

PyTorch Tutorial

May 21, 2025

Part 1: Tensor

dtype

- ① torch.bool
- ② torch.int32, torch.int64
- ③ torch.float32, torch.float64
- ④ torch.complex64

dtype

- ① torch.bool
- ② torch.int32, torch.int64
- ③ torch.float32, torch.float64
- ④ torch.complex64

Default floating dtype is torch.float32, which can be modified by

`torch.set_default_dtype(torch.float64)`

```
1 x = torch.ones(3, 1)
2 print (x.dtype)
3
4 torch.set_default_dtype(torch.float64)
5 x = torch.ones(3, 1)
6 print (x.dtype)
7
```

device

Data can be stored on CPU, GPU, ...

device

Data can be stored on CPU, GPU, ...

```
1 x = torch.ones(3, 1)
2
3 # show where the data is
4 print (x.device)
5
6 device = torch.device('cuda')
7
8 # move x to GPU
9 print (x.to(device))
10
11 # define a tensor on GPU
12 y = torch.ones(3, 1, device=device)
13
```

Part 2: Neural networks

Feedforward neural networks

$$f(x; \theta) = f_L(f_{L-1}(\cdots(f_1(x))\cdots)) \quad (1)$$

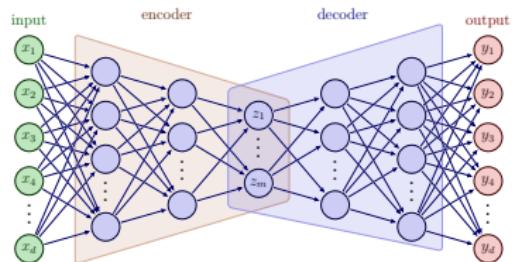
Recursive definition:

$$\begin{aligned} z_0 &= x \\ z_l &= f_l(z_{l-1}) = \varphi(b_l + W_l z_{l-1}), \quad l = 1, 2, \dots, L-1, \\ z_L &= f_L(z_{L-1}) = b_L + W_L z_{L-1}. \end{aligned} \quad (2)$$

where

- $z_l, b_l \in \mathbb{R}^{d_l}$, $W_l \in \mathbb{R}^{d_l \times d_{l-1}}$
- φ : activation function
- $\theta = (b_1, b_2, \dots, b_L, W_1, \dots, W_L)$

Feedforward neural networks



Building blocks

- ① Containers: Sequential, ModuleList
- ② nn.Linear
- ③ Activation: nn.ReLU, nn.Sigmoid, nn.Softplus, nn.Tanh
- ④ Convolution layers, dropout, recurrent¹

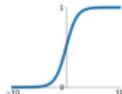
1. <https://docs.pytorch.org/docs/stable/nn.html>

Activation

Activation Functions

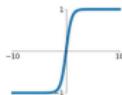
Sigmoid

$$\sigma(x) = \frac{1}{1+e^{-x}}$$



tanh

$$\tanh(x)$$



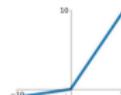
ReLU

$$\max(0, x)$$



Leaky ReLU

$$\max(0.1x, x)$$

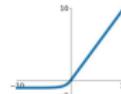


Maxout

$$\max(w_1^T x + b_1, w_2^T x + b_2)$$

ELU

$$\begin{cases} x & x \geq 0 \\ \alpha(e^x - 1) & x < 0 \end{cases}$$



torch.nn

```
1 class MyNet(nn.Module):  
2  
3     def __init__(self):  
4         super().__init__()  
5  
6         self.net = nn.Sequential(  
7             nn.Linear(1, 3),  
8             nn.Tanh(),  
9             nn.Linear(3, 4),  
10            nn.Tanh(),  
11            nn.Linear(4, 1),  
12        )  
13  
14    def forward(self, x):  
15        output = self.net(x)  
16        return output  
17  
18 model = MyNet()
```

A simpler way:

```
1 model = nn.Sequential(  
2     nn.Linear(1, 3),  
3     nn.Tanh(),  
4     nn.Linear(3, 4),  
5     nn.Tanh(),  
6     nn.Linear(4, 1),  
7 )
```

Parameters of the model are PyTorch tensors.

```
1 # print all (training) parameters of the model
2 for param in model.parameters():
3     print (param)
```

torch.nn

Parameters of the model are PyTorch tensors.

```
1 # print all (training) parameters of the model
2 for param in model.parameters():
3     print (param)
```

Output:

```
1 Parameter containing:
2 tensor([[0.1453],
3         [0.5682],
4         [0.7924]], requires_grad=True)
5 ...
```

Model evaluation

Model can be evaluated on mini-batches.

```
1 x = torch.ones(3, 1)  
2  
3 model(x)
```

Model evaluation

Input tensor and model tensors must have the same dtypes and devices.

```
1 # both x and tensors in the model have default dtype
  ↵ torch.float32
2 x = torch.ones(3, 1)
3 # works
4 model(x)
5
6 y = torch.ones(3, 1, dtype=torch.float64)
7 # doesn't work
8 model(y)
```

Example: positive function

Adding a positive activation to the output layer:

```
1 class MyNet(nn.Module):  
2  
3     def __init__(self):  
4         super().__init__()  
5  
6         self.net = nn.Sequential(  
7             nn.Linear(1, 3),  
8             nn.Tanh(),  
9             nn.Linear(3, 1),  
10            nn.Tanh(),  
11        )  
12  
13    def forward(self, x):  
14        output = self.net(x)  
15        return output  
16  
17 model = MyNet()
```

Example: features as input

Transform input tensor to features before passing to model.

```
1 class MyNet(nn.Module):
2
3     def __init__(self):
4         super().__init__()
5
6         self.net = nn.Sequential(
7             nn.Linear(1, 3),
8             nn.Tanh(),
9             nn.Linear(3, 1),
10            nn.Tanh(),
11        )
12
13    def forward(self, x):
14        output = self.net(torch.sin(x))
15        return output
16
17 model = MyNet()
```

Part 3: Automatic differentiation

computational graph

Whenever `require_grad=True`, PyTorch tracks the operations.

```
1 a = torch.tensor([2.0], requires_grad=True)
2
3 s = 0.5 * a ** 2
4
5 # requires_grad=True are set for parameters of model
6 model = nn.Sequential(
7     nn.Linear(1, 3),
8     nn.Tanh(),
9     nn.Linear(3, 4),
10    nn.Tanh(),
11    nn.Linear(4, 1),
12 )
```

computational graph

Backprogration:

```
1 a = torch.tensor([1.0], requires_grad=True)
2 b = torch.tensor([2.0])
3
4 s = 0.5 * a ** 2 + b
5
6 s.backward()
7 # output: 1 None
8 print (a.grad, b.grad)
```

computational graph

Gradient of **leaf nodes** will be accumulated in the grad field.

```
1 a = torch.tensor([1.0], requires_grad=True)
2
3 s = 0.5 * a ** 2
4 s.backward()
5 # gradient of s
6 print (a.grad)
7
8 s1 = 3 * a
9 # Gradients of s and s1 will be summed up.
10 # To avoid this, write:
11 # a.grad = None
12 s1.backward()
13
14 print (a.grad)
15
```

computational graph

Backward propagation will free intermediate values of the graph. To backward twice, add `retain_graph=True`.

```
1 a = torch.tensor([1.0], requires_grad=True)
2
3 s = 0.5 * a ** 2
4
5 # add retain_graph=True to avoid error below
6 s.backward(retain_graph=True)
7
8 s1 = 3 * s
9 s1.backward()
10
11 print (a.grad)
12
```

computational graph

Avoid gradient computation:

```
1 a = torch.tensor([1.0], requires_grad=True)
2
3 b = torch.tensor([1.5], requires_grad=True)
4 b.requires_grad = False
5
6 s = 0.5 * a ** 2 + b**2
7 s.backward()
```

computational graph

Detach: remove nodes from computational graph.

```
1 a = torch.tensor([1.0], requires_grad=True)
2
3 b = a
4 # quadratic in a
5 s = a * b
6 s.backward()
7
8 b = a.detach()
9 # linear in a
10 s = a * b
11 s.backward()
```

autograd.grad

autograd.grad allows to compute gradients with given tensors.

Example: compute $\int_0^1 |f'(x)|^2 dx$, where $f : \mathbb{R} \rightarrow \mathbb{R}$ is the model.

```
1 x = torch.linspace(0, 1, 101).reshape(-1,1)
  # we need to compute derivative wrt x
2 x.requires_grad_()
3
4
5 # f
6 y = model(x)
7
8 y_grad_vec = torch.autograd.grad(outputs=y.sum(), inputs=x,
  → retain_graph=True)[0]
9
10 int_val = torch.sum(y_grad_vec**2) * 0.01
```

Part 3: Dataset

different datasets

- ➊ training dataset: training models
- ➋ validation dataset: early stopping
- ➌ test dataset: evaluation of the trained model

dataset splitting

Method 1: torch.utils.data.random_split

```
1 x = torch.linspace(0, 1, 101).reshape(-1,1)
2 seed = 42
3
4 ratios = [0.8, 0.1, 0.1]
5
6 train_data, val_data, test_data =
    ↳ torch.utils.data.random_split(x, ratios,
    ↳ generator=torch.Generator().manual_seed(seed))
7
8 training_set = x[train_data.indices]
9 val_set = x[val_data.indices]
10 test_set = x[test_data.indices]
```

dataset splitting

Method 1: train_test_split from sklearn

```
1 import numpy as np
2 from sklearn.model_selection import train_test_split
3
4 X, y = np.arange(10).reshape((5, 2)), range(5)
5
6 X_train, X_test, y_train, y_test = train_test_split(X, y,
7     ↳ test_size=0.33, random_state=42)
8
9 print (X_train.shape, len(y_train))
```

DataLoader

Example 1:

```
1 from torch.utils.data import DataLoader  
2  
3 x = torch.linspace(0, 1, 101).reshape(-1,1)  
4 batch_size = 21  
5 data_loader = DataLoader(x, batch_size=batch_size,  
6                         shuffle=True, drop_last=True)  
7  
8 for data in data_loader:  
9     print (data.shape)
```

DataLoader

Example 2:

```
1 train_dataset = torchvision.datasets.MNIST(root='./data',
2     ↳ train=True, transform=transforms.ToTensor(),
2     ↳ download=True)
3
4 test_dataset = torchvision.datasets.MNIST(root='./data',
5     ↳ train=False, transform=transforms.ToTensor(),
5     ↳ download=True)
6
7 print (train_dataset.data.shape, test_dataset.data.shape)
8
9 train_loader = DataLoader(train_dataset, batch_size=1001,
10    ↳ shuffle=True)
11 test_loader = DataLoader(test_dataset, batch_size=1001,
11    ↳ shuffle=True)
12
13 for data in test_loader:
14     img, label = data
15     print (img.shape)
```

Part 3: Training

Optimizer

Optimization objective:

$$\mathcal{L}(\theta) := \frac{1}{N} \sum_{n=1}^N \ell(y_n, f(x_n; \theta)).$$

Full gradient: $\nabla_{\theta} \mathcal{L}(\theta) = \frac{1}{N} \sum_{n=1}^N \nabla_{\theta} \ell(y_n, f(x_n; \theta))$.

SGD:

$$\theta_{t+1} = \theta_t - \frac{\eta_t}{|\mathcal{B}_t|} \sum_{n \in \mathcal{B}_t} \nabla_{\theta} \ell(y_n, f(x_n; \theta_t))$$

Optimizer

Optimization objective:

$$\mathcal{L}(\theta) := \frac{1}{N} \sum_{n=1}^N \ell(y_n, f(x_n; \theta)).$$

Full gradient: $\nabla_{\theta} \mathcal{L}(\theta) = \frac{1}{N} \sum_{n=1}^N \nabla_{\theta} \ell(y_n, f(x_n; \theta))$.

SGD:

$$\begin{aligned}\theta_{t+1} &= \theta_t - \frac{\eta_t}{|\mathcal{B}_t|} \sum_{n \in \mathcal{B}_t} \nabla_{\theta} \ell(y_n, f(x_n; \theta_t)) \\ &= \theta_t - \eta_t \nabla_{\theta} \mathcal{L}(\theta_t) + \eta_t \left(\textcolor{blue}{\nabla_{\theta} \mathcal{L}(\theta_t)} - \frac{1}{|\mathcal{B}_t|} \sum_{n \in \mathcal{B}_t} \nabla_{\theta} \ell(y_n, f(x_n; \theta_t)) \right) \\ &= \theta_t - \eta_t \nabla_{\theta} \mathcal{L}(\theta_t) + \eta_t \textcolor{blue}{e_t}\end{aligned}$$

Optimizer

SGD:

```
1 import torch.optim as optim  
2  
3 model = nn.Sequential(nn.Linear(1, 3), nn.Tanh(),  
4                         nn.Linear(3, 4))  
5  
6 criterion = nn.CrossEntropyLoss()  
7 optimizer = optim.SGD(model.parameters(), lr=0.001,  
8                         momentum=0.9)
```

Adam:

```
1 criterion = nn.MSELoss(reduction='sum')  
2 optimizer = optim.Adam(model.parameters(), lr=0.001)
```

training loop

```
1   for epoch in range(total_epochs):      # for each epoch
2
3       # loop over all mini-batches
4       for data in data_loader:
5
6           y_pred = model(data)
7           y = fun(data)
8
9           loss = criterion(y_pred, y)
10
11          optimizer.zero_grad()
12          loss.backward()
13          optimizer.step()
```