

Contents

Figures	xix
Tables	xxvii
Examples	xxix
Foreword	xxxiii
Preface	xxxv
Acknowledgments	xxxix
About the Authors	xli

I Introducing Deep Learning 1

1 Biological and Machine Vision 3

Biological Vision	3
Machine Vision	8
The Neocognitron	8
LeNet-5	9
The Traditional Machine Learning Approach	12
ImageNet and the ILSVRC	13
AlexNet	14
TensorFlow Playground	17
Quick, Draw!	19
Summary	19

2 Human and Machine Language 21

Deep Learning for Natural Language Processing	21
Deep Learning Networks Learn Representations Automatically	22
Natural Language Processing	23
A Brief History of Deep Learning for NLP	24

Computational Representations of Language	25
One-Hot Representations of Words	25
Word Vectors	26
Word-Vector Arithmetic	29
word2viz	30
Localist Versus Distributed Representations	32
Elements of Natural Human Language	33
Google Duplex	35
Summary	37
3 Machine Art	39
A Boozy All-Nighter	39
Arithmetic on Fake Human Faces	41
Style Transfer: Converting Photos into Monet (and Vice Versa)	44
Make Your Own Sketches Photorealistic	45
Creating Photorealistic Images from Text	45
Image Processing Using Deep Learning	46
Summary	48
4 Game-Playing Machines	49
Deep Learning, AI, and Other Beasts	49
Artificial Intelligence	49
Machine Learning	50
Representation Learning	51
Artificial Neural Networks	51
Deep Learning	51
Machine Vision	52
Natural Language Processing	53
Three Categories of Machine Learning Problems	53
Supervised Learning	53
Unsupervised Learning	54
Reinforcement Learning	54
Deep Reinforcement Learning	56
Video Games	57
Board Games	59
AlphaGo	59

AlphaGo Zero	62
AlphaZero	65
Manipulation of Objects	67
Popular Deep Reinforcement Learning Environments	68
OpenAI Gym	68
DeepMind Lab	69
Unity ML-Agents	71
Three Categories of AI	71
Artificial Narrow Intelligence	72
Artificial General Intelligence	72
Artificial Super Intelligence	72
Summary	72

II Essential Theory Illustrated 73

5 The (Code) Cart Ahead of the (Theory) Horse 75

Prerequisites	75
Installation	76
A Shallow Network in Keras	76
The MNIST Handwritten Digits	76
A Schematic Diagram of the Network	77
Loading the Data	79
Reformatting the Data	81
Designing a Neural Network Architecture	83
Training a Deep Learning Model	83
Summary	84

6 Artificial Neurons Detecting Hot Dogs 85

Biological Neuroanatomy 101	85
The Perceptron	86
The Hot Dog / Not Hot Dog Detector	86
The Most Important Equation in This Book	90

Modern Neurons and Activation Functions	91
The Sigmoid Neuron	92
The Tanh Neuron	94
ReLU: Rectified Linear Units	94
Choosing a Neuron	96
Summary	96
Key Concepts	97
7 Artificial Neural Networks	99
The Input Layer	99
Dense Layers	99
A Hot Dog-Detecting Dense Network	101
Forward Propagation Through the First Hidden Layer	102
Forward Propagation Through Subsequent Layers	103
The Softmax Layer of a Fast Food-Classifying Network	106
Revisiting Our Shallow Network	108
Summary	110
Key Concepts	110
8 Training Deep Networks	111
Cost Functions	111
Quadratic Cost	112
Saturated Neurons	112
Cross-Entropy Cost	113
Optimization: Learning to Minimize Cost	115
Gradient Descent	115
Learning Rate	117
Batch Size and Stochastic Gradient Descent	119
Escaping the Local Minimum	122
Backpropagation	124
Tuning Hidden-Layer Count and Neuron Count	125
An Intermediate Net in Keras	127
Summary	129
Key Concepts	130

9 Improving Deep Networks 131

- Weight Initialization 131
 - Xavier Glorot Distributions 135
- Unstable Gradients 137
 - Vanishing Gradients 137
 - Exploding Gradients 138
 - Batch Normalization 138
- Model Generalization (Avoiding Overfitting) 140
 - L1 and L2 Regularization 141
 - Dropout 142
 - Data Augmentation 145
- Fancy Optimizers 145
 - Momentum 145
 - Nesterov Momentum 146
 - AdaGrad 146
 - AdaDelta and RMSProp 146
 - Adam 147
- A Deep Neural Network in Keras 147
- Regression 149
- TensorBoard 152
- Summary 154
- Key Concepts 155

III Interactive Applications of Deep Learning 157**10 Machine Vision 159**

- Convolutional Neural Networks 159
 - The Two-Dimensional Structure of Visual Imagery 159
 - Computational Complexity 160
 - Convolutional Layers 160
 - Multiple Filters 162
 - A Convolutional Example 163
 - Convolutional Filter
 - Hyperparameters 168

- Pooling Layers 169
- LeNet-5 in Keras 171
- AlexNet and VGGNet in Keras 176
- Residual Networks 179
 - Vanishing Gradients: The Bête Noire of Deep CNNs 179
 - Residual Connections 180
 - ResNet 182
- Applications of Machine Vision 182
 - Object Detection 183
 - Image Segmentation 186
 - Transfer Learning 188
 - Capsule Networks 192
- Summary 193
- Key Concepts 193

11 Natural Language Processing 195

- Preprocessing Natural Language Data 195
 - Tokenization 197
 - Converting All Characters to Lowercase 199
 - Removing Stop Words and Punctuation 200
 - Stemming 201
 - Handling n -grams 202
 - Preprocessing the Full Corpus 203
- Creating Word Embeddings with word2vec 206
 - The Essential Theory Behind word2vec 206
 - Evaluating Word Vectors 209
 - Running word2vec 209
 - Plotting Word Vectors 213
- The Area under the ROC Curve 217
 - The Confusion Matrix 218
 - Calculating the ROC AUC Metric 219
- Natural Language Classification with Familiar Networks 222
 - Loading the IMDb Film Reviews 222
 - Examining the IMDb Data 226
 - Standardizing the Length of the Reviews 228

Dense Network	229
Convolutional Networks	235
Networks Designed for Sequential Data	240
Recurrent Neural Networks	240
Long Short-Term Memory Units	244
Bidirectional LSTMs	247
Stacked Recurrent Models	248
Seq2seq and Attention	250
Transfer Learning in NLP	251
Non-sequential Architectures: The Keras Functional API	251
Summary	256
Key Concepts	257

12 Generative Adversarial Networks 259

Essential GAN Theory	259
The Quick, Draw! Dataset	263
The Discriminator Network	266
The Generator Network	269
The Adversarial Network	272
GAN Training	275
Summary	281
Key Concepts	282

13 Deep Reinforcement Learning 283

Essential Theory of Reinforcement Learning	283
The Cart-Pole Game	284
Markov Decision Processes	286
The Optimal Policy	288
Essential Theory of Deep Q-Learning Networks	290
Value Functions	291
Q-Value Functions	291
Estimating an Optimal Q-Value	291
Defining a DQN Agent	293
Initialization Parameters	295

Building the Agent's Neural Network Model	297
Remembering Gameplay	298
Training via Memory Replay	298
Selecting an Action to Take	299
Saving and Loading Model Parameters	300
Interacting with an OpenAI Gym Environment	300
Hyperparameter Optimization with SLM Lab	303
Agents Beyond DQN	306
Policy Gradients and the REINFORCE Algorithm	307
The Actor-Critic Algorithm	307
Summary	308
Key Concepts	309

IV You and AI 311

14 Moving Forward with Your Own Deep Learning Projects 313

Ideas for Deep Learning Projects	313
Machine Vision and GANs	313
Natural Language Processing	315
Deep Reinforcement Learning	316
Converting an Existing Machine Learning Project	316
Resources for Further Projects	317
Socially Beneficial Projects	318
The Modeling Process, Including Hyperparameter Tuning	318
Automation of Hyperparameter Search	321
Deep Learning Libraries	321
Keras and TensorFlow	321
PyTorch	323
MXNet, CNTK, Caffe, and So On	324
Software 2.0	324
Approaching Artificial General Intelligence	326
Summary	328

V Appendices 331**A Formal Neural Network Notation 333****B Backpropagation 335****C PyTorch 339**

PyTorch Features 339

Autograd System 339

Define-by-Run Framework 339

PyTorch Versus TensorFlow 340

PyTorch in Practice 341

PyTorch Installation 341

The Fundamental Units Within

PyTorch 341

Building a Deep Neural Network in

PyTorch 343

Index 345