

Deep Learning with Keras

Implementing deep learning models and neural networks with the power of Python

Antonio Gulli

Sujit Pal

Packt

BIRMINGHAM - MUMBAI

Table of Contents

Preface	1
Chapter 1: Neural Networks Foundations	9
Perceptron	11
The first example of Keras code	11
Multilayer perceptron — the first example of a network	12
Problems in training the perceptron and a solution	13
Activation function — sigmoid	14
Activation function — ReLU	15
Activation functions	15
A real example — recognizing handwritten digits	16
One-hot encoding — OHE	17
Defining a simple neural net in Keras	17
Running a simple Keras net and establishing a baseline	21
Improving the simple net in Keras with hidden layers	22
Further improving the simple net in Keras with dropout	25
Testing different optimizers in Keras	28
Increasing the number of epochs	34
Controlling the optimizer learning rate	34
Increasing the number of internal hidden neurons	35
Increasing the size of batch computation	37
Summarizing the experiments run for recognizing handwritten charts	37
Adopting regularization for avoiding overfitting	38
Hyperparameters tuning	40
Predicting output	40
A practical overview of backpropagation	40
Towards a deep learning approach	42
Summary	43
Chapter 2: Keras Installation and API	45
Installing Keras	46
Step 1 — install some useful dependencies	46
Step 2 — install Theano	47
Step 3 — install TensorFlow	47
Step 4 — install Keras	48
Step 5 — testing Theano, TensorFlow, and Keras	48

Configuring Keras	49
Installing Keras on Docker	50
Installing Keras on Google Cloud ML	53
Installing Keras on Amazon AWS	56
Installing Keras on Microsoft Azure	58
Keras API	60
Getting started with Keras architecture	60
What is a tensor?	60
Composing models in Keras	61
Sequential composition	61
Functional composition	61
An overview of predefined neural network layers	61
Regular dense	62
Recurrent neural networks — simple, LSTM, and GRU	62
Convolutional and pooling layers	63
Regularization	63
Batch normalization	64
An overview of predefined activation functions	64
An overview of losses functions	65
An overview of metrics	66
An overview of optimizers	66
Some useful operations	66
Saving and loading the weights and the architecture of a model	66
Callbacks for customizing the training process	67
Checkpointing	68
Using TensorBoard and Keras	69
Using Quiver and Keras	70
Summary	71
Chapter 3: Deep Learning with ConvNets	73
<hr/>	
Deep convolutional neural network — DCNN	74
Local receptive fields	74
Shared weights and bias	75
Pooling layers	76
Max-pooling	76
Average pooling	77
ConvNets summary	77
An example of DCNN — LeNet	78
LeNet code in Keras	78
Understanding the power of deep learning	85
Recognizing CIFAR-10 images with deep learning	86
Improving the CIFAR-10 performance with deeper a network	91

Improving the CIFAR-10 performance with data augmentation	93
Predicting with CIFAR-10	97
Very deep convolutional networks for large-scale image recognition	98
Recognizing cats with a VGG-16 net	99
Utilizing Keras built-in VGG-16 net module	100
Recycling pre-built deep learning models for extracting features	102
Very deep inception-v3 net used for transfer learning	103
Summary	106
Chapter 4: Generative Adversarial Networks and WaveNet	107
<hr/>	
What is a GAN?	108
Some GAN applications	110
Deep convolutional generative adversarial networks	114
Keras adversarial GANs for forging MNIST	117
Keras adversarial GANs for forging CIFAR	124
WaveNet — a generative model for learning how to produce audio	132
Summary	141
Chapter 5: Word Embeddings	143
<hr/>	
Distributed representations	144
word2vec	145
The skip-gram word2vec model	146
The CBOW word2vec model	150
Extracting word2vec embeddings from the model	152
Using third-party implementations of word2vec	155
Exploring GloVe	159
Using pre-trained embeddings	161
Learn embeddings from scratch	162
Fine-tuning learned embeddings from word2vec	167
Fine-tune learned embeddings from GloVe	171
Look up embeddings	172
Summary	176
Chapter 6: Recurrent Neural Network — RNN	179
<hr/>	
SimpleRNN cells	180
SimpleRNN with Keras — generating text	182
RNN topologies	187
Vanishing and exploding gradients	188
Long short term memory — LSTM	191
LSTM with Keras — sentiment analysis	193
Gated recurrent unit — GRU	200

GRU with Keras — POS tagging	202
Bidirectional RNNs	209
Stateful RNNs	210
Stateful LSTM with Keras — predicting electricity consumption	210
Other RNN variants	217
Summary	218
Chapter 7: Additional Deep Learning Models	219
Keras functional API	221
Regression networks	223
Keras regression example — predicting benzene levels in the air	224
Unsupervised learning — autoencoders	228
Keras autoencoder example — sentence vectors	230
Composing deep networks	239
Keras example — memory network for question answering	240
Customizing Keras	247
Keras example — using the lambda layer	248
Keras example — building a custom normalization layer	249
Generative models	252
Keras example — deep dreaming	252
Keras example — style transfer	261
Summary	267
Chapter 8: AI Game Playing	269
Reinforcement learning	270
Maximizing future rewards	271
Q-learning	272
The deep Q-network as a Q-function	273
Balancing exploration with exploitation	275
Experience replay, or the value of experience	276
Example - Keras deep Q-network for catch	276
The road ahead	289
Summary	291
Appendix: Conclusion	293
Keras 2.0 — what is new	295
Installing Keras 2.0	295
API changes	296
Index	299