
Table of Contents

Preface.....	ix
1. Getting Started with PyTorch.....	1
Building a Custom Deep Learning Machine	1
GPU	2
CPU/Motherboard	2
RAM	2
Storage	2
Deep Learning in the Cloud	3
Google Colaboratory	3
Cloud Providers	5
Which Cloud Provider Should I Use?	7
Using Jupyter Notebook	7
Installing PyTorch from Scratch	8
Download CUDA	8
Anaconda	9
Finally, PyTorch! (and Jupyter Notebook)	9
Tensors	10
Tensor Operations	11
Tensor Broadcasting	13
Conclusion	14
Further Reading	14
2. Image Classification with PyTorch.....	15
Our Classification Problem	15
Traditional Challenges	17
But First, Data	17
PyTorch and Data Loaders	18

Building a Training Dataset	18
Building Validation and Test Datasets	20
Finally, a Neural Network!	21
Activation Functions	22
Creating a Network	22
Loss Functions	23
Optimizing	24
Training	26
Making It Work on the GPU	27
Putting It All Together	27
Making Predictions	28
Model Saving	29
Conclusion	30
Further Reading	31
3. Convolutional Neural Networks.....	33
Our First Convolutional Model	33
Convolutions	34
Pooling	37
Dropout	38
History of CNN Architectures	39
AlexNet	39
Inception/GoogLeNet	40
VGG	41
ResNet	43
Other Architectures Are Available!	43
Using Pretrained Models in PyTorch	44
Examining a Model's Structure	44
BatchNorm	47
Which Model Should You Use?	48
One-Stop Shopping for Models: PyTorch Hub	48
Conclusion	49
Further Reading	49
4. Transfer Learning and Other Tricks.....	51
Transfer Learning with ResNet	51
Finding That Learning Rate	53
Differential Learning Rates	56
Data Augmentation	57
Torchvision Transforms	58
Color Spaces and Lambda Transforms	63
Custom Transform Classes	64

Start Small and Get Bigger!	65
Ensembles	66
Conclusion	67
Further Reading	67
5. Text Classification.....	69
Recurrent Neural Networks	69
Long Short-Term Memory Networks	71
Gated Recurrent Units	73
biLSTM	73
Embeddings	74
torchtext	76
Getting Our Data: Tweets!	77
Defining Fields	78
Building a Vocabulary	80
Creating Our Model	82
Updating the Training Loop	83
Classifying Tweets	84
Data Augmentation	84
Random Insertion	85
Random Deletion	85
Random Swap	86
Back Translation	86
Augmentation and torchtext	87
Transfer Learning?	88
Conclusion	88
Further Reading	89
6. A Journey into Sound.....	91
Sound	91
The ESC-50 Dataset	93
Obtaining the Dataset	93
Playing Audio in Jupyter	93
Exploring ESC-50	94
SoX and LibROSA	95
torchaudio	95
Building an ESC-50 Dataset	96
A CNN Model for ESC-50	98
This Frequency Is My Universe	100
Mel Spectrograms	100
A New Dataset	102
A Wild ResNet Appears	104

Finding a Learning Rate	105
Audio Data Augmentation	107
torchaudio Transforms	107
SoX Effect Chains	107
SpecAugment	109
Further Experiments	113
Conclusion	113
Further Reading	114
7. Debugging PyTorch Models.....	115
It's 3 a.m. What Is Your Data Doing?	115
TensorBoard	116
Installing TensorBoard	116
Sending Data to TensorBoard	117
PyTorch Hooks	120
Plotting Mean and Standard Deviation	121
Class Activation Mapping	122
Flame Graphs	125
Installing py-spy	127
Reading Flame Graphs	128
Fixing a Slow Transformation	129
Debugging GPU Issues	132
Checking Your GPU	132
Gradient Checkpointing	134
Conclusion	136
Further Reading	136
8. PyTorch in Production.....	137
Model Serving	137
Building a Flask Service	138
Setting Up the Model Parameters	140
Building the Docker Container	141
Local Versus Cloud Storage	144
Logging and Telemetry	145
Deploying on Kubernetes	147
Setting Up on Google Kubernetes Engine	147
Creating a k8s Cluster	148
Scaling Services	149
Updates and Cleaning Up	149
TorchScript	150
Tracing	150
Scripting	153

TorchScript Limitations	154
Working with libTorch	156
Obtaining libTorch and Hello World	156
Importing a TorchScript Model	158
Conclusion	159
Further Reading	160
9. PyTorch in the Wild.....	161
Data Augmentation: Mixed and Smoothed	161
mixup	161
Label Smoothing	165
Computer, Enhance!	166
Introduction to Super-Resolution	167
An Introduction to GANs	169
The Forger and the Critic	170
Training a GAN	171
The Dangers of Mode Collapse	172
ESRGAN	173
Further Adventures in Image Detection	173
Object Detection	173
Faster R-CNN and Mask R-CNN	175
Adversarial Samples	177
Black-Box Attacks	180
Defending Against Adversarial Attacks	180
More Than Meets the Eye: The Transformer Architecture	181
Paying Attention	181
Attention Is All You Need	182
BERT	183
FastBERT	183
GPT-2	185
Generating Text with GPT-2	185
ULMFiT	187
What to Use?	189
Conclusion	190
Further Reading	190
Index.....	193