

Table of Contents

Preface	V
Chapter 1: Introducing Advanced Deep Learning with Keras	1
Why is Keras the perfect deep learning library?	2
Installing Keras and TensorFlow	3
Implementing the core deep learning models - MLPs, CNNs and RNNs	4
The difference between MLPs, CNNs, and RNNs	5
Multilayer perceptrons (MLPs)	6
MNIST dataset	6
MNIST digits classifier model	8
Building a model using MLPs and Keras	12
Regularization	14
Output activation and loss function	15
Optimization	17
Performance evaluation	20
Model summary	21
Convolutional neural networks (CNNs)	23
Convolution	26
Pooling operations	27
Performance evaluation and model summary	28
Recurrent neural networks (RNNs)	31
Conclusion	37
Chapter 2: Deep Neural Networks	39
Functional API	40
Creating a two-input and one-output model	43

Deep residual networks (ResNet)	49
ResNet v2	59
Densely connected convolutional networks (DenseNet)	62
Building a 100-layer DenseNet-BC for CIFAR10	65
Conclusion	68
References	69
Chapter 3: Autoencoders	71
Principles of autoencoders	72
Building autoencoders using Keras	74
Denoising autoencoder (DAE)	84
Automatic colorization autoencoder	89
Conclusion	97
References	97
Chapter 4: Generative Adversarial Networks (GANs)	99
An overview of GANs	99
Principles of GANs	100
GAN implementation in Keras	105
Conditional GAN	114
Conclusion	124
References	124
Chapter 5: Improved GANs	125
Wasserstein GAN	126
Distance functions	126
Distance function in GANs	128
Use of Wasserstein loss	131
WGAN implementation using Keras	135
Least-squares GAN (LSGAN)	142
Auxiliary classifier GAN (ACGAN)	147
Conclusion	160
References	160
Chapter 6: Disentangled Representation GANs	161
Disentangled representations	162
InfoGAN	164
Implementation of InfoGAN in Keras	167
Generator outputs of InfoGAN	177
StackedGAN	179
Implementation of StackedGAN in Keras	181
Generator outputs of StackedGAN	197
Conclusion	200
Reference	201

Chapter 7: Cross-Domain GANs	203
Principles of CycleGAN	204
The CycleGAN Model	207
Implementing CycleGAN using Keras	211
Generator outputs of CycleGAN	225
CycleGAN on MNIST and SVHN datasets	227
Conclusion	234
References	235
Chapter 8: Variational Autoencoders (VAEs)	237
Principles of VAEs	238
Variational inference	239
Core equation	240
Optimization	241
Reparameterization trick	242
Decoder testing	243
VAEs in Keras	244
Using CNNs for VAEs	249
Conditional VAE (CVAE)	254
β -VAE: VAE with disentangled latent representations	264
Conclusion	268
References	268
Chapter 9: Deep Reinforcement Learning	271
Principles of reinforcement learning (RL)	272
The Q value	274
Q-Learning example	276
Q-Learning in Python	281
Nondeterministic environment	287
Temporal-difference learning	287
Q-Learning on OpenAI gym	288
Deep Q-Network (DQN)	293
DQN on Keras	296
Double Q-Learning (DDQN)	302
Conclusion	304
References	305
Chapter 10: Policy Gradient Methods	307
Policy gradient theorem	308
Monte Carlo policy gradient (REINFORCE) method	311
REINFORCE with baseline method	313
Actor-Critic method	315
Advantage Actor-Critic (A2C) method	317

Policy Gradient methods with Keras	318
Performance evaluation of policy gradient methods	335
Conclusion	341
References	341
Other Books You May Enjoy	343
Index	347