

# Table of Contents

## Preface

1

## Section 1: Core Concepts

<b>Chapter 1: The Nuts and Bolts of Neural Networks</b>	9
<b>The mathematical apparatus of NNs</b>	10
Linear algebra	10
Vector and matrix operations	12
Introduction to probability	16
Probability and sets	18
Conditional probability and the Bayes rule	20
Random variables and probability distributions	22
Probability distributions	25
Information theory	28
Differential calculus	31
<b>A short introduction to NNs</b>	34
Neurons	35
Layers as operations	36
NNs	38
Activation functions	40
The universal approximation theorem	45
<b>Training NNs</b>	47
Gradient descent	48
Cost functions	50
Backpropagation	52
Weight initialization	56
SGD improvements	58
<b>Summary</b>	60

## Section 2: Computer Vision

<b>Chapter 2: Understanding Convolutional Networks</b>	63
<b>Understanding CNNs</b>	64
Types of convolutions	70
Transposed convolutions	70
$1 \times 1$ convolutions	73
Depth-wise separable convolutions	74
Dilated convolutions	75
Improving the efficiency of CNNs	76
Convolution as matrix multiplication	77
Winograd convolutions	79
Visualizing CNNs	82

Guided backpropagation	82
Gradient-weighted class activation mapping	85
CNN regularization	87
<b>Introducing transfer learning</b>	90
Implementing transfer learning with PyTorch	91
Transfer learning with TensorFlow 2.0	97
<b>Summary</b>	102
<b>Chapter 3: Advanced Convolutional Networks</b>	103
<b>Introducing AlexNet</b>	104
<b>An introduction to Visual Geometry Group</b>	105
VGG with PyTorch and TensorFlow	107
<b>Understanding residual networks</b>	108
Implementing residual blocks	111
<b>Understanding Inception networks</b>	118
Inception v1	119
Inception v2 and v3	120
Inception v4 and Inception-ResNet	122
<b>Introducing Xception</b>	123
<b>Introducing MobileNet</b>	126
<b>An introduction to DenseNets</b>	128
<b>The workings of neural architecture search</b>	131
<b>Introducing capsule networks</b>	136
The limitations of convolutional networks	136
Capsules	138
Dynamic routing	140
The structure of the capsule network	142
<b>Summary</b>	143
<b>Chapter 4: Object Detection and Image Segmentation</b>	145
<b>Introduction to object detection</b>	146
Approaches to object detection	147
Object detection with YOLOv3	148
A code example of YOLOv3 with OpenCV	154
Object detection with Faster R-CNN	157
Region proposal network	159
Detection network	161
Implementing Faster R-CNN with PyTorch	163
<b>Introducing image segmentation</b>	166
Semantic segmentation with U-Net	167
Instance segmentation with Mask R-CNN	170
Implementing Mask R-CNN with PyTorch	172
<b>Summary</b>	174
<b>Chapter 5: Generative Models</b>	175
<b>Intuition and justification of generative models</b>	175

<b>Introduction to VAEs</b>	176
Generating new MNIST digits with VAE	181
<b>Introduction to GANs</b>	187
Training GANs	188
Training the discriminator	190
Training the generator	192
Putting it all together	193
Problems with training GANs	195
<b>Types of GAN</b>	196
Deep Convolutional GAN	196
Implementing DCGAN	197
Conditional GAN	202
Implementing CGAN	203
Wasserstein GAN	206
Implementing WGAN	209
Image-to-image translation with CycleGAN	212
Implementing CycleGAN	215
Building the generator and discriminator	216
Putting it all together	218
<b>Introducing artistic style transfer</b>	222
<b>Summary</b>	224

## Section 3: Natural Language and Sequence Processing

---

<b>Chapter 6: Language Modeling</b>	227
Understanding n-grams	228
Introducing neural language models	230
Neural probabilistic language model	232
Word2Vec	233
CBOW	234
Skip-gram	236
fastText	238
Global Vectors for Word Representation model	239
<b>Implementing language models</b>	243
Training the embedding model	243
Visualizing embedding vectors	246
<b>Summary</b>	250
<b>Chapter 7: Understanding Recurrent Networks</b>	251
<b>Introduction to RNNs</b>	251
RNN implementation and training	256
Backpropagation through time	258
Vanishing and exploding gradients	262
<b>Introducing long short-term memory</b>	264
Implementing LSTM	269
<b>Introducing gated recurrent units</b>	274

Implementing GRUs	275
<b>Implementing text classification</b>	278
Summary	282
<b>Chapter 8: Sequence-to-Sequence Models and Attention</b>	283
Introducing seq2seq models	284
Seq2seq with attention	287
Bahdanau attention	287
Luong attention	290
General attention	292
Implementing seq2seq with attention	294
Implementing the encoder	294
Implementing the decoder	295
Implementing the decoder with attention	296
Training and evaluation	298
<b>Understanding transformers</b>	302
The transformer attention	302
The transformer model	306
Implementing transformers	310
Multihead attention	310
Encoder	313
Decoder	315
Putting it all together	316
<b>Transformer language models</b>	319
Bidirectional encoder representations from transformers	319
Input data representation	321
Pretraining	322
Fine-tuning	324
Transformer-XL	326
Segment-level recurrence with state reuse	328
Relative positional encodings	329
XLNet	331
Generating text with a transformer language model	335
<b>Summary</b>	337
<b>Section 4: A Look to the Future</b>	
<b>Chapter 9: Emerging Neural Network Designs</b>	341
<b>Introducing Graph NNs</b>	342
Recurrent GNNs	345
Convolutional Graph Networks	348
Spectral-based convolutions	349
Spatial-based convolutions with attention	352
Graph autoencoders	355
Neural graph learning	359
Implementing graph regularization	361
<b>Introducing memory-augmented NNs</b>	365

Neural Turing machines	366
MANN*	373
<b>Summary</b>	374
<b>Chapter 10: Meta Learning</b>	375
<b>Introduction to meta learning</b>	376
Zero-shot learning	377
One-shot learning	378
Meta-training and meta-testing	380
<b>Metric-based meta learning</b>	382
Matching networks for one-shot learning	382
Siamese networks	385
Implementing Siamese networks	386
Prototypical networks	389
<b>Optimization-based learning</b>	393
<b>Summary</b>	400
<b>Chapter 11: Deep Learning for Autonomous Vehicles</b>	401
<b>Introduction to AVs</b>	402
Brief history of AV research	402
Levels of automation	405
<b>Components of an AV system</b>	406
Environment perception	408
Sensing	409
Localization	411
Moving object detection and tracking	411
Path planning	412
<b>Introduction to 3D data processing</b>	413
<b>Imitation driving policy</b>	417
Behavioral cloning with PyTorch	418
Generating the training dataset	421
Implementing the agent neural network	423
Training	424
Letting the agent drive	426
Putting it all together	427
<b>Driving policy with ChauffeurNet</b>	429
Input and output representations	429
Model architecture	432
Training	434
<b>Summary</b>	439
<b>Other Books You May Enjoy</b>	441
<b>Index</b>	445

---