

Preface	p. xiii
Acknowledgments	p. xv
About this book	p. xvi
About the authors	p. xx
About the cover	p. xxi
Fundamentals of Deep Learning	p. 1
What is deep learning?	p. 3
Artificial intelligence, machine learning, and deep learning	p. 4
Artificial intelligence	p. 4
Machine learning	p. 4
Learning representations from data	p. 6
The "deep" in deep learning	p. 8
Understanding how deep learning works, in three figures	p. 9
What deep learning has achieved so far	p. 11
Don't believe the short-term hype	p. 12
The promise of AI	p. 12
Before deep learning: a brief history of machine learning	p. 13
Probabilistic modeling	p. 14
Early neural networks	p. 14
Kernel methods	p. 15
Decision trees, random forests, and gradient boosting machines	p. 16
Back to neural networks	p. 17
What makes deep learning different	p. 17
The modern machine-learning landscape	p. 18
Why deep learning? Why now?	p. 19
Hardware	p. 19
Data	p. 20
Algorithms	p. 21
A new wave of investment	p. 21
The democratization of deep learning	p. 22
Will it last?	p. 22
Before we begin: the mathematical building blocks of neural networks	p. 24
A first look at a neural network	p. 25
Data representations for neural networks	p. 29
Scalars (0D tensors)	p. 29
Vectors (1D tensors)	p. 29
Matrices (2D tensors)	p. 30
3D tensors and higher-dimensional tensors	p. 30
Key attributes	p. 30
Manipulating tensors in R	p. 31
The notion of data batches	p. 32

Real-world examples of data tensors	p. 32
Vector data	p. 32
Timeseries data or sequence data	p. 33
Image data	p. 33
Video data	p. 34
The gears of neural networks: tensor operations	p. 34
Element-wise operations	p. 35
Operations involving tensors of different dimensions	p. 36
Tensor dot	p. 36
Tensor reshaping	p. 38
Geometric interpretation of tensor operations	p. 39
A geometric interpretation of deep learning	p. 40
The engine of neural networks: gradient-based optimization	p. 41
What's a derivative?	p. 42
Derivative of a tensor operation: the gradient	p. 43
Stochastic gradient descent	p. 44
Chaining derivatives: the Backpropagation algorithm	p. 46
Looking back at our first example	p. 47
Summary	p. 49
Getting started with neural networks	p. 50
Anatomy of a neural network	p. 51
Layers: the building blocks of deep learning	p. 52
Models: networks of layers	p. 52
Loss functions and optimizers: keys to configuring the learning process	p. 53
Introduction to Keras	p. 54
Kerns, TensorFlow, Theano, and CNTK	p. 54
Installing Keras	p. 56
Developing with Kerns: a quick overview	p. 56
Setting up a deep-learning workstation	p. 57
Getting Keras running: two options	p. 58
Running deep-learning jobs in the cloud: pros and cons	p. 58
What is the best GPU for deep learning?	p. 59
Classifying movie reviews: a binary classification example	p. 59
The IMDB dataset	p. 59
Preparing the data	p. 61
Building your network	p. 62
Validating your approach	p. 65
Using a trained network to generate predictions on new data	p. 68
Further experiments	p. 69
Wrapping up	p. 69
Classifying newswires: a multiclass classification example	p. 70

The Reuters dataset	p. 70
Preparing the data	p. 71
Building your network	p. 72
Validating your approach	p. 73
Generating predictions on new data	p. 74
A different way to handle the labels and the loss	p. 75
The importance of having sufficiently large intermediate layers	p. 75
Further experiments	p. 76
Wrapping up	p. 76
Predicting house prices: a regression example	p. 76
The Boston Housing Price dataset	p. 77
Preparing the data	p. 77
Building your network	p. 78
Validating your approach using K-fold validation	p. 79
Wrapping up	p. 83
Summary	p. 83
Fundamentals of machine learning	p. 84
Four branches of machine learning	p. 85
Supervised learning	p. 85
Unsupervised learning	p. 85
Self-supervised learning	p. 86
Reinforcement learning	p. 86
Evaluating machine-learning models	p. 87
Training, validation, and test sets	p. 88
Things to keep in mind	p. 91
Data preprocessing, feature engineering, and feature learning	p. 91
Data preprocessing for neural networks	p. 91
Feature engineering	p. 93
Overfitting and underfitting	p. 94
Reducing the network's size	p. 95
Adding weight regularization	p. 98
Adding dropout	p. 100
The universal workflow of machine learning	p. 102
Defining the problem and assembling a dataset	p. 102
Choosing a measure of success	p. 103
Deciding on an evaluation protocol	p. 104
Preparing your data	p. 104
Developing a model that does better than a baseline	p. 104
Scaling up: developing a model that overfits	p. 105
Regularizing your model and tuning your hyperparameters	p. 106
Summary	p. 107

Deep Learning in Practice	p. 109
Deep learning for computer vision	p. 111
Introduction to convnets	p. 111
The convolution operation	p. 114
The max-pooling operation	p. 119
Training a convnet from scratch on a small dataset	p. 121
The relevance of deep learning for small-data problems	p. 121
Downloading the data	p. 122
Building your network	p. 124
Data preprocessing	p. 126
Using data augmentation	p. 128
Using a pretrained convnet	p. 132
Feature extraction	p. 133
Fine-tuning	p. 142
Wrapping up	p. 146
Visualizing what convnets learn	p. 146
Visualizing intermediate activations	p. 146
Visualizing convnet filters	p. 153
Visualizing heatmaps of class activation	p. 159
Summary	p. 163
Deep learning for text and sequences	p. 164
Working with text data	p. 165
One-hot encoding of words and characters	p. 166
Using word embeddings	p. 169
Putting it all together: from raw text to word embeddings	p. 174
Wrapping up	p. 180
Understanding recurrent neural networks	p. 180
A recurrent layer in Keras	p. 182
Understanding the LSTM and GRU layers	p. 186
A concrete LSTM example in Keras	p. 188
Wrapping up	p. 190
Advanced use of recurrent neural networks	p. 190
A temperature-forecasting problem	p. 191
Preparing the data	p. 193
A common-sense, non-machine-learning baseline	p. 197
A basic machine-learning approach	p. 198
A first recurrent baseline	p. 199
Using recurrent dropout to fight overfitting	p. 201
Stacking recurrent layers	p. 202
Using bidirectional RNNs	p. 204
Going even further	p. 207

Wrapping up	p. 208
Sequence processing with convnets	p. 209
Understanding 1D convolution for sequence data	p. 209
1D pooling for sequence data	p. 210
Implementing a 1D convnet	p. 210
Combining CNNs and RNNs to process long sequences	p. 212
Wrapping up	p. 216
Summary	p. 216
Advanced deep-learning best practices	p. 218
Going beyond the sequential model: the Keras functional API	p. 219
Introduction to the functional API	p. 221
Multi-input models	p. 222
Multi-output models	p. 224
Directed acyclic graphs of layers	p. 227
Layer weight sharing	p. 231
Models as layers	p. 232
Wrapping up	p. 233
Inspecting and monitoring deep-learning models using Keras callbacks and TensorBoard	p. 233
Using callbacks to act on a model during training	p. 233
Introduction to TensorBoard: the TensorFlow visualization framework	p. 236
Wrapping up	p. 241
Getting the most out of your models	p. 241
Advanced architecture patterns	p. 241
Hyperparameter optimization	p. 245
Model ensembling	p. 246
Wrapping up	p. 248
Summary	p. 249
Generative deep learning	p. 250
Text generation with LSTM	p. 252
A brief history of generative recurrent networks	p. 252
How do you generate sequence data?	p. 253
The importance of the sampling strategy	p. 253
Implementing character-level LSTM text generation	p. 255
Wrapping up	p. 260
DeepDream	p. 260
Implementing DeepDream in Keras	p. 261
Wrapping up	p. 267
Neural style transfer	p. 267
The content loss	p. 268
The style loss	p. 268
Neural style transfer in Keras	p. 269

Wrapping up	p. 274
Generating images with variational autoencoders	p. 276
Sampling from latent spaces of images	p. 276
Concept vectors for image editing	p. 277
Variational autoencoders	p. 278
Wrapping up	p. 284
Introduction to generative adversarial networks	p. 284
A schematic GAN implementation	p. 286
A bag of tricks	p. 286
The generator	p. 287
The discriminator	p. 288
The adversarial network	p. 289
How to train your DCGAN	p. 290
Wrapping up	p. 292
Summary	p. 292
Conclusions	p. 293
Key concepts in review	p. 294
Various approaches to AI	p. 294
What makes deep learning special within the field of machine learning	p. 294
How to think about deep learning	p. 295
Key enabling technologies	p. 296
The universal machine-learning workflow	p. 297
Key network architectures	p. 298
The space of possibilities	p. 302
The limitations of deep learning	p. 303
The risk of anthropomorphizing machine-learning models	p. 304
Local generalization vs. extreme generalization	p. 306
Wrapping up	p. 307
The future of deep learning	p. 307
Models as programs	p. 308
Beyond backpropagation and differentiable layers	p. 310
Automated machine learning	p. 310
Lifelong learning and modular subroutine reuse	p. 311
The long-term vision	p. 313
Staying up to date in a fast-moving field	p. 313
Practice on real-world problems using Kaggle	p. 314
Read about the latest developments on arXiv	p. 314
Explore the Keras ecosystem	p. 315
Final words	p. 315
Installing Keras and its dependencies on Ubuntu	p. 316
Running RStudio Server on an EC2 GPU instance	p. 320

