this block.

class pl_bolts.models.self_supervised.amdim.networks.**ConvResNxN**(*n_in*, *n_out*, *width*, *stride*, *pad*, *use_bn=False*)

Bases: torch.nn.Module

forward(*x*)

init_weights(*init_scale=1.0*)

class pl_bolts.models.self_supervised.amdim.networks.**FakeRKHSConvNet**(*n_input*, *n_output*, *use_bn=False*)

Bases: torch.nn.Module

forward(*x*)

init_weights(*init_scale=1.0*)

class pl_bolts.models.self_supervised.amdim.networks.**MaybeBatchNorm2d**(*n_ftr*, *affine*, *use_bn*)

Bases: torch.nn.Module

forward(*x*)

class pl_bolts.models.self_supervised.amdim.networks.**NopNet**(*norm_dim=None*)

Bases: torch.nn.Module

forward(*x*)

pl_bolts.models.self_supervised.amdim.ssl_datasets module

class pl_bolts.models.self_supervised.amdim.ssl_datasets.**CIFAR10Mixed**(*root*, *split='val'*, *transform=None*, *target_transform=None*, *download=False*, *nb_labeled_per_class=None*, *val_pct=0.1*)

Bases: *pl_bolts.models.self_supervised.amdim.ssl_datasets.SSLDatasetMixin*, torchvision.datasets.CIFAR10

class pl_bolts.models.self_supervised.amdim.ssl_datasets.**SSLDatasetMixin**

Bases: abc.ABC

classmethod deterministic_shuffle(*x*, *y*)

classmethod generate_train_val_split(*examples*, *labels*, *pct_val*)
Splits dataset uniformly across classes :param _sphinx_paramlinks_pl_bolts.models.self_supervised.amdim.ssl_datasets.SSL
:param _sphinx_paramlinks_pl_bolts.models.self_supervised.amdim.ssl_datasets.SSLDatasetMixin.generate_train_val_split
:param _sphinx_paramlinks_pl_bolts.models.self_supervised.amdim.ssl_datasets.SSLDatasetMixin.generate_train_val_split
:return:

classmethod select_nb_imgs_per_class(*examples*, *labels*, *nb_imgs_in_val*)
Splits a dataset into two parts. The labeled split has nb_imgs_in_val per class :param _sphinx_paramlinks_pl_bolts.models.self_supervised.amdim.ssl_datasets.SSLDatasetMixin.select_nb_imgs_per_class.exam
:param _sphinx_paramlinks_pl_bolts.models.self_supervised.amdim.ssl_datasets.SSLDatasetMixin.select_nb_imgs_per_cla
:param _sphinx_paramlinks_pl_bolts.models.self_supervised.amdim.ssl_datasets.SSLDatasetMixin.select_nb_imgs_per_cla
:return:

pl_bolts.models.self_supervised.amdim.transforms module

class pl_bolts.models.self_supervised.amdim.transforms.**AMDIMEvalTransformsCIFAR10**
Bases: object

Transforms applied to AMDIM

Transforms:

```
transforms.ToTensor(),
normalize
```

Example:

```
x = torch.rand(5, 3, 32, 32)

transform = AMDIMEvalTransformsCIFAR10()
(view1, view2) = transform(x)
```

__call__(*inp*)
Call self as a function.

class pl_bolts.models.self_supervised.amdim.transforms.**AMDIMEvalTransformsImageNet128**(*height*
Bases: object

Transforms applied to AMDIM

Transforms:

```
transforms.Resize(height + 6, interpolation=3),
transforms.CenterCrop(height),
transforms.ToTensor(),
normalize
```

Example:

```
x = torch.rand(5, 3, 128, 128)

transform = AMDIMEvalTransformsImageNet128()
view1 = transform(x)
```

__call__(*inp*)
Call self as a function.

class pl_bolts.models.self_supervised.amdim.transforms.**AMDIMEvalTransformsSTL10**(*height=64*)
Bases: object

Transforms applied to AMDIM

Transforms:

```
transforms.Resize(height + 6, interpolation=3),
transforms.CenterCrop(height),
transforms.ToTensor(),
normalize
```

Example:

```
x = torch.rand(5, 3, 64, 64)

transform = AMDIMTrainTransformsSTL10()
view1 = transform(x)
```

__call__(*inp*)
Call self as a function.

class pl_bolts.models.self_supervised.amdim.transforms.**AMDIMTrainTransformsCIFAR10**
Bases: object

Transforms applied to AMDIM

Transforms:

```
img_jitter,
col_jitter,
rnd_gray,
transforms.ToTensor(),
normalize
```

Example:

```
x = torch.rand(5, 3, 32, 32)

transform = AMDIMTrainTransformsCIFAR10()
(view1, view2) = transform(x)
```

__call__(*inp*)
Call self as a function.

class pl_bolts.models.self_supervised.amdim.transforms.**AMDIMTrainTransformsImageNet128**(*heigh*
Bases: object

Transforms applied to AMDIM

Transforms:

```
img_jitter,
col_jitter,
rnd_gray,
transforms.ToTensor(),
normalize
```

Example:

```
x = torch.rand(5, 3, 128, 128)

transform = AMDIMTrainTransformsSTL10()
(view1, view2) = transform(x)
```

__call__(*inp*)
Call self as a function.

class pl_bolts.models.self_supervised.amdim.transforms.**AMDIMTrainTransformsSTL10**(*height=64*)
Bases: object

Transforms applied to AMDIM

Transforms:

```
img_jitter,
col_jitter,
rnd_gray,
transforms.ToTensor(),
normalize
```

Example:

```
x = torch.rand(5, 3, 64, 64)

transform = AMDIMTrainTransformsSTL10()
(view1, view2) = transform(x)
```

__call__(*inp*)
Call self as a function.

pl_bolts.models.self_supervised.cpc package

Submodules

pl_bolts.models.self_supervised.cpc.cpc_module module

CPC V2

class pl_bolts.models.self_supervised.cpc.cpc_module.**CPCV2**(*datamodule=None*, *encoder='cpc_encoder'*, *patch_size=8*, *patch_overlap=4*, *online_ft=True*, *task='cpc'*, *num_workers=4*, *learning_rate=0.0001*, *data_dir=''*, *batch_size=32*, *pretrained=None*, ***kwargs*)
Bases: pytorch_lightning.LightningModule

PyTorch Lightning implementation of Data-Efficient Image Recognition with Contrastive Predictive Coding

Paper authors: (Olivier J. Hénaff, Aravind Srinivas, Jeffrey De Fauw, Ali Razavi, Carl Doersch, S. M. Ali Eslami, Aaron van den Oord).

Model implemented by:

- William Falcon
- Tullie Murrell

Example

```
>>> from pl_bolts.models.self_supervised import CPCV2
...
>>> model = CPCV2()
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

CLI command:

```
# cifar10
python cpc_module.py --gpus 1

# imagenet
python cpc_module.py
    --gpus 8
    --dataset imagenet2012
    --data_dir /path/to/imagenet/
    --meta_dir /path/to/folder/with/meta.bin/
    --batch_size 32
```

Some uses:

```
# load resnet18 pretrained using CPC on imagenet
model = CPCV2(encoder='resnet18', pretrained=True)
resnet18 = model.encoder
renset18.freeze()

# it supportes any torchvision resnet
model = CPCV2(encoder='resnet50', pretrained=True)

# use it as a feature extractor
x = torch.rand(2, 3, 224, 224)
out = model(x)
```

Parameters

- **datamodule** (Optional[*LightningDataModule*]) – A Datamodule (optional). Otherwise set the dataloaders directly
- **encoder** (Union[str, Module, LightningModule]) – A string for any of the resnets in torchvision, or the original CPC encoder, or a custon nn.Module encoder
- **patch_size** (int) – How big to make the image patches

- **patch_overlap** (int) – How much overlap should each patch have.
- **online_ft** (int) – Enable a 1024-unit MLP to fine-tune online
- **task** (str) – Which self-supervised task to use ('cpc', 'amdim', etc...)
- **num_workers** (int) – num dataloader worksers
- **learning_rate** (int) – what learning rate to use
- **data_dir** (str) – where to store data
- **batch_size** (int) – batch size
- **pretrained** (Optional[str]) – If true, will use the weights pretrained (using CPC) on Imagenet

_CPCV2__compute_final_nb_c(*patch_size*)

_CPCV2__recover_z_shape(*Z*, *b*)

static add_model_specific_args(*parent_parser*)

configure_optimizers()

forward(*img_1*)

init_encoder()

load_pretrained(*encoder*)

prepare_data()

train_dataloader()

training_step(*batch*, *batch_nb*)

val_dataloader()

validation_epoch_end(*outputs*)

validation_step(*batch*, *batch_nb*)

pl_bolts.models.self_supervised.cpc.networks module

class pl_bolts.models.self_supervised.cpc.networks.**CPCResNet101**(*sample_batch*, *zero_init_residual=False*, *groups=1*, *width_per_group=64*, *replace_stride_with_dilation=None*, *norm_layer=None*)

Bases: torch.nn.Module

_make_layer(*sample_batch*, *block*, *planes*, *blocks*, *stride=1*, *dilate=False*, *expansion=4*)

flatten(*x*)

forward(*x*)

class pl_bolts.models.self_supervised.cpc.networks.**LNBottleneck**(*sample_batch*, *inplanes*, *planes*, *stride=1*, *downsample_conv=None*, *groups=1*, *base_width=64*, *dilation=1*, *norm_layer=None*, *expansion=4*)

Bases: torch.nn.Module

_LNBottleneck__init_layer_norms(*x*, *conv1*, *conv2*, *conv3*, *downsample_conv*)

forward(*x*)

pl_bolts.models.self_supervised.cpc.networks.**conv1x1**(*in_planes*, *out_planes*, *stride=1*)

1x1 convolution

pl_bolts.models.self_supervised.cpc.networks.**conv3x3**(*in_planes*, *out_planes*, *stride=1*, *groups=1*, *dilation=1*)

3x3 convolution with padding

pl_bolts.models.self_supervised.cpc.transforms module

class pl_bolts.models.self_supervised.cpc.transforms.**CPCEvalTransformsCIFAR10**(*patch_size=8*, *overlap=4*)

Bases: object

Transforms used for CPC:

Parameters

- **patch_size** – size of patches when cutting up the image into overlapping patches
- **overlap** – how much to overlap patches

Transforms:

```
random_flip
transforms.ToTensor()
normalize
Patchify(patch_size=patch_size, overlap_size=overlap)
```

Example:

```
# in a regular dataset
CIFAR10(..., transforms=CPCEvalTransformsCIFAR10())

# in a DataModule
module = CIFAR10DataModule(PATH)
train_loader = module.train_dataloader(batch_size=32,␣
↪transforms=CPCEvalTransformsCIFAR10())
```

__call__(*inp*)

Call self as a function.

class pl_bolts.models.self_supervised.cpc.transforms.**CPCEvalTransformsImageNet128**(*patch_size=*, *overlap=16*)

Bases: object

Transforms used for CPC:

Parameters

- **patch_size** – size of patches when cutting up the image into overlapping patches
- **overlap** – how much to overlap patches

Transforms:

```
random_flip
transforms.ToTensor()
normalize
Patchify(patch_size=patch_size, overlap_size=patch_size // 2)
```

Example:

```
# in a regular dataset
Imagenet(..., transforms=CPCEvalTransformsImageNet128())

# in a DataModule
module = ImagenetDataModule(PATH)
train_loader = module.train_dataloader(batch_size=32,␣
↪transforms=CPCEvalTransformsImageNet128())
```

__call__(*inp*)

Call self as a function.

class pl_bolts.models.self_supervised.cpc.transforms.**CPCEvalTransformsSTL10**(*patch_size=16*, *overlap=8*)

Bases: object

Transforms used for CPC:

Parameters

- **patch_size** – size of patches when cutting up the image into overlapping patches
- **overlap** – how much to overlap patches

Transforms:

```
random_flip
transforms.ToTensor()
normalize
Patchify(patch_size=patch_size, overlap_size=patch_size // 2)
```

Example:

```
# in a regular dataset
STL10(..., transforms=CPCEvalTransformsSTL10())

# in a DataModule
module = STL10DataModule(PATH)
train_loader = module.train_dataloader(batch_size=32,␣
↪transforms=CPCEvalTransformsSTL10())
```

__call__(*inp*)
Call self as a function.

class pl_bolts.models.self_supervised.cpc.transforms.**CPCTrainTransformsCIFAR10**(*patch_size=8*, *overlap=4*)

Bases: object

Transforms used for CPC:

Parameters

- **patch_size** – size of patches when cutting up the image into overlapping patches
- **overlap** – how much to overlap patches

Transforms:

```
random_flip
img_jitter
col_jitter
rnd_gray
transforms.ToTensor()
normalize
Patchify(patch_size=patch_size, overlap_size=patch_size // 2)
```

Example:

```
# in a regular dataset
CIFAR10(..., transforms=CPCTrainTransformsCIFAR10())

# in a DataModule
module = CIFAR10DataModule(PATH)
train_loader = module.train_dataloader(batch_size=32,␣
↪transforms=CPCTrainTransformsCIFAR10())
```

__call__(*inp*)
Call self as a function.

class pl_bolts.models.self_supervised.cpc.transforms.**CPCTrainTransformsImageNet128**(*patch_size=*, *overlap=16*)

Bases: object

Transforms used for CPC:

Parameters

- **patch_size** – size of patches when cutting up the image into overlapping patches
- **overlap** – how much to overlap patches

Transforms:

```
random_flip
transforms.ToTensor()
normalize
Patchify(patch_size=patch_size, overlap_size=patch_size // 2)
```

Example:

```
# in a regular dataset
Imagenet(..., transforms=CPCTrainTransformsImageNet128())

# in a DataModule
module = ImagenetDataModule(PATH)
train_loader = module.train_dataloader(batch_size=32,␣
↪transforms=CPCTrainTransformsImageNet128())
```

__call__(*inp*)
Call self as a function.

class pl_bolts.models.self_supervised.cpc.transforms.**CPCTrainTransformsSTL10**(*patch_size=16*, *overlap=8*)

Bases: object

Transforms used for CPC:

Parameters

- **patch_size** – size of patches when cutting up the image into overlapping patches
- **overlap** – how much to overlap patches

Transforms:

```
random_flip
img_jitter
col_jitter
rnd_gray
transforms.ToTensor()
normalize
Patchify(patch_size=patch_size, overlap_size=patch_size // 2)
```

Example:

```
# in a regular dataset
STL10(..., transforms=CPCTrainTransformsSTL10())

# in a DataModule
module = STL10DataModule(PATH)
train_loader = module.train_dataloader(batch_size=32,␣
↪transforms=CPCTrainTransformsSTL10())
```

__call__(*inp*)
Call self as a function.

pl_bolts.models.self_supervised.moco package

Submodules

pl_bolts.models.self_supervised.moco.callbacks module

class pl_bolts.models.self_supervised.moco.callbacks.**MocoLRScheduler**(*initial_lr=0.03*, *use_cosine_scheduler=False*, *schedule=(120, 160)*, *max_epochs=200*)

Bases: pytorch_lightning.Callback

on_epoch_start(*trainer*, *pl_module*)

pl_bolts.models.self_supervised.moco.moco2_module module

Adapted from: https://github.com/facebookresearch/moco

Original work is: Copyright (c) Facebook, Inc. and its affiliates. All Rights Reserved This implementation is: Copyright (c) PyTorch Lightning, Inc. and its affiliates. All Rights Reserved

class pl_bolts.models.self_supervised.moco.moco2_module.**MocoV2**(*base_encoder='resnet18'*, *emb_dim=128*, *num_negatives=65536*, *encoder_momentum=0.999*, *softmax_temperature=0.07*, *learning_rate=0.03*, *momentum=0.9*, *weight_decay=0.0001*, *datamodule=None*, *data_dir='./'*, *batch_size=256*, *use_mlp=False*, *num_workers=8*, **args*, ***kwargs*)

Bases: pytorch_lightning.LightningModule

PyTorch Lightning implementation of Moco

Paper authors: Xinlei Chen, Haoqi Fan, Ross Girshick, Kaiming He.

Code adapted from facebookresearch/moco to Lightning by:

- William Falcon

Example

```
>>> from pl_bolts.models.self_supervised import MocoV2
...
>>> model = MocoV2()
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

CLI command:

```
# cifar10
python moco2_module.py --gpus 1

# imagenet
python moco2_module.py
    --gpus 8
    --dataset imagenet2012
    --data_dir /path/to/imagenet/
    --meta_dir /path/to/folder/with/meta.bin/
    --batch_size 32
```

Parameters

- **base_encoder** (Union[str, Module]) – torchvision model name or torch.nn.Module
- **emb_dim** (int) – feature dimension (default: 128)
- **num_negatives** (int) – queue size; number of negative keys (default: 65536)
- **encoder_momentum** (float) – moco momentum of updating key encoder (default: 0.999)
- **softmax_temperature** (float) – softmax temperature (default: 0.07)
- **learning_rate** (float) – the learning rate
- **momentum** (float) – optimizer momentum
- **weight_decay** (float) – optimizer weight decay
- **datamodule** (Optional[*LightningDataModule*]) – the DataModule (train, val, test dataloaders)
- **data_dir** (str) – the directory to store data
- **batch_size** (int) – batch size
- **use_mlp** (bool) – add an mlp to the encoders
- **num_workers** (int) – workers for the loaders

_batch_shuffle_ddp(*x*)

Batch shuffle, for making use of BatchNorm. *** Only support DistributedDataParallel (DDP) model. ***

_batch_unshuffle_ddp(*x*, *idx_unshuffle*)

Undo batch shuffle. *** Only support DistributedDataParallel (DDP) model. ***

_dequeue_and_enqueue(*keys*)

_momentum_update_key_encoder()
Momentum update of the key encoder

static add_model_specific_args(*parent_parser*)

configure_optimizers()

forward(*img_q*, *img_k*)

Input: im_q: a batch of query images im_k: a batch of key images

Output: logits, targets

init_encoders(*base_encoder*)
Override to add your own encoders

prepare_data()

train_dataloader()

training_step(*batch*, *batch_idx*)

val_dataloader()

validation_epoch_end(*outputs*)

validation_step(*batch*, *batch_idx*)

pl_bolts.models.self_supervised.moco.moco2_module.**concat_all_gather**(*tensor*)
Performs all_gather operation on the provided tensors. * **Warning** *: torch.distributed.all_gather has no gradient.

pl_bolts.models.self_supervised.moco.transforms module

class pl_bolts.models.self_supervised.moco.transforms.**GaussianBlur**(*sigma=(0.1, 2.0)*)
Bases: object

Gaussian blur augmentation in SimCLR https://arxiv.org/abs/2002.05709

__call__(*x*)
Call self as a function.

class pl_bolts.models.self_supervised.moco.transforms.**Moco2EvalCIFAR10Transforms**(*height=32*)
Bases: object

Moco 2 augmentation: https://arxiv.org/pdf/2003.04297.pdf

__call__(*inp*)
Call self as a function.

class pl_bolts.models.self_supervised.moco.transforms.**Moco2EvalImagenetTransforms**(*height=128*)
Bases: object

Moco 2 augmentation: https://arxiv.org/pdf/2003.04297.pdf

__call__(*inp*)
Call self as a function.

class pl_bolts.models.self_supervised.moco.transforms.**Moco2EvalSTL10Transforms**(*height=64*)
Bases: object

Moco 2 augmentation: https://arxiv.org/pdf/2003.04297.pdf

__call__(*inp*)
Call self as a function.

class pl_bolts.models.self_supervised.moco.transforms.**Moco2TrainCIFAR10Transforms**(*height=32*)
Bases: object

Moco 2 augmentation: https://arxiv.org/pdf/2003.04297.pdf

__call__(*inp*)
Call self as a function.

class pl_bolts.models.self_supervised.moco.transforms.**Moco2TrainImagenetTransforms**(*height=128*
Bases: object

Moco 2 augmentation: https://arxiv.org/pdf/2003.04297.pdf

__call__(*inp*)
Call self as a function.

class pl_bolts.models.self_supervised.moco.transforms.**Moco2TrainSTL10Transforms**(*height=64*)
Bases: object

Moco 2 augmentation: https://arxiv.org/pdf/2003.04297.pdf

__call__(*inp*)
Call self as a function.

pl_bolts.models.self_supervised.simclr package

Submodules

pl_bolts.models.self_supervised.simclr.simclr_module module

class pl_bolts.models.self_supervised.simclr.simclr_module.**DensenetEncoder**
Bases: torch.nn.Module

forward(*x*)

class pl_bolts.models.self_supervised.simclr.simclr_module.**Projection**(*input_dim=1024*, *output_dim=128*)
Bases: torch.nn.Module

forward(*x*)

class pl_bolts.models.self_supervised.simclr.simclr_module.**SimCLR**(*datamodule=None*, *data_dir=''*, *learning_rate=6e-05*, *weight_decay=0.0005*, *input_height=32*, *batch_size=128*, *online_ft=False*, *num_workers=4*, *optimizer='lars'*, *lr_sched_step=30.0*, *lr_sched_gamma=0.5*, *lars_momentum=0.9*, *lars_eta=0.001*, *loss_temperature=0.5*, ***kwargs*)

Bases: pytorch_lightning.LightningModule

PyTorch Lightning implementation of SIMCLR

Paper authors: Ting Chen, Simon Kornblith, Mohammad Norouzi, Geoffrey Hinton.

Model implemented by:

- William Falcon
- Tullie Murrell

Example

```
>>> from pl_bolts.models.self_supervised import SimCLR
...
>>> model = SimCLR()
```

Train:

```
trainer = Trainer()
trainer.fit(model)
```

CLI command:

```
# cifar10
python simclr_module.py --gpus 1

# imagenet
python simclr_module.py
    --gpus 8
    --dataset imagenet2012
    --data_dir /path/to/imagenet/
    --meta_dir /path/to/folder/with/meta.bin/
    --batch_size 32
```

Parameters

- **datamodule** (Optional[*LightningDataModule*]) – The datamodule
- **data_dir** (str) – directory to store data
- **learning_rate** (float) – the learning rate
- **weight_decay** (float) – optimizer weight decay
- **input_height** (int) – image input height
- **batch_size** (int) – the batch size
- **online_ft** (bool) – whether to tune online or not
- **num_workers** (int) – number of workers
- **optimizer** (str) – optimizer name
- **lr_sched_step** (float) – step for learning rate scheduler
- **lr_sched_gamma** (float) – gamma for learning rate scheduler
- **lars_momentum** (float) – the mom param for lars optimizer
- **lars_eta** (float) – for lars optimizer
- **loss_temperature** (float) – float = 0.

static add_model_specific_args(*parent_parser*)

configure_optimizers()

forward(*x*)

init_encoder()

init_loss()

init_projection()

prepare_data()

train_dataloader()

training_step(*batch*, *batch_idx*)

val_dataloader()

validation_epoch_end(*outputs*)

validation_step(*batch*, *batch_idx*)

pl_bolts.models.self_supervised.simclr.simclr_transforms module

class pl_bolts.models.self_supervised.simclr.simclr_transforms.**GaussianBlur**(*kernel_size*, *min=0.1*, *max=2.0*)

Bases: object

__call__(*sample*)
Call self as a function.

class pl_bolts.models.self_supervised.simclr.simclr_transforms.**SimCLREvalDataTransform**(*input_s=1*)

Bases: object

Transforms for SimCLR

Transform:

```
Resize(input_height + 10, interpolation=3)
transforms.CenterCrop(input_height),
transforms.ToTensor()
```

Example:

```
from pl_bolts.models.self_supervised.simclr.transforms import␣
↪SimCLREvalDataTransform

transform = SimCLREvalDataTransform(input_height=32)
x = sample()
(xi, xj) = transform(x)
```

__call__(*sample*)
Call self as a function.

class pl_bolts.models.self_supervised.simclr.simclr_transforms.**SimCLRTrainDataTransform**(*inp*
s=1

Bases: object

Transforms for SimCLR

Transform:

```
RandomResizedCrop(size=self.input_height)
RandomHorizontalFlip()
RandomApply([color_jitter], p=0.8)
RandomGrayscale(p=0.2)
GaussianBlur(kernel_size=int(0.1 * self.input_height))
transforms.ToTensor()
```

Example:

```
from pl_bolts.models.self_supervised.simclr.transforms import␣
↪SimCLRTrainDataTransform

transform = SimCLRTrainDataTransform(input_height=32)
x = sample()
(xi, xj) = transform(x)
```

__call__(*sample*)
Call self as a function.

Submodules

pl_bolts.models.self_supervised.evaluator module

class pl_bolts.models.self_supervised.evaluator.**Flatten**
Bases: torch.nn.Module

forward(*input_tensor*)

class pl_bolts.models.self_supervised.evaluator.**SSLEvaluator**(*n_input*, *n_classes*, *n_hidden=512*, *p=0.1*)

Bases: torch.nn.Module

forward(*x*)

pl_bolts.models.self_supervised.resnets module

class pl_bolts.models.self_supervised.resnets.**ResNet**(*block*, *layers*, *num_classes=1000*, *zero_init_residual=False*, *groups=1*, *width_per_group=64*, *replace_stride_with_dilation=None*, *norm_layer=None*, *return_all_feature_maps=False*)

Bases: torch.nn.Module

_make_layer(*block*, *planes*, *blocks*, *stride=1*, *dilate=False*)

forward(*x*)

pl_bolts.models.self_supervised.resnets.**resnet18**(*pretrained=False*, *progress=True*, ***kwargs*)

ResNet-18 model from "Deep Residual Learning for Image Recognition" :param _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet18.pretrained: If True, returns a model pre-trained on ImageNet :type _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet18.pretrained: bool :param _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet18.progress: If True, displays a progress bar of the download to stderr :type _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet18.progress: bool

pl_bolts.models.self_supervised.resnets.**resnet34**(*pretrained=False*, *progress=True*, ***kwargs*)

ResNet-34 model from "Deep Residual Learning for Image Recognition" :param _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet34.pretrained: If True, returns a model pre-trained on ImageNet :type _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet34.pretrained: bool :param _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet34.progress: If True, displays a progress bar of the download to stderr :type _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet34.progress: bool

pl_bolts.models.self_supervised.resnets.**resnet50**(*pretrained=False*, *progress=True*, ***kwargs*)

ResNet-50 model from "Deep Residual Learning for Image Recognition" :param _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet50.pretrained: If True, returns a model pre-trained on ImageNet :type _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet50.pretrained: bool :param _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet50.progress: If True, displays a progress bar of the download to stderr :type _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet50.progress: bool

pl_bolts.models.self_supervised.resnets.**resnet101**(*pretrained=False*, *progress=True*, ***kwargs*)

ResNet-101 model from "Deep Residual Learning for Image Recognition" :param _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet101.pretrained: If True, returns a model pre-trained on ImageNet :type _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet101.pretrained: bool :param _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet101.progress: If True, displays a progress bar of the download to stderr :type _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet101.progress: bool

pl_bolts.models.self_supervised.resnets.**resnet152**(*pretrained=False*, *progress=True*, ***kwargs*)

ResNet-152 model from "Deep Residual Learning for Image Recognition" :param _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet152.pretrained: If True, returns a model

pre-trained on ImageNet :type _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet152.pretrained: bool :param _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet152.progress: If True, displays a progress bar of the download to stderr :type _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet152.progress: bool

pl_bolts.models.self_supervised.resnets.**resnext50_32x4d**(*pretrained=False*, *progress=True*, ***kwargs*)

ResNeXt-50 32x4d model from "Aggregated Residual Transformation for Deep Neural Networks" :param _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnext50_32x4d.pretrained: If True, returns a model pre-trained on ImageNet :type _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnext50_32x4d.pretrained: bool :param _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnext50_32x4d.progress: If True, displays a progress bar of the download to stderr :type _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnext50_32x4d.progress: bool

pl_bolts.models.self_supervised.resnets.**resnext101_32x8d**(*pretrained=False*, *progress=True*, ***kwargs*)

ResNeXt-101 32x8d model from "Aggregated Residual Transformation for Deep Neural Networks" :param _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnext101_32x8d.pretrained: If True, returns a model pre-trained on ImageNet :type _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnext101_32x8d.pretrained: bool :param _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnext101_32x8d.progress: If True, displays a progress bar of the download to stderr :type _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnext101_32x8d.progress: bool

pl_bolts.models.self_supervised.resnets.**wide_resnet50_2**(*pretrained=False*, *progress=True*, ***kwargs*)

Wide ResNet-50-2 model from "Wide Residual Networks" The model is the same as ResNet except for the bottleneck number of channels which is twice larger in every block. The number of channels in outer 1x1 convolutions is the same, e.g. last block in ResNet-50 has 2048-512-2048 channels, and in Wide ResNet-50-2 has 2048-1024-2048. :param _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.wide_resnet50_2.pretrained: If True, returns a model pre-trained on ImageNet :type _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.wide_resnet50_2.pretrained: bool :param _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.wide_resnet50_2.progress: If True, displays a progress bar of the download to stderr :type _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.wide_resnet50_2.progress: bool

pl_bolts.models.self_supervised.resnets.**wide_resnet101_2**(*pretrained=False*, *progress=True*, ***kwargs*)

Wide ResNet-101-2 model from "Wide Residual Networks" The model is the same as ResNet except for the bottleneck number of channels which is twice larger in every block. The number of channels in outer 1x1 convolutions is the same, e.g. last block in ResNet-50 has 2048-512-2048 channels, and in Wide ResNet-50-2 has 2048-1024-2048. :param _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.wide_resnet101_2.pretrained: If True, returns a model pre-trained on ImageNet :type _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.wide_resnet101_2.pretrained: bool :param _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.wide_resnet101_2.progress: If True, displays a progress bar of the download to stderr :type _sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.wide_resnet101_2.progress: bool

pl_bolts.models.vision package

Subpackages

pl_bolts.models.vision.image_gpt package

Submodules

pl_bolts.models.vision.image_gpt.gpt2 module

class pl_bolts.models.vision.image_gpt.gpt2.**Block**(*embed_dim*, *heads*)

Bases: torch.nn.Module

forward(*x*)

class pl_bolts.models.vision.image_gpt.gpt2.**GPT2**(*embed_dim*, *heads*, *layers*, *num_positions*, *vocab_size*, *num_classes*)

Bases: pytorch_lightning.LightningModule

GPT-2 from language Models are Unsupervised Multitask Learners

Paper by: Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever

Implementation contributed by:

- Teddy Koker

Example:

```
from pl_bolts.models import GPT2

seq_len = 17
batch_size = 32
vocab_size = 16
x = torch.randint(0, vocab_size, (seq_len, batch_size))
model = GPT2(embed_dim=32, heads=2, layers=2, num_positions=2, vocab_size=vocab_
↪size, num_classes=4)
results = model(x)
```

_init_embeddings()

_init_layers()

_init_sos_token()

forward(*x*, *classify=False*)

Expect input as shape [sequence len, batch] If classify, return classification logits

pl_bolts.models.vision.image_gpt.igpt_module module

class pl_bolts.models.vision.image_gpt.igpt_module.**ImageGPT**(*datamodule=None*, *embed_dim=16*, *heads=2*, *layers=2*, *pixels=28*, *vocab_size=16*, *num_classes=10*, *classify=False*, *batch_size=64*, *learning_rate=0.01*, *steps=25000*, *data_dir='.'*, *num_workers=8*, ***kwargs*)

Bases: pytorch_lightning.LightningModule

Paper: Generative Pretraining from Pixels [original paper code].

Paper by: Mark Che, Alec Radford, Rewon Child, Jeff Wu, Heewoo Jun, Prafulla Dhariwal, David Luan, Ilya Sutskever

Implementation contributed by:

- Teddy Koker

Original repo with results and more implementation details:

- https://github.com/teddykoker/image-gpt

Example Results (Photo credits: Teddy Koker):





Default arguments:

Table 1: Argument Defaults

Argument	Default	iGPT-S (Chen et al.)
–embed_dim	16	512
–heads	2	8
–layers	8	24
–pixels	28	32
–vocab_size	16	512
–num_classes	10	10
–batch_size	64	128
–learning_rate	0.01	0.01
–steps	25000	1000000

Example:

```
import pytorch_lightning as pl
from pl_bolts.models.vision import ImageGPT

dm = MNISTDataModule('.')
model = ImageGPT(dm)

pl.Trainer(gpu=4).fit(model)
```

As script:

```
cd pl_bolts/models/vision/image_gpt
python igpt_module.py --learning_rate 1e-2 --batch_size 32 --gpus 4
```

Parameters

- **datamodule** (Optional[*LightningDataModule*]) – LightningDataModule
- **embed_dim** (int) – the embedding dim
- **heads** (int) – number of attention heads
- **layers** (int) – number of layers
- **pixels** (int) – number of input pixels
- **vocab_size** (int) – vocab size
- **num_classes** (int) – number of classes in the input
- **classify** (bool) – true if should classify
- **batch_size** (int) – the batch size
- **learning_rate** (float) – learning rate
- **steps** (int) – number of steps for cosine annealing
- **data_dir** (str) – where to store data
- **num_workers** (int) – num_data workers

static add_model_specific_args(*parent_parser*)

configure_optimizers()

forward(*x*, *classify=False*)

prepare_data()

test_dataloader()

test_epoch_end(*outs*)

test_step(*batch*, *batch_idx*)

train_dataloader()

training_step(*batch*, *batch_idx*)

val_dataloader()

validation_epoch_end(*outs*)

validation_step(*batch*, *batch_idx*)

pl_bolts.models.vision.image_gpt.igpt_module.**_shape_input**(*x*)
shape batch of images for input into GPT2 model

Submodules

pl_bolts.models.vision.pixel_cnn module

PixelCNN Implemented by: William Falcon Reference: https://arxiv.org/pdf/1905.09272.pdf (page 15) Accessed: May 14, 2020

class pl_bolts.models.vision.pixel_cnn.**PixelCNN**(*input_channels*, *hidden_channels=256*, *num_blocks=5*)

Bases: torch.nn.Module

Implementation of Pixel CNN.

Paper authors: Aaron van den Oord, Nal Kalchbrenner, Oriol Vinyals, Lasse Espeholt, Alex Graves, Koray Kavukcuoglu

Implemented by:

- William Falcon

Example:

```
>>> from pl_bolts.models.vision import PixelCNN
>>> import torch
...
>>> model = PixelCNN(input_channels=3)
>>> x = torch.rand(5, 3, 64, 64)
>>> out = model(x)
...
>>> out.shape
torch.Size([5, 3, 64, 64])
```

conv_block(*input_channels*)

forward(*z*)

23.5.2 Submodules

pl_bolts.models.mnist_module module

class pl_bolts.models.mnist_module.**LitMNIST**(*hidden_dim=128*, *learning_rate=0.001*, *batch_size=32*, *num_workers=4*, *data_dir=''*)

Bases: pytorch_lightning.LightningModule

static add_model_specific_args(*parent_parser*)

configure_optimizers()

forward(*x*)

prepare_data()

test_dataloader()

test_epoch_end(*outputs*)

test_step(*batch*, *batch_idx*)

train_dataloader()

training_step(*batch*, *batch_idx*)

val_dataloader()

validation_epoch_end(*outputs*)

validation_step(*batch*, *batch_idx*)

23.6 pl_bolts.losses package

23.6.1 Submodules

pl_bolts.losses.rl module

Loss functions for the RL models

pl_bolts.losses.rl.**double_dqn_loss**(*batch*, *net*, *target_net*, *gamma=0.99*)

Calculates the mse loss using a mini batch from the replay buffer. This uses an improvement to the original DQN loss by using the double dqn. This is shown by using the actions of the train network to pick the value from the target network. This code is heavily commented in order to explain the process clearly

Parameters

- **batch** (Tuple[Tensor, Tensor]) – current mini batch of replay data
- **net** (Module) – main training network
- **target_net** (Module) – target network of the main training network
- **gamma** (float) – discount factor

Return type Tensor

Returns loss

pl_bolts.losses.rl.**dqn_loss**(*batch*, *net*, *target_net*, *gamma=0.99*)

Calculates the mse loss using a mini batch from the replay buffer

Parameters

- **batch** (Tuple[Tensor, Tensor]) – current mini batch of replay data
- **net** (Module) – main training network
- **target_net** (Module) – target network of the main training network
- **gamma** (float) – discount factor

Return type Tensor

Returns loss

pl_bolts.losses.rl.**per_dqn_loss**(*batch*, *batch_weights*, *net*, *target_net*, *gamma=0.99*)

Calculates the mse loss with the priority weights of the batch from the PER buffer

Parameters

- **batch** (Tuple[Tensor, Tensor]) – current mini batch of replay data
- **batch_weights** (List) – how each of these samples are weighted in terms of priority
- **net** (Module) – main training network
- **target_net** (Module) – target network of the main training network
- **gamma** (float) – discount factor

Return type Tuple[Tensor, ndarray]

Returns loss and batch_weights

pl_bolts.losses.self_supervised_learning module

class pl_bolts.losses.self_supervised_learning.**AmdimNCELoss**(*tclip*)

Bases: torch.nn.Module

forward(*anchor_representations*, *positive_representations*, *mask_mat*)

Compute the NCE scores for predicting r_src->r_trg. :param _sphinx_paramlinks_pl_bolts.losses.self_supervised_learning.AmdimNCELoss.forward.anchor_representations: (batch_size, emb_dim) :param _sphinx_paramlinks_pl_bolts.losses.self_supervised_learning.AmdimNCELoss.forward.posi (emb_dim, n_batch * w* h) (ie: nb_feat_vectors x embedding_dim) :param _sphinx_paramlinks_pl_bolts.losses.self_supervised_learning.AmdimNCELoss.forward.mask_mat: (n_batch_gpu, n_batch)

Output: raw_scores : (n_batch_gpu, n_locs) nce_scores : (n_batch_gpu, n_locs) lgt_reg : scalar

class pl_bolts.losses.self_supervised_learning.**CPCTask**(*num_input_channels*, *target_dim=64*, *embed_scale=0.1*)

Bases: torch.nn.Module

Loss used in CPC

compute_loss_h(*targets*, *preds*, *i*)

forward(*Z*)

class pl_bolts.losses.self_supervised_learning.**FeatureMapContrastiveTask**(*comparisons='00, 11', tclip=10.0, bidirectional=True*)

Bases: torch.nn.Module

Performs an anchor, positive negative pair comparison for each each tuple of feature maps passed.

```
# extract feature maps
pos_0, pos_1, pos_2 = encoder(x_pos)
anc_0, anc_1, anc_2 = encoder(x_anchor)

# compare only the 0th feature maps
task = FeatureMapContrastiveTask('00')
loss, regularizer = task((pos_0), (anc_0))

# compare (pos_0 to anc_1) and (pos_0, anc_2)
task = FeatureMapContrastiveTask('01, 02')
losses, regularizer = task((pos_0, pos_1, pos_2), (anc_0, anc_1, anc_2))
loss = losses.sum()

# compare (pos_1 vs a anc_random)
task = FeatureMapContrastiveTask('0r')
loss, regularizer = task((pos_0, pos_1, pos_2), (anc_0, anc_1, anc_2))
```

Parameters

- **comparisons** (str) – groupings of feature map indices to compare (zero indexed, 'r' means random) ex: '00, 1r'
- **tclip** (float) – stability clipping value
- **bidirectional** (bool) – if true, does the comparison both ways

```
# with bidirectional the comparisons are done both ways
task = FeatureMapContrastiveTask('01, 02')

# will compare the following:
# 01: (pos_0, anc_1), (anc_0, pos_1)
# 02: (pos_0, anc_2), (anc_0, pos_2)
```

_FeatureMapContrastiveTask__cache_dimension_masks(**args*)

_FeatureMapContrastiveTask__compare_maps(*m1*, *m2*)

_sample_src_ftr(*r_cnv*, *masks*)

feat_size_w_mask(*w*, *feature_map*)

forward(*anchor_maps*, *positive_maps*)

Takes in a set of tuples, each tuple has two feature maps with all matching dimensions

Example

```
>>> import torch
>>> from pytorch_lightning import seed_everything
>>> seed_everything(0)
0
>>> a1 = torch.rand(3, 5, 2, 2)
>>> a2 = torch.rand(3, 5, 2, 2)
>>> b1 = torch.rand(3, 5, 2, 2)
>>> b2 = torch.rand(3, 5, 2, 2)
...
>>> task = FeatureMapContrastiveTask('01, 11')
...
>>> losses, regularizer = task((a1, a2), (b1, b2))
>>> losses
tensor([2.2351, 2.1902])
>>> regularizer
tensor(0.0324)
```

static parse_map_indexes(*comparisons*)

Example:

```
>>> FeatureMapContrastiveTask.parse_map_indexes('11')
[(1, 1)]
>>> FeatureMapContrastiveTask.parse_map_indexes('11,59')
[(1, 1), (5, 9)]
>>> FeatureMapContrastiveTask.parse_map_indexes('11,59, 2r')
[(1, 1), (5, 9), (2, -1)]
```

pl_bolts.losses.self_supervised_learning.**nt_xent_loss**(*out_1*, *out_2*, *temperature*)

Loss used in SimCLR

pl_bolts.losses.self_supervised_learning.**tanh_clip**(*x*, *clip_val=10.0*)

soft clip values to the range [-clip_val, +clip_val]

23.7 pl_bolts.loggers package

Collection of PyTorchLightning loggers

class pl_bolts.loggers.**TrainsLogger**(*project_name=None*, *task_name=None*, *task_type='training'*, *reuse_last_task_id=True*, *output_uri=None*, *auto_connect_arg_parser=True*, *auto_connect_frameworks=True*, *auto_resource_monitoring=True*)

Bases: pytorch_lightning.loggers.base.LightningLoggerBase

Log using allegro.ai TRAINS. Install it with pip:

```
pip install trains
```

Example

```
>>> from pytorch_lightning import Trainer
>>> trains_logger = TrainsLogger(
...     project_name='pytorch lightning',
...     task_name='default',
...     output_uri='.',
... )
TRAINS Task: ...
TRAINS results page: ...
>>> trainer = Trainer(logger=trains_logger)
```

Use the logger anywhere in your `LightningModule` as follows:

```
>>> from pytorch_lightning import LightningModule
>>> class LitModel(LightningModule):
...     def training_step(self, batch, batch_idx):
...         # example
...         self.logger.experiment.whatever_trains_supports(...)
...
...     def any_lightning_module_function_or_hook(self):
...         self.logger.experiment.whatever_trains_supports(...)
```

Parameters

- **project_name** (Optional[str]) – The name of the experiment's project. Defaults to None.
- **task_name** (Optional[str]) – The name of the experiment. Defaults to None.
- **task_type** (str) – The name of the experiment. Defaults to 'training'.
- **reuse_last_task_id** (bool) – Start with the previously used task id. Defaults to True.
- **output_uri** (Optional[str]) – Default location for output models. Defaults to None.
- **auto_connect_arg_parser** (bool) – Automatically grab the ArgumentParser and connect it with the task. Defaults to True.
- **auto_connect_frameworks** (bool) – If True, automatically patch to trains backend. Defaults to True.
- **auto_resource_monitoring** (bool) – If True, machine vitals will be sent along side the task scalars. Defaults to True.

Examples

```
>>> logger = TrainsLogger("pytorch lightning", "default", output_uri=".")
TRAINS Task: ...
TRAINS results page: ...
>>> logger.log_metrics({"val_loss": 1.23}, step=0)
>>> logger.log_text("sample test")
sample test
>>> import numpy as np
>>> logger.log_artifact("confusion matrix", np.ones((2, 3)))
>>> logger.log_image("passed", "Image 1", np.random.randint(0, 255, (200, 150, 3),
↪ dtype=np.uint8))
```

classmethod bypass_mode()
Returns the bypass mode state.

Note: *GITHUB_ACTIONS* env will automatically set bypass_mode to `True` unless overridden specifically with `TrainsLogger.set_bypass_mode(False)`.

Return type bool

Returns If True, all outside communication is skipped.

finalize(*status=None*)

Return type None

log_artifact(*name*, *artifact*, *metadata=None*, *delete_after_upload=False*)
Save an artifact (file/object) in TRAINS experiment storage.

Parameters

- **name** (str) – Artifact name. Notice! it will override the previous artifact if the name already exists.
- **artifact** (Union[str, Path, Dict[str, Any], ndarray, Image]) – Artifact object to upload. Currently supports:
  - string / pathlib.Path are treated as path to artifact file to upload If a wildcard or a folder is passed, a zip file containing the local files will be created and uploaded.
  - dict will be stored as .json file and uploaded
  - pandas.DataFrame will be stored as .csv.gz (compressed CSV file) and uploaded
  - numpy.ndarray will be stored as .npz and uploaded
  - PIL.Image.Image will be stored to .png file and uploaded
- **metadata** (Optional[Dict[str, Any]]) – Simple key/value dictionary to store on the artifact. Defaults to None.
- **delete_after_upload** (bool) – If True, the local artifact will be deleted (only applies if artifact is a local file). Defaults to False.

Return type None

log_hyperparams(*params*)
Log hyperparameters (numeric values) in TRAINS experiments.

Parameters **params** (Union[Dict[str, Any], Namespace]) – The hyperparameters that passed through the model.

Return type None

log_image(*title*, *series*, *image*, *step=None*)
Log Debug image in TRAINS experiment

Parameters

- **title** (str) – The title of the debug image, i.e. “failed”, “passed”.
- **series** (str) – The series name of the debug image, i.e. “Image 0”, “Image 1”.
- **image** (Union[str, ndarray, Image, Tensor]) – Debug image to log. If numpy.ndarray or torch.Tensor, the image is assumed to be the following:

- shape: CHW
- color space: RGB
- value range: [0., 1.] (float) or [0, 255] (uint8)

- **step** (Optional[int]) – Step number at which the metrics should be recorded. Defaults to None.

Return type None

log_metric(*title*, *series*, *value*, *step=None*)
Log metrics (numeric values) in TRAINS experiments. This method will be called by the users.

Parameters

- **title** (str) – The title of the graph to log, e.g. loss, accuracy.
- **series** (str) – The series name in the graph, e.g. classification, localization.
- **value** (float) – The value to log.
- **step** (Optional[int]) – Step number at which the metrics should be recorded. Defaults to None.

Return type None

log_metrics(*metrics*, *step=None*)
Log metrics (numeric values) in TRAINS experiments. This method will be called by Trainer.

Parameters

- **metrics** (Dict[str, float]) – The dictionary of the metrics. If the key contains "/", it will be split by the delimiter, then the elements will be logged as "title" and "series" respectively.
- **step** (Optional[int]) – Step number at which the metrics should be recorded. Defaults to None.

Return type None

log_text(*text*)
Log console text data in TRAINS experiment.

Parameters **text** (str) – The value of the log (data-point).

Return type None

classmethod set_bypass_mode(*bypass*)
Will bypass all outside communication, and will drop all logs. Should only be used in "standalone mode", when there is no access to the *trains-server*.

Parameters **bypass** (bool) – If True, all outside communication is skipped.

Return type None

classmethod set_credentials(*api_host=None*, *web_host=None*, *files_host=None*, *key=None*, *secret=None*)
Set new default TRAINS-server host and credentials. These configurations could be overridden by either OS environment variables or trains.conf configuration file.

Note: Credentials need to be set *prior* to Logger initialization.

Parameters

- **api_host** (Optional[str]) – Trains API server url, example: host='http://localhost:8008'
- **web_host** (Optional[str]) – Trains WEB server url, example: host='http://localhost:8080'
- **files_host** (Optional[str]) – Trains Files server url, example: host='http://localhost:8081'
- **key** (Optional[str]) – user key/secret pair, example: key='thisisakey123'
- **secret** (Optional[str]) – user key/secret pair, example: secret='thisisseceret123'

Return type None

_bypass = None

property experiment

Actual TRAINS object. To use TRAINS features in your LightningModule do the following.

Example:

```
self.logger.experiment.some_trains_function()
```

Return type Task

property id

ID is a uuid (string) representing this specific experiment in the entire system.

Return type Optional[str]

property name

Name is a human readable non-unique name (str) of the experiment.

Return type Optional[str]

property version

Return type Optional[str]

23.7.1 Submodules

pl_bolts.loggers.trains module

TRAINS

class pl_bolts.loggers.trains.**TrainsLogger**(*project_name=None*, *task_name=None*, *task_type='training'*, *reuse_last_task_id=True*, *output_uri=None*, *auto_connect_arg_parser=True*, *auto_connect_frameworks=True*, *auto_resource_monitoring=True*)

Bases: pytorch_lightning.loggers.base.LightningLoggerBase

Log using allegro.ai TRAINS. Install it with pip:

```
pip install trains
```

Example

```
>>> from pytorch_lightning import Trainer
>>> trains_logger = TrainsLogger(
...     project_name='pytorch lightning',
...     task_name='default',
...     output_uri='.',
... )
TRAINS Task: ...
TRAINS results page: ...
>>> trainer = Trainer(logger=trains_logger)
```

Use the logger anywhere in your `LightningModule` as follows:

```
>>> from pytorch_lightning import LightningModule
>>> class LitModel(LightningModule):
...     def training_step(self, batch, batch_idx):
...         # example
...         self.logger.experiment.whatever_trains_supports(...)
...
...     def any_lightning_module_function_or_hook(self):
...         self.logger.experiment.whatever_trains_supports(...)
```

Parameters

- **project_name** (Optional[str]) – The name of the experiment's project. Defaults to None.
- **task_name** (Optional[str]) – The name of the experiment. Defaults to None.
- **task_type** (str) – The name of the experiment. Defaults to 'training'.
- **reuse_last_task_id** (bool) – Start with the previously used task id. Defaults to True.
- **output_uri** (Optional[str]) – Default location for output models. Defaults to None.
- **auto_connect_arg_parser** (bool) – Automatically grab the ArgumentParser and connect it with the task. Defaults to True.
- **auto_connect_frameworks** (bool) – If True, automatically patch to trains backend. Defaults to True.
- **auto_resource_monitoring** (bool) – If True, machine vitals will be sent along side the task scalars. Defaults to True.

Examples

```
>>> logger = TrainsLogger("pytorch lightning", "default", output_uri=".")
TRAINS Task: ...
TRAINS results page: ...
>>> logger.log_metrics({"val_loss": 1.23}, step=0)
>>> logger.log_text("sample test")
sample test
>>> import numpy as np
>>> logger.log_artifact("confusion matrix", np.ones((2, 3)))
>>> logger.log_image("passed", "Image 1", np.random.randint(0, 255, (200, 150, 3),
↪ dtype=np.uint8))
```

classmethod bypass_mode()
Returns the bypass mode state.

Note: *GITHUB_ACTIONS* env will automatically set bypass_mode to `True` unless overridden specifically with `TrainsLogger.set_bypass_mode(False)`.

Return type bool

Returns If True, all outside communication is skipped.

finalize(*status=None*)

Return type None

log_artifact(*name*, *artifact*, *metadata=None*, *delete_after_upload=False*)
Save an artifact (file/object) in TRAINS experiment storage.

Parameters

- **name** (str) – Artifact name. Notice! it will override the previous artifact if the name already exists.
- **artifact** (Union[str, Path, Dict[str, Any], ndarray, Image]) – Artifact object to upload. Currently supports:
  - string / pathlib.Path are treated as path to artifact file to upload If a wildcard or a folder is passed, a zip file containing the local files will be created and uploaded.
  - dict will be stored as .json file and uploaded
  - pandas.DataFrame will be stored as .csv.gz (compressed CSV file) and uploaded
  - numpy.ndarray will be stored as .npz and uploaded
  - PIL.Image.Image will be stored to .png file and uploaded
- **metadata** (Optional[Dict[str, Any]]) – Simple key/value dictionary to store on the artifact. Defaults to None.
- **delete_after_upload** (bool) – If True, the local artifact will be deleted (only applies if artifact is a local file). Defaults to False.

Return type None

log_hyperparams(*params*)
Log hyperparameters (numeric values) in TRAINS experiments.

Parameters **params** (Union[Dict[str, Any], Namespace]) – The hyperparameters that passed through the model.

Return type None

log_image(*title*, *series*, *image*, *step=None*)
Log Debug image in TRAINS experiment

Parameters

- **title** (str) – The title of the debug image, i.e. “failed”, “passed”.
- **series** (str) – The series name of the debug image, i.e. “Image 0”, “Image 1”.
- **image** (Union[str, ndarray, Image, Tensor]) – Debug image to log. If numpy.ndarray or torch.Tensor, the image is assumed to be the following:

- shape: CHW
- color space: RGB
- value range: [0., 1.] (float) or [0, 255] (uint8)

- **step** (Optional[int]) – Step number at which the metrics should be recorded. Defaults to None.

Return type None

log_metric(*title*, *series*, *value*, *step=None*)
Log metrics (numeric values) in TRAINS experiments. This method will be called by the users.

Parameters

- **title** (str) – The title of the graph to log, e.g. loss, accuracy.
- **series** (str) – The series name in the graph, e.g. classification, localization.
- **value** (float) – The value to log.
- **step** (Optional[int]) – Step number at which the metrics should be recorded. Defaults to None.

Return type None

log_metrics(*metrics*, *step=None*)
Log metrics (numeric values) in TRAINS experiments. This method will be called by Trainer.

Parameters

- **metrics** (Dict[str, float]) – The dictionary of the metrics. If the key contains "/", it will be split by the delimiter, then the elements will be logged as "title" and "series" respectively.
- **step** (Optional[int]) – Step number at which the metrics should be recorded. Defaults to None.

Return type None

log_text(*text*)
Log console text data in TRAINS experiment.

Parameters **text** (str) – The value of the log (data-point).

Return type None

classmethod set_bypass_mode(*bypass*)
Will bypass all outside communication, and will drop all logs. Should only be used in "standalone mode", when there is no access to the *trains-server*.

Parameters **bypass** (bool) – If True, all outside communication is skipped.

Return type None

classmethod set_credentials(*api_host=None*, *web_host=None*, *files_host=None*, *key=None*, *secret=None*)
Set new default TRAINS-server host and credentials. These configurations could be overridden by either OS environment variables or trains.conf configuration file.

Note: Credentials need to be set *prior* to Logger initialization.

Parameters

- **api_host** (Optional[str]) – Trains API server url, example: host='http://localhost:8008'
- **web_host** (Optional[str]) – Trains WEB server url, example: host='http://localhost:8080'
- **files_host** (Optional[str]) – Trains Files server url, example: host='http://localhost:8081'
- **key** (Optional[str]) – user key/secret pair, example: key='thisisakey123'
- **secret** (Optional[str]) – user key/secret pair, example: secret='thisisseceret123'

Return type None

_bypass = None

property experiment

Actual TRAINS object. To use TRAINS features in your LightningModule do the following.

Example:

```
self.logger.experiment.some_trains_function()
```

Return type Task

property id

ID is a uuid (string) representing this specific experiment in the entire system.

Return type Optional[str]

property name

Name is a human readable non-unique name (str) of the experiment.

Return type Optional[str]

property version

Return type Optional[str]

23.8 pl_bolts.optimizers package

23.8.1 Submodules

pl_bolts.optimizers.layer_adaptive_scaling module

Layer-wise adaptive rate scaling for SGD in PyTorch! Adapted from: https://github.com/noahgolmant/pytorch-lars/blob/master/lars.py

class pl_bolts.optimizers.layer_adaptive_scaling.**LARS**(*params*, *lr=<required parameter>*, *momentum=0.9*, *weight_decay=0.0005*, *eta=0.001*, *max_epoch=200*)

Bases: torch.optim.optimizer.Optimizer

Implements layer-wise adaptive rate scaling for SGD.

Parameters

- **params** (Iterable) – iterable of parameters to optimize or dicts defining parameter groups
- **lr** (float) – base learning rate (gamma_0)
- **momentum** (float) – momentum factor (default: 0) ("m")
- **weight_decay** (float) – weight decay (L2 penalty) (default: 0) ("beta")
- **eta** (float) – LARS coefficient
- **max_epoch** (int) – maximum training epoch to determine polynomial LR decay.

Reference: Based on Algorithm 1 of the following paper by You, Gitman, and Ginsburg. Large Batch Training of Convolutional Networks: https://arxiv.org/abs/1708.03888

Example

optimizer = LARS(model.parameters(), lr=0.1, eta=1e-3) optimizer.zero_grad() loss_fn(model(input), target).backward() optimizer.step()

step(*epoch=None*, *closure=None*)
Performs a single optimization step.

Parameters

- **closure** (Optional[Callable]) – A closure that reevaluates the model and returns the loss.
- **epoch** – current epoch to calculate polynomial LR decay schedule. if None, uses self.epoch and increments it.

class pl_bolts.optimizers.layer_adaptive_scaling.**_RequiredParameter**
Bases: object

Singleton class representing a required parameter for an Optimizer.

23.9 pl_bolts.transforms package

23.9.1 Subpackages

pl_bolts.transforms.self_supervised package

Submodules

pl_bolts.transforms.self_supervised.ssl_transforms module

class pl_bolts.transforms.self_supervised.ssl_transforms.**Patchify**(*patch_size*, *overlap_size*)
Bases: object

__call__(*x*)
Call self as a function.

class pl_bolts.transforms.self_supervised.ssl_transforms.**RandomTranslateWithReflect**(*max_tran*
Bases: object

Translate image randomly Translate vertically and horizontally by n pixels where n is integer drawn uniformly independently for each axis from [-max_translation, max_translation]. Fill the uncovered blank area with reflect padding.

__call__(*old_image*)
Call self as a function.

23.9.2 Submodules

pl_bolts.transforms.dataset_normalizations module

pl_bolts.transforms.dataset_normalizations.**cifar10_normalization**()

pl_bolts.transforms.dataset_normalizations.**imagenet_normalization**()

pl_bolts.transforms.dataset_normalizations.**stl10_normalization**()

23.10 pl_bolts.utils package

23.10.1 Submodules

pl_bolts.utils.pretrained_weights module

pl_bolts.utils.pretrained_weights.**load_pretrained**(*model*, *class_name=None*)

pl_bolts.utils.self_supervised module

pl_bolts.utils.self_supervised.**torchvision_ssl_encoder**(*name*, *pretrained=False*, *return_all_feature_maps=False*)

pl_bolts.utils.semi_supervised module

class pl_bolts.utils.semi_supervised.**Identity**
Bases: torch.nn.Module

An identity class to replace arbitrary layers in pretrained models

Example:

```
from pl_bolts.utils import Identity

model = resnet18()
model.fc = Identity()
```

forward(*x*)

pl_bolts.utils.semi_supervised.**balance_classes**(*X*, *Y*, *batch_size*)
Makes sure each batch has an equal amount of data from each class. Perfect balance

Parameters

- **X** (ndarray) – input features
- **Y** (list) – mixed labels (ints)
- **batch_size** (int) – the ultimate batch size

pl_bolts.utils.semi_supervised.**generate_half_labeled_batches**(*smaller_set_X*, *smaller_set_Y*, *larger_set_X*, *larger_set_Y*, *batch_size*)

Given a labeled dataset and an unlabeled dataset, this function generates a joint pair where half the batches are labeled and the other half is not

PYTHON MODULE INDEX

p

pl_bolts.callbacks, 130
pl_bolts.callbacks.printing, 130
pl_bolts.callbacks.variational, 131
pl_bolts.datamodules, 132
pl_bolts.datamodules.base_dataset, 132
pl_bolts.datamodules.cifar10_datamodule, 132
pl_bolts.datamodules.cifar10_dataset, 134
pl_bolts.datamodules.concat_dataset, 137
pl_bolts.datamodules.fashion_mnist_datamodule, 137
pl_bolts.datamodules.imagenet_datamodule, 138
pl_bolts.datamodules.imagenet_dataset, 140
pl_bolts.datamodules.lightning_datamodule, 141
pl_bolts.datamodules.mnist_datamodule, 144
pl_bolts.datamodules.sklearn_datamodule, 145
pl_bolts.datamodules.ssl_imagenet_datamodule, 149
pl_bolts.datamodules.stl10_datamodule, 151
pl_bolts.loggers, 220
pl_bolts.loggers.trains, 224
pl_bolts.losses, 217
pl_bolts.losses.rl, 217
pl_bolts.losses.self_supervised_learning, 218
pl_bolts.metrics, 152
pl_bolts.metrics.aggregation, 152
pl_bolts.models, 153
pl_bolts.models.autoencoders, 153
pl_bolts.models.autoencoders.basic_ae, 153
pl_bolts.models.autoencoders.basic_ae.basic_ae_module, 154
pl_bolts.models.autoencoders.basic_ae.components, 155
pl_bolts.models.autoencoders.basic_vae, 155
pl_bolts.models.autoencoders.basic_vae.basic_vae_mo 156
pl_bolts.models.autoencoders.basic_vae.components, 157
pl_bolts.models.gans, 158
pl_bolts.models.gans.basic, 158
pl_bolts.models.gans.basic.basic_gan_module, 158
pl_bolts.models.gans.basic.components, 159
pl_bolts.models.mnist_module, 217
pl_bolts.models.regression, 160
pl_bolts.models.regression.linear_regression, 160
pl_bolts.models.regression.logistic_regression, 161
pl_bolts.models.rl, 161
pl_bolts.models.rl.common, 161
pl_bolts.models.rl.common.agents, 162
pl_bolts.models.rl.common.cli, 163
pl_bolts.models.rl.common.experience, 163
pl_bolts.models.rl.common.memory, 165
pl_bolts.models.rl.common.networks, 167
pl_bolts.models.rl.common.wrappers, 170
pl_bolts.models.rl.double_dqn_model, 171
pl_bolts.models.rl.dqn_model, 174
pl_bolts.models.rl.dueling_dqn_model, 176
pl_bolts.models.rl.n_step_dqn_model, 179
pl_bolts.models.rl.noisy_dqn_model, 180
pl_bolts.models.rl.per_dqn_model, 182
pl_bolts.models.rl.reinforce_model, 185
pl_bolts.models.rl.vanilla_policy_gradient_model, 187
pl_bolts.models.self_supervised, 191
pl_bolts.models.self_supervised.amdim, 191
pl_bolts.models.self_supervised.amdim.amdim_module,

191
pl_bolts.models.self_supervised.amdim.datasets,
193
pl_bolts.models.self_supervised.amdim.networks,
193
pl_bolts.models.self_supervised.amdim.ssl_datasets,
194
pl_bolts.models.self_supervised.amdim.transforms,
195
pl_bolts.models.self_supervised.cpc, 197
pl_bolts.models.self_supervised.cpc.cpc_module,
197
pl_bolts.models.self_supervised.cpc.networks,
199
pl_bolts.models.self_supervised.cpc.transforms,
200
pl_bolts.models.self_supervised.evaluator,
210
pl_bolts.models.self_supervised.moco,
203
pl_bolts.models.self_supervised.moco.callbacks,
204
pl_bolts.models.self_supervised.moco.moco2_module,
204
pl_bolts.models.self_supervised.moco.transforms,
206
pl_bolts.models.self_supervised.resnets,
211
pl_bolts.models.self_supervised.simclr,
207
pl_bolts.models.self_supervised.simclr.simclr_module,
207
pl_bolts.models.self_supervised.simclr.simclr_transforms,
209
pl_bolts.models.vision, 213
pl_bolts.models.vision.image_gpt, 213
pl_bolts.models.vision.image_gpt.gpt2,
213
pl_bolts.models.vision.image_gpt.igpt_module,
214
pl_bolts.models.vision.pixel_cnn, 216
pl_bolts.optimizers, 228
pl_bolts.optimizers.layer_adaptive_scaling,
228
pl_bolts.transforms, 229
pl_bolts.transforms.dataset_normalizations,
230
pl_bolts.transforms.self_supervised, 229
pl_bolts.transforms.self_supervised.ssl_transforms,
229
pl_bolts.utils, 230
pl_bolts.utils.pretrained_weights, 230
pl_bolts.utils.self_supervised, 230
pl_bolts.utils.semi_supervised, 230

INDEX

Symbols

_CPCV2__compute_final_nb_c() (*pl_bolts.models.self_supervised.cpc.cpc_module.CPCV2 method*), 199

_CPCV2__recover_z_shape() (*pl_bolts.models.self_supervised.cpc.cpc_module.CPCV2 method*), 199

_FeatureMapContrastiveTask__cache_dimension_masks() (*pl_bolts.losses.self_supervised_learning.FeatureMapContrastiveTask method*), 219

_FeatureMapContrastiveTask__compare_maps() (*pl_bolts.losses.self_supervised_learning.FeatureMapContrastiveTask method*), 219

_LNBottleneck__init_layer_norms() (*pl_bolts.models.self_supervised.cpc.networks.LNBottleneck method*), 200

_RequiredParameter (*class in pl_bolts.optimizers.layer_adaptive_scaling*), 229

_VAE__init_system() (*pl_bolts.models.autoencoders.basic_vae.basic_vae_module.VAE method*), 156

_VAE__set_default_datamodule() (*pl_bolts.models.autoencoders.basic_vae.basic_vae_module.VAE method*), 157

_VAE__set_pretrained_dims() (*pl_bolts.models.autoencoders.basic_vae.basic_vae_module.VAE method*), 157

__call__() (*pl_bolts.models.rl.common.agents.Agent method*), 162

__call__() (*pl_bolts.models.rl.common.agents.PolicyAgent method*), 162

__call__() (*pl_bolts.models.rl.common.agents.ValueAgent method*), 162

__call__() (*pl_bolts.models.self_supervised.amdim.transforms.AMDIMEvalTransformsCIFAR10 method*), 195

__call__() (*pl_bolts.models.self_supervised.amdim.transforms.AMDIMEvalTransformsImageNet128 method*), 195

__call__() (*pl_bolts.models.self_supervised.amdim.transforms.AMDIMEvalTransformsSTL10 method*), 196

__call__() (*pl_bolts.models.self_supervised.amdim.transforms.AMDIMTrainTransformsCIFAR10 method*), 196

__call__() (*pl_bolts.models.self_supervised.amdim.transforms.AMDIM method*), 197

__call__() (*pl_bolts.models.self_supervised.amdim.transforms.AMDIM method*), 197

__call__() (*pl_bolts.models.self_supervised.cpc.transforms.CPCEvalTr method*), 200

__call__() (*pl_bolts.models.self_supervised.cpc.transforms.CPCEvalTr method*), 201

__call__() (*pl_bolts.models.self_supervised.cpc.transforms.CPCEvalTr method*), 201

__call__() (*pl_bolts.models.self_supervised.cpc.transforms.CPCTrainT method*), 202

__call__() (*pl_bolts.models.self_supervised.cpc.transforms.CPCTrainT method*), 203

__call__() (*pl_bolts.models.self_supervised.cpc.transforms.CPCTrainT method*), 203

__call__() (*pl_bolts.models.self_supervised.moco.transforms.Gaussian method*), 206

__call__() (*pl_bolts.models.self_supervised.moco.transforms.Moco2Ev method*), 206

__call__() (*pl_bolts.models.self_supervised.moco.transforms.Moco2Ev method*), 206

__call__() (*pl_bolts.models.self_supervised.moco.transforms.Moco2Ev method*), 206

__call__() (*pl_bolts.models.self_supervised.moco.transforms.Moco2Tr method*), 207

__call__() (*pl_bolts.models.self_supervised.moco.transforms.Moco2Tr method*), 207

__call__() (*pl_bolts.models.self_supervised.moco.transforms.Moco2Tr method*), 207

__call__() (*pl_bolts.models.self_supervised.simclr.simclr_transforms.C method*), 209

__call__() (*pl_bolts.models.self_supervised.simclr.simclr_transforms.S method*), 210

__call__() (*pl_bolts.models.self_supervised.simclr.simclr_transforms.S method*), 210

__call__() (*pl_bolts.transforms.self_supervised.ssl_transforms.Patchify method*), 229

__call__() (*pl_bolts.transforms.self_supervised.ssl_transforms.Random method*), 230

__asdict() (*pl_bolts.models.rl.common.memory.Experience method*), 165

_batch_shuffle_ddp() (*pl_bolts.models.self_supervised.moco.moco2_module.MocoV2 method*), 205
_batch_unshuffle_ddp() (*pl_bolts.models.self_supervised.moco.moco2_module.MocoV2 method*), 205
_bypass (*pl_bolts.loggers.TrainsLogger attribute*), 224
_bypass (*pl_bolts.loggers.trains.TrainsLogger attribute*), 228
_calculate_md5() (*in module pl_bolts.datamodules.imagenet_dataset*), 140
_calculate_output_dim() (*pl_bolts.models.autoencoders.basic_ae.components.AEEncoder method*), 155
_calculate_output_dim() (*pl_bolts.models.autoencoders.basic_vae.components.Encoder method*), 158
_calculate_output_size() (*pl_bolts.models.autoencoders.basic_vae.components.Decoder method*), 157
_check_exists() (*pl_bolts.datamodules.cifar10_dataset.CIFAR10 class method*), 135
_check_integrity() (*in module pl_bolts.datamodules.imagenet_dataset*), 140
_check_md5() (*in module pl_bolts.datamodules.imagenet_dataset*), 140
_config_modules() (*pl_bolts.models.self_supervised.amdim.networks.AMDIMEncoder method*), 193
_dataloader() (*pl_bolts.models.rl.reinforce_model.Reinforce method*), 185
_dataloader() (*pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradient method*), 188
_default_transforms() (*pl_bolts.datamodules.fashion_mnist_datamodule.FashionMNISTDataModule method*), 137
_default_transforms() (*pl_bolts.datamodules.mnist_datamodule.MNISTDataModule method*), 144
_default_transforms() (*pl_bolts.datamodules.ssl_imagenet_datamodule.SSLImagenetDataModule method*), 149
_dequeue_and_enqueue() (*pl_bolts.models.self_supervised.moco.moco2_module.MocoV2 method*), 205
_download_from_url() (*pl_bolts.datamodules.base_dataset.LightDataset method*), 132
_extract_archive_save_torch() (*pl_bolts.datamodules.cifar10_dataset.CIFAR10 method*), 135
_fields (*pl_bolts.models.rl.common.memory.Experience attribute*), 165
_fields_defaults (*pl_bolts.models.rl.common.memory.Experience attribute*), 165
_filter_acts() (*pl_bolts.models.self_supervised.amdim.networks.AM method*), 193
_get_conv_out() (*pl_bolts.models.rl.common.networks.CNN method*), 167
_get_conv_out() (*pl_bolts.models.rl.common.networks.DuelingCNN method*), 168
_get_conv_out() (*pl_bolts.models.rl.common.networks.NoisyCNN method*), 169
_init_datasets() (*pl_bolts.datamodules.sklearn_datamodule.Sklearn method*), 145
_init_embeddings() (*pl_bolts.models.vision.image_gpt.gpt2.GPT2 method*), 213
_init_layers() (*pl_bolts.models.vision.image_gpt.gpt2.GPT2 method*), 213
_init_sos_token() (*pl_bolts.models.vision.image_gpt.gpt2.GPT2 method*), 213
_is_gzip() (*in module pl_bolts.datamodules.imagenet_dataset*), 140
_is_tar() (*in module pl_bolts.datamodules.imagenet_dataset*), 140
_is_targz() (*in module pl_bolts.datamodules.imagenet_dataset*), 140
_is_tarxz() (*in module pl_bolts.datamodules.imagenet_dataset*), 140
_is_zip() (*in module pl_bolts.datamodules.imagenet_dataset*), 140
_make() (*pl_bolts.models.rl.common.memory.Experience class method*), 165
_make_layer() (*pl_bolts.models.self_supervised.cpc.networks.CPCRes method*), 199
_make_layer() (*pl_bolts.models.self_supervised.resnets.ResNet method*), 211
_momentum_update_key_encoder() (*pl_bolts.models.self_supervised.moco.moco2_module.MocoV2 method*), 205
_prepare_subset() (*pl_bolts.datamodules.base_dataset.LightDataset static method*), 132
_replace() (*pl_bolts.models.rl.common.memory.Experience method*), 165
_reset() (*pl_bolts.models.rl.common.experience.ExperienceSource method*), 164
_run_step() (*pl_bolts.models.autoencoders.basic_ae.basic_ae_module.*

method), 154
_run_step() (*pl_bolts.models.autoencoders.basic_vae.basic_vae_module.VAE method*), 157
_sample_src_ftr() (*pl_bolts.losses.self_supervised_learning.FeatureMapContrastiveTask method*), 219
_shape_input() (*in module pl_bolts.models.vision.image_gpt.igpt_module*), 216
_unpickle() (*pl_bolts.datamodules.cifar10_dataset.CIFAR10 method*), 135
_verify_archive() (*in module pl_bolts.datamodules.imagenet_dataset*), 140
_verify_splits() (*pl_bolts.datamodules.imagenet_datamodule.ImagenetDataModule method*), 138
_verify_splits() (*pl_bolts.datamodules.ssl_imagenet_datamodule.SSLImagenetDataModule method*), 149

A

accuracy() (*in module pl_bolts.metrics.aggregation*), 152
action() (*pl_bolts.models.rl.common.memory.Experience property*), 165
add() (*pl_bolts.models.rl.common.memory.MeanBuffer method*), 165
add_argparse_args() (*pl_bolts.datamodules.lightning_datamodule.LightningDataModule class method*), 141
add_base_args() (*in module pl_bolts.models.rl.common.cli*), 163
add_model_specific_args() (*pl_bolts.models.autoencoders.basic_ae.basic_ae_module.AE static method*), 154
add_model_specific_args() (*pl_bolts.models.autoencoders.basic_vae.basic_vae_module.VAE static method*), 157
add_model_specific_args() (*pl_bolts.models.gans.basic.basic_gan_module.GAN static method*), 159
add_model_specific_args() (*pl_bolts.models.mnist_module.LitMNIST static method*), 217
add_model_specific_args() (*pl_bolts.models.regression.linear_regression.LinearRegression static method*), 160
add_model_specific_args() (*pl_bolts.models.regression.logistic_regression.LogisticRegression static method*), 161
add_model_specific_args() (*pl_bolts.models.rl.dqn_model.DQN static method*), 175
add_model_specific_args() (*pl_bolts.models.rl.reinforce_model.Reinforce static method*), 186
add_model_specific_args() (*pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradient static method*), 188
add_model_specific_args() (*pl_bolts.models.self_supervised.amdim.amdim_module.AMDIM static method*), 192
add_model_specific_args() (*pl_bolts.models.self_supervised.cpc.cpc_module.CPCV2 static method*), 199
add_model_specific_args() (*pl_bolts.models.self_supervised.moco.moco2_module.MocoV2 static method*), 206
add_model_specific_args() (*pl_bolts.models.self_supervised.simclr.simclr_module.SimCLR static method*), 209
add_model_specific_args() (*pl_bolts.models.vision.image_gpt.igpt_module.ImageGPT static method*), 215
adv_val() (*pl_bolts.models.rl.common.networks.DuelingCNN method*), 168
adv_val() (*pl_bolts.models.rl.common.networks.DuelingMLP method*), 168
AE (*class in pl_bolts.models.autoencoders.basic_ae.basic_ae_module*), 154
AEEncoder (*class in pl_bolts.models.autoencoders.basic_ae.components*), 155
Agent (*class in pl_bolts.models.rl.common.agents*), 162
AMDIM (*class in pl_bolts.models.self_supervised.amdim.amdim_module*), 191
AMDIMEncoder (*class in pl_bolts.models.self_supervised.amdim.networks*), 193
AMDIMEvalTransformsCIFAR10 (*class in pl_bolts.models.self_supervised.amdim.transforms*), 195
AMDIMEvalTransformsImageNet128 (*class in pl_bolts.models.self_supervised.amdim.transforms*), 195
AMDIMEvalTransformsSTL10 (*class in pl_bolts.models.self_supervised.amdim.transforms*), 195
AmdimNCELoss (*class in pl_bolts.losses.self_supervised_learning*), 218
AMDIMPatchesPretraining (*class in pl_bolts.models.self_supervised.amdim.datasets*), 193
AMDIMPretraining (*class in pl_bolts.models.self_supervised.amdim.datasets*), 193
AMDIMTrainTransformsCIFAR10 (*class in pl_bolts.models.self_supervised.amdim.transforms*),

196
AMDIMTrainTransformsImageNet128 (*class in pl_bolts.models.self_supervised.amdim.transforms*), 196
AMDIMTrainTransformsSTL10 (*class in pl_bolts.models.self_supervised.amdim.transforms*), 197
append() (*pl_bolts.models.rl.common.memory.Buffer method*), 165
append() (*pl_bolts.models.rl.common.memory.MultiStepBuffer method*), 166
append() (*pl_bolts.models.rl.common.memory.PERBuffer method*), 166

B

balance_classes() (*in module pl_bolts.utils.semi_supervised*), 230
BASE_URL (*pl_bolts.datamodules.cifar10_dataset.CIFAR10 attribute*), 135
Block (*class in pl_bolts.models.vision.image_gpt.gpt2*), 213
Buffer (*class in pl_bolts.models.rl.common.memory*), 165
BufferWrapper (*class in pl_bolts.models.rl.common.wrappers*), 170
build_networks() (*pl_bolts.models.rl.dqn_model.DQN method*), 175
build_networks() (*pl_bolts.models.rl.dueling_dqn_model.DuelingDQN method*), 178
build_networks() (*pl_bolts.models.rl.noisy_dqn_model.NoisyDQN method*), 182
build_networks() (*pl_bolts.models.rl.reinforce_model.Reinforce method*), 186
build_networks() (*pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradient method*), 188
bypass_mode() (*pl_bolts.loggers.trains.TrainsLogger class method*), 225
bypass_mode() (*pl_bolts.loggers.TrainsLogger class method*), 221

C

cache_folder_name (*pl_bolts.datamodules.base_dataset.LightDataset attribute*), 132
cache_folder_name (*pl_bolts.datamodules.cifar10_dataset.CIFAR10 attribute*), 135
cached_folder_path() (*pl_bolts.datamodules.base_dataset.LightDataset property*), 132
calc_entropy_loss() (*pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradient method*), 188
calc_policy_loss() (*pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradient static method*), 189
calc_qvals() (*pl_bolts.models.rl.reinforce_model.Reinforce method*), 186
calc_qvals() (*pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradient method*), 189
CIFAR10 (*class in pl_bolts.datamodules.cifar10_dataset*), 134
cifar10() (*pl_bolts.models.self_supervised.amdim.datasets.AMDIMPretrainTransformCIFAR10 static method*), 193
cifar10() (*pl_bolts.models.self_supervised.amdim.datasets.AMDIMPretrainTransformSTL10 static method*), 193
cifar10_normalization() (*in module pl_bolts.transforms.dataset_normalizations*), 230
cifar10_tiny() (*pl_bolts.models.self_supervised.amdim.datasets.AMDIMPretrainTransformTinyImageNet static method*), 193
CIFAR10DataModule (*class in pl_bolts.datamodules.cifar10_datamodule*), 132
CIFAR10Mixed (*class in pl_bolts.models.self_supervised.amdim.ssl_datasets*), 194
CNN (*class in pl_bolts.models.rl.common.networks*), 167
compute_loss_h() (*pl_bolts.losses.self_supervised_learning.CPCTask method*), 218
concat_all_gather() (*in module pl_bolts.models.self_supervised.moco.moco2_module*), 206
ConcatDataset (*class in pl_bolts.datamodules.concat_dataset*), 137
configure_optimizers() (*pl_bolts.models.autoencoders.basic_ae.basic_ae_module.AE method*), 154
configure_optimizers() (*pl_bolts.models.autoencoders.basic_vae.basic_vae_module.VAE method*), 157
configure_optimizers() (*pl_bolts.models.gans.basic.basic_gan_module.GAN method*), 159
configure_optimizers() (*pl_bolts.models.mnist_module.LitMNIST method*), 217
configure_optimizers() (*pl_bolts.models.regression.linear_regression.LinearRegression method*), 160
configure_optimizers() (*pl_bolts.models.regression.logistic_regression.LogisticRegression method*), 161
configure_optimizers() (*pl_bolts.models.rl.dqn_model.DQN method*), 175
configure_optimizers()

(*pl_bolts.models.rl.reinforce_model.Reinforce method*), 186
configure_optimizers() (*pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradient method*), 189
configure_optimizers() (*pl_bolts.models.self_supervised.amdim.amdim_module.AMDIM method*), 192
configure_optimizers() (*pl_bolts.models.self_supervised.cpc.cpc_module.CPCV2 method*), 199
configure_optimizers() (*pl_bolts.models.self_supervised.moco.moco2_module.MocoV2 method*), 206
configure_optimizers() (*pl_bolts.models.self_supervised.simclr.simclr_module.SimCLR method*), 209
configure_optimizers() (*pl_bolts.models.vision.image_gpt.igpt_module.ImageGPT method*), 215
conv1x1() (*in module pl_bolts.models.self_supervised.cpc.networks*), 200
Conv3x3 (*class in pl_bolts.models.self_supervised.amdim.networks*), 193
conv3x3() (*in module pl_bolts.models.self_supervised.cpc.networks*), 200
conv_block() (*pl_bolts.models.vision.pixel_cnn.PixelCNN method*), 216
ConvResBlock (*class in pl_bolts.models.self_supervised.amdim.networks*), 193
ConvResNxN (*class in pl_bolts.models.self_supervised.amdim.networks*), 194
CPCEvalTransformsCIFAR10 (*class in pl_bolts.models.self_supervised.cpc.transforms*), 200
CPCEvalTransformsImageNet128 (*class in pl_bolts.models.self_supervised.cpc.transforms*), 200
CPCEvalTransformsSTL10 (*class in pl_bolts.models.self_supervised.cpc.transforms*), 201
CPCResNet101 (*class in pl_bolts.models.self_supervised.cpc.networks*), 199
CPCTask (*class in pl_bolts.losses.self_supervised_learning*), 218
CPCTrainTransformsCIFAR10 (*class in pl_bolts.models.self_supervised.cpc.transforms*), 202
CPCTrainTransformsImageNet128 (*class in pl_bolts.models.self_supervised.cpc.transforms*), 202
CPCTrainTransformsSTL10 (*class in pl_bolts.models.self_supervised.cpc.transforms*), 203
CPCV2 (*class in pl_bolts.models.self_supervised.cpc.cpc_module*), 197

D

data (*pl_bolts.datamodules.base_dataset.LightDataset attribute*), 132
data (*pl_bolts.datamodules.cifar10_dataset.CIFAR10 attribute*), 135
data (*pl_bolts.datamodules.cifar10_dataset.TrialCIFAR10 attribute*), 136
DataAugmentation (*class in pl_bolts.models.rl.common.wrappers*), 170
DATASET_NAME (*pl_bolts.datamodules.base_dataset.LightDataset attribute*), 132
DATASET_NAME (*pl_bolts.datamodules.cifar10_dataset.CIFAR10 attribute*), 135
Decoder (*class in pl_bolts.models.autoencoders.basic_vae.components*), 157
default_transforms() (*pl_bolts.datamodules.cifar10_datamodule.CIFAR10DataModule method*), 133
default_transforms() (*pl_bolts.datamodules.stl10_datamodule.STL10DataModule method*), 151
DenseBlock (*class in pl_bolts.models.autoencoders.basic_ae.components*), 155
DenseBlock (*class in pl_bolts.models.autoencoders.basic_vae.components*), 157
DensenetEncoder (*class in pl_bolts.models.self_supervised.simclr.simclr_module*), 207
deterministic_shuffle() (*pl_bolts.models.self_supervised.amdim.ssl_datasets.SSLDataset class method*), 194
dicts_to_table() (*in module pl_bolts.callbacks.printing*), 130
dir_path (*pl_bolts.datamodules.base_dataset.LightDataset attribute*), 132
dir_path (*pl_bolts.datamodules.cifar10_dataset.CIFAR10 attribute*), 135
dir_path (*pl_bolts.datamodules.cifar10_dataset.TrialCIFAR10 attribute*), 136
Discriminator (*class in pl_bolts.models.gans.basic.components*), 159
discriminator_loss() (*pl_bolts.models.gans.basic.basic_gan_module.GAN*

method), 159
discriminator_step() (*pl_bolts.models.gans.basic.basic_gan_module.GAN method*), 159
done() (*pl_bolts.models.rl.common.memory.Experience property*), 165
double_dqn_loss() (*in module pl_bolts.losses.rl*), 217
DoubleDQN (*class in pl_bolts.models.rl.double_dqn_model*), 171
download() (*pl_bolts.datamodules.cifar10_dataset.CIFAR10 method*), 135
DQN (*class in pl_bolts.models.rl.dqn_model*), 174
dqn_loss() (*in module pl_bolts.losses.rl*), 217
DuelingCNN (*class in pl_bolts.models.rl.common.networks*), 168
DuelingDQN (*class in pl_bolts.models.rl.dueling_dqn_model*), 176
DuelingMLP (*class in pl_bolts.models.rl.common.networks*), 168

E

elbo_loss() (*pl_bolts.models.autoencoders.basic_vae.basic_vae_module.VAE method*), 157
Encoder (*class in pl_bolts.models.autoencoders.basic_vae.components*), 157
EpisodicExperienceStream (*class in pl_bolts.models.rl.common.experience*), 163
Experience (*class in pl_bolts.models.rl.common.memory*), 165
ExperienceSource (*class in pl_bolts.models.rl.common.experience*), 163
experiment() (*pl_bolts.loggers.trains.TrainsLogger property*), 228
experiment() (*pl_bolts.loggers.TrainsLogger property*), 224
extra_args (*pl_bolts.datamodules.cifar10_datamodule.CIFAR10DataModule attribute*), 133
extract_archive() (*in module pl_bolts.datamodules.imagenet_dataset*), 140

F

FakeRKHSConvNet (*class in pl_bolts.models.self_supervised.amdim.networks*), 194
FashionMNISTDataModule (*class in pl_bolts.datamodules.fashion_mnist_datamodule*), 137
feat_size_w_mask() (*pl_bolts.losses.self_supervised_learning.FeatureMapContrastiveTask method*), 219
FeatureMapContrastiveTask (*class in pl_bolts.losses.self_supervised_learning*), 218
FILE_NAME (*pl_bolts.datamodules.cifar10_dataset.CIFAR10 attribute*), 135
finalize() (*pl_bolts.loggers.trains.TrainsLogger method*), 226
finalize() (*pl_bolts.loggers.TrainsLogger method*), 222
FireResetEnv (*class in pl_bolts.models.rl.common.wrappers*), 170
Flatten (*class in pl_bolts.models.self_supervised.evaluator*), 210
flatten() (*pl_bolts.models.self_supervised.cpc.networks.CPCResNet10 method*), 199
flatten_batch() (*pl_bolts.models.rl.reinforce_model.Reinforce static method*), 186
flatten_batch() (*pl_bolts.models.rl.vanilla_policy_gradient_model.P static method*), 189
forward() (*pl_bolts.losses.self_supervised_learning.AmdimNCELoss method*), 218
forward() (*pl_bolts.losses.self_supervised_learning.CPCTask method*), 218
forward() (*pl_bolts.losses.self_supervised_learning.FeatureMapContras method*), 219
forward() (*pl_bolts.models.autoencoders.basic_ae.basic_ae_module.AE method*), 154
forward() (*pl_bolts.models.autoencoders.basic_ae.components.AEEnco method*), 155
forward() (*pl_bolts.models.autoencoders.basic_ae.components.DenseBl method*), 155
forward() (*pl_bolts.models.autoencoders.basic_vae.basic_vae_module.V method*), 157
forward() (*pl_bolts.models.autoencoders.basic_vae.components.Decode method*), 157
forward() (*pl_bolts.models.autoencoders.basic_vae.components.DenseB method*), 157
forward() (*pl_bolts.models.autoencoders.basic_vae.components.Encode method*), 158
forward() (*pl_bolts.models.gans.basic.basic_gan_module.GAN method*), 159
forward() (*pl_bolts.models.gans.basic.components.Discriminator method*), 159
forward() (*pl_bolts.models.gans.basic.components.Generator method*), 159
forward() (*pl_bolts.models.mnist_module.LitMNIST method*), 217
forward() (*pl_bolts.models.regression.linear_regression.LinearRegressi method*), 160
forward() (*pl_bolts.models.regression.logistic_regression.LogisticRegre method*), 161
forward() (*pl_bolts.models.rl.common.networks.CNN method*), 167
forward() (*pl_bolts.models.rl.common.networks.DuelingCNN method*), 168
forward() (*pl_bolts.models.rl.common.networks.DuelingMLP*

method), 168
forward() (*pl_bolts.models.rl.common.networks.MLP method*), 169
forward() (*pl_bolts.models.rl.common.networks.NoisyCNN method*), 169
forward() (*pl_bolts.models.rl.common.networks.NoisyLinear method*), 169
forward() (*pl_bolts.models.rl.dqn_model.DQN method*), 175
forward() (*pl_bolts.models.rl.reinforce_model.Reinforce method*), 186
forward() (*pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradient method*), 189
forward() (*pl_bolts.models.self_supervised.amdim.amdim_module.AMDIM method*), 192
forward() (*pl_bolts.models.self_supervised.amdim.networks.AMDIMEncoder method*), 193
forward() (*pl_bolts.models.self_supervised.amdim.networks.Conv3x3 method*), 193
forward() (*pl_bolts.models.self_supervised.amdim.networks.ConvResBlock method*), 194
forward() (*pl_bolts.models.self_supervised.amdim.networks.ConvResNxN method*), 194
forward() (*pl_bolts.models.self_supervised.amdim.networks.FakeRKHSConvNet method*), 194
forward() (*pl_bolts.models.self_supervised.amdim.networks.MaybeBatchNorm2d method*), 194
forward() (*pl_bolts.models.self_supervised.amdim.networks.NopNet method*), 194
forward() (*pl_bolts.models.self_supervised.cpc.cpc_module.CPCV2 method*), 199
forward() (*pl_bolts.models.self_supervised.cpc.networks.CPCResNet101 method*), 199
forward() (*pl_bolts.models.self_supervised.cpc.networks.LNBottleneck method*), 200
forward() (*pl_bolts.models.self_supervised.evaluator.Flatten method*), 210
forward() (*pl_bolts.models.self_supervised.evaluator.SSLEvaluator method*), 210
forward() (*pl_bolts.models.self_supervised.moco.moco2_module.MocoV2 method*), 206
forward() (*pl_bolts.models.self_supervised.resnets.ResNet method*), 211
forward() (*pl_bolts.models.self_supervised.simclr.simclr_module.DensenetEncoder method*), 207
forward() (*pl_bolts.models.self_supervised.simclr.simclr_module.Projection method*), 207
forward() (*pl_bolts.models.self_supervised.simclr.simclr_module.SimCLR method*), 209
forward() (*pl_bolts.models.vision.image_gpt.gpt2.Block method*), 213
forward() (*pl_bolts.models.vision.image_gpt.gpt2.GPT2 method*), 213
forward() (*pl_bolts.models.vision.image_gpt.igpt_module.ImageGPT method*), 215
forward() (*pl_bolts.models.vision.pixel_cnn.PixelCNN method*), 216
forward() (*pl_bolts.utils.semi_supervised.Identity method*), 230
from_argparse_args() (*pl_bolts.datamodules.lightning_datamodule.LightningDataModu class method*), 141

G

GAN (*class in pl_bolts.models.gans.basic.basic_gan_module*), 158
GaussianBlur (*class in pl_bolts.models.self_supervised.moco.transforms*), 206
GaussianBlur (*class in pl_bolts.models.self_supervised.simclr.simclr_transforms*), 209
generate_half_labeled_batches() (*in module pl_bolts.utils.semi_supervised*), 231
generate_meta_bins() (*pl_bolts.datamodules.imagenet_dataset.UnlabeledImagenet class method*), 140
generate_train_val_split() (*pl_bolts.models.self_supervised.amdim.ssl_datasets.SSLDataset class method*), 195
Generator (*class in pl_bolts.models.gans.basic.components*), 159
generator_loss() (*pl_bolts.models.gans.basic.basic_gan_module.GA method*), 159
generator_step() (*pl_bolts.models.gans.basic.basic_gan_module.GA method*), 159
get_action() (*pl_bolts.models.rl.common.agents.ValueAgent method*), 162
get_approx_posterior() (*pl_bolts.models.autoencoders.basic_vae.basic_vae_module.VAE method*), 157
get_dataset() (*pl_bolts.models.self_supervised.amdim.datasets.AMDI static method*), 193
get_device() (*pl_bolts.models.rl.reinforce_model.Reinforce method*), 186
get_device() (*pl_bolts.models.rl.vanilla_policy_gradient_model.Policy method*), 190
get_init_arguments_and_types() (*pl_bolts.datamodules.lightning_datamodule.LightningDataModu class method*), 142
get_prior() (*pl_bolts.models.autoencoders.basic_vae.basic_vae_modu method*), 157
get_random_action() (*pl_bolts.models.rl.common.agents.ValueAgent method*), 163
get_transition_info() (*pl_bolts.models.rl.common.experience.NStepExperienceSource*

method), 164
get_transition_info() (*pl_bolts.models.rl.common.memory.MultiStepBuffer method*), 166
GPT2 (*class in pl_bolts.models.vision.image_gpt.gpt2*), 213

I

id() (*pl_bolts.loggers.trains.TrainsLogger property*), 228
id() (*pl_bolts.loggers.TrainsLogger property*), 224
Identity (*class in pl_bolts.utils.semi_supervised*), 230
ImageGenerator (*class in pl_bolts.models.gans.basic.basic_gan_module*), 159
ImageGPT (*class in pl_bolts.models.vision.image_gpt.igpt_module*), 214
imagenet() (*pl_bolts.models.self_supervised.amdim.datasets.AMDIMPatchesPretraining static method*), 193
imagenet() (*pl_bolts.models.self_supervised.amdim.datasets.AMDIMPretraining static method*), 193
imagenet_normalization() (*in module pl_bolts.transforms.dataset_normalizations*), 230
ImagenetDataModule (*class in pl_bolts.datamodules.imagenet_datamodule*), 138
ImageToPyTorch (*class in pl_bolts.models.rl.common.wrappers*), 170
init_decoder() (*pl_bolts.models.autoencoders.basic_ae.basic_ae_module.AE method*), 154
init_decoder() (*pl_bolts.models.autoencoders.basic_vae.basic_vae_module.VAE method*), 157
init_discriminator() (*pl_bolts.models.gans.basic.basic_gan_module.GAN method*), 159
init_encoder() (*pl_bolts.models.autoencoders.basic_ae.basic_ae_module.AE method*), 154
init_encoder() (*pl_bolts.models.autoencoders.basic_vae.basic_vae_module.VAE method*), 157
init_encoder() (*pl_bolts.models.self_supervised.amdim.amdim_module.AMDIM method*), 192
init_encoder() (*pl_bolts.models.self_supervised.cpc.cpc_module.CPCV2 method*), 199
init_encoder() (*pl_bolts.models.self_supervised.simclr.simclr_module.SimCLR method*), 209
init_encoders() (*pl_bolts.models.self_supervised.moco.moco2_module.MocoV2 method*), 206
init_generator() (*pl_bolts.models.gans.basic.basic_gan_module.GAN method*), 159
init_loss() (*pl_bolts.models.self_supervised.simclr.simclr_module.SimCLR method*), 209
init_projection() (*pl_bolts.models.self_supervised.simclr.simclr_module.SimCLR method*), 209
init_weights() (*pl_bolts.models.self_supervised.amdim.networks.AM method*), 193
init_weights() (*pl_bolts.models.self_supervised.amdim.networks.Cor method*), 194
init_weights() (*pl_bolts.models.self_supervised.amdim.networks.Cor method*), 194
init_weights() (*pl_bolts.models.self_supervised.amdim.networks.Fak method*), 194
interpolate_latent_space() (*pl_bolts.callbacks.variational.LatentDimInterpolator method*), 131

L

labels (*pl_bolts.datamodules.cifar10_dataset.CIFAR10 attribute*), 135
LARS (*class in pl_bolts.optimizers.layer_adaptive_scaling*), 128
LatentDimInterpolator (*class in pl_bolts.callbacks.variational*), 131
LightDataset (*class in pl_bolts.datamodules.base_dataset*), 132
LightningDataModule (*class in pl_bolts.datamodules.lightning_datamodule*), 141
LinearRegression (*class in pl_bolts.models.regression.linear_regression*), 160
LitMNIST (*class in pl_bolts.models.mnist_module*), 217
load_checkpoint() (*pl_bolts.models.self_supervised.cpc.networks.CPCResNet101 method*), 199
load_pretrained() (*in module pl_bolts.utils.pretrained_weights*), 230
load_pretrained() (*pl_bolts.models.autoencoders.basic_vae.basic_vae_module.VAE method*), 158
load_pretrained() (*pl_bolts.models.self_supervised.cpc.cpc_module.CPCV2 method*), 199
log_artifact() (*pl_bolts.loggers.trains.TrainsLogger method*), 226
log_artifact() (*pl_bolts.loggers.TrainsLogger method*), 222
log_hyperparams() (*pl_bolts.loggers.trains.TrainsLogger method*), 226
log_hyperparams() (*pl_bolts.loggers.TrainsLogger method*), 222
log_image() (*pl_bolts.loggers.trains.TrainsLogger method*), 226
log_image() (*pl_bolts.loggers.TrainsLogger method*), 222

log_metric() (*pl_bolts.loggers.trains.TrainsLogger method*), 227
log_metric() (*pl_bolts.loggers.TrainsLogger method*), 223
log_metrics() (*pl_bolts.loggers.trains.TrainsLogger method*), 227
log_metrics() (*pl_bolts.loggers.TrainsLogger method*), 223
log_text() (*pl_bolts.loggers.trains.TrainsLogger method*), 227
log_text() (*pl_bolts.loggers.TrainsLogger method*), 223
LogisticRegression (*class in pl_bolts.models.regression.logistic_regression*), 161
loss() (*pl_bolts.models.rl.reinforce_model.Reinforce method*), 186
loss() (*pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradient method*), 190

M

make_env() (*in module pl_bolts.models.rl.common.wrappers*), 171
MaxAndSkipEnv (*class in pl_bolts.models.rl.common.wrappers*), 170
MaybeBatchNorm2d (*class in pl_bolts.models.self_supervised.amdim.networks*), 194
mean() (*in module pl_bolts.metrics.aggregation*), 152
mean() (*pl_bolts.models.rl.common.memory.MeanBuffer method*), 166
MeanBuffer (*class in pl_bolts.models.rl.common.memory*), 165
MLP (*class in pl_bolts.models.rl.common.networks*), 168
MNISTDataModule (*class in pl_bolts.datamodules.mnist_datamodule*), 144
Moco2EvalCIFAR10Transforms (*class in pl_bolts.models.self_supervised.moco.transforms*), 206
Moco2EvalImagenetTransforms (*class in pl_bolts.models.self_supervised.moco.transforms*), 206
Moco2EvalSTL10Transforms (*class in pl_bolts.models.self_supervised.moco.transforms*), 206
Moco2TrainCIFAR10Transforms (*class in pl_bolts.models.self_supervised.moco.transforms*), 207
Moco2TrainImagenetTransforms (*class in pl_bolts.models.self_supervised.moco.transforms*), 207
Moco2TrainSTL10Transforms (*class in pl_bolts.models.self_supervised.moco.transforms*), 207
MocoLRScheduler (*class in pl_bolts.models.self_supervised.moco.callbacks*), 204
MocoV2 (*class in pl_bolts.models.self_supervised.moco.moco2_module*), 204
MultiStepBuffer (*class in pl_bolts.models.rl.common.memory*), 166

N

name (*pl_bolts.datamodules.cifar10_datamodule.CIFAR10DataModule attribute*), 133
name (*pl_bolts.datamodules.fashion_mnist_datamodule.FashionMNISTDa attribute*), 138
name (*pl_bolts.datamodules.imagenet_datamodule.ImagenetDataModule attribute*), 139
name (*pl_bolts.datamodules.lightning_datamodule.LightningDataModule attribute*), 143
name (*pl_bolts.datamodules.mnist_datamodule.MNISTDataModule attribute*), 145
name (*pl_bolts.datamodules.sklearn_datamodule.SklearnDataModule attribute*), 146
name (*pl_bolts.datamodules.ssl_imagenet_datamodule.SSLImagenetDataM attribute*), 150
name (*pl_bolts.datamodules.stl10_datamodule.STL10DataModule attribute*), 152
name() (*pl_bolts.loggers.trains.TrainsLogger property*), 228
name() (*pl_bolts.loggers.TrainsLogger property*), 224
new_state() (*pl_bolts.models.rl.common.memory.Experience property*), 165
NoisyCNN (*class in pl_bolts.models.rl.common.networks*), 169
NoisyDQN (*class in pl_bolts.models.rl.noisy_dqn_model*), 180
NoisyLinear (*class in pl_bolts.models.rl.common.networks*), 169
NopNet (*class in pl_bolts.models.self_supervised.amdim.networks*), 194
normalize (*pl_bolts.datamodules.base_dataset.LightDataset attribute*), 132
normalize (*pl_bolts.datamodules.cifar10_dataset.CIFAR10 attribute*), 135
normalize (*pl_bolts.datamodules.cifar10_dataset.TrialCIFAR10 attribute*), 136
NStepDQN (*class in pl_bolts.models.rl.n_step_dqn_model*), 179
NStepExperienceSource (*class in pl_bolts.models.rl.common.experience*), 164
nt_xent_loss() (*in module pl_bolts.losses.self_supervised_learning*), 220
num_classes() (*pl_bolts.datamodules.cifar10_datamodule.CIFAR10Da property*), 133

num_classes() (pl_bolts.datamodules.cifar10_datamodule.TinyCIFAR10DataModule property), 134
num_classes() (pl_bolts.datamodules.fashion_mnist_datamodule.FashionMNISTDataModule property), 138
num_classes() (pl_bolts.datamodules.imagenet_datamodule.ImagenetDataModule property), 139
num_classes() (pl_bolts.datamodules.mnist_datamodule.MNISTDataModule property), 145
num_classes() (pl_bolts.datamodules.ssl_imagenet_datamodule.SSLImagenetDataModule property), 150
num_classes() (pl_bolts.datamodules.stl10_datamodule.STL10DataModule property), 152

O

observation() (pl_bolts.models.rl.common.wrappers.BufferWrapper method), 170
observation() (pl_bolts.models.rl.common.wrappers.DataAugmentation method), 170
observation() (pl_bolts.models.rl.common.wrappers.ImageToPyTorch static method), 170
observation() (pl_bolts.models.rl.common.wrappers.ProcessFrame84 method), 171
observation() (pl_bolts.models.rl.common.wrappers.ScaledFloatFrame static method), 171
on_epoch_end() (pl_bolts.callbacks.printing.PrintTableMetricsCallback method), 130
on_epoch_end() (pl_bolts.callbacks.variational.LatentDimInterpolator method), 131
on_epoch_end() (pl_bolts.models.gans.basic.basic_gan_module.ImageGenerator method), 159
on_epoch_start() (pl_bolts.models.self_supervised.moco.callbacks.MocoLRScheduler method), 204
on_train_start() (pl_bolts.models.rl.noisy_dqn_model.NoisyDQN method), 182

P

parse_devkit_archive() (in module pl_bolts.datamodules.imagenet_dataset), 140
parse_map_indexes() (pl_bolts.losses.self_supervised_learning.FeatureMapContrastiveTask static method), 220
partition_train_set() (pl_bolts.datamodules.imagenet_dataset.UnlabeledImagenet method), 140
Patchify (class in pl_bolts.transforms.self_supervised.ssl_transforms), 229
per_dqn_loss() (in module pl_bolts.losses.rl), 218
PERBuffer (class in pl_bolts.models.rl.common.memory), 166
PERDQN (class in pl_bolts.models.rl.per_dqn_model), 182
PixelCNN (class in pl_bolts.models.vision.pixel_cnn), 216
pl_bolts (module), 130
pl_bolts.callbacks (module), 130
pl_bolts.callbacks.printing (module), 130
pl_bolts.callbacks.variational (module), 131
pl_bolts.datamodules (module), 132
pl_bolts.datamodules.base_dataset (module), 132
pl_bolts.datamodules.cifar10_datamodule (module), 132
pl_bolts.datamodules.cifar10_dataset (module), 134
pl_bolts.datamodules.concat_dataset (module), 137
pl_bolts.datamodules.fashion_mnist_datamodule (module), 137
pl_bolts.datamodules.imagenet_datamodule (module), 138
pl_bolts.datamodules.imagenet_dataset (module), 140
pl_bolts.datamodules.lightning_datamodule (module), 141
pl_bolts.datamodules.mnist_datamodule (module), 144
pl_bolts.datamodules.sklearn_datamodule (module), 145
pl_bolts.datamodules.ssl_imagenet_datamodule (module), 149
pl_bolts.datamodules.stl10_datamodule (module), 151
pl_bolts.loggers (module), 220
pl_bolts.loggers.trains (module), 224
pl_bolts.losses (module), 217
pl_bolts.losses.rl (module), 217
pl_bolts.losses.self_supervised_learning (module), 218
pl_bolts.metrics (module), 152
pl_bolts.metrics.aggregation (module), 152
pl_bolts.models (module), 153
pl_bolts.models.autoencoders (module), 153
pl_bolts.models.autoencoders.basic_ae (module), 153
pl_bolts.models.autoencoders.basic_ae.basic_ae_module (module), 154
pl_bolts.models.autoencoders.basic_ae.components (module), 155
pl_bolts.models.autoencoders.basic_vae (module), 155
pl_bolts.models.autoencoders.basic_vae.basic_vae_module (module), 156
pl_bolts.models.autoencoders.basic_vae.components (module), 157
pl_bolts.models.gans (module), 158
pl_bolts.models.gans.basic (module), 158
pl_bolts.models.gans.basic.basic_gan_module

(*module*), 158
pl_bolts.models.gans.basic.components (*module*), 159
pl_bolts.models.mnist_module (*module*), 217
pl_bolts.models.regression (*module*), 160
pl_bolts.models.regression.linear_regression (*module*), 160
pl_bolts.models.regression.logistic_regression (*module*), 161
pl_bolts.models.rl (*module*), 161
pl_bolts.models.rl.common (*module*), 161
pl_bolts.models.rl.common.agents (*module*), 162
pl_bolts.models.rl.common.cli (*module*), 163
pl_bolts.models.rl.common.experience (*module*), 163
pl_bolts.models.rl.common.memory (*module*), 165
pl_bolts.models.rl.common.networks (*module*), 167
pl_bolts.models.rl.common.wrappers (*module*), 170
pl_bolts.models.rl.double_dqn_model (*module*), 171
pl_bolts.models.rl.dqn_model (*module*), 174
pl_bolts.models.rl.dueling_dqn_model (*module*), 176
pl_bolts.models.rl.n_step_dqn_model (*module*), 179
pl_bolts.models.rl.noisy_dqn_model (*module*), 180
pl_bolts.models.rl.per_dqn_model (*module*), 182
pl_bolts.models.rl.reinforce_model (*module*), 185
pl_bolts.models.rl.vanilla_policy_gradient_model (*module*), 187
pl_bolts.models.self_supervised (*module*), 191
pl_bolts.models.self_supervised.amdim (*module*), 191
pl_bolts.models.self_supervised.amdim.amdim_module (*module*), 191
pl_bolts.models.self_supervised.amdim.datasets (*module*), 193
pl_bolts.models.self_supervised.amdim.networks (*module*), 193
pl_bolts.models.self_supervised.amdim.ssl_datasets (*module*), 194
pl_bolts.models.self_supervised.amdim.transforms (*module*), 195
pl_bolts.models.self_supervised.cpc (*module*), 197
pl_bolts.models.self_supervised.cpc.cpc_module (*module*), 197
pl_bolts.models.self_supervised.cpc.networks (*module*), 199
pl_bolts.models.self_supervised.cpc.transforms (*module*), 200
pl_bolts.models.self_supervised.evaluator (*module*), 210
pl_bolts.models.self_supervised.moco (*module*), 203
pl_bolts.models.self_supervised.moco.callbacks (*module*), 204
pl_bolts.models.self_supervised.moco.moco2_module (*module*), 204
pl_bolts.models.self_supervised.moco.transforms (*module*), 206
pl_bolts.models.self_supervised.resnets (*module*), 211
pl_bolts.models.self_supervised.simclr (*module*), 207
pl_bolts.models.self_supervised.simclr.simclr_module (*module*), 207
pl_bolts.models.self_supervised.simclr.simclr_transforms (*module*), 209
pl_bolts.models.vision (*module*), 213
pl_bolts.models.vision.image_gpt (*module*), 213
pl_bolts.models.vision.image_gpt.gpt2 (*module*), 213
pl_bolts.models.vision.image_gpt.igpt_module (*module*), 214
pl_bolts.models.vision.pixel_cnn (*module*), 216
pl_bolts.optimizers (*module*), 228
pl_bolts.optimizers.layer_adaptive_scaling (*module*), 228
pl_bolts.transforms (*module*), 229
pl_bolts.transforms.dataset_normalizations (*module*), 230
pl_bolts.transforms.self_supervised (*module*), 229
pl_bolts.transforms.self_supervised.ssl_transforms (*module*), 229
pl_bolts.utils (*module*), 230
pl_bolts.utils.pretrained_weights (*module*), 230
pl_bolts.utils.self_supervised (*module*), 230
pl_bolts.utils.semi_supervised (*module*), 230
PolicyAgent (*class in pl_bolts.models.rl.common.agents*), 162
PolicyGradient (*class in pl_bolts.models.rl.vanilla_policy_gradient_model*),

187
populate() (*pl_bolts.models.rl.dqn_model.DQN method*), 175
precision_at_k() (*in module pl_bolts.metrics.aggregation*), 152
prepare_data() (*pl_bolts.datamodules.cifar10_datamodule.CIFAR10DataModule method*), 133
prepare_data() (*pl_bolts.datamodules.cifar10_dataset.CIFAR10 method*), 135
prepare_data() (*pl_bolts.datamodules.cifar10_dataset.TrialCIFAR10 method*), 136
prepare_data() (*pl_bolts.datamodules.fashion_mnist_datamodule.FashionMNISTDataModule method*), 137
prepare_data() (*pl_bolts.datamodules.imagenet_datamodule.ImagenetDataModule method*), 138
prepare_data() (*pl_bolts.datamodules.lightning_datamodule.LightningDataModule method*), 142
prepare_data() (*pl_bolts.datamodules.mnist_datamodule.MNISTDataModule method*), 144
prepare_data() (*pl_bolts.datamodules.ssl_imagenet_datamodule.SSLImagenetDataModule method*), 149
prepare_data() (*pl_bolts.datamodules.stl10_datamodule.STL10DataModule method*), 151
prepare_data() (*pl_bolts.models.autoencoders.basic_ae.basic_ae_module.AE method*), 154
prepare_data() (*pl_bolts.models.autoencoders.basic_vae.basic_vae_module.VAE method*), 157
prepare_data() (*pl_bolts.models.gans.basic.basic_gan_module.GAN method*), 159
prepare_data() (*pl_bolts.models.mnist_module.LitMNIST method*), 217
prepare_data() (*pl_bolts.models.rl.dqn_model.DQN method*), 175
prepare_data() (*pl_bolts.models.rl.per_dqn_model.PERDQN method*), 184
prepare_data() (*pl_bolts.models.self_supervised.cpc.cpc_module.CPCV2 method*), 199
prepare_data() (*pl_bolts.models.self_supervised.moco.moco2_module.MocoV2 method*), 206
prepare_data() (*pl_bolts.models.self_supervised.simclr.simclr_module.SimCLR method*), 209
prepare_data() (*pl_bolts.models.vision.image_gpt.igpt_module.ImageGPT method*), 215
PrintTableMetricsCallback (*class in pl_bolts.callbacks.printing*), 130
PrioRLDataset (*class in pl_bolts.models.rl.common.experience*), 164
process() (*pl_bolts.models.rl.common.wrappers.ProcessFrame84 static method*), 171
process_batch() (*pl_bolts.models.rl.reinforce_model.Reinforce method*), 187
process_batch() (*pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradient method*), 190
ProcessFrame84 (*class in pl_bolts.models.rl.common.wrappers*), 171
Projection (*class in pl_bolts.models.self_supervised.simclr.simclr_module*), 207

R

RandomTranslateWithReflect (*class in pl_bolts.transforms.self_supervised.ssl_transforms*), 229
ReinForce (*class in pl_bolts.models.rl.reinforce_model*), 185
relabel() (*pl_bolts.datamodules.cifar10_dataset.CIFAR10 attribute*), 136
ReplayBuffer (*class in pl_bolts.models.rl.common.memory*), 167
reset() (*pl_bolts.models.rl.common.wrappers.BufferWrapper method*), 170
reset() (*pl_bolts.models.rl.common.wrappers.FireResetEnv method*), 170
reset() (*pl_bolts.models.rl.common.wrappers.MaxAndSkipEnv method*), 171
reset() (*pl_bolts.models.rl.common.wrappers.ToTensor method*), 171
reset_parameters() (*pl_bolts.models.rl.common.networks.NoisyLinear method*), 170
ResNet (*class in pl_bolts.models.self_supervised.resnets*), 211
resnet101() (*in module pl_bolts.models.self_supervised.resnets*), 211
resnet152() (*in module pl_bolts.models.self_supervised.resnets*), 211
resnet18() (*in module pl_bolts.models.self_supervised.resnets*), 211
resnet34() (*in module pl_bolts.models.self_supervised.resnets*), 211
resnet50() (*in module pl_bolts.models.self_supervised.resnets*), 211
resnext101_32x8d() (*in module pl_bolts.models.self_supervised.resnets*), 212
resnext50_32x4d() (*in module pl_bolts.models.self_supervised.resnets*), 212
retrieve() (*pl_bolts.models.rl.common.memory.Experience property*), 165
RLDataset (*class in pl_bolts.models.rl.common.experience*), 164

run_episode() (*pl_bolts.models.rl.common.experience.Experience method*), 164

S

sample() (*pl_bolts.models.rl.common.memory.Buffer method*), 165
sample() (*pl_bolts.models.rl.common.memory.MultiStepBuffer method*), 166
sample() (*pl_bolts.models.rl.common.memory.PERBuffer method*), 166
sample() (*pl_bolts.models.rl.common.memory.ReplayBuffer method*), 167
ScaledFloatFrame (*class in pl_bolts.models.rl.common.wrappers*), 171
select_nb_imgs_per_class() (*pl_bolts.models.self_supervised.amdim.ssl_datasets.SSLDatasetMixin class method*), 195
set_bypass_mode() (*pl_bolts.loggers.trains.TrainsLogger class method*), 227
set_bypass_mode() (*pl_bolts.loggers.TrainsLogger class method*), 223
set_credentials() (*pl_bolts.loggers.trains.TrainsLogger class method*), 227
set_credentials() (*pl_bolts.loggers.TrainsLogger class method*), 223
SimCLR (*class in pl_bolts.models.self_supervised.simclr.simclr_module*), 207
SimCLREvalDataTransform (*class in pl_bolts.models.self_supervised.simclr.simclr_transforms*), 209
SimCLRTrainDataTransform (*class in pl_bolts.models.self_supervised.simclr.simclr_transforms*), 210
single_step() (*pl_bolts.models.rl.common.experience.NStepExperienceSource method*), 164
size() (*pl_bolts.datamodules.lightning_datamodule.LightningDataModule method*), 142
SklearnDataModule (*class in pl_bolts.datamodules.sklearn_datamodule*), 145
SklearnDataset (*class in pl_bolts.datamodules.sklearn_datamodule*), 147
SSLDatasetMixin (*class in pl_bolts.models.self_supervised.amdim.ssl_datasets*), 194
SSLEvaluator (*class in pl_bolts.models.self_supervised.evaluator*), 210
SSLImagenetDataModule (*class in pl_bolts.datamodules.ssl_imagenet_datamodule*), 149
steps() (*pl_bolts.models.rl.common.memory.Experience property*), 165
step() (*pl_bolts.models.rl.common.experience.EpisodicExperienceStream method*), 163
step() (*pl_bolts.models.rl.common.experience.ExperienceSource method*), 164
step() (*pl_bolts.models.rl.common.experience.NStepExperienceSource method*), 164
step() (*pl_bolts.models.rl.common.wrappers.FireResetEnv method*), 170
step() (*pl_bolts.models.rl.common.wrappers.MaxAndSkipEnv method*), 171
step() (*pl_bolts.models.rl.common.wrappers.ToTensor method*), 171
step() (*pl_bolts.optimizers.layer_adaptive_scaling.LARS method*), 229
stl() (*pl_bolts.models.self_supervised.amdim.datasets.AMDIMPatchesPr static method*), 193
stl() (*pl_bolts.models.self_supervised.amdim.datasets.AMDIMPretrainin static method*), 193
stl10_normalization() (*in module pl_bolts.transforms.dataset_normalizations*), 230
STL10DataModule (*class in pl_bolts.datamodules.stl10_datamodule*), 151

T

tanh_clip() (*in module pl_bolts.losses.self_supervised_learning*), 220
targets (*pl_bolts.datamodules.base_dataset.LightDataset attribute*), 132
targets (*pl_bolts.datamodules.cifar10_dataset.CIFAR10 attribute*), 136
targets (*pl_bolts.datamodules.cifar10_dataset.TrialCIFAR10 attribute*), 136
TensorDataModule (*class in pl_bolts.datamodules.sklearn_datamodule*), 147
TensorDataset (*class in pl_bolts.datamodules.sklearn_datamodule*), 148
test_dataloader() (*pl_bolts.datamodules.cifar10_datamodule.CIFAR10DataModule method*), 133
test_dataloader() (*pl_bolts.datamodules.fashion_mnist_datamodule.FashionMNIST method*), 137
test_dataloader() (*pl_bolts.datamodules.imagenet_datamodule.ImagenetDataModu method*), 138
test_dataloader() (*pl_bolts.datamodules.lightning_datamodule.LightningDataModu*

method), 142
test_dataloader() (*pl_bolts.datamodules.mnist_datamodule.MNISTDataModule method*), 144
test_dataloader() (*pl_bolts.datamodules.sklearn_datamodule.SklearnDataModule method*), 145
test_dataloader() (*pl_bolts.datamodules.ssl_imagenet_datamodule.SSLImagenetDataModule method*), 149
test_dataloader() (*pl_bolts.datamodules.stl10_datamodule.STL10DataModule method*), 151
test_dataloader() (*pl_bolts.models.autoencoders.basic_ae.basic_ae_module.AE method*), 154
test_dataloader() (*pl_bolts.models.autoencoders.basic_vae.basic_vae_module.VAE method*), 157
test_dataloader() (*pl_bolts.models.mnist_module.LitMNIST method*), 217
test_dataloader() (*pl_bolts.models.rl.dqn_model.DQN method*), 175
test_dataloader() (*pl_bolts.models.vision.image_gpt.igpt_module.ImageGPT method*), 216
test_epoch_end() (*pl_bolts.models.autoencoders.basic_ae.basic_ae_module.AE method*), 154
test_epoch_end() (*pl_bolts.models.autoencoders.basic_vae.basic_vae_module.VAE method*), 157
test_epoch_end() (*pl_bolts.models.mnist_module.LitMNIST method*), 217
test_epoch_end() (*pl_bolts.models.regression.linear_regression.LinearRegression method*), 160
test_epoch_end() (*pl_bolts.models.regression.logistic_regression.LogisticRegression method*), 161
test_epoch_end() (*pl_bolts.models.rl.dqn_model.DQN method*), 175
test_epoch_end() (*pl_bolts.models.vision.image_gpt.igpt_module.ImageGPT method*), 216
TEST_FILE_NAME (*pl_bolts.datamodules.cifar10_dataset.CIFAR10 attribute*), 135
test_step() (*pl_bolts.models.autoencoders.basic_ae.basic_ae_module.AE method*), 154
test_step() (*pl_bolts.models.autoencoders.basic_vae.basic_vae_module.VAE method*), 157
test_step() (*pl_bolts.models.mnist_module.LitMNIST method*), 217
test_step() (*pl_bolts.models.regression.linear_regression.LinearRegression method*), 160
test_step() (*pl_bolts.models.regression.logistic_regression.LogisticRegression method*), 161
test_step() (*pl_bolts.models.rl.dqn_model.DQN method*), 176
test_step() (*pl_bolts.models.vision.image_gpt.igpt_module.ImageGPT method*), 216
test_transforms() (*pl_bolts.datamodules.lightning_datamodule.LightningDataModule property*), 143
TinyCIFAR10DataModule (*class in pl_bolts.datamodules.cifar10_datamodule*), 133
torchvision_ssl_encoder() (*in module pl_bolts.utils.self_supervised*), 230
ToTensor (*class in pl_bolts.models.rl.common.wrappers*), 171
train_dataloader() (*pl_bolts.datamodules.cifar10_datamodule.CIFAR10DataModule method*), 133
train_dataloader() (*pl_bolts.datamodules.fashion_mnist_datamodule.FashionMNIST method*), 137
train_dataloader() (*pl_bolts.datamodules.imagenet_datamodule.ImagenetDataModule method*), 139
train_dataloader() (*pl_bolts.datamodules.lightning_datamodule.LightningDataModule method*), 143
train_dataloader() (*pl_bolts.datamodules.mnist_datamodule.MNISTDataModule method*), 144
train_dataloader() (*pl_bolts.datamodules.sklearn_datamodule.SklearnDataModule method*), 146
train_dataloader() (*pl_bolts.datamodules.ssl_imagenet_datamodule.SSLImagenetDataModule method*), 150
train_dataloader() (*pl_bolts.datamodules.stl10_datamodule.STL10DataModule method*), 151
train_dataloader() (*pl_bolts.models.autoencoders.basic_ae.basic_ae_module.AE method*), 154
train_dataloader() (*pl_bolts.models.autoencoders.basic_vae.basic_vae_module.VAE method*), 157
train_dataloader() (*pl_bolts.models.gans.basic.basic_gan_module.GAN method*), 159
train_dataloader() (*pl_bolts.models.mnist_module.LitMNIST method*), 217
train_dataloader() (*pl_bolts.models.regression.linear_regression.LinearRegression method*), 160
train_dataloader() (*pl_bolts.models.rl.dqn_model.DQN method*), 176
train_dataloader() (*pl_bolts.models.regression.logistic_regression.LogisticRegression method*), 161
train_dataloader() (*pl_bolts.models.vision.image_gpt.igpt_module.ImageGPT method*), 216
train_transforms() (*pl_bolts.datamodules.lightning_datamodule.LightningDataModule property*), 143
transform_dataset() (*pl_bolts.datamodules.cifar10_dataset.CIFAR10 method*), 135
transform_dataset() (*pl_bolts.datamodules.mnist_dataset.MNIST method*), 140
transform_dataset() (*pl_bolts.datamodules.stl10_dataset.STL10 method*), 147
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 152
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 153
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 154
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 155
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 156
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 157
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 158
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 159
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 160
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 161
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 162
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 163
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 164
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 165
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 166
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 167
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 168
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 169
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 170
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 171
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 172
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 173
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 174
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 175
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 176
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 177
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 178
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 179
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 180
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 181
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 182
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 183
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 184
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 185
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 186
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 187
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 188
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 189
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 190
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 191
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 192
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 193
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 194
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 195
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 196
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 197
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 198
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 199
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 200
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 201
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 202
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 203
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 204
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 205
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 206
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 207
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 208
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 209
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 210
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 211
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 212
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 213
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 214
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 215
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 216
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 217
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 218
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 219
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 220
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 221
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 222
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 223
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 224
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 225
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 226
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 227
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 228
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 229
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 230
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 231
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 232
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 233
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 234
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 235
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 236
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 237
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 238
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 239
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 240
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 241
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 242
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 243
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 244
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 245
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 246
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 247
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 248
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 249
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 250
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 251
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 252
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 253
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 254
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 255
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 256
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 257
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 258
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 259
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 260
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 261
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 262
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 263
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 264
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 265
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 266
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 267
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 268
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 269
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 270
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 271
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 272
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 273
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 274
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 275
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 276
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 277
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 278
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 279
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 280
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 281
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 282
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 283
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 284
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 285
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 286
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 287
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 288
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 289
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 290
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 291
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 292
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 293
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 294
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 295
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 296
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 297
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 298
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 299
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 300
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 301
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 302
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 303
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 304
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 305
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 306
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 307
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 308
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 309
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 310
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 311
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 312
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 313
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 314
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 315
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 316
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 317
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 318
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 319
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 320
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 321
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 322
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 323
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 324
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 325
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 326
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 327
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 328
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 329
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 330
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 331
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 332
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 333
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 334
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 335
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 336
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 337
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 338
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 339
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 340
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 341
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 342
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 343
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 344
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 345
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 346
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 347
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 348
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 349
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 350
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 351
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 352
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 353
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 354
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 355
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 356
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 357
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 358
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 359
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 360
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 361
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 362
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 363
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 364
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 365
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 366
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 367
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 368
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 369
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 370
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 371
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 372
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 373
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 374
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 375
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 376
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 377
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 378
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 379
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 380
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 381
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 382
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 383
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 384
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 385
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 386
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 387
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 388
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 389
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 390
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 391
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 392
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 393
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 394
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 395
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 396
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 397
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 398
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 399
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 400
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 401
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 402
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 403
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 404
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 405
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 406
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 407
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 408
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 409
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 410
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 411
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 412
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 413
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 414
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 415
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 416
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 417
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 418
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 419
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 420
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 421
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 422
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 423
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 424
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 425
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 426
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 427
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 428
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 429
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 430
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 431
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 432
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 433
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 434
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 435
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 436
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 437
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 438
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 439
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 440
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 441
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 442
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 443
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 444
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 445
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 446
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 447
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 448
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 449
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 450
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 451
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 452
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 453
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 454
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 455
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 456
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 457
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 458
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 459
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 460
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 461
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 462
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 463
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 464
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 465
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 466
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 467
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 468
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 469
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 470
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 471
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 472
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 473
transform_dataset() (*pl_bolts.datamodules.vocdetection_dataset.VOCDetection method*), 474

(*pl_bolts.models.rl.reinforce_model.Reinforce method*), 187
train_dataloader() (*pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradient method*), 190
train_dataloader() (*pl_bolts.models.self_supervised.amdim.amdim_module.AMDIM method*), 192
train_dataloader() (*pl_bolts.models.self_supervised.cpc.cpc_module.CPCV2 method*), 199
train_dataloader() (*pl_bolts.models.self_supervised.moco.moco2_module.MocoV2 method*), 206
train_dataloader() (*pl_bolts.models.self_supervised.simclr.simclr_module.SimCLR method*), 209
train_dataloader() (*pl_bolts.models.vision.image_gpt.igpt_module.ImageGPT method*), 216
train_dataloader_mixed() (*pl_bolts.datamodules.stl10_datamodule.STL10DataModule method*), 151
TRAIN_FILE_NAME (*pl_bolts.datamodules.cifar10_dataset.CIFAR10 attribute*), 135
train_transform() (*pl_bolts.datamodules.imagenet_datamodule.ImagenetDataModule method*), 139
train_transforms() (*pl_bolts.datamodules.lightning_datamodule.LightningDataModule property*), 143
training_epoch_end() (*pl_bolts.models.gans.basic.basic_gan_module.GAN method*), 159
training_step() (*pl_bolts.models.autoencoders.basic_ae.basic_ae_module.AE method*), 154
training_step() (*pl_bolts.models.autoencoders.basic_vae.basic_vae_module.VAE method*), 157
training_step() (*pl_bolts.models.gans.basic.basic_gan_module.GAN method*), 159
training_step() (*pl_bolts.models.mnist_module.LitMNIST method*), 217
training_step() (*pl_bolts.models.regression.linear_regression.LinearRegression method*), 160
training_step() (*pl_bolts.models.regression.logistic_regression.LogisticRegression method*), 161
training_step() (*pl_bolts.models.rl.double_dqn_model.DoubleDQN method*), 173
training_step() (*pl_bolts.models.rl.dqn_model.DQN method*), 176
training_step() (*pl_bolts.models.rl.noisy_dqn_model.NoisyDQN method*), 182
training_step() (*pl_bolts.models.rl.per_dqn_model.PERDQN method*), 184
training_step() (*pl_bolts.models.rl.reinforce_model.Reinforce method*), 187
training_step() (*pl_bolts.models.rl.vanilla_policy_gradient_model.PolicyGradient method*), 190
training_step() (*pl_bolts.models.self_supervised.amdim.amdim_module.AMDIM method*), 192
training_step() (*pl_bolts.models.self_supervised.cpc.cpc_module.CPCV2 method*), 199
training_step() (*pl_bolts.models.self_supervised.moco.moco2_module.MocoV2 method*), 206
training_step() (*pl_bolts.models.self_supervised.simclr.simclr_module.SimCLR method*), 209
training_step() (*pl_bolts.models.vision.image_gpt.igpt_module.ImageGPT method*), 216
training_step_end() (*pl_bolts.models.self_supervised.amdim.amdim_module.AMDIM method*), 192
TrainsLogger (*class in pl_bolts.loggers*), 220
TrainsLogger (*class in pl_bolts.loggers.trains*), 224
TrialCIFAR10 (*class in pl_bolts.datamodules.cifar10_dataset*), 136

U

UnlabeledImagenet (*class in pl_bolts.datamodules.imagenet_dataset*), 140
update_beta() (*pl_bolts.models.rl.common.memory.PERBuffer method*), 166
update_epsilon() (*pl_bolts.models.rl.common.agents.ValueAgent method*), 163
update_priorities() (*pl_bolts.models.rl.common.memory.PERBuffer method*), 167

V

VAE (*class in pl_bolts.models.autoencoders.basic_vae.basic_vae_module*), 156
val_dataloader() (*pl_bolts.datamodules.cifar10_datamodule.CIFAR method*), 133
val_dataloader() (*pl_bolts.datamodules.fashion_mnist_datamodule.F method*), 137
val_dataloader() (*pl_bolts.datamodules.imagenet_datamodule.Image method*), 139
val_dataloader() (*pl_bolts.datamodules.lightning_datamodule.Light method*), 143
val_dataloader() (*pl_bolts.datamodules.mnist_datamodule.MNISTD method*), 144
val_dataloader() (*pl_bolts.datamodules.sklearn_datamodule.Sklearn method*), 146
val_dataloader() (*pl_bolts.datamodules.ssl_imagenet_datamodule.S method*), 150
val_dataloader() (*pl_bolts.datamodules.stl10_datamodule.STL10Da method*), 152

val_dataloader() (pl_bolts.models.autoencoders.basic_ae.basic_ae_module.AE method), 154
val_dataloader() (pl_bolts.models.autoencoders.basic_vae.basic_vae_module.VAE method), 157
val_dataloader() (pl_bolts.models.mnist_module.LitMNIST method), 217
val_dataloader() (pl_bolts.models.self_supervised.amdim.amdim_module.AMDIM method), 192
val_dataloader() (pl_bolts.models.self_supervised.cpc.cpc_module.CPCV2 method), 199
val_dataloader() (pl_bolts.models.self_supervised.moco.moco2_module.MocoV2 method), 206
val_dataloader() (pl_bolts.models.self_supervised.simclr.simclr_module.SimCLR method), 209
val_dataloader() (pl_bolts.models.vision.image_gpt.igpt_module.ImageGPT method), 216
val_dataloader_mixed() (pl_bolts.datamodules.stl10_datamodule.STL10DataModule method), 152
val_transform() (pl_bolts.datamodules.imagenet_datamodule.ImagenetDataModule method), 139
val_transforms() (pl_bolts.datamodules.lightning_datamodule.LightningDataModule property), 143
validation_epoch_end() (pl_bolts.models.autoencoders.basic_ae.basic_ae_module.AE method), 154
validation_epoch_end() (pl_bolts.models.autoencoders.basic_vae.basic_vae_module.VAE method), 157
validation_epoch_end() (pl_bolts.models.mnist_module.LitMNIST method), 217
validation_epoch_end() (pl_bolts.models.regression.linear_regression.LinearRegression method), 160
validation_epoch_end() (pl_bolts.models.regression.logistic_regression.LogisticRegression method), 161
validation_epoch_end() (pl_bolts.models.self_supervised.amdim.amdim_module.AMDIM method), 192
validation_epoch_end() (pl_bolts.models.self_supervised.cpc.cpc_module.CPCV2 method), 199
validation_epoch_end() (pl_bolts.models.self_supervised.moco.moco2_module.MocoV2 method), 206
validation_epoch_end() (pl_bolts.models.self_supervised.simclr.simclr_module.SimCLR method), 209
validation_epoch_end() (pl_bolts.models.vision.image_gpt.igpt_module.ImageGPT method), 216
validation_step() (pl_bolts.models.autoencoders.basic_ae.basic_ae_module.AE method), 154
validation_step() (pl_bolts.models.autoencoders.basic_vae.basic_vae_module.VAE method), 157
validation_step() (pl_bolts.models.mnist_module.LitMNIST method), 217
validation_step() (pl_bolts.models.regression.linear_regression.LinearRegression method), 160
validation_step() (pl_bolts.models.regression.logistic_regression.LogisticRegression method), 161
validation_step() (pl_bolts.models.self_supervised.amdim.amdim_module.AMDIM method), 192
validation_step() (pl_bolts.models.self_supervised.cpc.cpc_module.CPCV2 method), 199
validation_step() (pl_bolts.models.self_supervised.moco.moco2_module.MocoV2 method), 206
validation_step() (pl_bolts.models.self_supervised.simclr.simclr_module.SimCLR method), 209
validation_step() (pl_bolts.models.vision.image_gpt.igpt_module.ImageGPT method), 216
ValueAgent (class in pl_bolts.models.rl.common.agents), 162
version() (pl_bolts.loggers.trains.TrainsLogger property), 228
version (pl_bolts.loggers.TrainsLogger property), 224

W

wide_resnet101_2() (in module pl_bolts.models.self_supervised.resnets), 212
wide_resnet50_2() (in module pl_bolts.models.self_supervised.resnets), 212