



# PyTorch Introduction

## Training a network

David Rotermund & Udo Ernst

Recommendation

# Open Book: Dive into Deep Learning

- ASTON ZHANG
- ZACHARY C. LIPTON
- MU LI
- ALEXANDER J. SMOLA

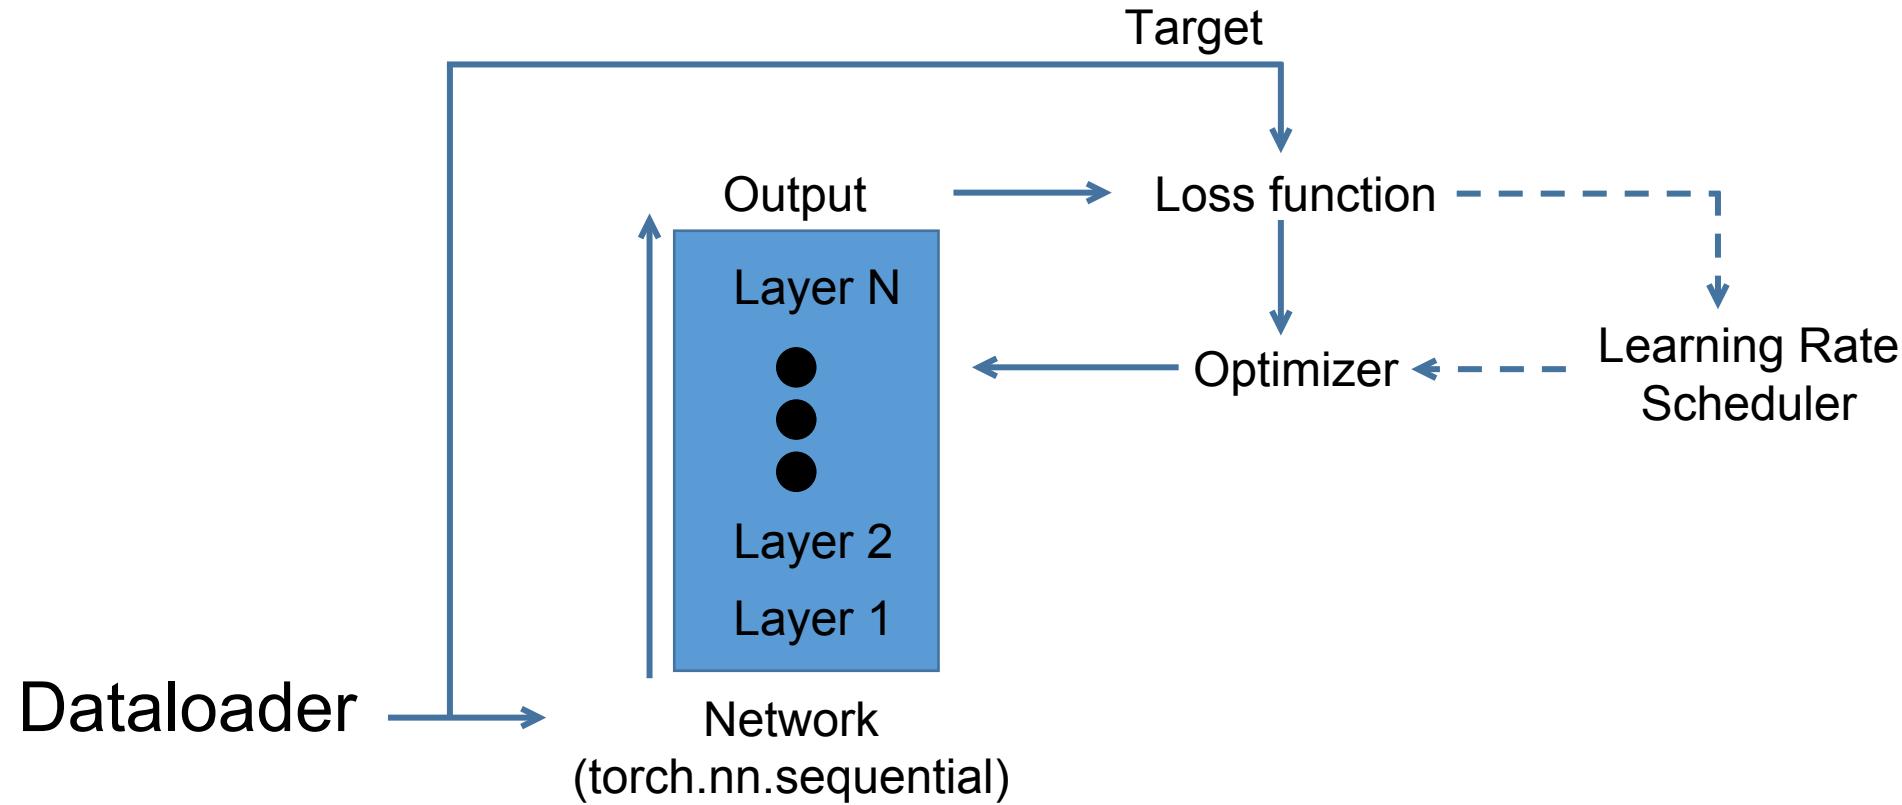
<https://d2l.ai/d2l-en.pdf>



Fig. 1.3.3 A donkey, a dog, a cat, and a rooster.

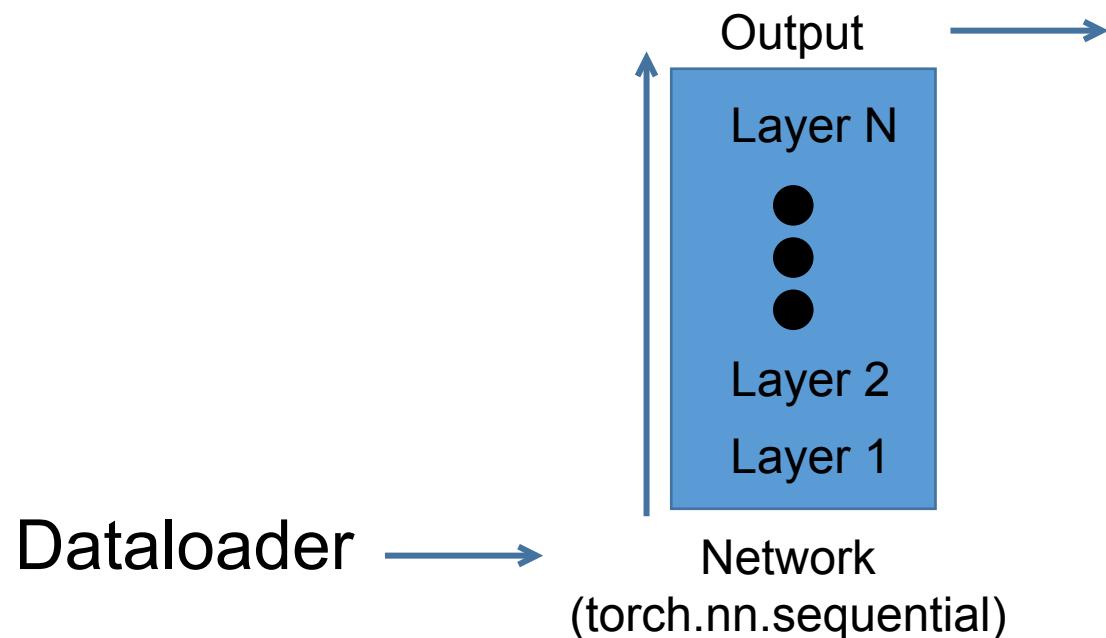
# Anatomy of a PyTorch Network + Support

## Training of the weights



# Anatomy of a PyTorch Network + Support

## Inference



`torch.tensor` is the `numpy.ndarray` of PyTorch



`== torch.tensor`  
`== stored data`

# Example network

I contains everything but it is not very well optimized.

Allowing **you** to improve on it.

(Or depending on your computer, switch it to the Fashion MNIST benchmark first...)

```
import torch
import torchvision # type: ignore
from torchvision.transforms import v2 # type: ignore
import time
import os
```

## The imports

```
number_of_epoch: int = 500
lr_parameter_max: float = 1e-9
```

```
ModelsPath: str = "Models"
os.makedirs(ModelsPath, exist_ok=True)
```

```
# Tensorboard
os.environ["TF_CPP_MIN_LOG_LEVEL"] = "3"
from torch.utils.tensorboard
import SummaryWriter
tb = SummaryWriter(log_dir="run")
```

## Setting up tensorboard

```
# GPU ?
device: torch.device =
    torch.device("cuda:0") if torch.cuda.is_available() else torch.device("cpu") GPU: Yes / No?
)
torch.set_default_dtype(torch.float32)

# Data augmentation
test_processing_chain = v2.Compose(
    transforms=[
        v2.ToImage(),
        v2.ToDtype(torch.float32, scale=True),
        v2.CenterCrop((28, 28)),
    ],
)

train_processing_chain = v2.Compose(
    transforms=[
        v2.ToImage(),
        v2.ToDtype(torch.float32, scale=True),
        v2.RandomCrop((28, 28)),
        v2.AutoAugment(),
    ],
)
```

**Data augmentation  
for test data**

**Data augmentation  
for training data**

```
# Data provider
tv_dataset_train = torchvision.datasets.CIFAR10(
    root="data",
    train=True,
    download=True,
    transform=train_processing_chain,
)
tv_dataset_test = torchvision.datasets.CIFAR10(
    root="data",
    train=False,
    download=True,
    transform=test_processing_chain,
)
```

**CIFAR10 training dataset  
(images, labels)**

**CIFAR10 test dataset  
(images, labels)**

```
# Data loader
train_data_load = torch.utils.data.DataLoader(
    tv_dataset_train, batch_size=100, shuffle=True)
```

## **data loader training data**

```
test_data_load = torch.utils.data.DataLoader(  
    tv_dataset_test, batch_size=100, shuffle=False)
```

# **data loader test data**

```
# Network
network = torch.nn.Sequential(
    torch.nn.Conv2d(
        in_channels=3,
        out_channels=32,
        kernel_size=5,
        stride=1,
        padding=0
```

# **network (layer in sequential container)**

**2D convolutional Layer  
(3 channel in, 32 out,  
5x5 kernel, 1x1 stride,  
no padding)**

## torch.nn.ReLU(), ReLU Layer

`torch.nn.BatchNorm2d(32)`, **BatchNorm2D Layer**

`torch.nn.MaxPool2d(kernel_size=2, stride=2, padding=0), no`

**max pooling  
( $2 \times 2$  kernel,  
 $2 \times 2$  stride,  
no padding)**

```
torch.nn.Conv2d(  
    in_channels=32,  
    out_channels=64,  
    kernel_size=5,  
    stride=1,  
    padding=0,  
),  
    torch.nn.ReLU(),  ReLU Layer  
    torch.nn.BatchNorm2d(64),   BatchNorm2D Layer  
    torch.nn.MaxPool2d(kernel_size=2, stride=2, padding=0),  
    torch.nn.Flatten(  
        start_dim=1,      Batch x 4 x 4 x 64 -> Batch x 1024  
    ),  
    torch.nn.Linear(  
        in_features=1024,  
        out_features=1024,  
        bias=True,  
    ),
```

**2D convolutional Layer  
(32 channel in, 64 out,  
5x5 kernel, 1x1 stride,  
no padding)**

**max pooling  
(2x2 kernel,  
2x2 stride,  
no padding)**

**Fully connected layer  
(1024 x 1024,  
with bias)**

```
torch.nn.ReLU(),          ReLU layer
torch.nn.Linear(  
    in_features=1024,  
    out_features=10,  
    bias=True,  
,           No ReLU / Softmax since CrossEntropyLoss contains Softmax  
).to(device)             Network to GPU if there is a GPU  
  
# Optimizer  
optimizer = torch.optim.Adam(network.parameters(), lr=0.001)  Adam optimizer  
  
# LR Scheduler  Learning rate scheduler  
lr_scheduler = torch.optim.lr_scheduler.ReduceLROnPlateau(optimizer)  
  
# Loss function Loss function  
loss_function = torch.nn.CrossEntropyLoss()
```

```
# Main loop
for epoch_id in range(0, number_of_epoch):
    print(f"Epoch: {epoch_id}")
    t_start: float = time.perf_counter()  Main loop over up to 500 epochs
                                            Getting a timestamp

    train_loss: float = 0.0
    train_correct: int = 0
    train_number: int = 0
    test_correct: int = 0
    test_number: int = 0

    # Switch the network into training mode
    network.train()  Setting the network into training mode
```

```
# This runs in total for one epoch split up into mini-batches  
for image, target in train_data_load:  
    # Clean the gradient  
    optimizer.zero_grad()
```

## Training mini batches

```
# Run data through network  
output = network(image.to(device))
```

## Pushing the mini batch through the network

```
# Measure the loss  
loss = loss_function(output, target.to(device))  
train_loss += loss.item()
```

```
# Classify  
train_correct += (  
    (output.argmax(dim=1) == target.to(device)).sum().detach().cpu().numpy()  
)  
train_number += target.shape[0]
```

```
# Calculate backprop  
loss.backward()
```

## Performing error back-propagation

```
# Update the parameter  
optimizer.step()
```

## Updating the parameter based on this mini batch

```
# Update the learning rate  
lr_scheduler.step(train_loss)
```

**Learning rate scheduler checks  
if learning rate needs adjustment**

```
t_training: float = time.perf_counter()
```

```
# Switch the network into evalution mode
```

```
network.eval()      Setting the network into evaluation mode
```

```
with torch.no_grad():    We don't need gradients
```

```
    for image, target in test_data_load:
```

```
# Run data through network
```

```
output = network(image.to(device))
```

**Pushing the mini batch  
through the network**

```
# Classifiy
```

```
test_correct += (    Comparing the argmax(output) with the target  
    (output.argmax(dim=1) == target.to(device)).sum().detach().cpu().numpy()  
)
```

```
test_number += target.shape[0]
```

```
t_testing = time.perf_counter()
```

```
perfomance_test_correct: float = 100.0 * test_correct / test_number
```

```
perfomance_train_correct: float = 100.0 * train_correct / train_number
```

### Store the data in Tensorboard

```
tb.add_scalar("Train Loss", train_loss, epoch_id)
```

```
tb.add_scalar("Train Number Correct", train_correct, epoch_id)
```

```
tb.add_scalar("Test Number Correct", test_correct, epoch_id)
```

```
tb.add_scalar("Error Test", 100.0 - perfomance_test_correct, epoch_id)
```

```
tb.add_scalar("Error Train", 100.0 - perfomance_train_correct, epoch_id)
```

```
tb.add_scalar("Learning Rate", optimizer.param_groups[-1][ "lr"], epoch_id)
```

```
tb.flush()
```

```
print(  
    f"Training: Loss={train_loss:.5f} Correct={perfomance_train_correct:.2f}% LR:{optimizer.param_groups[-1]['lr']}"  
)  
print(f"Testing: Correct={perfomance_test_correct:.2f}%")  
print(  
    f"Time: Training={(t_training - t_start):.1f}sec, Testing={(t_testing - t_training):.1f}sec"  
)
```

## Save the network

```
torch.save(network, os.path.join(ModelsPath, f"Model_MNIST_A_{epoch_id}.pt"))
```

```
print()
```


## Did we reach the final learning rate?

```
if optimizer.param_groups[-1]['lr'] < lr_parameter_max:  
    tb.close()  
    print("Done (lr_limit)")  
    exit()
```

```
tb.close() Close the Tensorboard connection
```

# VS Code Tensorboard extension


Observing the learning process...




```
# Tensorboard
os.environ["TF_CPP_MIN_LOG_LEVEL"] = "3"
▶ Launch TensorBoard Session
from torch.utils.tensorboard import SummaryWriter

tb = SummaryWriter(log_dir="run")
```

# A complete network... but not a good one



# A complete network... but not a good one



# A complete network... but not a good one

e.g. bad LR Scheduler settings

