

Intelligent Mobile Projects with TensorFlow

# TensorFlow

## 智能移动项目

(影印版)

Jeff Tang 著

Packt

[www.packtpub.com](http://www.packtpub.com)



东南大学出版社  
SOUTHEAST UNIVERSITY PRESS

TensorFlow 智能移动项目(影印版)  
Intelligent Mobile Projects  
with TensorFlow

Jeff Tang 著

南京 东南大学出版社

## 图书在版编目(CIP)数据

TensorFlow 智能移动项目:英文/(美)杰夫·唐(Jeff Tang)著. —影印本. —南京:东南大学出版社,2019.3

书名原文: Intelligent Mobile Projects with TensorFlow

ISBN 978-7-5641-8290-8

I. ①T… II. ①杰… 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 025341 号

图字:10-2018-490 号

© 2018 by PACKT Publishing Ltd.

Reprint of the English Edition, jointly published by PACKT Publishing Ltd and Southeast University Press, 2019. Authorized reprint of the original English edition, 2018 PACKT Publishing Ltd, the owner of all rights to publish and sell the same.

All rights reserved including the rights of reproduction in whole or in part in any form.

英文原版由 PACKT Publishing Ltd 出版 2018。

英文影印版由东南大学出版社出版 2019。此影印版的出版和销售得到出版权和销售权的所有者——PACKT Publishing Ltd 的许可。

版权所有,未得书面许可,本书的任何部分和全部不得以任何形式重制。

## TensorFlow 智能移动项目(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](mailto: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:25.25

字 数:494 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 978-7-5641-8290-8

定 价:98.00 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

*To Lisa and Wozi, who showed me that unconditional love and support  
can live in harmony with occasional need for attention.*



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](http://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](mailto:service@packtpub.com) for more details.

At [www.PacktPub.com](http://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

The past decade has seen the explosion of both machine learning and smartphones; today, these technologies are finally merging, and the result is an incredible variety of applications that you would have dismissed as far future Science Fiction just a few years ago. Think about it: you have already become accustomed to talking to your phone, asking it for directions, or telling it to schedule an appointment in your agenda. Your phone's camera tracks faces and recognizes objects. Games are becoming more interesting and challenging as the bots gets smarter and smarter. And countless apps use some form of artificial intelligence under the hood, in less obvious ways, such as recommending content that you will enjoy, anticipating your next trips to tell you when to leave, suggesting what to type next, and so on.

Until recently, all the intelligence happened on the server side, which meant that the user had to be connected to the internet, ideally with a fast and stable connection. The latency and service disruptions that this implied were show-stoppers for many applications. But today the intelligence is right there in the palm of your hand, thanks to tremendous hardware improvements and better Machine Learning libraries.

Most importantly, these technologies are now completely democratized: virtually any software engineer can learn to code an intelligent mobile application based on deep neural networks, using TensorFlow, Google's powerful and open source deep learning library. Jeff Tang's great and unique book will show you how to develop on-device TensorFlow-powered iOS, Android, and Raspberry Pi apps by guiding you through many concrete examples with step-by-step tutorials and hard-earned troubleshooting tips: from image classification, object detection, image captioning, and drawing recognition to speech recognition, forecasting time series, generative adversarial networks, reinforcement learning, and even building intelligent games using AlphaZero — the improved technology built on top of AlphaGo that beat Lee Sedol and Ke Jie, the world champions of the game of Go.

This is going to be a super popular book. It's such an important topic, and it's hard to get good reliable information. So roll up your sleeves, you have an exciting journey ahead of you! What intelligent mobile application will you build?

## **Aurélien Géron**

Former lead of YouTube's video classification team and author of the book *Hands-On Machine Learning with Scikit-Learn and TensorFlow* (O'Reilly, 2017)  
Paris, May 11th, 2018

# Contributors

## About the author

**Jeff Tang** fell in love with classical AI more than two decades ago. After his MS in CS, he worked on Machine Translation for 2 years and then, to survive the long AI winter, he worked on enterprise apps, voice apps, web apps, and mobile apps at startups, AOL, Baidu, and Qualcomm. He developed a top-selling iOS app with millions of downloads and was recognized by Google as a Top Android Market Developer. He reconnected with modern AI in 2015 and knew that AI will be his passion and commitment for the next two decades. One of his favorite topics is to make AI available anytime anywhere and hence the book.

*I'd like to thank Larissa Pinto for reaching out on the book idea, Flavian Vaz and Akhil Nair for all the feedback during content editing. Many thanks to Pete Warden, the TensorFlow mobile lead at Google, for his help before and after becoming a technical reviewer of the book, and to Amita Kapoor, another technical reviewer of the book, who also provided valuable feedback. Special thanks to Aurelien Géron, the best-seller author of the book Hands-On ML, for kindly responding to all my emails, sharing his insights, and writing the Foreword for the book from his packed agenda - Merci beaucoup, Aurelien. I also truly appreciate the understanding and support by my family, other than Lisa and Wozi, during the holiday season and all the months while I had to work like crazy day and night on the book - thanks Amy, Anna, Jenny, Sophia, Mark, Sandy, and Ben.*

## About the reviewers

**Pete Warden** is the technical lead of the mobile and embedded TensorFlow group in Google's Brain team.

**Amita Kapoor** is an associate professor in the department of electronics, SRCASW, University of Delhi since 1996. She is the recipient of the prestigious DAAD fellowship to pursue a part of her research work in Karlsruhe Institute of Technology, Germany. She had been awarded the Best Presentation Award at International Conference Photonics 2008. She is a member of different professional bodies and has more than 40 publications in international journals and conferences. Her present research areas include ML, AI, Neural Networks, Robotics, Buddhism, and Ethics in AI.

## Packt is searching for authors like you

If you're interested in becoming an author for Packt, please visit [authors.packtpub.com](http://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

<b>Preface</b>	1
<b>Chapter 1: Getting Started with Mobile TensorFlow</b>	11
<b>Setting up TensorFlow</b>	12
Setting up TensorFlow on MacOS	13
Setting up TensorFlow on GPU-powered Ubuntu	15
<b>Setting up Xcode</b>	20
<b>Setting up Android Studio</b>	20
<b>TensorFlow Mobile vs TensorFlow Lite</b>	23
<b>Running sample TensorFlow iOS apps</b>	23
<b>Running sample TensorFlow Android apps</b>	24
<b>Summary</b>	25
<b>Chapter 2: Classifying Images with Transfer Learning</b>	27
<b>Transfer learning – what and why</b>	28
<b>Retraining using the Inception v3 model</b>	29
<b>Retraining using MobileNet models</b>	38
<b>Using the retrained models in the sample iOS app</b>	40
<b>Using the retrained models in the sample Android app</b>	43
<b>Adding TensorFlow to your own iOS app</b>	44
Adding TensorFlow to your Objective-C iOS app	44
Adding TensorFlow to your Swift iOS app	52
<b>Adding TensorFlow to your own Android app</b>	58
<b>Summary</b>	63
<b>Chapter 3: Detecting Objects and Their Locations</b>	65
<b>Object detection—a quick overview</b>	66
<b>Setting up the TensorFlow Object Detection API</b>	68
Quick installation and example	68
Using pre-trained models	69
<b>Retraining SSD-MobileNet and Faster RCNN models</b>	75
<b>Using object detection models in iOS</b>	81
Building TensorFlow iOS libraries manually	81
Using TensorFlow iOS libraries in an app	82
Adding an object detection feature to an iOS app	86
<b>Using YOLO2—another object-detection model</b>	92
<b>Summary</b>	96
<b>Chapter 4: Transforming Pictures with Amazing Art Styles</b>	97

<b>Neural Style Transfer – a quick overview</b>	98
<b>Training fast neural-style transfer models</b>	99
<b>Using fast neural-style transfer models in iOS</b>	102
Adding and testing with fast neural transfer models	102
Looking back at the iOS code using fast neural transfer models	106
<b>Using fast neural-style transfer models in Android</b>	107
<b>Using the TensorFlow Magenta multi-style model in iOS</b>	114
<b>Using the TensorFlow Magenta multi-style model in Android</b>	122
<b>Summary</b>	126
<b>Chapter 5: Understanding Simple Speech Commands</b>	127
<b>Speech recognition – a quick overview</b>	128
<b>Training a simple commands recognition model</b>	130
<b>Using a simple speech recognition model in Android</b>	134
Building a new app using the model	135
Showing model-powered recognition results	140
<b>Using a simple speech recognition model in iOS with Objective-C</b>	143
Building a new app using the model	143
Fixing model-loading errors with tf_op_files.txt	150
<b>Using a simple speech recognition model in iOS with Swift</b>	151
<b>Summary</b>	155
<b>Chapter 6: Describing Images in Natural Language</b>	157
<b>Image captioning – how it works</b>	158
<b>Training and freezing an image captioning model</b>	160
Training and testing caption generation	160
Freezing the image captioning model	163
<b>Transforming and optimizing the image captioning model</b>	169
Fixing errors with transformed models	170
Optimizing the transformed model	173
<b>Using the image captioning model in iOS</b>	175
<b>Using the image captioning model in Android</b>	185
<b>Summary</b>	191
<b>Chapter 7: Recognizing Drawing with CNN and LSTM</b>	193
<b>Drawing classification – how it works</b>	194
<b>Training, predicting, and preparing the drawing classification model</b>	196
Training the drawing classification model	196
Predicting with the drawing classification model	197
Preparing the drawing classification model	200
<b>Using the drawing classification model in iOS</b>	206
Building custom TensorFlow library for iOS	206
Developing an iOS app to use the model	207
<b>Using the drawing classification model in Android</b>	215

---

Building custom TensorFlow library for Android	216
Developing an Android app to use the model	218
<b>Summary</b>	227
<b>Chapter 8: Predicting Stock Price with RNN</b>	229
<b>RNN and stock price prediction – what and how</b>	230
<b>Using the TensorFlow RNN API for stock price prediction</b>	232
Training an RNN model in TensorFlow	233
Testing the TensorFlow RNN model	237
<b>Using the Keras RNN LSTM API for stock price prediction</b>	239
Training an RNN model in Keras	240
Testing the Keras RNN model	244
<b>Running the TensorFlow and Keras models on iOS</b>	246
<b>Running the TensorFlow and Keras models on Android</b>	254
<b>Summary</b>	258
<b>Chapter 9: Generating and Enhancing Images with GAN</b>	259
<b>GAN – what and why</b>	260
<b>Building and training GAN models with TensorFlow</b>	262
Basic GAN model of generating handwritten digits	262
Advanced GAN model of enhancing image resolution	265
<b>Using the GAN models in iOS</b>	269
Using the basic GAN model	272
Using the advanced GAN model	274
<b>Using the GAN models in Android</b>	277
Using the basic GAN model	280
Using the advanced GAN model	282
<b>Summary</b>	284
<b>Chapter 10: Building an AlphaZero-like Mobile Game App</b>	285
<b>AlphaZero – how does it work?</b>	286
<b>Training and testing an AlphaZero-like model for Connect 4</b>	288
Training the model	288
Testing the model	292
Looking into the model-building code	295
Freezing the model	296
<b>Using the model in iOS to play Connect 4</b>	298
<b>Using the model in Android to play Connect 4</b>	311
<b>Summary</b>	322
<b>Chapter 11: Using TensorFlow Lite and Core ML on Mobile</b>	323
<b>TensorFlow Lite – an overview</b>	323
<b>Using TensorFlow Lite in iOS</b>	325
Running the example TensorFlow Lite iOS apps	325
Using a prebuilt TensorFlow Lite model in iOS	327

Using a retrained TensorFlow model for TensorFlow Lite in iOS	332
Using a custom TensorFlow Lite model in iOS	333
<b>Using TensorFlow Lite in Android</b>	335
<b>Core ML for iOS – an overview</b>	339
<b>Using Core ML with Scikit-Learn machine learning</b>	340
Building and converting the Scikit Learn models	340
Using the converted Core ML models in iOS	343
<b>Using Core ML with Keras and TensorFlow</b>	344
<b>Summary</b>	349
<b>Chapter 12: Developing TensorFlow Apps on Raspberry Pi</b>	351
<b>Setting up Raspberry Pi and making it move</b>	352
Setting up Raspberry Pi	353
Making Raspberry Pi move	356
<b>Setting up TensorFlow on Raspberry Pi</b>	359
<b>Image recognition and text to speech</b>	361
<b>Audio recognition and robot movement</b>	364
<b>Reinforcement learning on Raspberry Pi</b>	367
Understanding the CartPole simulated environment	368
Starting with basic intuitive policy	372
Using neural networks to build a better policy	374
<b>Summary</b>	381
<b>Final words</b>	381
<b>Other Books You May Enjoy</b>	383
<b>Index</b>	387

---

# Preface

**Artificial Intelligence (AI)**, the simulation of human intelligence in computers, has a long history. Since its official birth in 1956, AI has experienced several booms and busts. The ongoing AI resurgence, or the new AI revolution, started in 2012 with the breakthrough in deep learning, a branch of machine learning that is now the hottest branch of AI because of deep learning, when a **deep convolutional neural network (DCNN)** won the ImageNet Large-Scale Visual Recognition Challenge with an error rate of only 16.4%, compared to the second best non-DCNN entry with an error rate of 26.2%. Since 2012, improved DCNN-based entries have won the ImageNet challenge every year, and deep learning technology has been applied to many hard AI problems beyond computer vision, such as speech recognition, machine translation, and the game of Go, resulting in one breakthrough after another. In March 2016, Google DeepMind's AlphaGo, built with deep reinforcement learning, beat 18-time human world Go champion Lee Sedol 4:1. At Google I/O 2017, Google announced that they're shifting from mobile-first to AI-first world. Other leading companies such as Amazon, Apple, Facebook, and Microsoft have all invested heavily in AI and launched many AI-powered products.

TensorFlow is Google's open source framework for building machine learning AI apps. Since its initial release in November 2015, when there were already several popular open source deep learning frameworks, TensorFlow has quickly become the most popular open source deep learning framework in less than 2 years. New TensorFlow models to tackle all kinds of tasks that'd require human or even superhuman intelligence have been built on a weekly basis. Dozens of books on TensorFlow have been published. More online blogs, tutorials, courses, and videos on TensorFlow have been made available. It's obvious that AI and TensorFlow are hot, but why another book with the word "TensorFlow" in its title?

This is a unique book, and the first one that combines TensorFlow-powered AI with mobile, connecting the world of the brightest future with the world of most prosperous present. We have all witnessed and experienced the iOS and Android smart phone revolution in the past decade, and we're just starting the AI revolution that'll likely have an even deeper impact on the world around us. What can be better than a theme that integrates the best of the two worlds, a book that shows how to build TensorFlow AI apps on mobile, anytime, anywhere?

It's true that you can build AI apps using many cloud AI APIs that exist out there, and sometimes it makes sense. However, running AI apps completely on mobile devices have the benefits that you can run the apps even when no network connectivity is available, when you can't afford the round trip to a cloud server, or when users don't want to send the data on their phones to anyone else.

It's also true that there're already a few example TensorFlow iOS and Android apps in the TensorFlow open source project that can get you started with mobile TensorFlow. However, if you have ever tried to run a cool TensorFlow model that amazes you on your iOS or Android device, you'll most likely stumble upon many hiccups before you can see the model successfully running on your phone.

This book can save you a lot of time and effort by showing you how to solve all the common problems you may encounter when running TensorFlow models on mobile. You'll get to see more than 10 complete TensorFlow iOS and Android apps built from scratch in the book, running all kinds of cool TensorFlow models, including the latest and coolest **Generative Adversarial Network (GAN)** and AlphaZero-like models.

## Who this book is for

If you're an iOS and/or Android developer interested in building and retraining others' cool TensorFlow models and running them in your mobile apps, or if you're a TensorFlow developer and want to run your new amazing TensorFlow models on mobile devices, the book is for you. If you're interested in TensorFlow Lite, Core ML, or TensorFlow on Raspberry Pi, you'll also benefit from the book.

## What this book covers

*Chapter 1, Getting Started with Mobile TensorFlow*, discusses how to set up TensorFlow on Mac and Ubuntu and NVIDIA GPU on Ubuntu and how to set up Xcode and Android Studio. We'll also discuss the difference between TensorFlow Mobile and TensorFlow Lite and when you should use them. Finally, we'll show you how to run the sample TensorFlow iOS and Android apps.

*Chapter 2, Classifying Images with Transfer Learning*, covers what is transfer learning and why you should use it, how to retrain the Inception v3 and MobileNet models for more accurate and faster dog breed recognition, and how to use the retrained models in sample iOS and Android apps. Then, we'll show you how to add TensorFlow to your own iOS app, both in Objective-C and Swift, and your own Android app for dog breed recognition.

Chapter 3, *Detecting Objects and Their Locations*, gives a quick overview of Object Detection, and then shows you how to set up the TensorFlow Object Detection API and use it to retrain SSD-MobileNet and Faster RCNN models. We'll also show you how to use the models used in the example TensorFlow Android app in your iOS app by manually building the TensorFlow iOS library to support non-default TensorFlow operations. Finally, we'll show you how to train YOLO2, another popular object detection model, which is also used in the example TensorFlow Android app, and how to use it in your iOS app.

Chapter 4, *Transforming Pictures with Amazing Art Styles*, first gives an overview of neural style transfer with their rapid progress in the last few years. Then, it shows you how to train fast neural style transfer models and use them in iOS and Android apps. After that, we'll cover how to use the TensorFlow Magenta multi-style model in your own iOS and Android apps to easily create amazing art styles.

Chapter 5, *Understanding Simple Speech Commands*, outlines speech recognition and shows you how to train a simple speech commands recognition model. We'll then show you how to use the model in Android as well as in iOS using both Objective-C and Swift. We'll also cover more tips on how to fix possible model loading and running errors on mobile.

Chapter 6, *Describing Images in Natural Language*, describes how image captioning works, and then it covers how to train and freeze an image captioning model in TensorFlow. We'll further discuss how to transform and optimize the complicated model to get it ready for running on mobile. Finally, we'll offer complete iOS and Android apps using the model to generate natural language description of images.

Chapter 7, *Recognizing Drawing with CNN and LSTM*, explains how drawing classification works, and discusses how to train, predict, and prepare the model. Then, we'll show you how to build another custom TensorFlow iOS library to use the model in a fun iOS doodling app. Finally, we'll show you how to build a custom TensorFlow Android library to fix a new model loading error and then use the model in your own Android app.

Chapter 8, *Predicting Stock Price with RNN*, takes you through RNN and how to use it to predict stock prices. Then, we'll inform you of how to build an RNN model with the TensorFlow API to predict stock prices, and how to build a RNN LSTM model with the easier-to-use Keras API to achieve the same goal. We'll test and see whether such models can beat a random buy or sell strategy. Finally, we'll show you how to run the TensorFlow and Keras models in both iOS and Android apps.

Chapter 9, *Generating and Enhancing Images with GAN*, gives an overview of what GAN is and why it has such great potential. Then, it outlines how to build and train a basic GAN model that can be used to generate human-like handwritten digits and a more advanced model that can enhance low resolution images to high resolution ones. Finally, we'll cover how to use the two GAN models in your iOS and Android apps.

Chapter 10, *Building AlphaZero-like Mobile Game App*, begins with how the latest and coolest AlphaZero works, and how to train and test a AlphaZero-like model to play a simple but fun game called Connect 4 in Keras with TensorFlow as backend. We'll then show you the complete iOS and Android apps to use the model and play the game Connect 4 on your mobile devices.

Chapter 11, *Using TensorFlow Lite and Core ML on Mobile*, demonstrates TensorFlow Lite and then shows you how to use a prebuilt TensorFlow model, a retrained TensorFlow model for TensorFlow Lite, and a custom TensorFlow Lite model in iOS. We'll also show you how to use TensorFlow Lite in Android. After that, we'll give an overview of Apple's Core ML and show you how to use Core ML with standard machine learning using Scikit-Learn. Finally, we'll cover how to use Core ML with TensorFlow and Keras.

Chapter 12, *Developing TensorFlow Apps on Raspberry Pi*, first looks at how to set up Raspberry Pi and make it move, and how to set up TensorFlow on Raspberry Pi. Then, we'll cover how to use the TensorFlow image recognition and audio recognition models, along with text to speech and robot movement APIs, to build a Raspberry Pi robot that can move, see, listen, and speak. Finally, we'll discuss in detail how to use OpenAI Gym and TensorFlow to build and train a powerful neural network-based reinforcement learning policy model from scratch in a simulated environment to make the robot learn to keep its balance.

## To get the most out of this book

We recommend that you start with reading the first four chapters in order, along with running the accompanying iOS and Android apps available from the book's source code repository at <http://github.com/jeffxtang/mobiletfbook>. That'll help you ensure that you have the development environments all set up for TensorFlow mobile app development and that you know how to integrate TensorFlow into your own iOS and/or Android apps. If you're an iOS developer, you'll also learn how to use Objective-C or Swift with TensorFlow, and when and how to use the TensorFlow pod or the manual TensorFlow iOS library.

Then, if you need to build a custom TensorFlow Android library, go to Chapter 7, *Recognizing Drawing with CNN and LSTM*, and if you want to learn how to use a Keras model in your mobile app, check out Chapter 8, *Predicting Stock Price with RNN*, and Chapter 10, *Building an AlphaZero-like Mobile Game App*.

If you're more interested in TensorFlow Lite or Core ML, read Chapter 11, *Using TensorFlow Lite and Core ML on Mobile*, and if you're most interested in TensorFlow on Raspberry Pi, or reinforcement learning in TensorFlow, jump to Chapter 12, *Developing TensorFlow Apps on Raspberry Pi*.

Other than that, you can go through chapters 5 to 10 in order to see how to train different kinds of CNN, RNN, LSTM, GAN, and AlphaZero models and how to use them on mobile, maybe running the iOS and/or Android apps for each chapter before looking into the detailed implementation. Alternatively, you can jump directly to any chapter with the model you're most interested in; just be aware that a later chapter may refer to an earlier chapter for some duplicated details, such as steps of adding a TensorFlow custom iOS library to your iOS app, or fixing some model loading or running errors by building a TensorFlow custom library. However, rest assured that you won't be lost, or at least we've done our best to provide user-friendly and step-by-step tutorials, with occasional references to some steps of previous tutorials, to help you avoid all possible pitfalls you may encounter when building mobile TensorFlow apps, while also avoiding repeating ourselves.

## When to read the book

AI, or its hottest branch machine learning, or its hottest subbranch deep learning, has enjoyed rapid progress in recent years. New releases of TensorFlow, backed by Google and with the most popular developer community of all open source machine learning frameworks, has also been launched at a faster speed. When we started writing the book in December 2017, the latest TensorFlow release was 1.4.0, released on November 2, 2017, and after that 1.5.0 was released on January 26, 2018, 1.6.0 on February 28, 2018, 1.7.0 on March 29, 2018, and 1.8.0 on April 27, 2018. All the iOS, Android, and Python code in the book has been tested with all those TensorFlow versions. Still, by the time you read the book, the latest TensorFlow version likely will be later than 1.8.0.

It turns out you don't need to worry too much about new releases of TensorFlow; the code in the book will most likely run seamlessly on the latest TensorFlow releases. During our testing run the apps on TensorFlow 1.4, 1.5, 1.6, 1.7, and 1.8, we've made no code changes at all. It's likely that in a later version, more TensorFlow operations will be supported by default, so you won't need to build a custom TensorFlow library, or you'll be able to build the custom TensorFlow library in a simpler way.