Armando Fandango 著

精通TensorFlow 1.x

（影印版）

Mastering TensorFlow 1.x

东南大学出版社
SOUTHEAST UNIVERSITY PRESS

Packt

精通 TensorFlow 1.x(影印版)

Mastering TensorFlow 1.x

Armando Fandango 著

南京　东南大学出版社

图书在版编目(CIP)数据

精通 TensorFlow 1.x：英文/(美)阿曼多·范丹戈(Armando Fandango)著. —影印本. —南京：东南大学出版社，2019.3

书名原文：Mastering TensorFlow 1.x

ISBN 978-7-5641-8292-2

Ⅰ.①精… Ⅱ.①阿… Ⅲ.①人工智能-算法-英文 Ⅳ.①TP18

中国版本图书馆 CIP 数据核字(2019)第 025336 号

图字：10-2018-493 号

精通 TensorFlow 1.x(影印版)

出版发行：东南大学出版社
地　　址：南京四牌楼 2 号　　邮编：210096
出 版 人：江建中
网　　址：http://www.seupress.com
电子邮件：press@seupress.com
印　　刷：常州市武进第三印刷有限公司
开　　本：787 毫米×980 毫米　　16 开本
印　　张：29.75
字　　数：583 千字
版　　次：2019 年 3 月第 1 版
印　　次：2019 年 3 月第 1 次印刷
书　　号：ISBN 978-7-5641-8292-2
定　　价：108.00 元

本社图书若有印装质量问题，请直接与营销部联系。电话(传真)：025-83791830

mapt.io

Mapt is an online digital library that gives you full access to over 5,000 books and videos, as well as industry leading tools to help you plan your personal development and advance your career. For more information, please visit our website.

Why subscribe?

- Spend less time learning and more time coding with practical eBooks and Videos from over 4,000 industry professionals
- Improve your learning with Skill Plans built especially for you
- Get a free eBook or video every month
- Mapt is fully searchable
- Copy and paste, print, and bookmark content

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.

Foreword

TensorFlow and Keras are a key part of the "Data Science for Internet of Things" course, which I teach at the University of Oxford. My TensorFlow journey started with Keras. Over time, in our course, we increasingly gravitated towards core TensorFlow in addition to Keras. I believe many people's 'TensorFlow journey' will follow this trajectory.

Armando Fandango's book "Mastering TensorFlow 1.x" provides a road map for this journey. The book is an ambitious undertaking, interweaving Keras and core TensorFlow libraries. It delves into complex themes and libraries such as Sonnet, distributed TensorFlow with TF Clusters, deploying production models with TensorFlow Serving, TensorFlow mobile, and TensorFlow for embedded devices.

In that sense, this is an advanced book. But the author covers deep learning models such as RNN, CNN, autoencoders, generative adversarial models, and deep reinforcement learning through Keras. Armando has clearly drawn upon his experience to make this complex journey easier for readers.

I look forward to increased adoption of this book and learning from it.

Ajit Jaokar

Data Science for IoT Course Creator and Lead Tutor at the University of Oxford / Principal Data Scientist.

Contributors

About the author

Armando Fandango creates AI-empowered products by leveraging his expertise in deep learning, computational methods, and distributed computing. He advises Owen.ai Inc on AI product strategy. He founded NeuraSights Inc. with the goal of creating insights using neural networks. He is the founder of Vets2Data Inc., a non-profit organization assisting US military veterans in building AI skills.

Armando has authored books titled Python Data Analysis - 2nd Edition and Mastering TensorFlow and published research in international journals and conferences.

I would like to thank Dr. Paul Wiegand (UCF), Dr. Brian Goldiez (UCF), Tejas Limkar (Packt), and Tushar Gupta (Packt) for being able to complete this book. This work would not be possible without their inspiration.

About the reviewer

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. He has also authored TensorFlow Machine Learning Cookbook by Packt.

He has a passion for learning and advocating for analytics, machine learning, and artificial intelligence. he occasionally puts his thoughts and musings on his blog, `fromdata.org`, or through his Twitter account at `@nfmcclure`.

Packt is searching for authors like you

If you're interested in becoming an author for Packt, please visit `authors.packtpub.com` and apply today. We have worked with thousands of developers and tech professionals, just like you, to help them share their insight with the global tech community. You can make a general application, apply for a specific hot topic that we are recruiting an author for, or submit your own idea.

Table of Contents

Preface

Google's TensorFlow has become a major player and a go-to tool for developers to bring smart processing within an application. TensorFlow has become a major research and engineering tool in every organization. Thus, there is a need to learn advanced use cases of TensorFlow that can be implemented in all kinds of software and devices to build intelligent systems. TensorFlow is one of its kind, with lots of new updates and bug fixes to bring smart automation into your projects. So in today's world, it becomes a necessity to master TensorFlow in order to create advanced machine learning and deep learning applications. *Mastering TensorFlow* will help you learn all the advanced features TensorFlow has to offer. This book funnels down the key information to provide the required expertise to the readers to enter the world of artificial intelligence, thus extending the knowledge of intermediate TensorFlow users to the next level. From implementing advanced computations to trending real-world research areas, this book covers it all. Get to the grips with this highly comprehensive guide to make yourself well established in the developer community, and you'll have a platform to contribute to research works or projects.

Who this book is for

This book is for anyone who wants to build or upgrade their skills in applying TensorFlow to deep learning problems. Those who are looking for an easy-to-follow guide that underlines the intricacies and complex use cases of deep learning will find this book useful. A basic understanding of TensorFlow and Python is required to get the most out of the book.

What this book covers

`Chapter 1`, *TensorFlow 101*, recaps the basics of TensorFlow, such as how to create tensors, constants, variables, placeholders, and operations. We learn about computation graphs and how to place computation graph nodes on various devices such as GPU. We also learn how to use TensorBoard to visualize various intermediate and final output values.

`Chapter 2`, *High-Level Libraries for TensorFLow*, covers several high-level libraries such as TF Contrib Learn, TF Slim, TFLearn, Sonnet, and Pretty Tensor.

`Chapter 3`, *Keras 101*, gives a detailed overview of the high-level library Keras, which is now part of the TensorFlow core.

`Chapter 4`, *Classical Machine Learning with TensorFlow*, teaches us to use TensorFlow to implement classical machine learning algorithms, such as linear regression and classification with logistic regression.

`Chapter 5`, *Neural Networks and MLP with TensorFlow and Keras*, introduces the concept of neural networks and shows how to build simple neural network models. We also cover how to build deep neural network models known as MultiLayer Perceptrons.

`Chapter 6`, *RNNs with TensorFlow and Keras*, covers how to build Recurrent Neural Networks with TensorFlow and Keras. We cover the internal architecture of RNN, Long Short-Term Networks (LSTM), and Gated Recurrent Units (GRU). We provide a brief overview of the API functions and classes provided by TensorFlow and Keras to implement RNN models.

`Chapter 7`, *RNN for Time Series Data with TensorFlow and Keras*, shows how to build and train RNN models for time series data and provide examples in TensorFlow and Keras libraries.

`Chapter 8`, *RNN for Text Data with TensorFlow and Keras*, teaches us how to build and train RNN models for text data and provides examples in TensorFlow and Keras libraries. We learn to build word vectors and embeddings with TensorFlow and Keras, followed by LSTM models for using embeddings to generate text from sample text data.

`Chapter 9`, *CNN with TensorFlow and Keras*, covers CNN models for image data and provides examples in TensorFlow and Keras libraries. We implement the LeNet architecture pattern for our example.

`Chapter 10`, *Autoencoder with TensorFlow and Keras*, illustrates the Autoencoder models for image data and again provides examples in TensorFlow and Keras libraries. We show the implementation of Simple Autoencoder, Denoising Autoencoder, and Variational Autoencoders.

`Chapter 11`, *TensorFlow Models in Production with TF Serving*, teaches us to deploy the models with TensorFlow Serving. We learn how to deploy using TF Serving in Docker containers and Kubernetes clusters.

`Chapter 12`, *Transfer Learning and Pre-Trained Models*, shows the use of pretrained models for predictions. We learn how to retrain the models on a different dataset. We provide examples to apply the VGG16 and Inception V3 models, pretrained on the ImageNet dataset, to predict images in the COCO dataset. We also show examples of retraining only the last layer of the models with the COCO dataset to improve the predictions.

`Chapter 13`, *Deep Reinforcement Learning*, covers reinforcement learning and the OpenAI gym. We build and train several models using various reinforcement learning strategies, including deep Q networks.

`Chapter 14`, *Generative Adversarial Networks*, shows how to build and train generative adversarial models in TensorFLow and Keras. We provide examples of SimpleGAN and DCGAN.

`Chapter 15`, *Distributed Models with TensorFlow Clusters*, covers distributed training for TensorFLow models using TensorFLow clusters. We provide examples of asynchronous and synchronous update methods for training models in data-parallel fashion.

`Chapter 16`, *TensorFlow Models on Mobile and Embedded Platforms*, shows how to deploy TensorFlow models on mobile devices running on iOS and Android platforms. We cover both TF Mobile and TF Lite APIs of the TensorFlow Library.

`Chapter 17`, *TensorFlow and Keras in R*, covers how to build and train TensorFlow models in R statistical software. We learn about the three packages provided by R Studio that implement the TF Core, TF Estimators, and Keras API in R.

`Chapter 18`, *Debugging TensorFlow Models*, tells us strategies and techniques to find problem hotspots when the models do not work as expected. We cover TensorFlow debugger, along with other methods.

`Appendix`, *Tensor Processing Units*, gives a brief overview of Tensor Processing Units. TPUs are futuristic platforms optimized to train and run TensorFlow models. Although not widely available yet, they are available on the Google Cloud Platform and slated to be available soon outside the GCP.

To get the most out of this book

1. We assume that you are familiar with coding in Python and the basics of TensorFlow and Keras.
2. If you haven't done already, then install Jupyter Notebooks, TensorFlow, and Keras.
3. Download the code bundle for this book that contains the Python, R, and notebook code files.