
PyTorch-Lightning-Bolts Documentation

Release 0.2.1

PyTorchLightning et al.

Sep 13, 2020

START HERE

1	Introduction Guide	1
2	Model quality control	11
3	Build a Callback	15
4	Info Callbacks	17
5	Self-supervised Callbacks	19
6	Variational Callbacks	21
7	Vision Callbacks	23
8	DataModules	27
9	Sklearn Datamodule	29
10	Vision DataModules	33
11	AsynchronousLoader	47
12	DummyDataset	49
13	Losses	51
14	How to use models	53
15	Autoencoders	61
16	Classic ML Models	67
17	Convolutional Architectures	71
18	GANs	75
19	Self-supervised Learning	79
20	Self-supervised learning Transforms	93
21	Self-supervised learning	103
22	Semi-supervised learning	105

23 Self-supervised Learning Contrastive tasks	107
24 Indices and tables	111
Python Module Index	187
Index	189

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='st110').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 = 150528
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 first 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
```


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.

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.

INFO CALLBACKS

These callbacks give all sorts of useful information during training.

4.1 Print Table Metrics

This callback 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
```


SELF-SUPERVISED CALLBACKS

Useful callbacks for self-supervised learning models

5.1 BYOLMAWeightUpdate

The exponential moving average weight-update rule from Bring Your Own Latent (BYOL).

```
class pl_bolts.callbacks.self_supervised.BYOLMAWeightUpdate (initial_tau=0.996)  
    Bases: pytorch_lightning.Callback
```

Weight update rule from BYOL.

Your model should have a:

- self.online_network.
- self.target_network.

Updates the target_network params using an exponential moving average update rule weighted by tau. BYOL claims this keeps the online_network from collapsing.

Note: Automatically increases tau from *initial_tau* to 1.0 with every training step

Example:

```
from pl_bolts.callbacks.self_supervised import BYOLMAWeightUpdate  
  
# model must have 2 attributes  
model = Model()  
model.online_network = ...  
model.target_network = ...  
  
# make sure to set max_steps in Trainer  
trainer = Trainer(callbacks=[BYOLMAWeightUpdate()], max_steps=1000)
```

Parameters *initial_tau* – starting tau. Auto-updates with every training step

5.2 SSLOnlineEvaluator

Appends a MLP for fine-tuning to the given model. Callback has its own mini-inner loop.

```
class pl_bolts.callbacks.self_supervised.SSLOnlineEvaluator(drop_p=0.2,    hid-  
                                                         den_dim=1024,  
                                                         z_dim=None,  
                                                         num_classes=None)
```

Bases: `pytorch_lightning.Callback`

Attaches a MLP for finetuning using the standard self-supervised protocol.

Example:

```
from pl_bolts.callbacks.self_supervised import SSLOnlineEvaluator  
  
# your model must have 2 attributes  
model = Model()  
model.z_dim = ... # the representation dim  
model.num_classes = ... # the num of classes in the model
```

Parameters

- **drop_p** (`float`) – (0.2) dropout probability
- **hidden_dim** (`int`) –

(1024) the hidden dimension for the finetune MLP

get_representations (`pl_module, x`)

Override this to customize for the particular model :`param_sphinx_paramlinks_pl_bolts.callbacks.self_supervised.SSLOnlineEvaluator.get_representations`:`param_sphinx_paramlinks_pl_bolts.callbacks.self_supervised.SSLOnlineEvaluator.get_representations.x`:


VARIATIONAL CALLBACKS

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

6.1 Latent Dim Interpolator

Interpolates latent dims.

Example output:



```
class pl_bolts.callbacks.variational.LatentDimInterpolator(interpolate_epoch_interval=20,  
                                                         range_start=-  
                                                         5,      range_end=5,  
                                                         num_samples=2)
```

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 $[-5, 5]$ $(-5, -4, -3, \dots, 3, 4, 5)$

Example:

```
from pl_bolts.callbacks import LatentDimInterpolator

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

Parameters

- **interpolate_epoch_interval** –
- **range_start** – default -5
- **range_end** – default 5
- **num_samples** – default 2


VISION CALLBACKS

Useful callbacks for vision models

7.1 Confused Logit

Shows how the input would have to change to move the prediction from one logit to the other

Example outputs:



```
class pl_bolts.callbacks.vision.confused_logit.ConfusedLogitCallback(top_k,  
                                                                    projec-  
                                                                    tion_factor=3,  
                                                                    min_logit_value=5.0,  
                                                                    log-  
                                                                    ging_batch_interval=20,  
                                                                    max_logit_difference=0.1)
```

Bases: `pytorch_lightning.Callback`

Takes the logit predictions of a model and when the probabilities of two classes are very close, the model doesn't have high certainty that it should pick one vs the other class.

This callback shows how the input would have to change to swing the model from one label prediction to the other.

In this case, the network predicts a 5... but gives almost equal probability to an 8. The images show what about the original 5 would have to change to make it more like a 5 or more like an 8.

For each confused logit the confused images are generated by taking the gradient from a logit wrt an input for the top two closest logits.

Example:

```
from pl_bolts.callbacks.vision import ConfusedLogitCallback
trainer = Trainer(callbacks=[ConfusedLogitCallback()])
```

Note: whenever called, this model will look for `self.last_batch` and `self.last_logits` in the `LightningModule`

Note: this callback supports tensorboard only right now

Parameters

- **top_k** – How many “offending” images we should plot
- **projection_factor** – How much to multiply the input image to make it look more like this logit label
- **min_logit_value** – Only consider logit values above this threshold
- **logging_batch_interval** – how frequently to inspect/potentially plot something
- **max_logit_difference** – when the top 2 logits are within this threshold we consider them confused

Authored by:

- Alfredo Canziani
-

7.2 Tensorboard Image Generator

Generates images from a generative model and plots to tensorboard

class `pl_bolts.callbacks.vision.image_generation.TensorboardGenerativeModelImageSampler` (*num*
Bases: `pytorch_lightning.Callback`

Generates images and logs to tensorboard. Your model must implement the forward function for generation

Requirements:

```
# model must have img_dim arg
model.img_dim = (1, 28, 28)

# model forward must work for sampling
z = torch.rand(batch_size, latent_dim)
img_samples = your_model(z)
```

Example:

```
from pl_bolts.callbacks import TensorboardGenerativeModelImageSampler

trainer = Trainer(callbacks=[TensorboardGenerativeModelImageSampler()])
```


DATAMODULES

DataModules (introduced in PyTorch Lightning 0.9.0) decouple the data from a model. A DataModule is simply a collection of a training dataloader, val dataloader and test dataloader. In addition, it specifies how to:

- Downloading/preparing data.
- Train/val/test splits.
- Transforms

Then you can use it like this:

Example:

```
dm = MNISTDataModule('path/to/data')
model = LitModel()

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

Or use it manually with plain PyTorch

Example:

```
dm = MNISTDataModule('path/to/data')
for batch in dm.train_dataloader():
    ...
for batch in dm.val_dataloader():
    ...
for batch in dm.test_dataloader():
    ...
```

Please visit the PyTorch Lightning documentation for more details on DataModules

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)
```

9.1 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
```

9.2 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,
                                                                ran-
                                                                dom_state=1234,
                                                                shuffle=True,
                                                                *args,
                                                                **kwargs)
```

Bases: `pytorch_lightning.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
```


VISION DATAMODULES

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


10.1 Supervised learning

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

10.1.1 BinaryMNIST

```
class pl_bolts.datamodules.binary_mnist_datamodule.BinaryMNISTDataModule(data_dir,  
                                                                           val_split=5000,  
                                                                           num_workers=16,  
                                                                           nor-  
                                                                           mal-  
                                                                           ize=False,  
                                                                           seed=42,  
                                                                           *args,  
                                                                           **kwargs)
```

Bases: `pytorch_lightning.LightningDataModule`



Specs:

- 10 classes (1 per digit)
- Each image is (1 x 28 x 28)

Binary MNIST, train, val, test splits and transforms

Transforms:

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

Example:

```
from pl_bolts.datamodules import BinaryMNISTDataModule

dm = BinaryMNISTDataModule('.')
model = LitModel()

Trainer().fit(model, 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

10.1.2 CityScapes

```
class pl_bolts.datamodules.cityscapes_datamodule.CityscapesDataModule(data_dir,
                                                                    val_split=5000,
                                                                    num_workers=16,
                                                                    batch_size=32,
                                                                    seed=42,
                                                                    *args,
                                                                    **kwargs)
```

Bases: `pytorch_lightning.LightningDataModule`



Standard Cityscapes, train, val, test splits and transforms

Specs:

- 30 classes (road, person, sidewalk, etc...)
- (image, target) - image dims: (3 x 32 x 32), target dims: (3 x 32 x 32)

Transforms:

```
transforms = transform_lib.Compose([
    transform_lib.ToTensor(),
    transform_lib.Normalize(
        mean=[0.28689554, 0.32513303, 0.28389177],
        std=[0.18696375, 0.19017339, 0.18720214]
    )
])
```

Example:

```
from pl_bolts.datamodules import CityscapesDataModule

dm = CityscapesDataModule(PATH)
model = LitModel()

Trainer().fit(model, 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
- **batch_size** – number of examples per training/eval step

prepare_data()

Saves Cityscapes files to data_dir

test_dataloader()

Cityscapes test set uses the test split

train_dataloader()

Cityscapes train set with removed subset to use for validation

val_dataloader()


Cityscapes val set uses a subset of the training set for validation

property num_classes

Return: 30

10.1.3 CIFAR-10

```
class pl_bolts.datamodules.cifar10_datamodule.CIFAR10DataModule(data_dir=None,  
                                                                val_split=5000,  
                                                                num_workers=16,  
                                                                batch_size=32,  
                                                                seed=42,  
                                                                *args,  
                                                                **kwargs)
```

Bases: `pytorch_lightning.LightningDataModule`**Specs:**

- 10 classes (1 per class)
- Each image is (3 x 32 x 32)

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()

Trainer().fit(model, dm)
```

Or you can set your own transforms

Example:

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

Parameters

- **data_dir** (Optional[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
- **batch_size** (int) – number of examples per training/eval step

prepare_data()

Saves CIFAR10 files to data_dir

test_dataloader()

CIFAR10 test set uses the test split

train_dataloader()

CIFAR train set removes a subset to use for validation

val_dataloader()

CIFAR10 val set uses a subset of the training set for validation


property num_classes

Return: 10

10.1.4 FashionMNIST

```
class pl_bolts.datamodules.fashion_mnist_datamodule.FashionMNISTDataModule(data_dir,
                                                                    val_split=5000,
                                                                    num_workers=16,
                                                                    seed=42,
                                                                    *args,
                                                                    **kwargs)

Bases: pytorch_lightning.LightningDataModule
```



Specs:

- 10 classes (1 per type)
- Each image is (1 x 28 x 28)

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()

Trainer().fit(model, 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

10.1.5 Imagenet

```
class pl_bolts.datamodules.imagenet_datamodule.ImagenetDataModule (data_dir,
                                                                meta_dir=None,
                                                                num_imgs_per_val_class=50,
                                                                im-
                                                                age_size=224,
                                                                num_workers=16,
                                                                batch_size=32,
                                                                *args,
                                                                **kwargs)
```

Bases: `pytorch_lightning.LightningDataModule`


Specs:

- 1000 classes
- Each image is (3 x varies x varies) (here we default to 3 x 224 x 224)

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

dm = ImagenetDataModule(IMGNET_PATH)
model = LitModel()

Trainer().fit(model, dm)
```

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
- **batch_size** (`int`) – batch_size

`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()`

Uses the validation split of imagenet2012 for testing

train_dataloader()

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

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()

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


10.1.6 MNIST

```
class pl_bolts.datamodules.mnist_datamodule.MNISTDataModule(data_dir='/',
                                                             val_split=5000,
                                                             num_workers=16,
                                                             normalize=False,
                                                             seed=42,
                                                             batch_size=32,
                                                             *args, **kwargs)
```

Bases: `pytorch_lightning.LightningDataModule`

Specs:

- 10 classes (1 per digit)
- Each image is (1 x 28 x 28)



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()

Trainer().fit(model, 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

10.2 Semi-supervised learning

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

10.2.1 Imagenet (ssl)

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

Bases: `pytorch_lightning.LightningDataModule`

10.2.2 STL-10

```
class pl_bolts.datamodules.stl10_datamodule.STL10DataModule (data_dir=None,  
unla-  
beled_val_split=5000,  
train_val_split=500,  
num_workers=16,  
batch_size=32,  
seed=42, *args,  
**kwargs)
```

Bases: `pytorch_lightning.LightningDataModule`

Specs:

- 10 classes (1 per type)
- Each image is (3 x 96 x 96)

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()

Trainer().fit(model, dm)
```

Parameters

- **data_dir** (Optional[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
- **batch_size** (int) – the batch size

prepare_data()

Downloads the unlabeled, train and test split

test_dataloader()

Loads the test split of STL10

Parameters

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

train_dataloader()

Loads the ‘unlabeled’ split minus a portion set aside for validation via *unlabeled_val_split*.

train_dataloader_mixed()

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()

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()

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

$\text{unlabeled_val} = (\text{unlabeled} - \text{train_val_split})$

$\text{labeled_val} = (\text{train} - \text{train_val_split})$

$\text{full_val} = \text{unlabeled_val} + \text{labeled_val}$

Parameters

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

ASYNCHRONOUSLOADER

This dataloader behaves identically to the standard pytorch dataloader, but will transfer data asynchronously to the GPU with training. You can also use it to wrap an existing dataloader.

Example:: `dataloader = AsynchronousLoader(DataLoader(ds, batch_size=16), device=device)`

for b in dataloader: ...

```
class pl_bolts.datamodules.async_dataloader.AsynchronousLoader(data, device=torch.device,
                                                                q_size=10,
                                                                num_batches=None,
                                                                **kwargs)
```

Bases: `object`

Class for asynchronously loading from CPU memory to device memory with DataLoader.

Note that this only works for single GPU training, multiGPU uses PyTorch's DataParallel or DistributedDataParallel which uses its own code for transferring data across GPUs. This could just break or make things slower with DataParallel or DistributedDataParallel.

Parameters

- **data** – The PyTorch Dataset or DataLoader we're using to load.
- **device** – The PyTorch device we are loading to
- **q_size** – Size of the queue used to store the data loaded to the device
- **num_batches** – Number of batches to load. This must be set if the dataloader doesn't have a finite `__len__`. It will also override `DataLoader.__len__` if set and `DataLoader` has a `__len__`. Otherwise it can be left as `None`
- ****kwargs** – Any additional arguments to pass to the dataloader if we're constructing one here

DUMMYDATASET

```
class pl_bolts.datamodules.dummy_dataset.DummyDataset(*shapes,  
                                                       num_samples=10000)
```

Bases: `torch.utils.data.Dataset`

Generate a dummy dataset

Parameters

- ***shapes** – list of shapes
- **num_samples** – how many samples to use in this dataset

Example:

```
from pl_bolts.datamodules import DummyDataset

# mnist dims
>>> ds = DummyDataset((1, 28, 28), (1,))
>>> dl = DataLoader(ds, batch_size=7)
...
>>> batch = next(iter(dl))
>>> x, y = batch
>>> x.size()
torch.Size([7, 1, 28, 28])
>>> y.size()
torch.Size([7, 1])
```


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

13.1 Your Loss

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

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

14.1 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='st110')
```

14.2 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.

14.2.1 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

14.2.2 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()
```

14.2.3 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)
```

14.3 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!

14.4 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.

14.4.1 Extending work

Perhaps a research project requires modifying a part of a known approach. In this case, you're better off only changing that part of a system that is already known 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
```

14.4.2 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
```

14.4.3 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):
        suocer().__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
```


AUTOENCODERS

This section houses autoencoders and variational autoencoders.

15.1 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
```

You can use the pretrained models present in bolts.

CIFAR-10 pretrained model:


```
from pl_bolts.models.autoencoders import AE

ae = AE(input_height=32)
print(AE.pretrained_weights_available())
ae = ae.from_pretrained('cifar10-resnet18')

ae.freeze()
```

- [Tensorboard for AE on CIFAR-10](#)

Training:



```
class pl_bolts.models.autoencoders.AE(input_height, enc_type='resnet18', first_conv=False,
                                       maxpool1=False, enc_out_dim=512, kl_coeff=0.1, la-
                                       tent_dim=256, lr=0.0001, **kwargs)
```

Bases: `pytorch_lightning.LightningModule`

Standard AE

Model is available pretrained on different datasets:

Example:

```
# not pretrained
ae = AE()

# pretrained on imagenet
ae = AE.from_pretrained('resnet50-imagenet')

# pretrained on cifar10
ae = AE.from_pretrained('resnet18-cifar10')
```

Parameters

- **input_height** – height of the images
- **enc_type** – option between resnet18 or resnet50
- **first_conv** – use standard kernel_size 7, stride 2 at start or replace it with kernel_size 3, stride 1 conv
- **maxpool1** – use standard maxpool to reduce spatial dim of feat by a factor of 2
- **enc_out_dim** – set according to the out_channel count of encoder used (512 for resnet18, 2048 for resnet50)
- **latent_dim** – dim of latent space
- **lr** – learning rate for Adam

15.1.1 Variational Autoencoders

15.2 Basic VAE

Use the VAE like so.

```
from pl_bolts.models.autoencoders import VAE

model = VAE()
```

(continues on next page)

(continued from previous page)

```
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
```

You can use the pretrained models present in bolts.

CIFAR-10 pretrained model:


```
from pl_bolts.models.autoencoders import VAE

vae = VAE(input_height=32)
print(VAE.pretrained_weights_available())
vae = vae.from_pretrained('cifar10-resnet18')

vae.freeze()
```

- Tensorboard for VAE on CIFAR-10

Training:



STL-10 pretrained model:

```
from pl_bolts.models.autoencoders import VAE

vae = VAE(input_height=96, first_conv=True)
```

(continues on next page)


(continued from previous page)

```
print(VAE.pretrained_weights_available())
vae = vae.from_pretrained('cifar10-resnet18')

vae.freeze()
```

- Tensorboard for VAE on STL-10

Training:



```
class pl_bolts.models.autoencoders.VAE(input_height, enc_type='resnet18', first_conv=False,
                                       maxpool1=False, enc_out_dim=512, kl_coeff=0.1,
                                       latent_dim=256, lr=0.0001, **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.from_pretrained('resnet50-imagenet')

# pretrained on cifar10
vae = VAE.from_pretrained('resnet18-cifar10')
```

Parameters

- **input_height** – height of the images
- **enc_type** – option between resnet18 or resnet50
- **first_conv** – use standard kernel_size 7, stride 2 at start or replace it with kernel_size 3, stride 1 conv
- **maxpool1** – use standard maxpool to reduce spatial dim of feat by a factor of 2

- **enc_out_dim** – set according to the out_channel count of encoder used (512 for resnet18, 2048 for resnet50)
- **kl_coeff** – coefficient for kl term of the loss
- **latent_dim** – dim of latent space
- **lr** – learning rate for Adam

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.

16.1 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 that minimizes the mean squared error between the predicted and true target values.

We formulate the linear regression model as a single-layer neural network. By default we include only one neuron in the output layer, although you can specify the `output_dim` yourself.

Add either L1 or L2 regularization, or both, by specifying the regularization strength (default 0).

```
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,
                                                                    out-
                                                                    put_dim=1,
                                                                    bias=True,
                                                                    learn-
                                                                    ing_rate=0.0001,
                                                                    opti-
                                                                    mizer=torch.optim.Adam,
                                                                    l1_strength=0.0,
                                                                    l2_strength=0.0,
                                                                    **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+)
 - **output_dim** (`int`) – number of dimensions of the output (default=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** (`float`) – L1 regularization strength (default=None)
 - **l2_strength** (`float`) – L2 regularization strength (default=None)
-

16.2 Logistic Regression

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

In the binary case, we formulate the logistic regression model as a one-layer neural network with one neuron in the output layer and a sigmoid activation function. In the multi-class case, we use a single-layer neural network but now with k neurons in the output, where k is the number of classes. This is also referred to as multinomial logistic regression.

Add either L1 or L2 regularization, or both, by specifying the regularization strength (default 0).

```
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 first 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
```

(continues on next page)

(continued from previous page)

```

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)

CONVOLUTIONAL ARCHITECTURES

This package lists contributed convolutional architectures.

17.1 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=seq_len, 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

17.2 Image GPT

```
class pl_bolts.models.vision.image_gpt.igpt_module.ImageGPT (datamodule=None,
                                                             embed_dim=16,
                                                             heads=2,          lay-
                                                             ers=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.)
<code>-embed_dim</code>	16	512
<code>-heads</code>	2	8
<code>-layers</code>	8	24
<code>-pixels</code>	28	32
<code>-vocab_size</code>	16	512
<code>-num_classes</code>	10	10
<code>-batch_size</code>	64	128
<code>-learning_rate</code>	0.01	0.01
<code>-steps</code>	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

17.3 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])
```

GANs

Collection of Generative Adversarial Networks


18.1 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`(*input_channels*, *input_height*, *input_width*, *latent_dim*=32, *learning_rate*=0.0002, ***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** – the datamodule (train, val, test splits)
- **latent_dim** (`int`) – emb dim for encoder
- **batch_size** – the batch size
- **learning_rate** (`float`) – the learning rate
- **data_dir** – where to store data

- **num_workers** – 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)
```


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

19.1 Use cases

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

19.1.1 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)
```

19.1.2 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),
)
```

19.1.3 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)
```

19.2 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.

19.2.1 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

19.2.2 BYOL

```
class pl_bolts.models.self_supervised.BYOL (num_classes,          learning_rate=0.2,
                                             weight_decay=1.5e-06,   input_height=32,
                                             batch_size=32,          num_workers=0,
                                             warmup_epochs=10,        max_epochs=1000,
                                             **kwargs)
```

Bases: `pytorch_lightning.LightningModule`

PyTorch Lightning implementation of [Bootstrap Your Own Latent \(BYOL\)](#)

Paper authors: Jean-Bastien Grill ,Florian Strub, Florent Alth  , Corentin Tallec, Pierre H. Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, R  mi Munos, Michal Valko.

Model implemented by:

- [Annika Brundyn](#)

Warning: Work in progress. This implementation is still being verified.

TODOs:

- verify on CIFAR-10
- verify on STL-10
- pre-train on imagenet

Example:

```
import pytorch_lightning as pl
from pl_bolts.models.self_supervised import BYOL
from pl_bolts.datamodules import CIFAR10DataModule
from pl_bolts.models.self_supervised.simclr.simclr_transforms import (
    SimCLREvalDataTransform, SimCLRTrainDataTransform)

# model
model = BYOL(num_classes=10)

# data
dm = CIFAR10DataModule(num_workers=0)
dm.train_transforms = SimCLRTrainDataTransform(32)
dm.val_transforms = SimCLREvalDataTransform(32)

trainer = pl.Trainer()
trainer.fit(model, dm)
```

Train:

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

CLI command:

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

(continues on next page)

(continued from previous page)

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

Parameters

- **datamodule** – The datamodule
- **learning_rate** (*float*) – the learning rate
- **weight_decay** (*float*) – optimizer weight decay
- **input_height** (*int*) – image input height
- **batch_size** (*int*) – the batch size
- **num_workers** (*int*) – number of workers
- **warmup_epochs** (*int*) – num of epochs for scheduler warm up
- **max_epochs** (*int*) – max epochs for scheduler

19.2.3 CPC (V2)

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](#)

To Train:

```
import pytorch_lightning as pl
from pl_bolts.models.self_supervised import CPCV2
from pl_bolts.datamodules import CIFAR10DataModule
from pl_bolts.models.self_supervised.cpc import (
    CPCTrainTransformsCIFAR10, CPCEvalTransformsCIFAR10)

# data
dm = CIFAR10DataModule(num_workers=0)
dm.train_transforms = CPCTrainTransformsCIFAR10()
dm.val_transforms = CPCEvalTransformsCIFAR10()

# model
model = CPCV2()

# fit
trainer = pl.Trainer()
trainer.fit(model, dm)
```

To finetune:

```
python cpc_finetuner.py
  --ckpt_path path/to/checkpoint.ckpt
  --dataset cifar10
  --gpus 1
```

CIFAR-10 and STL-10 baselines

CPCv2 does not report baselines on CIFAR-10 and STL-10 datasets. Results in table are reported from the [YADIM](#) paper.

Table 1: CPCv2 implementation results

Dataset	test acc	Encoder	Optimizer	Batch	Epochs	Hardware	LR
CIFAR-10	84.52	CPCresnet101	Adam	64	1000 (upto 24 hours)	1 V100 (32GB)	4e-5
STL-10	78.36	CPCresnet101	Adam	144	1000 (upto 72 hours)	4 V100 (32GB)	1e-4
ImageNet	54.82	CPCresnet101	Adam	3072	1000 (upto 21 days)	64 V100 (32GB)	4e-5

CIFAR-10 pretrained model:


```
from pl_bolts.models.self_supervised import CPCV2

weight_path = 'https://pl-bolts-weights.s3.us-east-2.amazonaws.com/cpc/cpc-cifar10-v4-
  exp3/epoch%3D474.ckpt'
cpc_v2 = CPCV2.load_from_checkpoint(weight_path, strict=False)


cpc_v2.freeze()
```

- Tensorboard for CIFAR10

Pre-training:



Fine-tuning:



STL-10 pretrained model:


```
from pl_bolts.models.self_supervised import CPCV2

weight_path = 'https://pl-bolts-weights.s3.us-east-2.amazonaws.com/cpc/cpc-stl10-v0-
exp3/epoch%3D624.ckpt'
cpc_v2 = CPCV2.load_from_checkpoint(weight_path, strict=False)


cpc_v2.freeze()
```

- Tensorboard for STL10

Pre-training:



Fine-tuning:



ImageNet pretrained model:

```

from pl_bolts.models.self_supervised import CPCV2

weight_path = 'https://pl-bolts-weights.s3.us-east-2.amazonaws.com/cpc/cpcv2_weights/
↳checkpoints/epoch%3D526.ckpt'
cpc_v2 = CPCV2.load_from_checkpoint(weight_path, strict=False)

cpc_v2.freeze()


```

- Tensorboard for ImageNet

Pre-training:



Fine-tuning:



CPCV2 API

```

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

```

Bases: `pytorch_lightning.LightningModule`

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 custom 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 workers
- **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

19.2.4 Moco (V2)

```
class pl_bolts.models.self_supervised.MocoV2 (base_encoder='resnet18', emb_dim=128,
                                             num_negatives=65536,          en-
                                             coder_momentum=0.999,        soft-
                                             max_temperature=0.07,         learn-
                                             ing_rate=0.03,                momentum=0.9,
                                             weight_decay=0.0001,         datamod-
                                             ule=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
```

(continues on next page)

(continued from previous page)

```
--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: `img_q`: a batch of query images `img_k`: a batch of key images

Output: logits, targets

init_encoders (`base_encoder`)

Override to add your own encoders

19.2.5 SimCLR

PyTorch Lightning implementation of SimCLR

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

Model implemented by:

- William Falcon
- Tullie Murrell

To Train:

```
import pytorch_lightning as pl
from pl_bolts.models.self_supervised import SimCLR
from pl_bolts.datamodules import CIFAR10DataModule
from pl_bolts.models.self_supervised.simclr.simclr_transforms import (
    SimCLREvalDataTransform, SimCLRTrainDataTransform)

# data
dm = CIFAR10DataModule(num_workers=0)
dm.train_transforms = SimCLRTrainDataTransform(32)
dm.val_transforms = SimCLREvalDataTransform(32)

# model
model = SimCLR(num_samples=dm.num_samples, batch_size=dm.batch_size)

# fit
trainer = pl.Trainer()
trainer.fit(model, dm)
```

CIFAR-10 baseline

Table 2: Cifar-10 implementation results

Implementation	test acc	Encoder	Optimizer	Batch	Epochs	Hardware	LR
Original	92.00?	resnet50	LARS	512	1000	1 V100 (32GB)	1.0
Ours	85.68	resnet50	LARS	512	960 (12 hr)	1 V100 (32GB)	1e-6


CIFAR-10 pretrained model:

```
from pl_bolts.models.self_supervised import SimCLR



weight_path = 'https://pl-bolts-weights.s3.us-east-2.amazonaws.com/simclr/simclr-
→cifar10-v1-exp12_87_52/epoch%3D960.ckpt'
simclr = SimCLR.load_from_checkpoint(weight_path, strict=False)

simclr.freeze()
```

Pre-training:



Fine-tuning (Single layer MLP, 1024 hidden units):



To reproduce:

```
# pretrain
python simclr_module.py
    --gpus 1
    --dataset cifar10
    --batch_size 512
    --learning_rate 1e-06
```

(continues on next page)

(continued from previous page)

```
--num_workers 8

# finetune
python simclr_finetuner.py
    --ckpt_path path/to/epoch=xyz.ckpt
    --gpus 1
```

SimCLR API

```
class pl_bolts.models.self_supervised.SimCLR(batch_size, num_samples,
                                              warmup_epochs=10, lr=0.0001,
                                              opt_weight_decay=1e-06,
                                              loss_temperature=0.5, **kwargs)
```

Bases: `pytorch_lightning.LightningModule`

Parameters

- **batch_size** – the batch size
- **num_samples** – num samples in the dataset
- **warmup_epochs** – epochs to warmup the lr for
- **lr** – the optimizer learning rate
- **opt_weight_decay** – the optimizer weight decay
- **loss_temperature** – the loss temperature

SELF-SUPERVISED LEARNING TRANSFORMS

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

20.1 CPC transforms

Transforms used for CPC

20.1.1 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.
```

20.1.2 CIFAR-10 Eval (c)

```
class pl_bolts.models.self_supervised.cpc.transforms.CPCEvalTransformsCIFAR10(patch_size=8,
                                                                              over-
                                                                              lap=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.
```

20.1.3 Imagenet Train (c)

```
class pl_bolts.models.self_supervised.cpc.transforms.CPCTrainTransformsImageNet128(patch_size=
                                                                                       over-
                                                                                       lap=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.

20.1.4 Imagenet Eval (c)

`class pl_bolts.models.self_supervised.cpc.transforms.CPCEvalTransformsImageNet128` (*patch_size=*
over-
lap=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.

20.1.5 STL-10 Train (c)

class pl_bolts.models.self_supervised.cpc.transforms.**CPCTrainTransformsSTL10** (*patch_size=16, over-lap=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.

20.1.6 STL-10 Eval (c)

class pl_bolts.models.self_supervised.cpc.transforms.**CPCEvalTransformsSTL10** (*patch_size=16, over-lap=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.

20.2 AMDIM transforms

Transforms used for AMDIM

20.2.1 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.

20.2.2 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.

20.2.3 Imagenet Train (a)

class `pl_bolts.models.self_supervised.amdim.transforms.AMDIMTrainTransformsImageNet128` (*height*)

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.

20.2.4 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.

20.2.5 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.

20.2.6 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 = AMDIMEvalTransformsSTL10()  
view1 = transform(x)
```

`__call__` (*inp*)

Call self as a function.

20.3 MOCO V2 transforms

Transforms used for MOCO V2

20.3.1 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.
```

20.3.2 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.
```

20.3.3 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.
```

20.3.4 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.
```


20.3.5 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.
```

20.3.6 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.
```

20.4 SimCLR transforms

Transforms used for SimCLR

20.4.1 Train (sc)

```
class pl_bolts.models.self_supervised.simclr.simclr_transforms.SimCLRTrainDataTransform (input_height, input_width, 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.
```

20.4.2 Eval (sc)

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.

SELF-SUPERVISED LEARNING

Collection of useful functions for self-supervised learning

21.1 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()
```

21.2 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, re-
    turn_all_feature_maps=False)
```

21.3 SSL backbone finetuner

```
class pl_bolts.models.self_supervised.ssl_finetuner.SSLFineTuner (backbone,  
                                                             in_features,  
                                                             num_classes,  
                                                             hid-  
                                                             den_dim=1024)
```

Bases: `pytorch_lightning.LightningModule`

Finetunes a self-supervised learning backbone using the standard evaluation protocol of a singler layer MLP with 1024 units

Example:

```
from pl_bolts.utils.self_supervised import SSLFineTuner
from pl_bolts.models.self_supervised import CPCV2
from pl_bolts.datamodules import CIFAR10DataModule
from pl_bolts.models.self_supervised.cpc.transforms import _
↳ CPCEvalTransformsCIFAR10,
                                     ↳
↳ CPCTrainTransformsCIFAR10

# pretrained model
backbone = CPCV2.load_from_checkpoint(PATH, strict=False)

# dataset + transforms
dm = CIFAR10DataModule(data_dir='.')
dm.train_transforms = CPCTrainTransformsCIFAR10()
dm.val_transforms = CPCEvalTransformsCIFAR10()

# finetuner
finetuner = SSLFineTuner(backbone, in_features=backbone.z_dim, num_
↳ classes=backbone.num_classes)

# train
trainer = pl.Trainer()
trainer.fit(finetuner, dm)

# test
trainer.test(datamodule=dm)
```

Parameters

- **backbone** – a pretrained model
- **in_features** – feature dim of backbone outputs
- **num_classes** – classes of the dataset
- **hidden_dim** – dim of the MLP (1024 default used in self-supervised literature)

SEMI-SUPERVISED LEARNING

Collection of utilities for semi-supervised learning where some part of the data is labeled and the other part is not.

22.1 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

22.2 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

SELF-SUPERVISED LEARNING CONTRASTIVE TASKS

This section implements popular contrastive learning tasks used in self-supervised learning.

23.1 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 $B \times B$ 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,
                                                                    bidi-
                                                                    rec-
                                                                    tional=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)]

```

23.2 Context prediction tasks

The following tasks aim to predict a target using a context representation.

23.2.1 CPCTask

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, em-
bed_scale=0.1)

Bases: `torch.nn.Module`

Loss used in CPC

INDICES AND TABLES

- [genindex](#)
- [modindex](#)
- [search](#)

Logo

24.1 PyTorch Lightning Bolts

Pretrained SOTA Deep Learning models, callbacks and more for research and production with PyTorch Lightning and PyTorch

24.1.1 Trending contributors

24.1.2 Continuous Integration

| System / PyTorch ver. | 1.4 (min. req.) | 1.6 (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 |

24.1.3 Install

Simple installation from PyPI

```
pip install pytorch-lightning-bolts
```

Install bleeding-edge (no guarantees)

```
pip install git+https://github.com/PytorchLightning/pytorch-lightning-bolts.  
↪git@master --upgrade
```

24.1.4 Docs

- master
- stable
- 0.2.0
- 0.1.1

24.1.5 What is Bolts

Bolts is a Deep learning research and production toolbox of:

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

24.1.6 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
        # -----

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

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

```

24.1.7 Who is Bolts for?

- Corporate production teams
- Professional researchers
- Ph.D. students
- Linear + Logistic regression heroes

24.1.8 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())

```

24.1.9 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.

24.1.10 Team

Bolts is supported by the PyTorch Lightning team and the PyTorch Lightning community!

24.2 pl_bolts.callbacks package

Collection of PyTorchLightning callbacks

24.2.1 Subpackages

pl_bolts.callbacks.vision package

Submodules

pl_bolts.callbacks.vision.confused_logit module

```
class pl_bolts.callbacks.vision.confused_logit.ConfusedLogitCallback(top_k,  
                                                                    projec-  
                                                                    tion_factor=3,  
                                                                    min_logit_value=5.0,  
                                                                    log-  
                                                                    ging_batch_interval=20,  
                                                                    max_logit_difference=0.1)
```

Bases: `pytorch_lightning.Callback`

Takes the logit predictions of a model and when the probabilities of two classes are very close, the model doesn't have high certainty that it should pick one vs the other class.

This callback shows how the input would have to change to swing the model from one label prediction to the other.

In this case, the network predicts a 5... but gives almost equal probability to an 8. The images show what about the original 5 would have to change to make it more like a 5 or more like an 8.

For each confused logit the confused images are generated by taking the gradient from a logit wrt an input for the top two closest logits.

Example:

```
from pl_bolts.callbacks.vision import ConfusedLogitCallback  
trainer = Trainer(callbacks=[ConfusedLogitCallback()])
```

Note: whenever called, this model will look for `self.last_batch` and `self.last_logits` in the `LightningModule`

Note: this callback supports tensorboard only right now

Parameters

- **top_k** – How many “offending” images we should plot
- **projection_factor** – How much to multiply the input image to make it look more like this logit label
- **min_logit_value** – Only consider logit values above this threshold
- **logging_batch_interval** – how frequently to inspect/potentially plot something
- **max_logit_difference** – when the top 2 logits are within this threshold we consider them confused

Authored by:

- Alfredo Canziani

```
static _ConfusedLogitCallback__draw_sample (fig, axarr, row_idx, col_idx, img, title)
__plot (confusing_x, confusing_y, trainer, model, mask_idxs)
on_train_batch_end (trainer, pl_module, batch, batch_idx, dataloader_idx)
```

pl_bolts.callbacks.vision.image_generation module

class `pl_bolts.callbacks.vision.image_generation.TensorboardGenerativeModelImageSampler` (*num*
 Bases: `pytorch_lightning.Callback`

Generates images and logs to tensorboard. Your model must implement the forward function for generation

Requirements:

```
# model must have img_dim arg
model.img_dim = (1, 28, 28)

# model forward must work for sampling
z = torch.rand(batch_size, latent_dim)
img_samples = your_model(z)
```

Example:

```
from pl_bolts.callbacks import TensorboardGenerativeModelImageSampler

trainer = Trainer(callbacks=[TensorboardGenerativeModelImageSampler()])

on_epoch_end (trainer, pl_module)
```

24.2.2 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.self_supervised module

class pl_bolts.callbacks.self_supervised.BYOLMAWeightUpdate (*initial_tau=0.996*)
 Bases: pytorch_lightning.Callback

Weight update rule from BYOL.

Your model should have a:

- self.online_network.
- self.target_network.

Updates the target_network params using an exponential moving average update rule weighted by tau. BYOL claims this keeps the online_network from collapsing.

Note: Automatically increases tau from *initial_tau* to 1.0 with every training step

Example:

```
from pl_bolts.callbacks.self_supervised import BYOLMAWeightUpdate

# model must have 2 attributes
model = Model()
model.online_network = ...
model.target_network = ...

# make sure to set max_steps in Trainer
trainer = Trainer(callbacks=[BYOLMAWeightUpdate()], max_steps=1000)
```

Parameters *initial_tau* – starting tau. Auto-updates with every training step

on_train_batch_end (*trainer, pl_module, batch, batch_idx, dataloader_idx*)

update_tau (*pl_module, trainer*)

update_weights (*online_net, target_net*)

class pl_bolts.callbacks.self_supervised.SSLOnlineEvaluator (*drop_p=0.2, hidden_dim=1024, z_dim=None, num_classes=None*)
 Bases: pytorch_lightning.Callback

Attaches a MLP for finetuning using the standard self-supervised protocol.

Example:

```
from pl_bolts.callbacks.self_supervised import SSLOnlineEvaluator

# your model must have 2 attributes
model = Model()
model.z_dim = ... # the representation dim
model.num_classes = ... # the num of classes in the model
```

Parameters

- **drop_p**(float) – (0.2) dropout probability
- **hidden_dim**(int) –

(1024) the hidden dimension for the finetune MLP

get_representations(*pl_module*, *x*)

Override this to customize for the particular model :param_sphinx_paramlinks_pl_bolts.callbacks.self_supervised.SSLOnlineEvaluator.get_representations.x:

on_pretrainRoutine_start(*trainer*, *pl_module*)

on_train_batch_end(*trainer*, *pl_module*, *batch*, *batch_idx*, *dataloader_idx*)

to_device(*batch*, *device*)

pl_bolts.callbacks.variational module

```
class pl_bolts.callbacks.variational.LatentDimInterpolator(interpolate_epoch_interval=20,
                                                           range_start=-5,
                                                           range_end=5,
                                                           num_samples=2)
```

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
- **num_samples** – default 2

interpolate_latent_space(*pl_module*, *latent_dim*)

on_epoch_end(*trainer*, *pl_module*)

24.3 pl_bolts.datamodules package

24.3.1 Submodules

pl_bolts.datamodules.async_dataloader module

```
class pl_bolts.datamodules.async_dataloader.AsynchronousLoader (data, device=torch.device,
q_size=10,
num_batches=None,
**kwargs)
```

Bases: `object`

Class for asynchronously loading from CPU memory to device memory with DataLoader.

Note that this only works for single GPU training, multiGPU uses PyTorch's DataParallel or DistributedDataParallel which uses its own code for transferring data across GPUs. This could just break or make things slower with DataParallel or DistributedDataParallel.

Parameters

- **data** – The PyTorch Dataset or DataLoader we're using to load.
- **device** – The PyTorch device we are loading to
- **q_size** – Size of the queue used to store the data loaded to the device
- **num_batches** – Number of batches to load. This must be set if the dataloader doesn't have a finite `__len__`. It will also override `DataLoader.__len__` if set and `DataLoader` has a `__len__`. Otherwise it can be left as `None`
- ****kwargs** – Any additional arguments to pass to the dataloader if we're constructing one here

`load_instance` (*sample*)

`load_loop` ()

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.binary_mnist_datamodule` module

```
class pl_bolts.datamodules.binary_mnist_datamodule.BinaryMNIST(*args,
                                                                **kwargs)
    Bases: torchvision.datasets.MNIST
class pl_bolts.datamodules.binary_mnist_datamodule.BinaryMNISTDataModule(data_dir,
                                                                            val_split=5000,
                                                                            num_workers=16,
                                                                            normalize=False,
                                                                            seed=42,
                                                                            *args,
                                                                            **kwargs)
    Bases: pytorch_lightning.LightningDataModule
```



Specs:

- 10 classes (1 per digit)
- Each image is (1 x 28 x 28)

Binary MNIST, train, val, test splits and transforms

Transforms:

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

Example:

```

from pl_bolts.datamodules import BinaryMNISTDataModule

dm = BinaryMNISTDataModule('.')
model = LitModel()

Trainer().fit(model, 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 = 'mnist'


property num_classes

Return: 10

pl_bolts.datamodules.cifar10_datamodule module

```
class pl_bolts.datamodules.cifar10_datamodule.CIFAR10DataModule(data_dir=None,  
                                                                val_split=5000,  
                                                                num_workers=16,  
                                                                batch_size=32,  
                                                                seed=42,  
                                                                *args,  
                                                                **kwargs)
```

Bases: `pytorch_lightning.LightningDataModule`

**Specs:**

- 10 classes (1 per class)
- Each image is (3 x 32 x 32)

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()  
  
Trainer().fit(model, dm)
```

Or you can set your own transforms

Example:

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

Parameters

- **data_dir** (*Optional[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
- **batch_size** (*int*) – number of examples per training/eval step

default_transforms ()

prepare_data ()

Saves CIFAR10 files to data_dir

test_dataloader ()

CIFAR10 test set uses the test split

train_dataloader ()

CIFAR train set removes a subset to use for validation

val_dataloader ()

CIFAR10 val set uses a subset of the training set for validation

extra_args = {}

name = '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])
>>> data, label = dataset[0]
>>> data.shape
torch.Size([3, 32, 32])
```

(continues on next page)

(continued from previous page)

```
>>> 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), rela-
                                                         bel=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.cityscapes_datamodule` module

```
class pl_bolts.datamodules.cityscapes_datamodule.CityscapesDataModule(data_dir,
                                                                    val_split=5000,
                                                                    num_workers=16,
                                                                    batch_size=32,
                                                                    seed=42,
                                                                    *args,
                                                                    **kwargs)
```

Bases: `pytorch_lightning.LightningDataModule`

Standard Cityscapes, train, val, test splits and transforms

Specs:

- 30 classes (road, person, sidewalk, etc...)
- (image, target) - image dims: (3 x 32 x 32), target dims: (3 x 32 x 32)



Transforms:

```
transforms = transform_lib.Compose([
    transform_lib.ToTensor(),
    transform_lib.Normalize(
        mean=[0.28689554, 0.32513303, 0.28389177],
        std=[0.18696375, 0.19017339, 0.18720214]
    )
])
```

Example:

```
from pl_bolts.datamodules import CityscapesDataModule

dm = CityscapesDataModule(PATH)
model = LitModel()

Trainer().fit(model, 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
- **batch_size** – number of examples per training/eval step

default_transforms()

prepare_data()

Saves Cityscapes files to data_dir

test_dataloader()

Cityscapes test set uses the test split

```
train_dataloader()
    Cityscapes train set with removed subset to use for validation

val_dataloader()
    Cityscapes val set uses a subset of the training set for validation

extra_args = {}

name = 'Cityscapes'

property num_classes
    Return: 30
```

pl_bolts.datamodules.concat_dataset module

```
class pl_bolts.datamodules.concat_dataset.ConcatDataset(*datasets)
    Bases: torch.utils.data.Dataset
```

pl_bolts.datamodules.dummy_dataset module

```
class pl_bolts.datamodules.dummy_dataset.DummyDataset(*shapes,
                                                         num_samples=10000)
```

Bases: torch.utils.data.Dataset

Generate a dummy dataset

Parameters

- ***shapes** – list of shapes
- **num_samples** – how many samples to use in this dataset

Example:

```
from pl_bolts.datamodules import DummyDataset

# mnist dims
>>> ds = DummyDataset((1, 28, 28), (1,))
>>> dl = DataLoader(ds, batch_size=7)
...
>>> batch = next(iter(dl))
>>> x, y = batch
>>> x.size()
torch.Size([7, 1, 28, 28])
>>> y.size()
torch.Size([7, 1])
```

```
class pl_bolts.datamodules.dummy_dataset.DummyDetectionDataset(img_shape=(3,
                                                                              256,
                                                                              256),
                                                                num_boxes=1,
                                                                num_classes=2,
                                                                num_samples=10000)


    Bases: torch.utils.data.Dataset

    _random_bbox()
```

pl_bolts.datamodules.fashion_mnist_datamodule module

```
class pl_bolts.datamodules.fashion_mnist_datamodule.FashionMNISTDataModule(data_dir,
                                                                           val_split=5000,
                                                                           num_workers=16,
                                                                           seed=42,
                                                                           *args,
                                                                           **kwargs)
```

Bases: `pytorch_lightning.LightningDataModule`

**Specs:**

- 10 classes (1 per type)
- Each image is (1 x 28 x 28)

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()

Trainer().fit(model, 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 = '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,  
                                                                    im-  
                                                                    age_size=224,  
                                                                    num_workers=16,  
                                                                    batch_size=32,  
                                                                    *args,  
                                                                    **kwargs)
```


Bases: `pytorch_lightning.LightningDataModule`

Specs:

- 1000 classes
- Each image is (3 x varies x varies) (here we default to 3 x 224 x 224)

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

dm = ImagenetDataModule(IMGNET_PATH)
model = LitModel()

Trainer().fit(model, dm)
```

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
- **batch_size** (`int`) – batch_size

`_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()

Uses the validation split of imagenet2012 for testing

train_dataloader()

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

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()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 = '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.mnist_datamodule module

```
class pl_bolts.datamodules.mnist_datamodule.MNISTDataModule(data_dir='/',
                                                             val_split=5000,
                                                             num_workers=16,
                                                             normalize=False,
                                                             seed=42,
                                                             batch_size=32,
                                                             *args, **kwargs)
```

Bases: `pytorch_lightning.LightningDataModule`



Specs:

- 10 classes (1 per digit)
- Each image is (1 x 28 x 28)

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()

Trainer().fit(model, 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 = '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: pytorch_lightning.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)
```

```
train_dataloader(batch_size=16)
```

```
val_dataloader(batch_size=16)
```

```
name = '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)
...
>>> # 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: pytorch_lightning.LightningDataModule

_default_transforms()
_verify_splits(data_dir, split)
prepare_data()
test_dataloader(batch_size, num_images_per_class, add_normalize=False)
train_dataloader(batch_size, num_images_per_class=-1, add_normalize=False)
val_dataloader(batch_size, num_images_per_class=50, add_normalize=False)
name = 'imagenet'
property num_classes

```

pl_bolts.datamodules.stl10_datamodule module

```

class pl_bolts.datamodules.stl10_datamodule.STL10DataModule(data_dir=None,
                                                             unla-
                                                             beled_val_split=5000,
                                                             train_val_split=500,
                                                             num_workers=16,
                                                             batch_size=32,
                                                             seed=42, *args,
                                                             **kwargs)

Bases: pytorch_lightning.LightningDataModule

```

**Specs:**

- 10 classes (1 per type)
- Each image is (3 x 96 x 96)

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()

Trainer().fit(model, dm)
```

Parameters

- **data_dir** (*Optional*[*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
- **batch_size** (*int*) – the batch size

default_transforms ()

prepare_data ()

Downloads the unlabeled, train and test split

test_dataloader ()

Loads the test split of STL10

Parameters

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

train_dataloader ()

Loads the ‘unlabeled’ split minus a portion set aside for validation via *unlabeled_val_split*.

train_dataloader_labeled ()

train_dataloader_mixed ()

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()

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_labeled()

val_dataloader_mixed()

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 = 'st110'

property num_classes

pl_bolts.datamodules.vocdetection_datamodule module

class pl_bolts.datamodules.vocdetection_datamodule.**Compose**(transforms)

Bases: `object`

Like `torchvision.transforms.compose` but works for (image, target)

__call__(image, target)

Call self as a function.

class pl_bolts.datamodules.vocdetection_datamodule.**VOCDetectionDataModule**(data_dir, year='2012', num_workers=16, normalize=False, *args, **kwargs)

Bases: `pytorch_lightning.LightningDataModule`

TODO(teddykoker) docstring

_default_transforms()

prepare_data()

Saves VOCDetection files to data_dir

train_dataloader(batch_size=1, transforms=None)

VOCDetection train set uses the *train* subset

Parameters

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

val_dataloader (*batch_size=1, transforms=None*)
VOCDetection val set uses the *val* subset

Parameters

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

name = 'vocdetection'

property num_classes

Return: 21

`pl_bolts.datamodules.vocdetection_datamodule._collate_fn(batch)`

`pl_bolts.datamodules.vocdetection_datamodule._prepare_voc_instance(image, target)`

Prepares VOC dataset into appropriate target for fasterrcnn

<https://github.com/pytorch/vision/issues/1097#issuecomment-508917489>

24.4 pl_bolts.metrics package

24.4.1 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

24.5 pl_bolts.models package

Collection of PyTorchLightning models

24.5.1 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 CIFAR10 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 is a resnet18 backbone followed by linear layers which map representations to latent space. The default decoder mirrors the encoder architecture and is similar to an inverted resnet18.

```
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(input_height,
                                                                enc_type='resnet18',
                                                                first_conv=False,
                                                                max-
                                                                pool1=False,
                                                                enc_out_dim=512,
                                                                kl_coeff=0.1, la-
                                                                tent_dim=256,
                                                                lr=0.0001,
                                                                **kwargs)
```

Bases: `pytorch_lightning.LightningModule`

Standard AE

Model is available pretrained on different datasets:

Example:

```
# not pretrained
ae = AE()

# pretrained on imagenet
ae = AE.from_pretrained('resnet50-imagenet')

# pretrained on cifar10
ae = AE.from_pretrained('resnet18-cifar10')
```

Parameters

- **input_height** – height of the images
- **enc_type** – option between resnet18 or resnet50
- **first_conv** – use standard kernel_size 7, stride 2 at start or replace it with kernel_size 3, stride 1 conv
- **maxpool1** – use standard maxpool to reduce spatial dim of feat by a factor of 2
- **enc_out_dim** – set according to the out_channel count of encoder used (512 for resnet18, 2048 for resnet50)
- **latent_dim** – dim of latent space
- **lr** – learning rate for Adam

```
static add_model_specific_args (parent_parser)
configure_optimizers ()
forward (z)
from_pretrained (checkpoint_name)
static pretrained_weights_available ()
step (batch, batch_idx)
training_step (batch, batch_idx)
validation_step (batch, batch_idx)
pretrained_urls = {'cifar10-resnet18': 'https://pl-bolts-weights.s3.us-east-2.amazonaws.com/cifar10-resnet18.pth'}
pl_bolts.models.autoencoders.basic_ae.basic_ae_module.cli_main (args=None)
```

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 CIFAR10 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 is a resnet18 backbone followed by linear layers which map representations to μ and σ . The default decoder mirrors the encoder architecture and is similar to an inverted resnet18. 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(input_height,
                                                                    enc_type='resnet18',
                                                                    first_conv=False,
                                                                    max-
                                                                    pool1=False,
                                                                    enc_out_dim=512,
                                                                    kl_coeff=0.1,
                                                                    la-
                                                                    tent_dim=256,
                                                                    lr=0.0001,
                                                                    **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.from_pretrained('resnet50-imagenet')

# pretrained on cifar10
vae = VAE.from_pretrained('resnet18-cifar10')
```

Parameters

- **input_height** – height of the images
- **enc_type** – option between resnet18 or resnet50
- **first_conv** – use standard kernel_size 7, stride 2 at start or replace it with kernel_size 3, stride 1 conv
- **maxpool1** – use standard maxpool to reduce spatial dim of feat by a factor of 2
- **enc_out_dim** – set according to the out_channel count of encoder used (512 for resnet18, 2048 for resnet50)
- **kl_coeff** – coefficient for kl term of the loss
- **latent_dim** – dim of latent space
- **lr** – learning rate for Adam

`_run_step(x)`

`static add_model_specific_args(parent_parser)`

`configure_optimizers()`

`forward(z)`

`from_pretrained(checkpoint_name)`

```
static pretrained_weights_available()
sample(mu, log_var)
step(batch, batch_idx)
training_step(batch, batch_idx)
validation_step(batch, batch_idx)

pretrained_urls = {'cifar10-resnet18': 'https://pl-bolts-weights.s3.us-east-2.amazonaws.com/cifar10-resnet18.pth'}

pl_bolts.models.autoencoders.basic_vae.basic_vae_module.cli_main(args=None)
```

Submodules

pl_bolts.models.autoencoders.components module

```
class pl_bolts.models.autoencoders.components.DecoderBlock(inplanes, planes,
                                                            scale=1, upsam-
                                                            ple=None)
```

Bases: `torch.nn.Module`

ResNet block, but convs replaced with resize convs, and channel increase is in second conv, not first

forward(x)

expansion = 1

```
class pl_bolts.models.autoencoders.components.DecoderBottleneck(inplanes, planes,
                                                                scale=1, upsam-
                                                                ple=None)
```

Bases: `torch.nn.Module`

ResNet bottleneck, but convs replaced with resize convs

forward(x)

expansion = 4

```
class pl_bolts.models.autoencoders.components.EncoderBlock(inplanes, planes,
                                                            stride=1, downsam-
                                                            ple=None)
```

Bases: `torch.nn.Module`

ResNet block, copied from <https://github.com/pytorch/vision/blob/master/torchvision/models/resnet.py#L35>

forward(x)

expansion = 1

```
class pl_bolts.models.autoencoders.components.EncoderBottleneck(inplanes, planes,
                                                                stride=1, downsam-
                                                                ple=None)
```

Bases: `torch.nn.Module`

ResNet bottleneck, copied from <https://github.com/pytorch/vision/blob/master/torchvision/models/resnet.py#L75>

```

    forward (x)

    expansion = 4
class pl_bolts.models.autoencoders.components.Interpolate (size=None,
                                                            scale_factor=None)

    Bases: torch.nn.Module

    nn.Module wrapper for F.interpolate

    forward (x)

class pl_bolts.models.autoencoders.components.ResNetDecoder (block,          layers,
                                                                latent_dim,          in-
                                                                put_height,
                                                                first_conv=False,
                                                                maxpool1=False)

    Bases: torch.nn.Module

    Resnet in reverse order

    _make_layer (block, planes, blocks, scale=1)

    forward (x)

class pl_bolts.models.autoencoders.components.ResNetEncoder (block,          layers,
                                                                first_conv=False,
                                                                maxpool1=False)

    Bases: torch.nn.Module

    _make_layer (block, planes, blocks, stride=1)

    forward (x)

pl_bolts.models.autoencoders.components.conv1x1 (in_planes, out_planes, stride=1)
    1x1 convolution

pl_bolts.models.autoencoders.components.conv3x3 (in_planes, out_planes, stride=1)
    3x3 convolution with padding

pl_bolts.models.autoencoders.components.resize_conv1x1 (in_planes,          out_planes,
                                                            scale=1)
    upsample + 1x1 convolution with padding to avoid checkerboard artifact

pl_bolts.models.autoencoders.components.resize_conv3x3 (in_planes,          out_planes,
                                                            scale=1)
    upsample + 3x3 convolution with padding to avoid checkerboard artifact

pl_bolts.models.autoencoders.components.resnet18_decoder (latent_dim, input_height,
                                                            first_conv, maxpool1)

pl_bolts.models.autoencoders.components.resnet18_encoder (first_conv, maxpool1)

pl_bolts.models.autoencoders.components.resnet50_decoder (latent_dim, input_height,
                                                            first_conv, maxpool1)

pl_bolts.models.autoencoders.components.resnet50_encoder (first_conv, maxpool1)

```

pl_bolts.models.detection package

Submodules

pl_bolts.models.detection.faster_rcnn module

```
class pl_bolts.models.detection.faster_rcnn.FasterRCNN (learning_rate=0.0001,
                                                         num_classes=91,      pre-
                                                         trained=False,      pre-
                                                         trained_backbone=True,
                                                         train-
                                                         able_backbone_layers=3,
                                                         replace_head=True,
                                                         **kwargs)
```

Bases: `pytorch_lightning.LightningModule`

PyTorch Lightning implementation of [Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks](#).

Paper authors: Shaoqing Ren, Kaiming He, Ross Girshick, Jian Sun

Model implemented by:

- *Teddy Koker* <<https://github.com/teddykoker>>

During training, the model expects both the input tensors, as well as targets (list of dictionary), containing:

- `boxes` (`FloatTensor[N, 4]`): the ground truth boxes in `[x1, y1, x2, y2]` format.
- `labels` (`Int64Tensor[N]`): the class label for each ground truth box

CLI command:

```
# PascalVOC
python faster_rcnn.py --gpus 1 --pretrained True
```

Parameters

- **learning_rate** (`float`) – the learning rate
- **num_classes** (`int`) – number of detection classes (including background)
- **pretrained** (`bool`) – if true, returns a model pre-trained on COCO train2017
- **pretrained_backbone** (`bool`) – if true, returns a model with backbone pre-trained on Imagenet
- **trainable_backbone_layers** (`int`) – number of trainable resnet layers starting from final block

static `add_model_specific_args` (`parent_parser`)

`configure_optimizers` ()

`forward` (`x`)

`training_step` (`batch`, `batch_idx`)

`validation_epoch_end` (`outs`)

`validation_step` (`batch`, `batch_idx`)


```
pl_bolts.models.detection.faster_rcnn._evaluate_iou(target, pred)
    Evaluate intersection over union (IOU) for target from dataset and output prediction from model

pl_bolts.models.detection.faster_rcnn.run_cli()
```

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(input_channels, input_height,
                                                       input_width, latent_dim=32,
                                                       learning_rate=0.0002,
                                                       **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** – the datamodule (train, val, test splits)
- **latent_dim** (`int`) – emb dim for encoder
- **batch_size** – the batch size
- **learning_rate** (`float`) – the learning rate
- **data_dir** – where to store data
- **num_workers** – 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*)

training_step (*batch, batch_idx, optimizer_idx*)

`pl_bolts.models.gans.basic.basic_gan_module.cli_main(args=None)`

pl_bolts.models.gans.basic.components module

class `pl_bolts.models.gans.basic.components.Discriminator` (*img_shape, hid-*
den_dim=1024)

Bases: `torch.nn.Module`

forward (*img*)

class `pl_bolts.models.gans.basic.components.Generator` (*latent_dim, img_shape, hid-*
den_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,*
out-
put_dim=1,
bias=True,
learn-
ing_rate=0.0001,
opti-
mizer=torch.optim.Adam,
l1_strength=0.0,
l2_strength=0.0,
***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+)
- **output_dim** (*int*) – number of dimensions of the output (default=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** (*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.regression.linear_regression.cli_main()`

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.regression.logistic_regression.cli_main()
```

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',
    en-
    coder='amdim_encoder',
    con-
    trastive_task=torch.nn.Module,
    im-
    age_channels=3,
    im-
    age_height=32,
    en-
    coder_feature_dim=320,
    embed-
    ding_fx_dim=1280,
    conv_block_depth=10,
    use_bn=False,
    tclip=20.0,
    learn-
    ing_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.amdim_module.cli_main ()
```

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,
                                                                en-
                                                                coder_feature_dim=64,
                                                                embed-
                                                                ding_fx_dim=512,
                                                                conv_block_depth=3,
                                                                en-
                                                                coder_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',
                                                                    trans-
                                                                    form=None,
                                                                    tar-
                                                                    get_transform=None,
                                                                    down-
                                                                    load=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_spli
        :param _sphinx_paramlinks_pl_bolts.models.self_supervised.amdim.ssl_datasets.SSLDatasetMixin.generate_train_val_spli
        :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.exa
        :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` (*height*

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.byol package

Submodules

pl_bolts.models.self_supervised.byol.byol_module module

```
class pl_bolts.models.self_supervised.byol.byol_module.BYOL(num_classes, learning_rate=0.2,
weight_decay=1.5e-06, input_height=32, batch_size=32, num_workers=0,
warmup_epochs=10, max_epochs=1000, **kwargs)
```

Bases: `pytorch_lightning.LightningModule`

PyTorch Lightning implementation of [Bootstrap Your Own Latent \(BYOL\)](#)

Paper authors: Jean-Bastien Grill ,Florian Strub, Florent Alth  , Corentin Tallec, Pierre H. Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, R  mi Munos, Michal Valko.

Model implemented by:

- [Annika Brundyn](#)

Warning: Work in progress. This implementation is still being verified.

TODOs:

- verify on CIFAR-10
- verify on STL-10
- pre-train on imagenet

Example:

```
import pytorch_lightning as pl
from pl_bolts.models.self_supervised import BYOL
from pl_bolts.datamodules import CIFAR10DataModule
from pl_bolts.models.self_supervised.simclr.simclr_transforms import (
    SimCLREvalDataTransform, SimCLRTrainDataTransform)

# model
model = BYOL(num_classes=10)

# data
dm = CIFAR10DataModule(num_workers=0)
dm.train_transforms = SimCLRTrainDataTransform(32)
dm.val_transforms = SimCLREvalDataTransform(32)

trainer = pl.Trainer()
trainer.fit(model, dm)
```

Train:

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

CLI command:

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

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

Parameters

- **datamodule** – The datamodule
- **learning_rate** (*float*) – the learning rate
- **weight_decay** (*float*) – optimizer weight decay
- **input_height** (*int*) – image input height
- **batch_size** (*int*) – the batch size
- **num_workers** (*int*) – number of workers
- **warmup_epochs** (*int*) – num of epochs for scheduler warm up
- **max_epochs** (*int*) – max epochs for scheduler

static `add_model_specific_args` (*parent_parser*)

configure_optimizers ()

cosine_similarity (*a*, *b*)

forward (*x*)

on_train_batch_end (*batch*, *batch_idx*, *dataloader_idx*)

Return type None

shared_step (*batch, batch_idx*)

training_step (*batch, batch_idx*)

validation_step (*batch, batch_idx*)

`pl_bolts.models.self_supervised.byol.byol_module.cli_main()`

pl_bolts.models.self_supervised.byol.models module

class `pl_bolts.models.self_supervised.byol.models.MLP` (*input_dim=2048, den_size=4096, put_dim=256*, *hid-out-*)

Bases: `torch.nn.Module`

forward (*x*)

class `pl_bolts.models.self_supervised.byol.models.SiameseArm` (*encoder=None*)

Bases: `torch.nn.Module`

forward (*x*)

pl_bolts.models.self_supervised.cpc package

Submodules

pl_bolts.models.self_supervised.cpc.cpc_finetuner module

`pl_bolts.models.self_supervised.cpc.cpc_finetuner.cli_main()`

pl_bolts.models.self_supervised.cpc.cpc_module module

CPC V2

class `pl_bolts.models.self_supervised.cpc.cpc_module.CPCV2` (*datamodule=None, en-coder='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`

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 custom 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 workers
- **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_I)
init_encoder()
load_pretrained(encoder)
shared_step(batch)
training_step(batch, batch_nb)
validation_step(batch, batch_nb)
```

`pl_bolts.models.self_supervised.cpc.networks` module

```
class pl_bolts.models.self_supervised.cpc.networks.CPCResNet(sample_batch,
                                                             block,          layers,
                                                             zero_init_residual=False,
                                                             groups=1,
                                                             width_per_group=64,
                                                             re-
                                                             place_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,
                                                                downsam-
                                                                ple_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, dila-
                                                       tion=1)

    3x3 convolution with padding

pl_bolts.models.self_supervised.cpc.networks.cpc_resnet101 (sample_batch,
                                                            **kwargs)

pl_bolts.models.self_supervised.cpc.networks.cpc_resnet50 (sample_batch,
                                                            **kwargs)

```

pl_bolts.models.self_supervised.cpc.transforms module

```

class pl_bolts.models.self_supervised.cpc.transforms.CPCEvalTransformsCIFAR10 (patch_size=8,
                                                                                over-
                                                                                lap=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

```

(continues on next page)

(continued from previous page)

```

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=*
over-
lap=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,*
over-
lap=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, over-lap=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=8, over-lap=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, over-lap=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,
                                                                    sched-
                                                                    ule=(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,
                                                                    en-
                                                                    coder_momentum=0.999,
                                                                    soft-
                                                                    max_temperature=0.07,
                                                                    learn-
                                                                    ing_rate=0.03,
                                                                    momentum=0.9,
                                                                    weight_decay=0.0001,
                                                                    datamod-
                                                                    ule=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](https://github.com/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: `img_q`: a batch of query images `img_k`: a batch of key images

Output: logits, targets

`init_encoders(base_encoder)`

Override to add your own encoders

`training_step(batch, batch_idx)`

`validation_epoch_end(outputs)`

`validation_step(batch, batch_idx)`

`pl_bolts.models.self_supervised.moco.moco2_module.cli_main()`

`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_finetuner module

```
pl_bolts.models.self_supervised.simclr.simclr_finetuner.cli_main()
```

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=2048,
                                                                    hid-
                                                                    den_dim=2048,
                                                                    out-
                                                                    put_dim=128)
    Bases: torch.nn.Module

    forward (x)

class pl_bolts.models.self_supervised.simclr.simclr_module.SimCLR (batch_size,
                                                                    num_samples,
                                                                    warmup_epochs=10,
                                                                    lr=0.0001,
                                                                    opt_weight_decay=1e-06,
                                                                    loss_temperature=0.5,
                                                                    **kwargs)
    Bases: pytorch_lightning.LightningModule
```

Parameters

- **batch_size** – the batch size
- **num_samples** – num samples in the dataset
- **warmup_epochs** – epochs to warmup the lr for
- **lr** – the optimizer learning rate
- **opt_weight_decay** – the optimizer weight decay
- **loss_temperature** – the loss temperature

```
static add_model_specific_args (parent_parser)
```

```
configure_optimizers ()
```

```
exclude_from_wt_decay (named_params, weight_decay, skip_list=['bias', 'bn'])
```

```
forward (x)
```

```
init_encoder ()
```

```
setup (stage)
```

```
shared_step (batch, batch_idx)
```

```
training_step (batch, batch_idx)
```

```
validation_step (batch, batch_idx)
```

```
pl_bolts.models.self_supervised.simclr.simclr_module.cli_main ()
```

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(input_height, 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.resnet50_bn` (*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_bn.pretrained`: If True, returns a model pre-trained on ImageNet :type `_sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet50_bn.pretrained`: bool :param `_sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet50_bn.progress`: If True, displays a progress bar of the download to stderr :type `_sphinx_paramlinks_pl_bolts.models.self_supervised.resnets.resnet50_bn.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.self_supervised.ssl_finetuner module

```
class pl_bolts.models.self_supervised.ssl_finetuner.SSLFineTuner(backbone,
                                                                in_features,
                                                                num_classes,
                                                                hid-
                                                                den_dim=1024)
```

Bases: `pytorch_lightning.LightningModule`

Finetunes a self-supervised learning backbone using the standard evaluation protocol of a singler layer MLP with 1024 units

Example:

```
from pl_bolts.utils.self_supervised import SSLFineTuner
from pl_bolts.models.self_supervised import CPCV2
from pl_bolts.datamodules import CIFAR10DataModule
from pl_bolts.models.self_supervised.cpc.transforms import _
    CPCEvalTransformsCIFAR10,
    CPCTrainTransformsCIFAR10

# pretrained model
backbone = CPCV2.load_from_checkpoint(PATH, strict=False)

# dataset + transforms
dm = CIFAR10DataModule(data_dir='.')
dm.train_transforms = CPCTrainTransformsCIFAR10()
dm.val_transforms = CPCEvalTransformsCIFAR10()

# finetuner
finetuner = SSLFineTuner(backbone, in_features=backbone.z_dim, num_
    classes=backbone.num_classes)

# train
trainer = pl.Trainer()
trainer.fit(finetuner, dm)

# test
trainer.test(datamodule=dm)
```

Parameters

- **backbone** – a pretrained model
- **in_features** – feature dim of backbone outputs
- **num_classes** – classes of the dataset
- **hidden_dim** – dim of the MLP (1024 default used in self-supervised literature)

configure_optimizers ()

on_train_epoch_start ()

Return type None

shared_step (batch)

test_step (batch, batch_idx)

training_step (batch, batch_idx)

validation_step (batch, batch_idx)

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=seq_len, 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,          lay-
                                                             ers=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.)
<code>-embed_dim</code>	16	512
<code>-heads</code>	2	8
<code>-layers</code>	8	24
<code>-pixels</code>	28	32
<code>-vocab_size</code>	16	512
<code>-num_classes</code>	10	10
<code>-batch_size</code>	64	128
<code>-learning_rate</code>	0.01	0.01
<code>-steps</code>	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)
```

```

test_epoch_end (outs)
test_step (batch, batch_idx)
training_step (batch, batch_idx)
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

pl_bolts.models.vision.image_gpt.igpt_module.cli_main()

```

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, hid-
                                                    den_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)

24.5.2 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="", **kwargs)  
  
    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)  
  
pl_bolts.models.mnist_module.cli_main ()
```

24.6 pl_bolts.losses package

24.6.1 Submodules

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, em-  
                                                         bed_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,
                                                                           bidi-
                                                                           rec-
                                                                           tional=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_ftn (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]

24.7 pl_bolts.optimizers package

24.7.1 Submodules

pl_bolts.optimizers.lars_scheduling module

References

- <https://github.com/NVIDIA/apex/blob/master/apex/parallel/LARC.py>
- <https://arxiv.org/pdf/1708.03888.pdf>
- <https://github.com/noahgolmant/pytorch-lars/blob/master/lars.py>

class `pl_bolts.optimizers.lars_scheduling.LARSWrapper` (*optimizer*, *eta*=0.02,
clip=True, *eps*=1e-08)

Bases: `object`

Wrapper that adds LARS scheduling to any optimizer. This helps stability with huge batch sizes.

Parameters

- **optimizer** – torch optimizer
- **eta** – LARS coefficient (trust)
- **clip** – True to clip LR
- **eps** – adaptive_lr stability coefficient

step()

update_p(*p, group, weight_decay*)

property param_groups

property state

pl_bolts.optimizers.lr_scheduler module

```
class pl_bolts.optimizers.lr_scheduler.LinearWarmupCosineAnnealingLR(optimizer,
                                                                    warmup_epochs,
                                                                    max_epochs,
                                                                    warmup_start_lr=0.0,
                                                                    eta_min=0.0,
                                                                    last_epoch=-1)
```

Bases: torch.optim.lr_scheduler._LRScheduler

Sets the learning rate of each parameter group to follow a linear warmup schedule between warmup_start_lr and base_lr followed by a cosine annealing schedule between base_lr and eta_min.

Warning: It is recommended to call `step()` for `LinearWarmupCosineAnnealingLR` after each iteration as calling it after each epoch will keep the starting lr at warmup_start_lr for the first epoch which is 0 in most cases.

Warning: passing epoch to `step()` is being deprecated and comes with an EPOCH_DEPRECATION_WARNING. It calls the `_get_closed_form_lr()` method for this scheduler instead of `get_lr()`. Though this does not change the behavior of the scheduler, when passing epoch param to `step()`, the user should call the `step()` function before calling train and validation methods.

Parameters

- **optimizer** (*Optimizer*) – Wrapped optimizer.
- **warmup_epochs** (*int*) – Maximum number of iterations for linear warmup
- **max_epochs** (*int*) – Maximum number of iterations
- **warmup_start_lr** (*float*) – Learning rate to start the linear warmup. Default: 0.
- **eta_min** (*float*) – Minimum learning rate. Default: 0.
- **last_epoch** (*int*) – The index of last epoch. Default: -1.

Example

```
>>> layer = nn.Linear(10, 1)
>>> optimizer = Adam(layer.parameters(), lr=0.02)
>>> scheduler = LinearWarmupCosineAnnealingLR(optimizer, warmup_epochs=10, max_
↳ epochs=40)
>>> #
>>> # the default case
>>> for epoch in range(40):
...     # train(...)
...     # validate(...)
...     scheduler.step()
>>> #
>>> # passing epoch param case
>>> for epoch in range(40):
...     scheduler.step(epoch)
...     # train(...)
...     # validate(...)
```

`_get_closed_form_lr()`

Called when epoch is passed as a param to the *step* function of the scheduler.

Return type `List[float]`

`get_lr()`

Compute learning rate using chainable form of the scheduler

Return type `List[float]`

24.8 pl_bolts.transforms package

24.8.1 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*,
over-
lap_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.

24.8.2 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()
```

24.9 pl_bolts.utils package

24.9.1 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

`pl_bolts.utils.shaping` module

`pl_bolts.utils.shaping.tile` (*a*, *dim*, *n_tile*)

PYTHON MODULE INDEX

p

`pl_bolts.callbacks`, 114
`pl_bolts.callbacks.printing`, 116
`pl_bolts.callbacks.self_supervised`, 117
`pl_bolts.callbacks.variational`, 118
`pl_bolts.callbacks.vision`, 114
`pl_bolts.callbacks.vision.confused_logit`, 114
`pl_bolts.callbacks.vision.image_generation`, 115
`pl_bolts.datamodules`, 119
`pl_bolts.datamodules.async_dataloader`, 119
`pl_bolts.datamodules.base_dataset`, 119
`pl_bolts.datamodules.binary_mnist_datamodule`, 120
`pl_bolts.datamodules.cifar10_datamodule`, 122
`pl_bolts.datamodules.cifar10_dataset`, 124
`pl_bolts.datamodules.cityscapes_datamodule`, 126
`pl_bolts.datamodules.concat_dataset`, 128
`pl_bolts.datamodules.dummy_dataset`, 128
`pl_bolts.datamodules.fashion_mnist_datamodule`, 129
`pl_bolts.datamodules.imagenet_datamodule`, 130
`pl_bolts.datamodules.imagenet_dataset`, 133
`pl_bolts.datamodules.mnist_datamodule`, 134
`pl_bolts.datamodules.sklearn_datamodule`, 135
`pl_bolts.datamodules.ssl_imagenet_datamodule`, 139
`pl_bolts.datamodules.stl10_datamodule`, 139
`pl_bolts.datamodules.vocdetection_datamodule`, 141
`pl_bolts.losses`, 180
`pl_bolts.losses.self_supervised_learning`, 180
`pl_bolts.metrics`, 142
`pl_bolts.metrics.aggregation`, 142
`pl_bolts.models`, 142
`pl_bolts.models.autoencoders`, 142
`pl_bolts.models.autoencoders.basic_ae`, 143
`pl_bolts.models.autoencoders.basic_ae.basic_ae_module`, 143
`pl_bolts.models.autoencoders.basic_vae`, 144
`pl_bolts.models.autoencoders.basic_vae.basic_vae_module`, 145
`pl_bolts.models.autoencoders.components`, 146
`pl_bolts.models.detection`, 148
`pl_bolts.models.detection.faster_rcnn`, 148
`pl_bolts.models.gans`, 149
`pl_bolts.models.gans.basic`, 149
`pl_bolts.models.gans.basic.basic_gan_module`, 149
`pl_bolts.models.gans.basic.components`, 150
`pl_bolts.models.mnist_module`, 180
`pl_bolts.models.regression`, 150
`pl_bolts.models.regression.linear_regression`, 150
`pl_bolts.models.regression.logistic_regression`, 151
`pl_bolts.models.self_supervised`, 152
`pl_bolts.models.self_supervised.amdim`, 152
`pl_bolts.models.self_supervised.amdim.amdim_module`, 153
`pl_bolts.models.self_supervised.amdim.datasets`, 154
`pl_bolts.models.self_supervised.amdim.networks`, 155
`pl_bolts.models.self_supervised.amdim.ssl_datasets`, 156
`pl_bolts.models.self_supervised.amdim.transforms`, 156

157
pl_bolts.models.self_supervised.byol, 159
pl_bolts.models.self_supervised.byol.byol_module, 159
pl_bolts.models.self_supervised.byol.models, 161
pl_bolts.models.self_supervised.cpc, 161
pl_bolts.models.self_supervised.cpc.cpc_finetuner, 161
pl_bolts.models.self_supervised.cpc.cpc_module, 161
pl_bolts.models.self_supervised.cpc.networks, 162
pl_bolts.models.self_supervised.cpc.transforms, 163
pl_bolts.models.self_supervised.evaluator, 172
pl_bolts.models.self_supervised.moco, 167
pl_bolts.models.self_supervised.moco.callbacks, 167
pl_bolts.models.self_supervised.moco.moco2_module, 167
pl_bolts.models.self_supervised.moco.transforms, 169
pl_bolts.models.self_supervised.resnets, 172
pl_bolts.models.self_supervised.simclr, 170
pl_bolts.models.self_supervised.simclr.simclr_finetuner, 170
pl_bolts.models.self_supervised.simclr.simclr_module, 170
pl_bolts.models.self_supervised.simclr.simclr_transforms, 171
pl_bolts.models.self_supervised.ssl_finetuner, 174
pl_bolts.models.vision, 176
pl_bolts.models.vision.image_gpt, 176
pl_bolts.models.vision.image_gpt.gpt2, 176
pl_bolts.models.vision.image_gpt.igpt_module, 177
pl_bolts.models.vision.pixel_cnn, 179
pl_bolts.optimizers, 182
pl_bolts.optimizers.lars_scheduling, 182
pl_bolts.optimizers.lr_scheduler, 183
pl_bolts.transforms, 184
pl_bolts.transforms.dataset_normalizations, 185
pl_bolts.transforms.self_supervised, 184
pl_bolts.transforms.self_supervised.ssl_transforms, 184
pl_bolts.utils, 185
pl_bolts.utils.pretrained_weights, 185
pl_bolts.utils.self_supervised, 185
pl_bolts.utils.semi_supervised, 185
pl_bolts.utils.shaping, 186

INDEX

Symbols

Symbols

<code>__CPCV2__compute_final_nb_c()</code>	<code>__call__()</code> (<code>pl_bolts.models.self_supervised.cpc.transforms.CPCTrainTransform</code>), 166
<code>(pl_bolts.models.self_supervised.cpc.cpc_module.CPCV2</code>	<code>method)</code> , 166
<code>method)</code> , 162	<code>__call__()</code> (<code>pl_bolts.models.self_supervised.moco.transforms.GaussianMixture</code>), 169
<code>__CPCV2__recover_z_shape()</code>	<code>method)</code> , 169
<code>(pl_bolts.models.self_supervised.cpc.cpc_module.CPCV2</code>	<code>__call__()</code> (<code>pl_bolts.models.self_supervised.moco.transforms.Moco2EvalTransform</code>), 169
<code>method)</code> , 162	<code>method)</code> , 169
<code>__ConfusedLogitCallback__draw_sample()</code>	<code>__call__()</code> (<code>pl_bolts.models.self_supervised.moco.transforms.Moco2EvalTransform</code>), 169
<code>(pl_bolts.callbacks.vision.confused_logit.ConfusedLogitCallback</code>	<code>method)</code> , 169
<code>static method)</code> , 115	<code>__call__()</code> (<code>pl_bolts.models.self_supervised.moco.transforms.Moco2EvalTransform</code>), 169
<code>__FeatureMapContrastiveTask__cache_dimension_masks()</code>	<code>method)</code> , 169
<code>(pl_bolts.losses.self_supervised_learning.FeatureMapContrastiveTask</code>	<code>__call__()</code> (<code>pl_bolts.models.self_supervised.moco.transforms.Moco2TrainTransform</code>), 170
<code>method)</code> , 181	<code>method)</code> , 170
<code>__FeatureMapContrastiveTask__compare_maps()</code>	<code>__call__()</code> (<code>pl_bolts.models.self_supervised.moco.transforms.Moco2TrainTransform</code>), 170
<code>(pl_bolts.losses.self_supervised_learning.FeatureMapContrastiveTask</code>	<code>method)</code> , 170
<code>method)</code> , 181	<code>__call__()</code> (<code>pl_bolts.models.self_supervised.moco.transforms.Moco2TrainTransform</code>), 170
<code>__LNBottleneck__init_layer_norms()</code>	<code>method)</code> , 170
<code>(pl_bolts.models.self_supervised.cpc.networks.LNBottleneck</code>	<code>__call__()</code> (<code>pl_bolts.models.self_supervised.simclr.simclr_transforms.CIFAR10</code>), 171
<code>method)</code> , 163	<code>method)</code> , 171
<code>__call__()</code> (<code>pl_bolts.datamodules.vocdetetection_datamodule.Compose</code>	<code>__call__()</code> (<code>pl_bolts.models.self_supervised.simclr.simclr_transforms.SimCLR</code>), 171
<code>method)</code> , 141	<code>method)</code> , 171
<code>__call__()</code> (<code>pl_bolts.models.self_supervised.amdim.transforms.AMDIMTrainTransform</code>	<code>__call__()</code> (<code>pl_bolts.models.self_supervised.simclr.simclr_transforms.SimCLR</code>), 172
<code>method)</code> , 157	<code>method)</code> , 172
<code>__call__()</code> (<code>pl_bolts.models.self_supervised.amdim.transforms.AMDIMEvalTransform</code>	<code>__call__()</code> (<code>pl_bolts.models.self_supervised.ssl_transforms.Patchify</code>), 178
<code>method)</code> , 157	<code>method)</code> , 184
<code>__call__()</code> (<code>pl_bolts.models.self_supervised.amdim.transforms.AMDIMTrainTransform</code>	<code>__call__()</code> (<code>pl_bolts.models.self_supervised.ssl_transforms.Random</code>), 184
<code>method)</code> , 158	<code>method)</code> , 184
<code>__call__()</code> (<code>pl_bolts.models.self_supervised.amdim.transforms.AMDIMTrainTransform</code>	<code>batch_shuffle_ddp()</code>
<code>method)</code> , 158	<code>(pl_bolts.models.self_supervised.moco.moco2_module.MocoV2</code>
<code>__call__()</code> (<code>pl_bolts.models.self_supervised.amdim.transforms.AMDIMTrainTransform</code>	<code>method)</code> , 168
<code>method)</code> , 158	<code>__batch_unshuffle_ddp()</code>
<code>__call__()</code> (<code>pl_bolts.models.self_supervised.amdim.transforms.AMDIMTrainTransform</code>	<code>(pl_bolts.models.self_supervised.moco.moco2_module.MocoV2</code>
<code>method)</code> , 159	<code>method)</code> , 168
<code>__call__()</code> (<code>pl_bolts.models.self_supervised.cpc.transforms.CPCEvalTransform</code>	<code>calculate_md5()</code>
<code>method)</code> , 164	<code>CIFAR10 (in module</code>
<code>__call__()</code> (<code>pl_bolts.models.self_supervised.cpc.transforms.CPCEvalTransform</code>	<code>pl_bolts.datamodules.imagenet_dataset)</code> , 133
<code>method)</code> , 164	<code>__check_exists()</code> (<code>pl_bolts.datamodules.cifar10_dataset.CIFAR10</code>
<code>__call__()</code> (<code>pl_bolts.models.self_supervised.cpc.transforms.CPCEvalTransform</code>	<code>class method)</code> , 135
<code>method)</code> , 165	<code>__check_integrity()</code> (in module
<code>__call__()</code> (<code>pl_bolts.models.self_supervised.cpc.transforms.CPCTrainTransform</code>	<code>pl_bolts.datamodules.imagenet_dataset)</code> , 133
<code>method)</code> , 165	<code>__check_md5()</code> (in module
<code>__call__()</code> (<code>pl_bolts.models.self_supervised.cpc.transforms.CPCTrainTransform</code>	<code>CIFAR10 (in module</code>

`pl_bolts.datamodules.imagenet_dataset),` 133
`_collate_fn()` (in module `pl_bolts.datamodules.imagenet_dataset),`
`pl_bolts.datamodules.vocdetection_datamodule),` 133
`_config_modules()` (in module `pl_bolts.datamodules.imagenet_dataset),`
`(pl_bolts.models.self_supervised.amdim.networks.AMDIMEncoder`
`method), 155` `_is_zip()` (in module
`pl_bolts.datamodules.imagenet_dataset),`
`_default_transforms()` (in module `pl_bolts.datamodules.binary_mnist_datamodule.BinaryMNISTDataModule`
`method), 121` `_make_layer()` (`pl_bolts.models.autoencoders.components.ResNetDecoder`
`method), 147`
`_default_transforms()` (in module `pl_bolts.datamodules.fashion_mnist_datamodule.FashionMNISTDataModule`
`method), 130` `_make_layer()` (`pl_bolts.models.autoencoders.components.ResNetEncoder`
`method), 147`
`_default_transforms()` (in module `pl_bolts.datamodules.mnist_datamodule.MNISTDataModule`
`method), 135` `_make_layer()` (`pl_bolts.models.self_supervised.cpc.networks.CPCResNet`
`method), 172`
`_default_transforms()` (in module `pl_bolts.datamodules.ssl_imagenet_datamodule.SSLImageNetDataModule`
`method), 139` `_make_layer()` (`pl_bolts.models.self_supervised.resnets.ResNet`
`method), 168`
`_default_transforms()` (in module `pl_bolts.datamodules.vocdetection_datamodule.VOCDetectionDataModule`
`method), 141` `_prepare_subset()`
`_dequeue_and_enqueue()` (in module `pl_bolts.models.self_supervised.moco.moco2_module.MocoV2`
`method), 168` `static method), 119`
`_download_from_url()` (in module `pl_bolts.datamodules.base_dataset.LightDataset`
`method), 119` `_prepare_voc_instance()` (in module
`pl_bolts.datamodules.vocdetection_datamodule),`
`_evaluate_iou()` (in module `pl_bolts.datamodules.dummy_dataset.DummyDetectionDataset`
`method), 128`
`_extract_archive_save_torch()` (in module `pl_bolts.models.autoencoders.basic_vae.basic_vae_module`
`method), 145`
`_forward_acts()` (`pl_bolts.models.self_supervised.amdim.networks.AMDIMEncoder`
`method), 155` `_sample_src_ftr()`
`_get_closed_form_lr()` (in module `pl_bolts.optimizers.lr_scheduler.LinearWarmupCosineAnnealingLR`
`method), 184` `_shape_input()` (in module
`pl_bolts.models.vision.image_gpt.igpt_module),`
`_init_datasets()` (`pl_bolts.datamodules.sklearn_datamodule.SklearnDataModule`
`method), 136` `_verify_archive()` (in module
`pl_bolts.datamodules.imagenet_dataset),`
`_init_embeddings()` (in module `pl_bolts.models.vision.image_gpt.gpt2.GPT2`
`method), 176` `_verify_splits()` (`pl_bolts.datamodules.imagenet_datamodule.ImageNetDataModule`
`method), 131`
`_init_layers()` (`pl_bolts.models.vision.image_gpt.gpt2.GPT2`
`method), 176` `_verify_splits()` (`pl_bolts.datamodules.ssl_imagenet_datamodule.SSLImageNetDataModule`
`method), 139`
`_init_sos_token()` (in module `pl_bolts.models.vision.image_gpt.gpt2.GPT2`
`method), 176`
`_is_gzip()` (in module `pl_bolts.datamodules.imagenet_dataset),` 133
`_is_tar()` (in module `pl_bolts.datamodules.imagenet_dataset),` 133

A

`accuracy()` (in module `pl_bolts.metrics.aggregation),` 142
`add_model_specific_args()` (`pl_bolts.models.autoencoders.basic_ae.basic_ae_module.AE`

```

static method), 144
add_model_specific_args()
(pl_bolts.models.autoencoders.basic_vae.basic_vae_module.BasicVAEPretraining
static method), 145
add_model_specific_args()
(pl_bolts.models.detection.faster_rcnn.FasterRCNNAMDIMPRetraining
static method), 148
add_model_specific_args()
(pl_bolts.models.gans.basic.basic_gan_module.GANAMDIMTrainTransformsCIFAR10
static method), 149
add_model_specific_args()
(pl_bolts.models.mnist_module.LitMNIST
static method), 180
add_model_specific_args()
(pl_bolts.models.mnist_module.LitMNIST
static method), 180
add_model_specific_args()
(pl_bolts.models.regression.linear_regression.LinearRegressionAMDIMPRetraining
static method), 151
add_model_specific_args()
(pl_bolts.models.regression.logistic_regression.LogisticRegressionAMDIMPRetraining
static method), 152
add_model_specific_args()
(pl_bolts.models.self_supervised.amdim.amdim_module.AMDIM
static method), 154
add_model_specific_args()
balance_classes() (in module
(pl_bolts.models.self_supervised.byol.byol_module.BYOL
static method), 160
add_model_specific_args()
BASE_URL(pl_bolts.datamodules.cifar10_dataset.CIFAR10
attribute), 125
(pl_bolts.models.self_supervised.cpc.cpc_module.GPCV2
static method), 162
add_model_specific_args()
(pl_bolts.models.self_supervised.moco.moco2_module.MocoV2
static method), 169
add_model_specific_args()
(pl_bolts.models.self_supervised.moco.moco2_module.MocoV2
static method), 169
add_model_specific_args()
(pl_bolts.models.self_supervised.simclr.simclr_module.SimCLR
static method), 171
add_model_specific_args()
BYOL(class in pl_bolts.models.self_supervised.byol.byol_module),
(pl_bolts.models.vision.image_gpt.igpt_module.ImageGPT
static method), 178
AE(class in pl_bolts.models.autoencoders.basic_ae.basic_ae_module)
143
AMDIM(class in pl_bolts.models.self_supervised.amdim.amdim_module),
153
AMDIMEncoder(class in
cache_folder_name
(pl_bolts.datamodules.base_dataset.LightDataset
attribute), 119
155
AMDIMEvalTransformsCIFAR10(class in
cache_folder_name
(pl_bolts.datamodules.cifar10_dataset.CIFAR10
attribute), 125
157
AMDIMEvalTransformsImageNet128(class in
cached_folder_path()
(pl_bolts.datamodules.base_dataset.LightDataset
property), 119
157
AMDIMEvalTransformsSTL10(class in
CIFAR10(class in pl_bolts.datamodules.cifar10_dataset),
124
pl_bolts.models.self_supervised.amdim.transforms),
157
AmdimNCELoss(class in
cifar10() (pl_bolts.models.self_supervised.amdim.datasets.AMDIMPat
static method), 154

```

`cifar10()` (`pl_bolts.models.self_supervised.amdim.datasets.AMDIMDataset` in `pl_bolts.datamodules.vocdetection_datamodule`), static method), 154

`cifar10_normalization()` (`in module compute_loss_h()` (`pl_bolts.losses.self_supervised_learning.CPCTask` in `pl_bolts.transforms.dataset_normalizations`), method), 180

`cifar10_tiny()` (`pl_bolts.models.self_supervised.amdim.datasets.AMDIMDataset` in `pl_bolts.datamodules.vocdetection_datamodule`), static method), 154

`CIFAR10DataModule` (class in `ConcatDataset` (class in `pl_bolts.datamodules.concat_dataset`), 122

`CIFAR10Mixed` (class in `(pl_bolts.models.autoencoders.basic_ae.basic_ae_module.AE` in `pl_bolts.models.self_supervised.amdim.ssl_datasets`), method), 144

`CityscapesDataModule` (class in `(pl_bolts.models.autoencoders.basic_vae.basic_vae_module.VAE` in `pl_bolts.datamodules.cityscapes_datamodule`), method), 145

`cli_main()` (`in module (pl_bolts.models.detection.faster_rcnn.FasterRCNN` in `pl_bolts.models.autoencoders.basic_ae.basic_ae_module`), method), 148

`cli_main()` (`in module (pl_bolts.models.gans.basic.basic_gan_module.GAN` in `pl_bolts.models.autoencoders.basic_vae.basic_vae_module`), method), 149

`cli_main()` (`in module (pl_bolts.models.mnist_module.LitMNIST` in `pl_bolts.models.gans.basic.basic_gan_module`), method), 180

`cli_main()` (`in module (pl_bolts.models.regression.linear_regression.LinearRegression` in `pl_bolts.models.mnist_module`), 180

`cli_main()` (`in module (pl_bolts.models.regression.logistic_regression.LogisticRegression` in `pl_bolts.models.regression.linear_regression`), method), 152

`cli_main()` (`in module (pl_bolts.models.self_supervised.amdim.amdim_module.AMDIM` in `pl_bolts.models.regression.logistic_regression`), method), 154

`cli_main()` (`in module (pl_bolts.models.self_supervised.byol.byol_module.BYOL` in `pl_bolts.models.self_supervised.amdim.amdim_module`), method), 160

`cli_main()` (`in module (pl_bolts.models.self_supervised.cpc.cpc_module.CPCV2` in `pl_bolts.models.self_supervised.byol.byol_module`), method), 162

`cli_main()` (`in module (pl_bolts.models.self_supervised.moco.moco2_module.MocoV2` in `pl_bolts.models.self_supervised.cpc.cpc_finetuner`), method), 169

`cli_main()` (`in module (pl_bolts.models.self_supervised.simclr.simclr_module.SimCLR` in `pl_bolts.models.self_supervised.moco.moco2_module`), method), 171

`cli_main()` (`in module (pl_bolts.models.self_supervised.ssl_finetuner.SSLSFineTuner` in `pl_bolts.models.self_supervised.simclr.simclr_finetuner`), method), 175

`cli_main()` (`in module (pl_bolts.models.vision.image_gpt.igpt_module.ImageGPT` in `pl_bolts.models.self_supervised.simclr.simclr_module`), method), 178

`cli_main()` (`in module ConfusedLogitCallback` (class in `pl_bolts.callbacks.vision.confused_logit`), 114

conv1x1 () (in module CPCV2 (class in *pl_bolts.models.self_supervised.cpc.cpc_module*),
pl_bolts.models.autoencoders.components), 161
 147

conv1x1 () (in module **D**
pl_bolts.models.self_supervised.cpc.networks), data (*pl_bolts.datamodules.base_dataset.LightDataset*
 163 attribute), 119

Conv3x3 (class in *pl_bolts.models.self_supervised.amdim.networks*), (*pl_bolts.datamodules.cifar10_dataset.CIFAR10*
 155 attribute), 125

conv3x3 () (in module data (*pl_bolts.datamodules.cifar10_dataset.TrialCIFAR10*
pl_bolts.models.autoencoders.components), attribute), 126
 147

conv3x3 () (in module DATASET_NAME (*pl_bolts.datamodules.base_dataset.LightDataset*
pl_bolts.models.self_supervised.cpc.networks), attribute), 119
 163 DATASET_NAME (*pl_bolts.datamodules.cifar10_dataset.CIFAR10*
 attribute), 125

conv_block () (*pl_bolts.models.vision.pixel_cnn.PixelCNN*EncoderBlock (class in
 method), 179 *pl_bolts.models.autoencoders.components*),

ConvResBlock (class in 146
pl_bolts.models.self_supervised.amdim.networks)DecoderBottleneck (class in
 155 *pl_bolts.models.autoencoders.components*),

ConvResNxN (class in 146
pl_bolts.models.self_supervised.amdim.networks)default_transforms ()
 155 (*pl_bolts.datamodules.cifar10_datamodule.CIFAR10DataModule*
 method), 123

cosine_similarity ()
 (*pl_bolts.models.self_supervised.byol.byol_module.BYOL*default_transforms ()
 method), 160 (*pl_bolts.datamodules.cityscapes_datamodule.CityscapesDataModule*
 method), 127

cpc_resnet101 () (in module default_transforms ()
pl_bolts.models.self_supervised.cpc.networks), (*pl_bolts.datamodules.stl10_datamodule.STL10DataModule*
 163 method), 140

cpc_resnet50 () (in module DensenetEncoder (class in
pl_bolts.models.self_supervised.cpc.networks), *pl_bolts.models.self_supervised.simclr.simclr_module*),
 163 170

CPCEvalTransformsCIFAR10 (class in deterministic_shuffle ()
pl_bolts.models.self_supervised.cpc.transforms), (*pl_bolts.models.self_supervised.amdim.ssl_datasets.SSLDataset*
 163 class method), 156

CPCEvalTransformsImageNet128 (class in dicts_to_table () (in module
pl_bolts.models.self_supervised.cpc.transforms), *pl_bolts.callbacks.printing*), 116
 164

CPCEvalTransformsSTL10 (class in dir_path (*pl_bolts.datamodules.base_dataset.LightDataset*
pl_bolts.models.self_supervised.cpc.transforms), attribute), 119
 164 dir_path (*pl_bolts.datamodules.cifar10_dataset.CIFAR10*
 attribute), 125

CPCResNet (class in dir_path (*pl_bolts.datamodules.cifar10_dataset.TrialCIFAR10*
pl_bolts.models.self_supervised.cpc.networks), attribute), 126
 162

CPCTask (class in *pl_bolts.losses.self_supervised_learning*)Discriminator (class in
 180 *pl_bolts.models.gans.basic.components*),

CPCTrainTransformsCIFAR10 (class in 150
pl_bolts.models.self_supervised.cpc.transforms), discriminator_loss ()
 165 (*pl_bolts.models.gans.basic.basic_gan_module.GAN*
 method), 149

CPCTrainTransformsImageNet128 (class in discriminator_step ()
pl_bolts.models.self_supervised.cpc.transforms), (*pl_bolts.models.gans.basic.basic_gan_module.GAN*
 165 method), 150

CPCTrainTransformsSTL10 (class in download () (*pl_bolts.datamodules.cifar10_dataset.CIFAR10*
pl_bolts.models.self_supervised.cpc.transforms), method), 125
 166

DummyDataset (class in `pl_bolts.datamodules.dummy_dataset`), 128

DummyDetectionDataset (class in `pl_bolts.datamodules.dummy_dataset`), 128

E

EncoderBlock (class in `pl_bolts.models.autoencoders.components`), 146

EncoderBottleneck (class in `pl_bolts.models.autoencoders.components`), 146

`exclude_from_wt_decay()` (`pl_bolts.models.self_supervised.simclr.simclr_module.SimCLRModule` method), 171

`expansion` (`pl_bolts.models.autoencoders.components.DecoderBlock` attribute), 146

`expansion` (`pl_bolts.models.autoencoders.components.DecoderBottleneck` attribute), 146

`expansion` (`pl_bolts.models.autoencoders.components.EncoderBlock` attribute), 146

`expansion` (`pl_bolts.models.autoencoders.components.EncoderBottleneck` attribute), 147

`extra_args` (`pl_bolts.datamodules.cifar10_datamodule.CIFAR10DataModule` attribute), 123

`extra_args` (`pl_bolts.datamodules.cityscapes_datamodule.CityscapesDataModule` attribute), 128

`extract_archive()` (in module `pl_bolts.datamodules.imagenet_dataset`), 133

F

FakeRKHSConvNet (class in `pl_bolts.models.self_supervised.amdim.networks`), 155

FashionMNISTDataModule (class in `pl_bolts.datamodules.fashion_mnist_datamodule`), 129

FasterRCNN (class in `pl_bolts.models.detection.faster_rcnn`), 148

`feat_size_w_mask()` (`pl_bolts.losses.self_supervised_learning.FeatureMapContrastiveTask` method), 181

FeatureMapContrastiveTask (class in `pl_bolts.losses.self_supervised_learning`), 181

FILE_NAME (`pl_bolts.datamodules.cifar10_dataset.CIFAR10Dataset` attribute), 125

Flatten (class in `pl_bolts.models.self_supervised.evaluator`), 172

`flatten()` (`pl_bolts.models.self_supervised.cpc.networks.CPCResNet` method), 162

`forward()` (`pl_bolts.losses.self_supervised_learning.AmdimNCELoss` method), 180

`forward()` (`pl_bolts.losses.self_supervised_learning.CPCTask` method), 181

`forward()` (`pl_bolts.losses.self_supervised_learning.FeatureMapContrastiveTask` method), 181

`forward()` (`pl_bolts.models.autoencoders.basic_ae.basic_ae_module.AE` method), 144

`forward()` (`pl_bolts.models.autoencoders.basic_vae.basic_vae_module.BasicVAE` method), 145

`forward()` (`pl_bolts.models.autoencoders.components.DecoderBlock` method), 146

`forward()` (`pl_bolts.models.autoencoders.components.DecoderBottleneck` method), 146

`forward()` (`pl_bolts.models.autoencoders.components.EncoderBlock` method), 146

`forward()` (`pl_bolts.models.autoencoders.components.EncoderBottleneck` method), 147

`forward()` (`pl_bolts.models.autoencoders.components.Interpolate` method), 147

`forward()` (`pl_bolts.models.autoencoders.components.ResNetDecoder` method), 147

`forward()` (`pl_bolts.models.autoencoders.components.ResNetEncoder` method), 147

`forward()` (`pl_bolts.models.detection.faster_rcnn.FasterRCNN` method), 148

`forward()` (`pl_bolts.models.gans.basic.basic_gan_module.GAN` method), 150

`forward()` (`pl_bolts.models.gans.basic.components.Discriminator` method), 150

`forward()` (`pl_bolts.models.gans.basic.components.Generator` method), 150

`forward()` (`pl_bolts.models.mnist_module.LitMNIST` method), 180

`forward()` (`pl_bolts.models.regression.linear_regression.LinearRegression` method), 151

`forward()` (`pl_bolts.models.regression.logistic_regression.LogisticRegression` method), 152

`forward()` (`pl_bolts.models.self_supervised.amdim.amdim_module.AMDIM` method), 154

`forward()` (`pl_bolts.models.self_supervised.amdim.networks.AMDIMEnsemble` method), 155

`forward()` (`pl_bolts.models.self_supervised.amdim.networks.Conv3x3` method), 155

`forward()` (`pl_bolts.models.self_supervised.amdim.networks.ConvResBlock` method), 155

`forward()` (`pl_bolts.models.self_supervised.amdim.networks.ConvResNext` method), 155

`forward()` (`pl_bolts.models.self_supervised.amdim.networks.FakeRKHS` method), 155

`forward()` (`pl_bolts.models.self_supervised.amdim.networks.MaybeBatchNorm` method), 156

`forward()` (`pl_bolts.models.self_supervised.amdim.networks.NopNet` method), 156

`forward()` (`pl_bolts.models.self_supervised.byol.byol_module.BYOL` method), 160

`forward()` (`pl_bolts.models.self_supervised.byol.models.MLP` class method), 133
`method`), 161 `generate_train_val_split()`
`forward()` (`pl_bolts.models.self_supervised.byol.models.SiameseA` `pl_bolts.models.self_supervised.amdim.ssl_datasets.SSLDataset`
`method`), 161 class method), 156
`forward()` (`pl_bolts.models.self_supervised.cpc.cpc_module.CPCV2r` (class in
`method`), 162 `pl_bolts.models.gans.basic.components`),
`forward()` (`pl_bolts.models.self_supervised.cpc.networks.CPCResNet` `generator_loss()` (`pl_bolts.models.gans.basic.basic_gan_module.GAN`
`method`), 162 `method`), 150
`forward()` (`pl_bolts.models.self_supervised.cpc.networks.LNBottleneck` `generator_step()` (`pl_bolts.models.gans.basic.basic_gan_module.GAN`
`method`), 163 `method`), 150
`forward()` (`pl_bolts.models.self_supervised.evaluator.Flatten` `method`), 150
`method`), 172 `get_dataset()` (`pl_bolts.models.self_supervised.amdim.datasets.AMDIMPa`
`forward()` (`pl_bolts.models.self_supervised.evaluator.SSLEvaluator` static method), 154
`method`), 172 `get_lr()` (`pl_bolts.optimizers.lr_scheduler.LinearWarmupCosineAnneali`
`forward()` (`pl_bolts.models.self_supervised.moco.moco2_module.MathV2` `method`), 184
`method`), 169 `get_representations()`
`forward()` (`pl_bolts.models.self_supervised.resnets.ResNet` (`pl_bolts.callbacks.self_supervised.SSLOnlineEvaluator`
`method`), 173 `method`), 118
`forward()` (`pl_bolts.models.self_supervised.simclr.simclr_module.Densimclr` `pl_bolts.models.vision.image_gpt.gpt2`),
`method`), 170 `method`), 176
`forward()` (`pl_bolts.models.self_supervised.simclr.simclr_module.Projection`
`method`), 170 `method`), 171
`forward()` (`pl_bolts.models.self_supervised.simclr.simclr_module.SimCLR` (class in `pl_bolts.utils.semi_supervised`), 185
`method`), 171 `ImageGPT` (class in `pl_bolts.models.vision.image_gpt.igpt_module`),
`forward()` (`pl_bolts.models.vision.image_gpt.gpt2.Block` `imagenet()` (`pl_bolts.models.self_supervised.amdim.datasets.AMDIMPa`
`method`), 176 `method`), 154
`forward()` (`pl_bolts.models.vision.image_gpt.gpt2.GPT2` static method), 154
`method`), 176 `imagenet()` (`pl_bolts.models.self_supervised.amdim.datasets.AMDIMPa`
`forward()` (`pl_bolts.models.vision.image_gpt.igpt_module.ImageGPT` static method), 154
`method`), 178 `imagenet_normalization()` (in module
`forward()` (`pl_bolts.models.vision.pixel_cnn.PixelCNN` `pl_bolts.transforms.dataset_normalizations`),
`method`), 179 `method`), 185
`forward()` (`pl_bolts.utils.semi_supervised.Identity` `ImagenetDataModule` (class in
`method`), 185 `pl_bolts.datamodules.imagenet_datamodule`),
`from_pretrained()` `method`), 130
`(pl_bolts.models.autoencoders.basic_ae.basic_ae_module.AE` `Discriminator()`
`method`), 144 `(pl_bolts.models.gans.basic.basic_gan_module.GAN`
`from_pretrained()` `method`), 150
`(pl_bolts.models.autoencoders.basic_vae.basic_vae_module.VAE` `encoder()` (`pl_bolts.models.self_supervised.amdim.amdim_modu`
`method`), 145 `method`), 154
`init_encoder()` (`pl_bolts.models.self_supervised.cpc.cpc_module.CPC`
`method`), 162 `method`), 171
`GAN` (class in `pl_bolts.models.gans.basic.basic_gan_module`), `init_encoder()` (`pl_bolts.models.self_supervised.simclr.simclr_modu`
149 `method`), 171
`GaussianBlur` (class in `init_encoders()` (`pl_bolts.models.self_supervised.moco.moco2_modu`
`pl_bolts.models.self_supervised.moco.transforms`), `method`), 169
169 `init_generator()` (`pl_bolts.models.gans.basic.basic_gan_module.GAN`
`GaussianBlur` (class in `method`), 150
`pl_bolts.models.self_supervised.simclr.simclr_transforms` `init_weights()` (`pl_bolts.models.self_supervised.amdim.networks.AM`
171 `method`), 155
`generate_half_labeled_batches()` (in mod- `init_weights()` (`pl_bolts.models.self_supervised.amdim.networks.Con`
ule `pl_bolts.utils.semi_supervised`), 185 `method`), 155
`generate_meta_bins()` `init_weights()` (`pl_bolts.models.self_supervised.amdim.networks.Con`
(`pl_bolts.datamodules.imagenet_dataset.UnlabeledImagenet` `method`), 155

`init_weights()` (*pl_bolts.models.self_supervised.amdim.networks.fake_kl_hsc_nets* (class in *pl_bolts.models.self_supervised.amdim.networks*), 155)
`Interpolate` (class in *pl_bolts.models.autoencoders.components*), 147
`interpolate_latent_space()` (*pl_bolts.callbacks.variational.LatentDimInterpolator* (class in *pl_bolts.callbacks.variational*), 118)
L
`labels` (*pl_bolts.datamodules.cifar10_dataset.CIFAR10* (class in *pl_bolts.datamodules.cifar10_dataset*), 125)
`LARSWrapper` (class in *pl_bolts.optimizers.lars_scheduling*), 182
`LatentDimInterpolator` (class in *pl_bolts.callbacks.variational*), 118
`LightDataset` (class in *pl_bolts.datamodules.base_dataset*), 119
`LinearRegression` (class in *pl_bolts.models.regression.linear_regression*), 150
`LinearWarmupCosineAnnealingLR` (class in *pl_bolts.optimizers.lr_scheduler*), 183
`LitMNIST` (class in *pl_bolts.models.mnist_module*), 180
`LN Bottleneck` (class in *pl_bolts.models.self_supervised.cpc.networks*), 162
`load_instance()` (*pl_bolts.datamodules.async_dataloader.AsynchronousLoader* (class in *pl_bolts.datamodules.async_dataloader*), 119)
`load_loop()` (*pl_bolts.datamodules.async_dataloader.AsynchronousLoader* (class in *pl_bolts.datamodules.async_dataloader*), 119)
`load_pretrained()` (in module *pl_bolts.utils.pretrained_weights*), 185
`load_pretrained()` (*pl_bolts.models.self_supervised.cpc.cpc_module.CPCV2* (class in *pl_bolts.models.self_supervised.cpc.cpc_module*), 162)
`LogisticRegression` (class in *pl_bolts.models.regression.logistic_regression*), 151
M
`MaybeBatchNorm2d` (class in *pl_bolts.models.self_supervised.amdim.networks*), 156
`mean()` (in module *pl_bolts.metrics.aggregation*), 142
`MLP` (class in *pl_bolts.models.self_supervised.byol.models*), 161
`MNISTDataModule` (class in *pl_bolts.datamodules.mnist_datamodule*), 134
`Moco2EvalCIFAR10Transforms` (class in *pl_bolts.models.self_supervised.moco.transforms*), 169
`Moco2EvalSTL10Transforms` (class in *pl_bolts.models.self_supervised.moco.transforms*), 169
`Moco2TrainCIFAR10Transforms` (class in *pl_bolts.models.self_supervised.moco.transforms*), 169
`Moco2TrainImagenetTransforms` (class in *pl_bolts.models.self_supervised.moco.transforms*), 170
`Moco2TrainSTL10Transforms` (class in *pl_bolts.models.self_supervised.moco.transforms*), 170
`MocoLRScheduler` (class in *pl_bolts.models.self_supervised.moco.callbacks*), 167
`MocoV2` (class in *pl_bolts.models.self_supervised.moco.moco2_module*), 167
N
`name` (*pl_bolts.datamodules.binary_mnist_datamodule.BinaryMNISTDataModule* (class in *pl_bolts.datamodules.binary_mnist_datamodule*), 121)
`name` (*pl_bolts.datamodules.cifar10_datamodule.CIFAR10DataModule* (class in *pl_bolts.datamodules.cifar10_datamodule*), 123)
`name` (*pl_bolts.datamodules.cityscapes_datamodule.CityscapesDataModule* (class in *pl_bolts.datamodules.cityscapes_datamodule*), 128)
`name` (*pl_bolts.datamodules.fashion_mnist_datamodule.FashionMNISTDataModule* (class in *pl_bolts.datamodules.fashion_mnist_datamodule*), 130)
`name` (*pl_bolts.datamodules.imagenet_datamodule.ImagenetDataModule* (class in *pl_bolts.datamodules.imagenet_datamodule*), 132)
`name` (*pl_bolts.datamodules.mnist_datamodule.MNISTDataModule* (class in *pl_bolts.datamodules.mnist_datamodule*), 135)
`name` (*pl_bolts.datamodules.sklearn_datamodule.SklearnDataModule* (class in *pl_bolts.datamodules.sklearn_datamodule*), 136)
`name` (*pl_bolts.datamodules.ssl_imagenet_datamodule.SSLImagenetDataModule* (class in *pl_bolts.datamodules.ssl_imagenet_datamodule*), 139)
`name` (*pl_bolts.datamodules.stl10_datamodule.STL10DataModule* (class in *pl_bolts.datamodules.stl10_datamodule*), 141)
`name` (*pl_bolts.datamodules.vocdetection_datamodule.VOCDetectionDataModule* (class in *pl_bolts.datamodules.vocdetection_datamodule*), 142)
`NopNet` (class in *pl_bolts.models.self_supervised.amdim.networks*), 156
`normalize` (*pl_bolts.datamodules.base_dataset.LightDataset* (class in *pl_bolts.datamodules.base_dataset*), 119)
`normalize` (*pl_bolts.datamodules.cifar10_dataset.CIFAR10DataModule* (class in *pl_bolts.datamodules.cifar10_dataset*), 125)
`normalize` (*pl_bolts.datamodules.cifar10_dataset.TrialCIFAR10DataModule* (class in *pl_bolts.datamodules.cifar10_dataset*), 126)
`nt_xent_loss()` (in module *pl_bolts.losses.self_supervised_learning*), 182

`num_classes()` (`pl_bolts.datamodules.binary_mnist_datamodule.BinaryMNISTDataModule`
`property`), 121 `parse_map_indexes()`
`num_classes()` (`pl_bolts.datamodules.cifar10_datamodule.CIFAR10DataModule`, `pl_bolts.models.self_supervised_learning.FeatureMapContrastive`
`property`), 123 `static method`), 182
`num_classes()` (`pl_bolts.datamodules.cifar10_datamodule.pl_bolts.datamodules.cifar10_datamodule`)
`property`), 124 (`pl_bolts.datamodules.imagenet_dataset.UnlabeledImagenet`
`property`), 124 `method`), 124
`num_classes()` (`pl_bolts.datamodules.cityscapes_datamodule.CityscapesDataModule`
`property`), 128 `Patchify` (class in `pl_bolts.transforms.self_supervised.ssl_transforms`),
`num_classes()` (`pl_bolts.datamodules.fashion_mnist_datamodule.FashionMNISTDataModule`
`property`), 130 `PixelCNN` (class in `pl_bolts.models.vision.pixel_cnn`),
`num_classes()` (`pl_bolts.datamodules.imagenet_datamodule.ImagenetDataModule`
`property`), 132 `pl_bolts.callbacks` (module), 114
`num_classes()` (`pl_bolts.datamodules.mnist_datamodule.MNISTDataModule`, `pl_bolts.callbacks.printing` (module), 116
`property`), 135 `pl_bolts.callbacks.self_supervised` (mod-
`num_classes()` (`pl_bolts.datamodules.ssl_imagenet_datamodule.SSLImagenetDataModule`
`property`), 139 `pl_bolts.callbacks.variational` (module),
`num_classes()` (`pl_bolts.datamodules.stl10_datamodule.STL10DataModule`
`property`), 141 `pl_bolts.callbacks.vision` (module), 114
`num_classes()` (`pl_bolts.datamodules.vocdetection_datamodule.VOCDetectionDataModule`, `pl_bolts.callbacks.vision.confused_logit`
`property`), 142 (module), 114
`pl_bolts.callbacks.vision.image_generation`
`(module)`, 115
O
`on_epoch_end()` (`pl_bolts.callbacks.printing.PrintTableMetricsCallback`, `pl_bolts.datamodules` (module), 119
`method`), 116 `pl_bolts.datamodules.async_dataloader`
`on_epoch_end()` (`pl_bolts.callbacks.variational.LatentDimInterpolation` (module), 119
`method`), 118 `pl_bolts.datamodules.base_dataset` (mod-
`on_epoch_end()` (`pl_bolts.callbacks.vision.image_generation.TensorBoardGenerativeModelImageSampler`
`method`), 115 `pl_bolts.datamodules.binary_mnist_datamodule`
`on_epoch_start()` (`pl_bolts.models.self_supervised.moco.callbacks.MocoLRScheduler` (module), 128
`method`), 167 `pl_bolts.datamodules.cifar10_datamodule`
`on_pretrain_routine_start()` (module), 122
`(pl_bolts.callbacks.self_supervised.SSLOnlineEvaluator`, `pl_bolts.datamodules.cifar10_dataset`
`method`), 118 (module), 124
`on_train_batch_end()` `pl_bolts.datamodules.cityscapes_datamodule`
`(pl_bolts.callbacks.self_supervised.BYOLMAWeightUpdate` (module), 126
`method`), 117 `pl_bolts.datamodules.concat_dataset`
`on_train_batch_end()` (module), 128
`(pl_bolts.callbacks.self_supervised.SSLOnlineEvaluator`, `pl_bolts.datamodules.dummy_dataset` (mod-
`method`), 118 `ule`), 128
`on_train_batch_end()` `pl_bolts.datamodules.fashion_mnist_datamodule`
`(pl_bolts.callbacks.vision.confused_logit.ConfusedLogitCallback` (module), 129
`method`), 115 `pl_bolts.datamodules.imagenet_datamodule`
`on_train_batch_end()` (module), 130
`(pl_bolts.models.self_supervised.byol.byol_module.BYOL`, `pl_bolts.datamodules.imagenet_dataset`
`method`), 160 (module), 133
`on_train_epoch_start()` `pl_bolts.datamodules.mnist_datamodule`
`(pl_bolts.models.self_supervised.ssl_finetuner.SSLFineTuner` (module), 134
`method`), 175 `pl_bolts.datamodules.sklearn_datamodule`
`(module)`, 135
P `pl_bolts.datamodules.ssl_imagenet_datamodule`
`param_groups()` (`pl_bolts.optimizers.lars_scheduling.LARSWrapper` (module), 139
`property`), 183 `pl_bolts.datamodules.stl10_datamodule`
`parse_devkit_archive()` (in module (module), 139
`pl_bolts.datamodules.imagenet_dataset`), `pl_bolts.datamodules.vocdetection_datamodule`

(module), 141
 pl_bolts.losses (module), 180
 pl_bolts.losses.self_supervised_learning (module), 180
 pl_bolts.metrics (module), 142
 pl_bolts.metrics.aggregation (module), 142
 pl_bolts.models (module), 142
 pl_bolts.models.autoencoders (module), 142
 pl_bolts.models.autoencoders.basic_ae (module), 143
 pl_bolts.models.autoencoders.basic_ae.basic_ae (module), 143
 pl_bolts.models.autoencoders.basic_vae (module), 144
 pl_bolts.models.autoencoders.basic_vae.basic_vae (module), 145
 pl_bolts.models.autoencoders.components (module), 146
 pl_bolts.models.detection (module), 148
 pl_bolts.models.detection.faster_rcnn (module), 148
 pl_bolts.models.gans (module), 149
 pl_bolts.models.gans.basic (module), 149
 pl_bolts.models.gans.basic.basic_gan_module (module), 149
 pl_bolts.models.gans.basic.components (module), 150
 pl_bolts.models.mnist_module (module), 180
 pl_bolts.models.regression (module), 150
 pl_bolts.models.regression.linear_regression (module), 150
 pl_bolts.models.regression.logistic_regression (module), 151
 pl_bolts.models.self_supervised (module), 152
 pl_bolts.models.self_supervised.amdim (module), 152
 pl_bolts.models.self_supervised.amdim.amdim_module (module), 153
 pl_bolts.models.self_supervised.amdim.datasets (module), 154
 pl_bolts.models.self_supervised.amdim.nets (module), 155
 pl_bolts.models.self_supervised.amdim.ssp_datasets (module), 156
 pl_bolts.models.self_supervised.amdim.transforms (module), 157
 pl_bolts.models.self_supervised.byol (module), 159
 pl_bolts.models.self_supervised.byol.byol_module (module), 159
 pl_bolts.models.self_supervised.byol.models (module), 161
 pl_bolts.models.self_supervised.cpc (module), 161
 pl_bolts.models.self_supervised.cpc.cpc_finetuner (module), 161
 pl_bolts.models.self_supervised.cpc.cpc_module (module), 161
 pl_bolts.models.self_supervised.cpc.networks (module), 162
 pl_bolts.models.self_supervised.cpc.transforms (module), 163
 pl_bolts.models.self_supervised.evaluator (module), 167
 pl_bolts.models.self_supervised.moco (module), 167
 pl_bolts.models.self_supervised.moco.callbacks (module), 167
 pl_bolts.models.self_supervised.moco.moco2_module (module), 167
 pl_bolts.models.self_supervised.moco.transforms (module), 169
 pl_bolts.models.self_supervised.resnets (module), 172
 pl_bolts.models.self_supervised.simclr (module), 170
 pl_bolts.models.self_supervised.simclr.simclr_finetuner (module), 170
 pl_bolts.models.self_supervised.simclr.simclr_module (module), 170
 pl_bolts.models.self_supervised.simclr.simclr_transforms (module), 171
 pl_bolts.models.self_supervised.ssl_finetuner (module), 174
 pl_bolts.models.vision (module), 176
 pl_bolts.models.vision.image_gpt (module), 176
 pl_bolts.models.vision.image_gpt.gpt2 (module), 176
 pl_bolts.models.vision.image_gpt.igpt_module (module), 177
 pl_bolts.models.vision.pixel_cnn (module), 179
 pl_bolts.optimizers (module), 182
 pl_bolts.optimizers.lars_scheduling (module), 182
 pl_bolts.optimizers.lr_scheduler (module), 183
 pl_bolts.transforms (module), 184
 pl_bolts.transforms.dataset_normalizations (module), 185
 pl_bolts.transforms.self_supervised (module), 184
 pl_bolts.transforms.self_supervised.ssl_transforms (module), 184
 pl_bolts.utils (module), 185
 pl_bolts.utils.pretrained_weights (module), 185

[ule](#)), 185
[pl_bolts.utils.self_supervised](#) (module), [pl_bolts.models.autoencoders.components](#), 185
[pl_bolts.utils.semi_supervised](#) (module), [resize_conv3x3\(\)](#) (in module [pl_bolts.models.autoencoders.components](#)), 185
[pl_bolts.utils.shaping](#) (module), 186
[precision_at_k\(\)](#) (in module [ResNet](#) (class in [pl_bolts.models.self_supervised.resnets](#)), [pl_bolts.metrics.aggregation](#)), 142
[prepare_data\(\)](#) ([pl_bolts.datamodules.binary_mnist_datamodule.BinaryMNISTDataModule](#) module method), 121
[prepare_data\(\)](#) ([pl_bolts.datamodules.cifar10_datamodule.CIFAR10DataModule](#) module method), 123
[prepare_data\(\)](#) ([pl_bolts.datamodules.cifar10_dataset.CIFAR10Dataset](#) module method), 125
[prepare_data\(\)](#) ([pl_bolts.datamodules.cifar10_dataset.TrainCIFAR10](#) (in module [pl_bolts.models.self_supervised.resnets](#)), module method), 126
[prepare_data\(\)](#) ([pl_bolts.datamodules.cityscapes_datamodule.CityscapesDataModule](#) module method), 127
[prepare_data\(\)](#) ([pl_bolts.datamodules.fashion_mnist_datamodule.FashionMNISTDataModule](#) module method), 130
[prepare_data\(\)](#) ([pl_bolts.datamodules.imagenet_datamodule.ImageNetDataModule](#) (in module [pl_bolts.models.autoencoders.components](#)), module method), 131
[prepare_data\(\)](#) ([pl_bolts.datamodules.mnist_datamodule.MNISTDataModule](#) module method), 135
[prepare_data\(\)](#) ([pl_bolts.datamodules.ssl_imagenet_datamodule.SSLImageNetDataModule](#) (in module [pl_bolts.models.self_supervised.resnets](#)), module method), 139
[prepare_data\(\)](#) ([pl_bolts.datamodules.stl10_datamodule.STL10DataModule](#) (in module [pl_bolts.models.self_supervised.resnets](#)), module method), 140
[prepare_data\(\)](#) ([pl_bolts.datamodules.vocdetection_datamodule.VOCDetectionDataModule](#) module method), 141
[prepare_data\(\)](#) ([pl_bolts.models.mnist_module.LitMNIST](#) [pl_bolts.models.self_supervised.resnets](#), module method), 180
[pretrained_urls](#) ([pl_bolts.models.autoencoders.basic_vae_basic_vae_module.VAE](#) (in module [pl_bolts.models.autoencoders.components](#)), attribute), 144
[pretrained_urls](#) ([pl_bolts.models.autoencoders.basic_vae_basic_vae_module.VAE](#) (in module [pl_bolts.models.autoencoders.components](#)), attribute), 146
[pretrained_weights_available\(\)](#) ([pl_bolts.models.autoencoders.basic_vae_basic_vae_module.VAE](#) static method), 144
[pretrained_weights_available\(\)](#) ([pl_bolts.models.autoencoders.basic_vae_basic_vae_module.VAE](#) static method), 145
[PrintTableMetricsCallback](#) (class in [pl_bolts.models.autoencoders.components](#), [pl_bolts.callbacks.printing](#)), 116
[Projection](#) (class in [resnext101_32x8d\(\)](#) (in module [pl_bolts.models.self_supervised.simclr.simclr_module](#)), [pl_bolts.models.self_supervised.resnets](#), 170
[resnext50_32x4d\(\)](#) (in module [pl_bolts.models.self_supervised.resnets](#), 173
[RandomTranslateWithReflect](#) (class in [pl_bolts.transforms.self_supervised.ssl_transforms](#)), 184
[relabel](#) ([pl_bolts.datamodules.cifar10_dataset.CIFAR10](#) attribute), 125
[run_cli\(\)](#) (in module [pl_bolts.models.detection.faster_rcnn](#)), 149

S

[sample\(\)](#) (`pl_bolts.models.autoencoders.basic_vae.basic_vae_module.VAE` static method), 146
[select_nb_imgs_per_class\(\)](#) (`pl_bolts.models.self_supervised.amdim.ssl_datasets.SSLDatasetMixin` class method), 156
[setup\(\)](#) (`pl_bolts.models.self_supervised.simclr.simclr_module.SimCLR` method), 171
[shared_step\(\)](#) (`pl_bolts.models.self_supervised.byol.byol_module.BYOE` method), 161
[shared_step\(\)](#) (`pl_bolts.models.self_supervised.cpc.cpc_module.CPCV2` method), 162
[shared_step\(\)](#) (`pl_bolts.models.self_supervised.simclr.simclr_module.SimCLR` method), 171
[shared_step\(\)](#) (`pl_bolts.models.self_supervised.ssl_finetuner.SSLFinetuner` method), 175
[SiameseArm](#) (class in `pl_bolts.models.self_supervised.byol.models`), 161
[SimCLR](#) (class in `pl_bolts.models.self_supervised.simclr.simclr_module`), 170
[SimCLREvalDataTransform](#) (class in `pl_bolts.models.self_supervised.simclr.simclr_transforms`), 171
[SimCLRTrainDataTransform](#) (class in `pl_bolts.models.self_supervised.simclr.simclr_transforms`), 172
[SklearnDataModule](#) (class in `pl_bolts.datamodules.sklearn_datamodule`), 135
[SklearnDataset](#) (class in `pl_bolts.datamodules.sklearn_datamodule`), 136
[SSLDatasetMixin](#) (class in `pl_bolts.models.self_supervised.amdim.ssl_datasets`), 156
[SSLEvaluator](#) (class in `pl_bolts.models.self_supervised.evaluator`), 172
[SSLFinetuner](#) (class in `pl_bolts.models.self_supervised.ssl_finetuner`), 174
[SSLImagenetDataModule](#) (class in `pl_bolts.datamodules.ssl_imagenet_datamodule`), 139
[SSLOnlineEvaluator](#) (class in `pl_bolts.callbacks.self_supervised`), 117
[state\(\)](#) (`pl_bolts.optimizers.lars_scheduling.LARSWrapper` property), 183
[step\(\)](#) (`pl_bolts.models.autoencoders.basic_ae.basic_ae_module.AE` method), 144
[step\(\)](#) (`pl_bolts.models.autoencoders.basic_vae.basic_vae_module.VAE` method), 146
[step\(\)](#) (`pl_bolts.optimizers.lars_scheduling.LARSWrapper` method), 183
[stl\(\)](#) (`pl_bolts.models.self_supervised.amdim.datasets.AMDIMPatchesPretrain` static method), 154
[stl\(\)](#) (`pl_bolts.models.self_supervised.amdim.datasets.AMDIMPatchesPretrain` static method), 154
[stl10_normalization\(\)](#) (in `pl_bolts.transforms.dataset_normalizations`), 185
[STL10DataModule](#) (class in `pl_bolts.datamodules.stl10_datamodule`), 139
[TanhClip](#) (in `pl_bolts.losses.self_supervised_learning`), 182
[targets](#) (`pl_bolts.datamodules.base_dataset.LightDataset` attribute), 120
[targets](#) (`pl_bolts.datamodules.cifar10_dataset.CIFAR10` attribute), 125
[targets](#) (`pl_bolts.datamodules.cifar10_dataset.TrialCIFAR10` attribute), 126
[TensorBoardGenerativeModelImageSampler](#) (class in `pl_bolts.callbacks.vision.image_generation`), 115
[TensorDataModule](#) (class in `pl_bolts.datamodules.sklearn_datamodule`), 137
[TensorDataset](#) (class in `pl_bolts.datamodules.sklearn_datamodule`), 138
[test_dataloader\(\)](#) (`pl_bolts.datamodules.binary_mnist_datamodule.BinaryMNISTDataModule` method), 121
[test_dataloader\(\)](#) (`pl_bolts.datamodules.cifar10_datamodule.CIFAR10DataModule` method), 123
[test_dataloader\(\)](#) (`pl_bolts.datamodules.cityscapes_datamodule.CityscapesDataModule` method), 127
[test_dataloader\(\)](#) (`pl_bolts.datamodules.fashion_mnist_datamodule.FashionMNISTDataModule` method), 130
[test_dataloader\(\)](#) (`pl_bolts.datamodules.imagenet_datamodule.ImagenetDataModule` method), 132
[test_dataloader\(\)](#) (`pl_bolts.datamodules.mnist_datamodule.MNISTDataModule` method), 135
[test_dataloader\(\)](#) (`pl_bolts.datamodules.sklearn_datamodule.SklearnDataModule` method), 136
[test_dataloader\(\)](#) (`pl_bolts.datamodules.ssl_imagenet_datamodule.SSLImagenetDataModule` method), 139

`method`), 139
`test_dataloader()` (`pl_bolts.datamodules.stl10_datamodule.STL10DataModule`), 140
`test_dataloader()` (`pl_bolts.models.mnist_module.LitMNIST`), 180
`test_epoch_end()` (`pl_bolts.models.mnist_module.LitMNIST`), 180
`test_epoch_end()` (`pl_bolts.models.regression.linear_regression.LinearRegression`), 151
`test_epoch_end()` (`pl_bolts.models.regression.logistic_regression.LogisticRegression`), 152
`test_epoch_end()` (`pl_bolts.models.vision.image_gpt.igpt_module.ImageGPT`), 178
`TEST_FILE_NAME` (`pl_bolts.datamodules.cifar10_dataset.CIFAR10`), 125
`test_step()` (`pl_bolts.models.mnist_module.LitMNIST`), 180
`test_step()` (`pl_bolts.models.regression.linear_regression.LinearRegression`), 151
`test_step()` (`pl_bolts.models.regression.logistic_regression.LogisticRegression`), 152
`test_step()` (`pl_bolts.models.self_supervised.ssl_finetuner.SSLFinetuner`), 175
`test_step()` (`pl_bolts.models.vision.image_gpt.igpt_module.ImageGPT`), 179
`tile()` (in module `pl_bolts.utils.shaping`), 186
`TinyCIFAR10DataModule` (class in `pl_bolts.datamodules.cifar10_datamodule`), 123
`to_device()` (`pl_bolts.callbacks.self_supervised.SSLOnlineEvaluation`), 118
`torchvision_ssl_encoder()` (in module `pl_bolts.utils.self_supervised`), 185
`train_dataloader()` (`pl_bolts.datamodules.binary_mnist_datamodule.BinaryMNISTDataModule`), 121
`train_dataloader()` (`pl_bolts.datamodules.cifar10_datamodule.CIFAR10DataModule`), 123
`train_dataloader()` (`pl_bolts.datamodules.cityscapes_datamodule.CityscapesDataModule`), 127
`train_dataloader()` (`pl_bolts.datamodules.fashion_mnist_datamodule.FashionMNISTDataModule`), 130
`train_dataloader()` (`pl_bolts.datamodules.imagenet_datamodule.ImagenetDataModule`), 132
`train_dataloader()` (`pl_bolts.datamodules.mnist_datamodule.MNISTDataModule`), 135
`train_dataloader()` (`pl_bolts.datamodules.sklearn_datamodule.SklearnDataModule`), 136
`train_dataloader()` (`pl_bolts.datamodules.ssl_imagenet_datamodule.SSLImagenetDataModule`), 139
`train_dataloader()` (`pl_bolts.datamodules.stl10_datamodule.STL10DataModule`), 140
`train_dataloader()` (`pl_bolts.datamodules.voc_detection_datamodule.VOCDetectionDataModule`), 141
`train_dataloader()` (`pl_bolts.models.mnist_module.LitMNIST`), 180
`train_dataloader_labeled()` (`pl_bolts.datamodules.stl10_datamodule.STL10DataModule`), 140
`train_dataloader_mixed()` (`pl_bolts.datamodules.stl10_datamodule.STL10DataModule`), 140
`TRAIN_FILE_NAME` (`pl_bolts.datamodules.cifar10_dataset.CIFAR10`), 125
`train_form()` (`pl_bolts.datamodules.imagenet_datamodule.ImagenetDataModule`), 132
`training_step()` (`pl_bolts.models.autoencoders.basic_ae.basic_ae_module.BasicAE`), 144
`training_step()` (`pl_bolts.models.autoencoders.basic_vae.basic_vae_module.BasicVAE`), 146
`training_step()` (`pl_bolts.models.detection.faster_rcnn.FasterRCNN`), 148
`training_step()` (`pl_bolts.models.gans.basic.basic_gan_module.BasicGAN`), 150
`training_step()` (`pl_bolts.models.mnist_module.LitMNIST`), 180
`training_step()` (`pl_bolts.models.regression.linear_regression.LinearRegression`), 151
`training_step()` (`pl_bolts.models.regression.logistic_regression.LogisticRegression`), 152
`training_step()` (`pl_bolts.models.self_supervised.amdim.amdim_module.AMDIM`), 154
`training_step()` (`pl_bolts.models.self_supervised.byol.byol_module.Byol`), 162
`training_step()` (`pl_bolts.models.self_supervised.cpc.cpc_module.CPC`), 169
`training_step()` (`pl_bolts.models.self_supervised.moco.moco2_module.MOCO2`), 171
`training_step()` (`pl_bolts.models.self_supervised.simclr.simclr_module.SimCLR`), 171
`training_step()` (`pl_bolts.models.self_supervised.ssl_finetuner.SSLFinetuner`), 175


```
validation_step()  
    (pl_bolts.models.self_supervised.ssl_finetuner.SSLFineTuner  
     method), 175  
validation_step()  
    (pl_bolts.models.vision.image_gpt.igpt_module.ImageGPT  
     method), 179  
VOCDetectionDataModule      (class      in  
    pl_bolts.datamodules.vocdetection_datamodule),  
141
```

W

```
wide_resnet101_2()          (in      module  
    pl_bolts.models.self_supervised.resnets),  
174  
wide_resnet50_2()          (in      module  
    pl_bolts.models.self_supervised.resnets),  
174
```