

Pytorch with Lightning ⚡

Why ⚡

- **Full flexibility**
 - Try any ideas using raw PyTorch without the boilerplate.
- **Reproducible + Readable**
 - Decoupled research and engineering code enable reproducibility and better readability.
- **Simple multi-GPU training**
 - Use multiple GPUs/TPUs/HPUs etc... without code changes.
- **Built-in testing**

Torch Model

```
1 class Net(nn.Module):
2     def __init__(self):
3         super(Net, self).__init__()
4         self.conv1 = nn.Conv2d(3, 6, 5)
5         self.pool = nn.MaxPool2d(2, 2)
6         self.conv2 = nn.Conv2d(6, 16, 5)
7         self.fc1 = nn.Linear(16 * 5 * 5, 120)
8         self.fc2 = nn.Linear(120, 10)
9     def forward(self, x):
10        x = self.pool(F.relu(self.conv1(x)))
11        x = self.pool(F.relu(self.conv2(x)))
12        x = x.view(-1, 16 * 5 * 5)
13        x = F.relu(self.fc1(x))
14        return self.fc2(x)
```

Torch Training Loop

```
1 for epoch in range(EPOCHS):
2     for i, data in enumerate(trainloader, 0):
3         inputs, labels = data
4
5         optimizer.zero_grad()
6
7         # forward + backward + optimize
8         outputs = net(inputs)
9         loss = criterion(outputs, labels)
10        loss.backward()
11        optimizer.step()
12
13        # print statistics
14        ...
```

Ok... Someone told me LR scheduler helps training

```
1 scheduler = lr_scheduler.StepLR(optimizer, step_size=7, gamma=0.1)
2 for epoch in range(EPOCHS):
3     for i, data in enumerate(trainloader, 0):
4         inputs, labels = data
5         optimizer.zero_grad()
6
7         outputs = net(inputs)
8         loss = criterion(outputs, labels)
9         loss.backward()
10        optimizer.step()
11        scheduler.step()
```

Perhaps I should clip my gradients

```
1 scheduler = lr_scheduler.StepLR(optimizer_ft, step_size=7,  
2 for epoch in range(EPOCHS):  
3     for i, data in enumerate(trainloader, 0):  
4         inputs, labels = data  
5         optimizer.zero_grad()  
6  
7         outputs = net(inputs)  
8         loss = criterion(outputs, labels)  
9         loss.backward()  
10        torch.nn.utils.clip_grad_norm_(model.parameters(),  
11        optimizer.step()  
12        scheduler.step()
```

Let's speed up the training and accumulate gradients

```
1 scheduler = lr_scheduler.StepLR(optimizer_ft, step_size=7,
2 for epoch in range(EPOCHS):
3     loss = 0
4     for i, data in enumerate(trainloader, 0):
5         inputs, labels = data
6         optimizer.zero_grad()
7
8         outputs = net(inputs)
9         loss += criterion(outputs, labels) # Compute loss
10        if (i+1) % accumulation_steps == 0:
11            loss = loss / accumulation_steps
12            loss.backward()
13            torch.nn.utils.clip_grad_norm_(model.parameters)
14            optimizer.step()
15            scheduler.step()
```

Oh.. I found this github repo with something interesting

```
1 from suspicious_not_tested_github_repo import EarlyStopping  
2  
3 es = EarlyStopping(no_documentation_whatsoever=True)
```


Ups, I forgot to log my metrics and losses

```
1 scheduler = lr_scheduler.StepLR(optimizer_ft, step_size=7,
2 writer = SummaryWriter()
3
4 for epoch in range(EPOCHS):
5     loss = 0
6     for i, data in enumerate(trainloader, 0):
7         inputs, labels = data
8         optimizer.zero_grad()
9
10        outputs = net(inputs)
11        loss += criterion(outputs, labels) # Compute loss
12
13        if phase == "training":
14            if (i+1) % accumulation_steps == 0:
15                loss = loss / accumulation_steps
```

Wait my training loop isn't working...

Where is the bug

```
1 for epoch in range(EPOCHS):
2     loss = 0
3     for i, data in enumerate(trainloader, 0):
4         inputs, labels = data
5         #optimizer.zero_grad()
6
7         outputs = net(inputs)
8         loss += criterion(outputs, labels) # Compute loss
9
10        if phase == "training":
11            if (i+1) % accumulation_steps == 0:
12                loss = loss / accumulation_steps
13                loss.backward()
14                torch.nn.utils.clip_grad_norm_(model.parameters())
15                optimizer.step()
```

Nice we got a GPU, lets update the code!

```
1 device = torch.device("cuda:0" if torch.cuda.is_available(  
2 model = model.to(device)  
3  
4 for epoch in range(EPOCHS):  
5     for i, data in enumerate(trainloader, 0):  
6         inputs, labels = data  
7         inputs = inputs.to(device)  
8         labels = labels.to(device)  
9  
10        optimizer.zero_grad()  
11  
12        outputs = net(inputs)  
13        loss += criterion(outputs, labels) # Compute loss  
14        ...
```

But wait... What if I want to use more than 1 GPU? 😞

Lets simplify it with ⚡

The main principle

You do the cool stuff

Lightning takes care of the boilerplate

```
1 for epoch in range(EPOCHS):
2     for i, data in enumerate(trainloader, 0):
3         inputs, labels = data
4         optimizer.zero_grad()
5         outputs = net(inputs)
6         loss = criterion(outputs, labels)
7         loss.backward()
8         optimizer.step()
9         ...
```

```
1 class CustomClassifier(pl.LightningModule):
2
3     def __init__(self):
4         super().__init__()
5         self.criterion = nn.CrossEntropyLoss()
6         self.fc1 = nn.Linear(784, 10)
7
8     def forward(self, x):
9         return self.fc1(x)
10
11    def training_step(self, batch, batch_idx):
12        x, y = batch
13        logits = self.forward(x)
14
15        loss = self.criterion(logits, y)
```


We have it all, lets train!

```
1 model = CustomClassifier()  
2 dm = CustomDatamodule(batch_size=21)  
3  
4 trainer = pl.Trainer()  
5 trainer.fit(model, dm, epochs=37)
```


Can we finally use multiple GPUs? Or even a TPU?

```
1 # dp = DataParallel
2 trainer = Trainer(gpus=2, distributed_backend='dp')
3
4 # ddp = DistributedDataParallel
5 trainer = Trainer(gpus=2, num_nodes=2, distributed_backend=
6
7 # ddp2 = DistributedDataParallel + dp
8 trainer = Trainer(gpus=2, num_nodes=2, distributed_backend=

1 trainer = pl.Trainer(tpu_cores=8)
2 trainer.fit(model, dm)
```

Lets see a full example

Package to interact with the Operating System

Pytorch package

Pytorch submodule with the Neural Network layers

Pytorch submodule for tensor manipulation

Torchvision submodule with dataset transformations

Torchvision submodule with predefined datasets

Class used to load a dataset




```
import os
import torch
from torch import nn
import torch.nn.functional as F
from torchvision import transforms
from torchvision.datasets import MNIST
from torch.utils.data import DataLoader
import lightning as L
```

How Lightning is organised

A `LightningModule` will help us organise our code into 6 sections:

- Initialisation
- Train Loop
- Validation Loop
- Test Loop
- Prediction Loop
- Optimizers and LR Schedulers

Encoder Decoder in Pytorch



```
class Encoder(nn.Module):
    def __init__(self):
        super().__init__()
        self.l1 = nn.Sequential(nn.Linear(28 * 28, 64), nn.ReLU(), nn.Linear(64, 3))

    def forward(self, x):
        return self.l1(x)

class Decoder(nn.Module):
    def __init__(self):
        super().__init__()
        self.l1 = nn.Sequential(nn.Linear(3, 64), nn.ReLU(), nn.Linear(64, 28 * 28))

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

Lets wrap it up with ⚡

For this example we only need the **training step** and the **optimizers**

```
class LitAutoEncoder(L.LightningModule):
    def __init__(self, encoder, decoder):
        super().__init__()
        self.encoder = encoder
        self.decoder = decoder

    def training_step(self, batch, batch_idx):
        # training_step defines the train loop.
        x, y = batch
        x = x.view(x.size(0), -1)
        z = self.encoder(x)
        x_hat = self.decoder(z)
        loss = F.mse_loss(x_hat, x)
        return loss

    def configure_optimizers(self):
        optimizer = torch.optim.Adam(self.parameters(), lr=1e-3)
        return optimizer
```

Lets define the dataset and the training

```
dataset = MNIST(os.getcwd(), download=True, transform=transforms.ToTensor())
train_loader = DataLoader(dataset)
```

```
# model
autoencoder = LitAutoEncoder(Encoder(), Decoder())

# train model
trainer = L.Trainer()
trainer.fit(model=autoencoder, train_dataloaders=train_loader)
```

What is happening under the hood

```
autoencoder = LitAutoEncoder(Encoder(), Decoder())
optimizer = autoencoder.configure_optimizers()

for batch_idx, batch in enumerate(train_loader):
    loss = autoencoder.training_step(batch, batch_idx)

    loss.backward()
    optimizer.step()
    optimizer.zero_grad()
```

Validation and Testing

Imports

```
import torch.utils.data as data
from torchvision import datasets
import torchvision.transforms as transforms

# Load data sets
transform = transforms.ToTensor()
train_set = datasets.MNIST(root="MNIST", download=True, train=True, transform=transform)
test_set = datasets.MNIST(root="MNIST", download=True, train=False, transform=transform)
```

test_step function

```
class LitAutoEncoder(L.LightningModule):
    def training_step(self, batch, batch_idx):
        ...

    def test_step(self, batch, batch_idx):
        # this is the test loop
        x, y = batch
        x = x.view(x.size(0), -1)
        z = self.encoder(x)
        x_hat = self.decoder(z)
        test_loss = F.mse_loss(x_hat, x)
        self.log("test_loss", test_loss)
```

How to test the model?

```
from torch.utils.data import DataLoader

# initialize the Trainer
trainer = Trainer()

# test the model
trainer.test(model, dataloaders=DataLoader(test_set))
```

Validation and Testing

validation_step function

```
class LitAutoEncoder(L.LightningModule):
    def training_step(self, batch, batch_idx):
        ...

    def validation_step(self, batch, batch_idx):
        # this is the validation loop
        x, y = batch
        x = x.view(x.size(0), -1)
        z = self.encoder(x)
        x_hat = self.decoder(z)
        val_loss = F.mse_loss(x_hat, x)
        self.log("val_loss", val_loss)
```

How to use the validation split?

```
from torch.utils.data import DataLoader

train_loader = DataLoader(train_set)
valid_loader = DataLoader(valid_set)
model = LitAutoEncoder(...)

# train with both splits
trainer = L.Trainer()
trainer.fit(model, train_loader, valid_loader)
```


Saving and Loading Checkpoints

Defining the root directory for the checkpoints

```
# saves checkpoints to 'some/path/' at every epoch end  
trainer = Trainer(default_root_dir="some/path/")
```

How to load from a checkpoint

```
model = MyLightningModule.load_from_checkpoint("/path/to/checkpoint.ckpt")  
  
# disable randomness, dropout, etc...  
model.eval()  
  
# predict with the model  
y_hat = model(x)
```