


🔥 Deep Learning Workshop



Aron van de Pol

Leiden University Centre For Digital Humanities

2024-03-05





Images as data



 **Today**

1. Loading images into tensors
2. Batching & Dataloaders
3. Creating model architectures for image data.



Images as data

- Humanities we have lots of potential data
- Most often DH works with text
- Digitized texts not always *machine-readable*



Images as data

- Art Historians
- Visual history
- Print shops
- Advertisements
- Etc



FashionMNIST

- Very common dataset to practice on
- Produced from Zalando product images
- 70,000 grayscale images
 - 28*28px
- 9 Classes

 **FashionMNIST**

0. T-shirt/top
1. Trouser
2. Pullover
3. Dress
4. Coat
5. Sandal
6. Shirt
7. Sneaker
8. Bag
9. Ankle boot



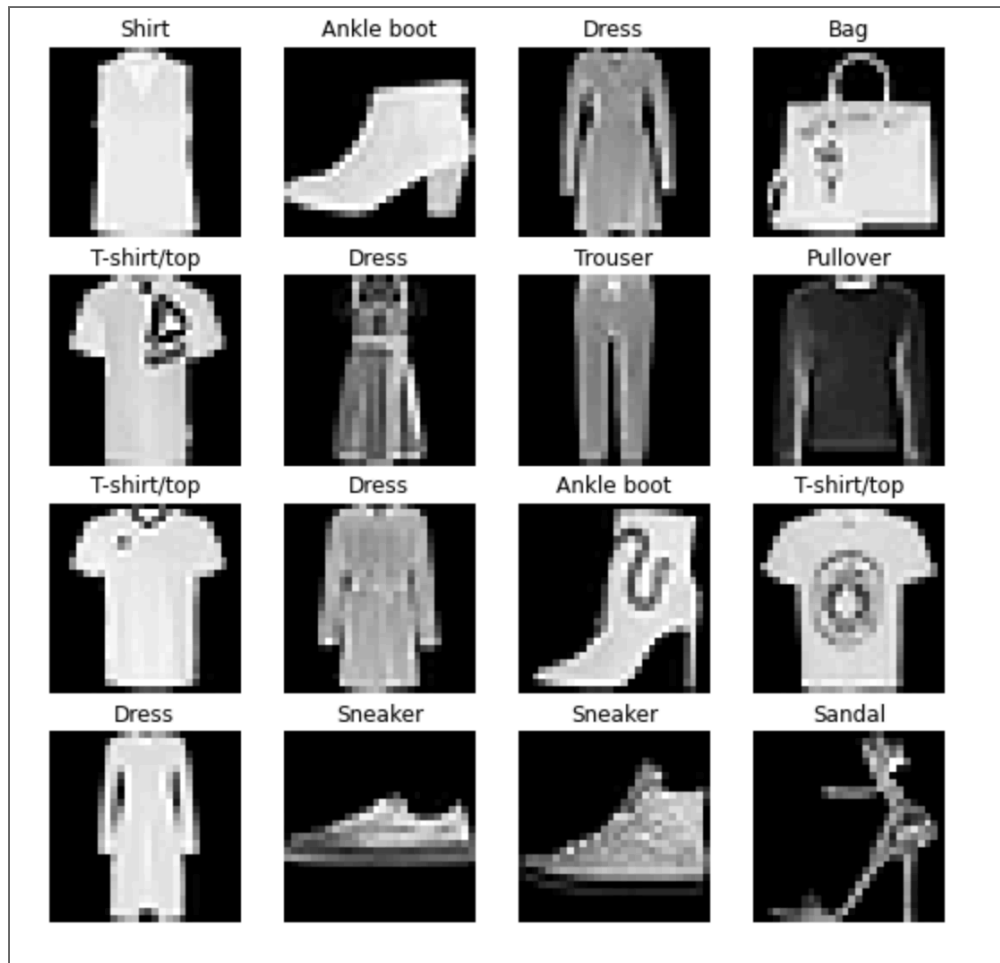
 **FashionMNIST**

Figure 1: Random sampled images from the FashionMNIST dataset



FashionMNIST

We can load it from pytorch:

```
1 train_data = datasets.FashionMNI
2     root="data", # where to down
3     train=True, # get training d
4     download=True, # download da
5     transform=ToTensor(), # turn
6     target_transform=None # you
7 )
8
9 # Setup testing data
10 test_data = datasets.FashionMNIS
11     root="data",
12     train=False, # get test data
13     download=True,
14     transform=ToTensor()
15 )
```



Batches and Dataloaders





Batches

- Single number data points are light.
- Images as data are heavy.
- Computationally inefficient to push the entire dataset through the model.





Batches

- Batching as a solution
- Split up the data into *batches*
- push batches through the model and every batch update parameters.





Batches

1. **Sample** - smallest unit, represents one data point.
2. **Batch size** - The hyperparameter that defines the number of samples per batch.
 - Often this is a power of 2 (2, 4, 8 16 etc...)
3. **Batch** - N amount of samples.
4. **Epoch** - Full pass over all the data (thus all batches)



Batches in PyTorch

```
1 from torch.utils.data import Data
2
3 # Setup the batch size hyperpara
4 BATCH_SIZE = 32
5
6 # Turn datasets into iterables (
7 train_dataloader = DataLoader(tr
8                               ba
9                               sh
10 )
11
12 test_dataloader = DataLoader(tes
13                       bat
14                       shu
15 )
```



Model Architectures for image data





Baseline model

- We start of with a baseline model
 - `nn.Linear()`
 - `nn.ReLU()`
- But, how can we input data?

■ `nn.Flatten()`

- Flatten changes our square image tensors into vectors.

Figure 2: `nn.Flatten()` visualized

■ nn.Flatten()

```
1 # Create a flatten layer
2 flatten_model = nn.Flatten()
3
4 # Get a single sample
5 x = train_features_batch[0]
6
7 # Flatten the sample
8 output = flatten_model(x)
9
10 # Print out what happened
11 print(f"Shape before flattening:")
12 print(f"Shape after flattening:")
13
14 #returns
15 Shape before flattening: torch.S
16 Shape after flattening: torch.Si
```



Baseline model

```
1 # Create a model with non-linear
2 class FashionModelV0(nn.Module):
3     def __init__(self):
4         super().__init__()
5         self.layer_stack = nn.Se
6             nn.Flatten(), # flat
7             nn.Linear(in_feature
8             nn.ReLU(),
9             nn.Linear(in_feature
10            nn.ReLU()
11        )
12
13    def forward(self, x: torch.T
14        return self.layer_stack(
```



Training loop differences

```
1 epochs = 10
2
3 for epoch in range(epochs): # loop
4     print(f'epoch:{epoch}')
5     train_loss = 0 # tracking the
6     for batch, (X, y) in enumerate
7         X, y = X.to(device), y.to
8         model_1.train()
9         # 1. Forward pass
10        y_pred = model_1(X)
11        loss = loss_fn(y_pred, y)
12        train_loss += loss # accumulate
13        # rest of train loop
14
15        # (average loss per batch per
16        train_loss /= len(train_data)
17        # rest of the test loop...
```



Let's code!





Convolutional Layers





Convolutional Layers


- Also known as CNN or ConvNet
- Good at finding patterns in visual data
- Makes use of so-called ‘convolutions’





A Convolution

1. We have our image
2. We have a kernel that applies some calculation to the pixels it 'slides' over. In this case the kernel is 3×3 .
3. The result of the calculation is a single pixel value for a 'new' image.
4. Slide over all the pixels to get a 'compressed' image.





A Convolution

CNN-explainer website





TinyVGG Network

1. 1st Conv2d layer (the convolution layer)
2. ReLU activation function
3. 2nd Conv2d layer
4. ReLU activation function
5. MaxPool2d layer (compressing the image even further)



TinyVGG Network

```
1 # Create a convolutional neural network
2 class FashionModelV1(nn.Module):
3     """
4     Model architecture copying VGG
5     """
6     def __init__(self):
7         super().__init__()
8         self.block_1 = nn.Sequential(
9             nn.Conv2d(in_channels=3, out_channels=64,
10                      kernel_size=3, stride=1, padding=1),
11             nn.ReLU(),
12             nn.Conv2d(in_channels=64, out_channels=64,
13                      kernel_size=3, stride=1, padding=1),
14             nn.ReLU(),
15             nn.Conv2d(in_channels=64, out_channels=128,
16                      kernel_size=3, stride=1, padding=1),
17             nn.ReLU(),
18             nn.Conv2d(in_channels=128, out_channels=128,
```

Let's Code



Confusion Matrix

Figure 3: Confusion Matrix of TinyVGG architecture on FashionMNIST



Works Cited

