

The logo for Carnegie Mellon University, featuring a dark blue background with a grid of colorful lines (red, green, yellow, blue) forming a diamond pattern.

**Carnegie
Mellon
University**

Introduction to PyTorch

10-301/10-601 Introduction to Machine learning

Nov 10, 2023

(Credit to 16-824 Spring '22, Fall '23 TAs - modified by ML Fall '23 TAs)

Configure Colab with PyTorch

- Covered at the end of HW6 recitation (recording in canvas) ([colab link](#))

```
[ ] import torch
```

The following command will return `True` if we have a device that supports CUDA, which in our case is the T4 GPU.

```
[ ] torch.cuda.is_available()
```

```
True
```

This command tells us how many GPUs we have available.

```
[ ] torch.cuda.device_count()
```

```
1
```

This command tells us the name of the GPU that we are using.

```
[ ] torch.cuda.get_device_name(0)
```

```
'Tesla T4'
```

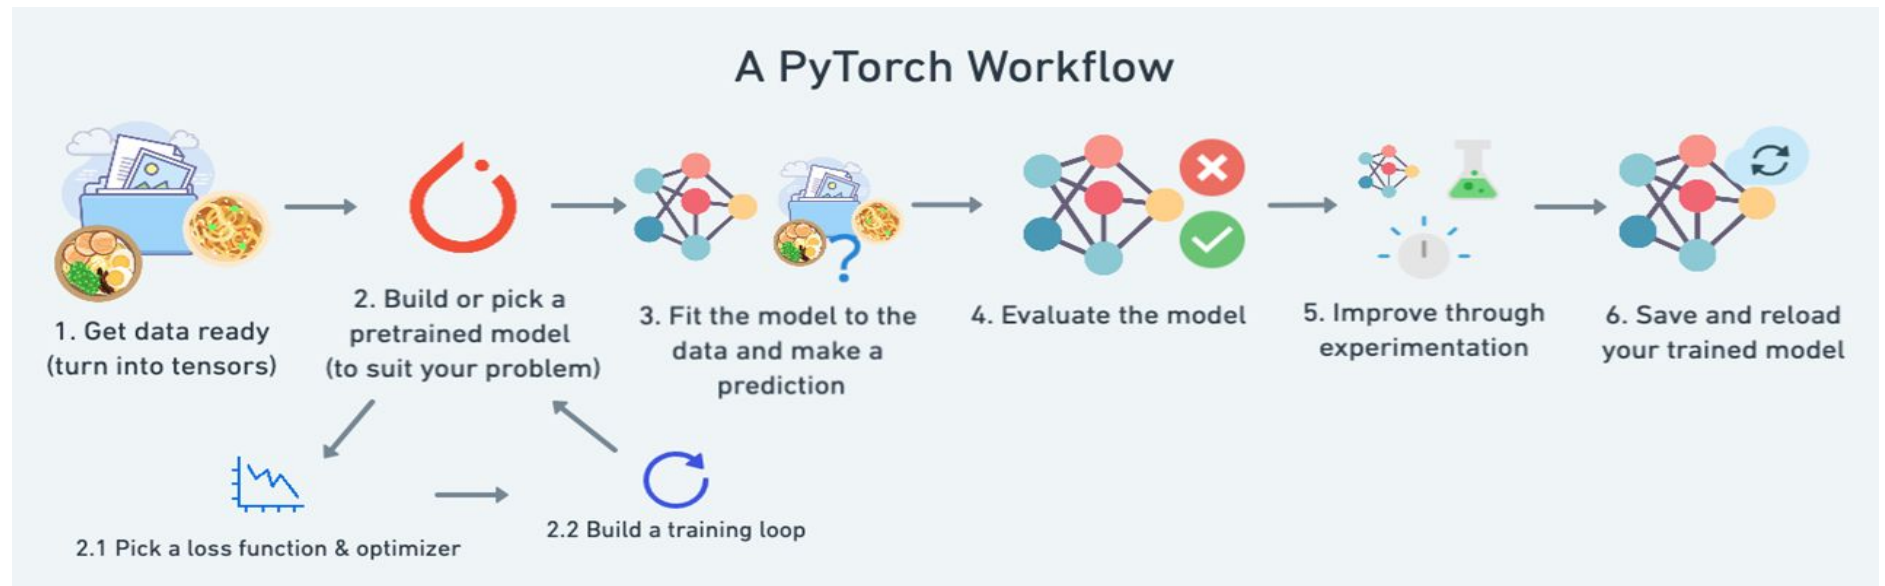


Contents

- What's PyTorch
- Tensors in PyTorch
- Model definition: nn Module
- Training
- Validation
- Key Metrics
- Summary of NN


What is PyTorch?

- An open source machine learning framework that accelerates the path from research prototyping to production deployment.
- A library of deep learning modules/functions/losses/optimizers, etc.



Tensors: Backbone of PyTorch

- Multi-dimensional array, same as numpy array
- Biggest difference: Tensors can be run on CPU/GPU



```
import torch
import numpy as np
a_np = np.zeros((32,32))
a_torch = torch.from_numpy(a_np)
a_np = a_torch.numpy()
```

Basic elements: Tensor

- torch.Tensor
- Just like numpy array

Create from list

```
1 import torch
2
3 data = [[1, 2],[3, 4]]
4 x_data = torch.tensor(data)
```

Create from numpy array

```
1 import torch
2 import numpy as np
3
4 np_array = np.array(data)
5 x_np = torch.from_numpy(np_array)
```

Basic elements: Tensor

- Create tensor

```
1 import torch
2
3 a = torch.ones(3, 3)
4 a = torch.zeros(3, 3)
5 a = torch.randn(3, 3)
```

Basic elements: Tensor

- Indexing and slicing

```
1 import torch
2
3 tensor = torch.ones(4, 4)
4 tensor[:, 1] = 0
5 tensor[1:2, 3:-1] = 2
```


Basic elements: Tensor

- device and type

```
1 import torch
2
3 a = torch.ones(3,3, dtype=torch.float32)
4 a.device      # cpu
5 a.dtype       # torch.float32
6 b = torch.ones_like(a)
7 b.to("cuda")
8 b.to(torch.int32)
9 b.to(a.device)
10 b.to(a.dtype)
```

Basic elements: Tensor


- Tensor operation
- Add/subtract/multiply/matrix multiply/transpose/expand ...
- Support operation in batch

```
1 import torch
2
3 a = torch.ones(2, 3, 1)
4 b = 2*torch.ones(2, 3, 1)
5 c1 = a + b # [[[3], [3], [3]], [[3], [3], [3]]]
6 c2 = a - b # [[[-1], [-1], [-1]], [[-1], [-1], [-1]]]
7 d = a * b # [[[2], [2], [2]], [[2], [2], [2]]]
8 e = a.transpose(1, 2) # e.size(): [2, 1, 3]
9 f = b @ e # 2*torch.ones(2, 3, 3)
10 g = a.expand(-1, -1, 3) # torch.ones(2, 3, 3)
```

Basic elements: Tensor

- `Tensor.Contiguous()`

```
arr = np.arange(12).reshape(3,4)
```




arr

`arr.contiguous()`



arr.T



Model Definition: nn.Module

- Defining a MLP
- Define basic operations
- Define forward functions
- A `nn.Module.parameters()` returns trainable parameters

```
1 import torch
2 import torch.nn as nn
3 import torch.nn.functional as F
4
5 class Net(nn.Module):
6
7     def __init__(self):
8         super(Net, self).__init__()
9         self.fc1 = nn.Linear(128, 128)
10        self.fc2 = nn.Linear(128, 64)
11        self.fc3 = nn.Linear(64, 10)
12
13    def forward(self, x):
14        x = F.relu(self.fc1(x))
15        x = F.relu(self.fc2(x))
16        x = self.fc3(x)
17
18        return x
19
20 net = Net()
21 print(net)
```

Model Definition: nn.Module

- Customize your layer(network)
- Declare param
- Initialize params
- Define operations

```
1 import torch
2 import torch.nn as nn
3 import torch.nn.functional as F
4
5 class MyLinear(nn.Module):
6
7     def __init__(self, in_ch, out_ch):
8         super(MyLinear, self).__init__()
9         self.W = nn.Parameter(torch.zeros(in_ch, out_ch), requires_grad=True)
10        self.b = nn.Parameter(torch.zeros(out_ch), requires_grad=True)
11
12        ## initialize weights
13        nn.init.xavier_uniform_(self.W)
14        nn.init.zeros_(self.b)
15
16    def forward(self, x):
17        return (x[:, :, None] @ self.W[None, :, :])[..., 0] + self.b[None]
```

Training Loop

```
1 running_loss = 0.0
2 for i, data in enumerate(trainloader, 0):
3     # get the inputs; data is a list of [inputs, labels]
4     inputs, labels = data
5
6     # zero the parameter gradients
7     optimizer.zero_grad()
8
9     # forward + backward + optimize
10    outputs = net(inputs)
11    loss = criterion(outputs, labels)
12    loss.backward()
13    optimizer.step()
```

Validation

- On the test dataset periodically run to look at accuracy/loss
- Set `model.eval()` to deactivate all the layers from updating
- Run with `torch.no_grad` to deactivate gradients

```
@torch.no_grad()  
def add_ab(a,b):  
    return a+b
```



Key metrics to track

- Train loss / accuracy
- Validation loss / accuracy

General rule of thumb:

- Train loss low, Validation loss low: things are working
- Train loss low, Validation loss high: overfitting
- Train loss high, Validation loss high: underfitting
- Train loss high, Validation loss low: some bug in evaluation



Summary of training a NN

1. Load data
 - DataLoader for batching, shuffling
2. Define Forward (of the neural network)
 - Implement nn.Module
3. Define loss
 - Pytorch provides a lot of these if needed
4. Define Backward
 - PyTorch automatically computes gradients
5. Optimizer to update the given parameters
 - torch.optim
6. Track key metrics