

Antonio Gulli, Amita Kapoor

TensorFlow 1.x Deep Learning Cookbook

Over 90 unique recipes to solve artificial-intelligence
driven problems with Python



Packt>

TensorFlow 1.x Deep Learning Cookbook

Deep neural networks (DNNs) have achieved a lot of success in the field of computer vision, speech recognition, and natural language processing. This exciting recipe-based guide will take you from the realm of DNN theory to implementing them practically to solve real-life problems in the artificial intelligence domain.

In this book, you will learn how to efficiently use TensorFlow, Google's open source framework for deep learning. You will implement different deep learning networks, such as Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), Deep Q-learning Networks (DQNs), and Generative Adversarial Networks (GANs), with easy-to-follow standalone recipes. You will learn how to use TensorFlow with Keras as the backend.

You will learn how different DNNs perform on some popularly used datasets, such as MNIST, CIFAR-10, and Youtube8m. You will not only learn about the different mobile and embedded platforms supported by TensorFlow, but also how to set up cloud platforms for deep learning applications. You will also get a sneak peek at TPU architecture and how it will affect the future of DNNs.

By using crisp, no-nonsense recipes, you will become an expert in implementing deep learning techniques in growing real-world applications and research areas such as reinforcement learning, GANs, and autoencoders.

Things you will learn:

- Leverage different data sets such as MNIST, CIFAR-10, and Youtube8m with TensorFlow and learn how to access and use them in your code
- Use TensorBoard to understand neural network architectures, optimize the learning process, and peek inside the neural network black box
- Use different regression techniques for prediction and classification problems
- Build single and multilayer perceptrons in TensorFlow
- Implement a CNN and a RNN in TensorFlow, and use them to solve real-world problems
- Learn how Restricted Boltzmann Machines can be used to recommend movies
- Understand the implementation of autoencoders and deep belief networks, and use them for emotion detection
- Master the different reinforcement learning methods in order to implement game playing agents

TensorFlow 1.x Deep Learning Cookbook

Antonio Gulli, Amrita Kapoor



TensorFlow 1.x Deep Learning Cookbook

Over 90 unique recipes to solve artificial-intelligence driven problems with Python

Antonio Gulli

Amita Kapoor

Packt>

BIRMINGHAM - MUMBAI

TensorFlow 1.x Deep Learning Cookbook

Copyright © 2017 Packt Publishing

All rights reserved. No part of this book may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, without the prior written permission of the publisher, except in the case of brief quotations embedded in critical articles or reviews.

Every effort has been made in the preparation of this book to ensure the accuracy of the information presented. However, the information contained in this book is sold without warranty, either express or implied. Neither the authors, nor Packt Publishing, and its dealers and distributors will be held liable for any damages caused or alleged to be caused directly or indirectly by this book.

Packt Publishing has endeavored to provide trademark information about all of the companies and products mentioned in this book by the appropriate use of capitals. However, Packt Publishing cannot guarantee the accuracy of this information.

First published: December 2017

Production reference: 1081217

Published by Packt Publishing Ltd.

Livery Place

35 Livery Street

Birmingham

B3 2PB, UK.

ISBN 978-1-78829-359-4

www.packtpub.com

Credits

Authors

Antonio Gulli
Amita Kapoor

Copy Editors

Safis Editing
Vikrant Phadkay

Reviewers

Nick McClure
Narotam Singh
Corrado Zoccolo

Project Coordinator

Manthan Patel

Commissioning Editor

Sunith Shetty

Proofreader

Safis Editing

Acquisition Editor

Tushar Gupta

Indexer

Rekha Nair

Content Development Editor

Tejas Limkar

Graphics

Tania Dutta

Technical Editor

Danish Shaikh

Production Coordinator

Deepika Naik

About the Authors

Antonio Gulli is a transformational software executive and business leader with a passion for establishing and managing global technological talent for innovation and execution. He is an expert in search engines, online services, machine learning, information retrieval, analytics, and cloud computing. So far, he has been lucky enough to gain professional experience in four different countries in Europe and manage teams in six different countries in Europe and America. Currently, he works as site lead and director of cloud in Google Warsaw, driving European efforts for Serverless, Kubernetes, and Google Cloud UX. Previously, Antonio helped to innovate academic search as the vice president for Elsevier, a worldwide leading publisher. Before that, he drove query suggestions and news search as a principal engineer for Microsoft. Earlier, he served as the CTO for Ask.com, driving multimedia and news search. Antonio has filed for 20+ patents, published multiple academic papers, and served as a senior PC member in multiple international conferences. He truly believes that to be successful, you must have a great combination of management, research skills, just-get-it-done, and selling attitude.

I thank every reader of this book for your attention and for the trust. I am humbled by the number of comments received on LinkedIn and Facebook: You, the reader, provided immense help in making this book better. I would also like to thank various people for providing support during the process of writing the book. In no order: Susana, Ewa, Ignacy, Dawid, Max, Jarek, Jerzy, Nina, Laura, Antonella, Eric, Ettore, Francesco, Liubov, Marco, Fabio, Giacomo, Saskia, Christina, Wieland, and Yossi. I am very grateful to my coauthor, Amita, for her valuable comments and suggestions. I am extremely thankful to the reviewers of this book, Eric Brewer, Corrado Zoccolo, and Sujit Pal, for going through the entire book content. Special thanks to my manager, Eyal, for supporting me during the writing process and for the trust constantly offered. Part of this book has been written in Charlotte Menora (<http://bistrocharlotte.pl/>), a pub in Warsaw, where I found myself writing pages after work. This is an inspirational place, which I definitively recommend if you are visiting Poland. Modern and cool as the city of Warsaw is these days. Last and not the least, I am grateful to the entire editorial team of Packt, especially Tushar Gupta and Tejas Limkar for all the support, constant reminders regarding the schedule, and continuous motivation. Thanks for your patience.

Amita Kapoor is an associate professor in the Department of Electronics, SRCASW, University of Delhi. She has been actively teaching neural networks for the last 20 years. She did her master's in electronics in 1996 and PhD in 2011. During her PhD, she was awarded the prestigious DAAD fellowship to pursue a part of her research work in Karlsruhe Institute of Technology, Karlsruhe, Germany. She had been awarded the best presentation award at International Conference Photonics 2008 for her paper. She is a member of professional bodies such as OSA (Optical Society of America), IEEE (Institute of Electrical and Electronics Engineers), INNS (International Neural Network Society), and ISBS (Indian Society for Buddhist Studies). Amita has more than 40 publications in international journals and conferences to her credit. Her present research areas include machine learning, artificial intelligence, neural networks, robotics, Buddhism (philosophy and psychology) and ethics in AI.

This book is an attempt to summarize what all I had learned in the field of deep neural networks. I have presented it in a manner that readers find easy to understand and apply, and so the prime motivation of this book comes from you, the readers. I thank every reader of this book for being consistently present at the back of my mind, especially when I felt lazy. I would also like to thank professor Parongama Sen, University of Calcutta, for introducing me to the subject in 1994, and my friends Nirjara Jain and Shubha Swaminathan for the hours spent in college library discussing Asimov, his stories, and the future that neural networks behold for our society. I am very grateful to my coauthor, Antonio Guili, for his valuable comments and suggestions and the reviewers of this book, Narotam Singh and Nick McClure, for painstakingly going through the entire content and rechecking the codes. Last and not the least, I am grateful to the entire editorial team of Packt, especially Tushar Gupta and Tejas Limkar for all the support, constant reminders regarding schedule, and continuous motivation.

About the Reviewers

Narotam Singh has been with India Meteorological Department, Ministry of Earth Sciences, India, since 1996. He has been actively involved with various technical programs and training of officers of GoI in the field of information technology and communication. He did his post-graduation in the field of electronics in 1996 and both diploma and post-graduate diploma in the field of computer engineering in 1994 and 1997 respectively. He is currently working in the enigmatic field of neural networks, machine learning, and deep learning.

Nick McClure is currently a senior data scientist at PayScale Inc in Seattle, Washington, USA. Previously, he worked at Zillow and Caesar's Entertainment. He has degrees in applied mathematics from the University of Montana and the College of Saint Benedict and Saint John's University. Nick has also authored *TensorFlow Machine Learning Cookbook* by Packt Publishing.

He has a passion for learning and advocating for analytics, machine learning, and artificial intelligence. Nick occasionally puts his thoughts and musings on his blog, fromdata.org, or through his Twitter account at [@nfmccclure](https://twitter.com/nfmccclure).

Corrado Zoccolo is a Senior Staff Software Engineer at Google, with 10+ years of experience in distributed indexing and information retrieval systems.

I would like to show my gratitude to the many people who helped me succeed in my studies and my career, in particular: my wife Ermelinda, that supported me over the many years I dedicated to computer science; Prof. Marco Vanneschi who introduced me to the beautiful world of Distributed Systems; my first manager at Google, Peter Dickman, who set me on the right track for my career, and all my colleagues from which I keep learning every day.

www.PacktPub.com

For support files and downloads related to your book, please visit www.PacktPub.com. Did you know that Packt offers eBook versions of every book published, with PDF and ePub files available? You can upgrade to the eBook version at www.PacktPub.com and as a print book customer, you are entitled to a discount on the eBook copy. Get in touch with us at service@packtpub.com for more details. At www.PacktPub.com, you can also read a collection of free technical articles, sign up for a range of free newsletters and receive exclusive discounts and offers on Packt books and eBooks.



<https://www.packtpub.com/mapt>

Get the most in-demand software skills with Mapt. Mapt gives you full access to all Packt books and video courses, as well as industry-leading tools to help you plan your personal development and advance your career.

Why subscribe?

- Fully searchable across every book published by Packt
- Copy and paste, print, and bookmark content
- On demand and accessible via a web browser

Customer Feedback

Thanks for purchasing this Packt book. At Packt, quality is at the heart of our editorial process. To help us improve, please leave us an honest review on this book's Amazon page at www.amazon.com/dp/1788293592.

If you'd like to join our team of regular reviewers, you can email us at customerreviews@packtpub.com. We award our regular reviewers with free eBooks and videos in exchange for their valuable feedback. Help us be relentless in improving our products!

This book is dedicated to my son Lorenzo - for your growing passion for algorithms - to my son Leonardo - for your curiosity to explore the world -, and to my daughter Aurora - for your love. A special thought to Elio and Maria, my father and my mother.

All my royalties will go to theschoolfund.org a non-profit organization devoted to helping fund scholarships for high potential students who cannot afford to attend school.

- Antonio Gulli

This book is dedicated to the loving memory of my mother Late Smt Swarnlata Kapoor, who instilled in me the value of education, and motivated me to continue learning despite odds, and my father Mr Anil Mohan Kapoor for all his love and support.

A part of my royalties will go to smilefoundation.org a non-profit organization based in India working in the welfare projects on education, healthcare, livelihood, and women empowerment in remote villages and slums across different states of India.

- Amita Kapoor

Table of Contents

Preface	1
Chapter 1: TensorFlow - An Introduction	9
Introduction	10
Installing TensorFlow	12
Getting ready	12
How to do it...	13
How it works...	15
There's more...	16
Hello world in TensorFlow	17
How to do it...	17
How it works...	19
Understanding the TensorFlow program structure	20
How to do it...	20
How it works...	21
There's more...	22
Working with constants, variables, and placeholders	23
How to do it...	24
How it works...	28
There's more...	29
Performing matrix manipulations using TensorFlow	30
How to do it...	30
How it works...	32
There's more...	32
Using a data flow graph	32
How to do it...	32
Migrating from 0.x to 1.x	34
How to do it...	34
There's more...	35
Using XLA to enhance computational performance	35
Getting ready	36
How to do it...	37
Invoking CPU/GPU devices	38
How to do it...	38
How it works...	41

TensorFlow for Deep Learning	41
How to do it...	43
There's more	47
Different Python packages required for DNN-based problems	47
How to do it...	48
See also	49
Chapter 2: Regression	51
Introduction	51
Choosing loss functions	54
Getting ready	54
How to do it...	55
How it works...	56
There's more...	57
Optimizers in TensorFlow	57
Getting ready	57
How to do it...	58
There's more...	60
See also	60
Reading from CSV files and preprocessing data	61
Getting ready	61
How to do it...	62
There's more...	64
House price estimation-simple linear regression	64
Getting ready	64
How to do it...	64
How it works...	67
There's more...	69
House price estimation-multiple linear regression	70
How to do it...	70
How it works...	74
There's more...	74
Logistic regression on the MNIST dataset	75
How to do it...	75
How it works...	78
See also	81
Chapter 3: Neural Networks - Perceptron	83
Introduction	83
Activation functions	87

Getting ready	87
How to do it...	88
How it works...	93
There's more...	94
See also	94
Single layer perceptron	94
Getting ready	95
How to do it...	95
There's more...	97
Calculating gradients of backpropagation algorithm	97
Getting ready	98
How to do it...	99
How it works...	102
There's more...	102
See also	102
MNIST classifier using MLP	102
Getting ready	102
How to do it...	103
How it works...	105
Function approximation using MLP-predicting Boston house prices	106
Getting ready	106
How to do it...	106
How it works...	110
There's more...	111
Tuning hyperparameters	112
How to do it...	112
There's more...	114
See also	114
Higher-level APIs-Keras	114
How to do it...	115
There's more...	116
See also	116
Chapter 4: Convolutional Neural Networks	117
Introduction	117
Local receptive fields	118
Shared weights and bias	118
A mathematical example	118
ConvNets in TensorFlow	119

Pooling layers	120
Max pooling	120
Average pooling	121
ConvNets summary	122
Creating a ConvNet to classify handwritten MNIST numbers	122
Getting ready	122
How to do it...	123
How it works...	129
Creating a ConvNet to classify CIFAR-10	129
Getting ready	129
How to do it...	130
How it works...	132
There's more...	133
Transferring style with VGG19 for image repainting	133
Getting ready	133
How to do it...	134
How it works...	144
There's more...	144
Using a pretrained VGG16 net for transfer learning	145
Getting ready	145
How to do it...	146
How it works...	149
There's more...	149
Creating a DeepDream network	150
Getting ready	150
How to do it...	151
How it works...	154
There's more...	155
See also	155
Chapter 5: Advanced Convolutional Neural Networks	157
Introduction	157
Creating a ConvNet for Sentiment Analysis	158
Getting ready	158
How to do it...	159
How it works...	161
There is more...	162
Inspecting what filters a VGG pre-built network has learned	163
Getting ready	163