Quan Hua, Shams Ul Azeem, Saif Ahmed 著

TensorFlow 1.x机器学习

（影印版）

Machine Learning with TensorFlow 1.x

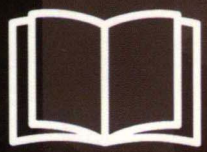

东南大学出版社
SOUTHEAST UNIVERSITY PRESS

Packt

TensorFlow 1.x 机器学习（影印版）

Machine Learning with TensorFlow 1.x

Quan Hua, Shams Ul Azeem, Saif Ahmed 著

南京　东南大学出版社

图书在版编目(CIP)数据

TensorFlow 1.x 机器学习：英文/(越)全华(Quan Hua),(巴)夏姆斯·乌尔·阿齐姆(Shams Ul Azeem),(美)赛义夫·阿哈迈德(Saif Ahmed)著. —影印本. —南京：东南大学出版社，2018.8

书名原文：Machine Learning with TensorFlow 1.x

ISBN 978-7-5641-7754-6

Ⅰ.①T… Ⅱ.①全… ②夏… ③赛… Ⅲ.①人工智能—算法—英文 Ⅳ.①TP18

中国版本图书馆 CIP 数据核字(2018)第 098984 号

图字：10-2018-098 号

TensorFlow 1.x 机器学习(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：19
字　　数：372 千字
版　　次：2018 年 8 月第 1 版
印　　次：2018 年 8 月第 1 次印刷
书　　号：ISBN 978-7-5641-7754-6
定　　价：82.00 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

Credits

Authors
Quan Hua
Shams Ul Azeem
Saif Ahmed

Reviewer

Nathan Lintz

Commissioning Editor

Kunal Parikh

Acquisition Editor

Tushar Gupta

Content Development Editor

Siddhi Chavan

Technical Editor

Mehul Singh

Copy Editor

Zainab Bootwala

Project Coordinator

Prajakta Naik

Proofreader

Safis Editing

Indexer

Rekha Nair

Graphics

Jason Monteiro

Production Coordinator

Deepika Naik

About the Authors

Quan Hua is a Computer Vision and Machine Learning Engineer at BodiData, a data platform for body measurements, where he focuses on developing computer vision and machine learning applications for a handheld technology capable of acquiring a body avatar while a person is fully clothed. He earned a bachelor of science degree from the University of Science, Vietnam, specializing in Computer Vision. He has been working in the field of computer vision and machine learning for about 3 years at start-ups.

Quan has been writing for Packt since 2015 for a Computer Vision book, *OpenCV 3 Blueprints*.

I wish to thank everyone who has encouraged me on the way while writing this book.

I want to express my sincere gratitude to my co-authors, editors, and reviewers for their advice and assistance.

I would like to thank the members of my family and my wife, Kim Ngoc, who supported and encouraged me in spite of all the time it took me away from them. They all kept me going, and this book would not have been possible without them.

I would also like to thank my teachers who gave me knowledge of Computer Vision and Machine Learning.

Shams Ul Azeem is an undergraduate in electrical engineering from NUST Islamabad, Pakistan. He has a great interest in the computer science field, and he started his journey with Android development. Now, he's pursuing his career in Machine Learning, particularly in deep learning, by doing medical-related freelancing projects with different companies.
He was also a member of the RISE lab, NUST, and he has a publication credit at the IEEE International Conference, ROBIO as a co-author of *Designing of motions for humanoid goalkeeper robots.*

Saif Ahmed is an accomplished quantitative analyst and data scientist with 15 years of industry experience. His career started in management consulting at Accenture and lead him to quantitative and senior management roles at Goldman Sachs and AIG Investments. Most recently, he co-founded and runs a start-up focused on applying Deep Learning to automating medical imaging. He obtained his bachelor's degree in computer science from Cornell University and is currently pursuing a graduate degree in data science at U.C. Berkeley.

About the Reviewer

Nathan Lintz is a Machine Learning researcher, focusing on text classification. When he began with Machine Learning, he primarily used Theano but quickly switched to TensorFlow when it was released. TensorFlow has greatly reduced the time it takes to build Machine Learning systems thanks to its intuitive and powerful neural network utilities.

I want to thank my family and professors for all the help they have given me. Without them, I would have never been able to pursue my passion for software engineering and Machine Learning.

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 `https://www.amazon.com/dp/1787123421`. 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!

Table of Contents

Preface

Machine Learning has revolutionized the modern world. Many machine learning algorithms, especially deep learning, have been used worldwide, ranging from mobile devices to cloud-based services. TensorFlow is one of the leading open source software libraries and helps you build, train, and deploy your Machine Learning system for a variety of applications. This practical book is designed to bring you the best of TensorFlow and help you build real-world Machine Learning systems.

By the end of this book, you will have a deep understanding of TensorFlow and be able to apply Machine Learning techniques to your application.

What this book covers

`Chapter 1`, *Getting Started with TensorFlow*, shows how to install Tensorflow and get started on Ubuntu, macOS, and Windows.

`Chapter 2`, *Your First Classifier*, guides you through your first journey with a handwriting recognizer.

`Chapter 3`, *The TensorFlow Toolbox*, gives you an overview of the tools that Tensorflow provides to work more effectively and easily.

`Chapter 4`, *Cats and Dogs*, teaches you how to build an image classifier using Convolutional Neural Networks in TensorFlow.

`Chapter 5`, *Sequence to Sequence Models—Parlez-vous Français?*, discusses how to build an English to French translator using sequence-to-sequence models.

`Chapter 6`, *Finding Meaning*, explores the ways to find the meaning in the text by using sentiment analysis, entity extraction, keyword extraction, and word-relation extraction.

`Chapter 7`, *Making Money with Machine Learning*, dives into an area with copious amounts of data: the financial world. You will learn how to work with the time series data to solve the financial problems.

`Chapter 8`, *The Doctor Will See You Now*, investigates ways to tackle an *enterprise-grade* problem—medical diagnosis—using deep neural networks.

`Chapter 9`, *Cruise Control - Automation*, teaches you how to create a production system, ranging from training to serving a model. The system can also receive user feedbacks and automatically train itself every day.

`Chapter 10`, *Go Live and Go Big*, guides you through the world of Amazon Web Services and shows you how to take advantage of a multiple GPUs system on Amazon servers.

`Chapter 11`, *Going Further - 21 Problems*, introduces 21 real-life problems that you can use in deep learning—TensorFlow to solve after reading this book.

`Appendix`, *Advanced Installation*, discusses GPUs and focuses on a step-by-step CUDA setup and a GPU-based TensorFlow installation.

What you need for this book

For software, the whole book is based on TensorFlow. You can use either Linux, Windows, or macOS.

For hardware, you will need a computer or laptop that runs Ubuntu, macOS, or Windows. As authors, we encourage you to have an NVIDIA graphics card if you want to work with deep neural networks, especially when you want to work with large-scale datasets.

Who this book is for

This book is ideal for you if you aspire to build Machine Learning systems that are smart and practical enough for real-world applications. You should be comfortable with Machine Learning concepts, Python programming, IDEs, and the command line. This book will be useful to people who program professionally as part of their job, or those who are working as scientists and engineers and need to learn about Machine Learning and TensorFlow in support of their work.

Conventions

In this book, you will find a number of text styles that distinguish between different kinds of information. Here are some examples of these styles and an explanation of their meaning.

Code words in text, database table names, folder names, filenames, file extensions, pathnames, dummy URLs, user input, and Twitter handles are shown as follows: "We can include other contexts through the use of the `include` directive."

A block of code is set as follows:

```
batch_size = 128
num_steps = 10000
learning_rate = 0.3
data_showing_step = 500
```

When we wish to draw your attention to a particular part of a code block, the relevant lines or items are set in bold:

```
Layer 1 CONV (32, 28, 28, 4)
Layer 2 CONV (32, 14, 14, 4)
Layer 3 CONV (32, 7, 7, 4)
```

Any command-line input or output is written as follows:

```
sudo apt-get install python-pip python-dev
```

New terms and **important words** are shown in bold.

Warnings or important notes appear like this.

Tips and tricks appear like this.

Reader feedback

Feedback from our readers is always welcome. Let us know what you think about this book—what you liked or disliked. Reader feedback is important for us as it helps us develop titles that you will really get the most out of.

To send us general feedback, simply e-mail `feedback@packtpub.com`, and mention the book's title on the subject of your message.