
Tensorflow

Predict valuable insights of your data with TensorFlow

Create powerful machine learning algorithms with TensorFlow

Build advanced computer vision applications using machine learning and deep learning techniques

Solution for Edge Computing Applications

Learning TensorFlow

TensorFlow 2 Pocket Primer

Intelligent Mobile Projects with TensorFlow

Teach language to machines using Python's deep learning library

Solve computer vision problems with modeling in TensorFlow and Python

Machine Learning with TensorFlow

Hands-On Computer Vision with TensorFlow 2

A hands-on guide to building deep learning models from scratch using real-world datasets

Mastering Computer Vision with TensorFlow 2.x

Natural Language Processing with TensorFlow

Pro Deep Learning with TensorFlow

TensorFlow 2.0 Quick Start Guide

Get up to speed with the newly introduced features of TensorFlow 2.0

Build, manage, and scale machine learning workloads seamlessly using Google's TensorFlow Enterprise

Hands-On Neural Networks with TensorFlow 2.0

Create images, text, and music with VAEs, GANs, LSTMs, GPT models and more

Deep Learning with TensorFlow

Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow

TensorFlow For Dummies

TensorFlow for Deep Learning

TensorFlow 1.x Deep Learning Cookbook

A Mathematical Approach to Advanced Artificial Intelligence in Python

Deep Learning with TensorFlow 2 and Keras - Second Edition

Hands-On Convolutional Neural Networks with TensorFlow

From Linear Regression to Reinforcement Learning

Learn TensorFlow Enterprise

Second generation machine learning with Google's brainchild - TensorFlow 1.x

Concepts, Tools, and Techniques to Build Intelligent Systems

Learning TensorFlow.js

Implement Machine Learning and Deep Learning Models with Python

Generative AI with Python and TensorFlow 2

10 real-world projects on computer vision, machine translation, chatbots, and reinforcement learning

Hands-On Neural Networks with TensorFlow 2.0

Learning Tensorflow

Learn TensorFlow 2.0

API Oriented Deep Learning with Python

Tensorflow

*Downloaded from
archive.imba.com by
guest*

TRAVIS ATKINSON

Predict valuable insights of your data with TensorFlow Packt Publishing Ltd
Apply neural network architectures to build state-of-the-art computer vision applications using the Python programming language
Key Features
Gain a fundamental understanding of advanced computer vision and neural

network models in use today
Cover tasks such as low-level vision, image classification, and object detection
Develop deep learning models on cloud platforms and optimize them using TensorFlow Lite and the OpenVINO toolkit
Book Description
Computer vision allows machines to gain human-level understanding to visualize, process, and analyze images and videos. This book focuses on using TensorFlow to help you learn advanced computer vision tasks

such as image acquisition, processing, and analysis. You'll start with the key principles of computer vision and deep learning to build a solid foundation, before covering neural network architectures and understanding how they work rather than using them as a black box. Next, you'll explore architectures such as VGG, ResNet, Inception, R-CNN, SSD, YOLO, and MobileNet. As you advance, you'll learn to use visual search methods using transfer learning. You'll also cover advanced computer vision concepts such as semantic segmentation, image inpainting with GAN's, object tracking, video segmentation, and action recognition. Later, the book focuses on how machine learning and deep learning concepts can be used to perform tasks

such as edge detection and face recognition. You'll then discover how to develop powerful neural network models on your PC and on various cloud platforms. Finally, you'll learn to perform model optimization methods to deploy models on edge devices for real-time inference. By the end of this book, you'll have a solid understanding of computer vision and be able to confidently develop models to automate tasks. What you will learn

Explore methods of feature extraction and image retrieval and visualize different layers of the neural network model

Use TensorFlow for various visual search methods for real-world scenarios

Build neural networks or adjust parameters to optimize the performance of models

Understand TensorFlow DeepLab to perform

semantic segmentation on images and DCGAN for image inpainting Evaluate your model and optimize and integrate it into your application to operate at scale Get up to speed with techniques for performing manual and automated image annotation Who this book is for This book is for computer vision professionals, image processing professionals, machine learning engineers and AI developers who have some knowledge of machine learning and deep learning and want to build expert-level computer vision applications. In addition to familiarity with TensorFlow, Python knowledge will be required to get started with this book. *Create powerful machine learning algorithms with TensorFlow* Packt Publishing Ltd

Explore machine learning concepts using the latest numerical computing library — TensorFlow — with the help of this comprehensive cookbook About This Book Your quick guide to implementing TensorFlow in your day-to-day machine learning activities Learn advanced techniques that bring more accuracy and speed to machine learning Upgrade your knowledge to the second generation of machine learning with this guide on TensorFlow Who This Book Is For This book is ideal for data scientists who are familiar with C++ or Python and perform machine learning activities on a day-to-day basis. Intermediate and advanced machine learning implementers who need a quick guide they can easily navigate will find it useful. What You Will Learn Become familiar with the basics of

the TensorFlow machine learning library
Get to know Linear Regression
techniques with TensorFlow Learn SVMs
with hands-on recipes Implement neural
networks and improve predictions Apply
NLP and sentiment analysis to your data
Master CNN and RNN through practical
recipes Take TensorFlow into production
In Detail TensorFlow is an open source
software library for Machine Intelligence.
The independent recipes in this book will
teach you how to use TensorFlow for
complex data computations and will let
you dig deeper and gain more insights
into your data than ever before. You'll
work through recipes on training models,
model evaluation, sentiment analysis,
regression analysis, clustering analysis,
artificial neural networks, and deep
learning – each using Google's machine

learning library TensorFlow. This guide
starts with the fundamentals of the
TensorFlow library which includes
variables, matrices, and various data
sources. Moving ahead, you will get
hands-on experience with Linear
Regression techniques with TensorFlow.
The next chapters cover important high-
level concepts such as neural networks,
CNN, RNN, and NLP. Once you are
familiar and comfortable with the
TensorFlow ecosystem, the last chapter
will show you how to take it to
production. Style and approach This
book takes a recipe-based approach
where every topic is explicated with the
help of a real-world example.
*Build advanced computer vision
applications using machine learning and
deep learning techniques* Apress

This book is designed to guide you through TensorFlow and how to use it effectively. Throughout the book, you will work through recipes and get hands-on experience to perform complex data computations, gain insights into your data, and more.

Solution for Edge Computing

Applications Packt Publishing Ltd

Learn how to solve real life problems using different methods like logic regression, random forests and SVM's with TensorFlow. Key Features Understand predictive analytics along with its challenges and best practices Embedded with assessments that will help you revise the concepts you have learned in this book Book Description Predictive analytics discovers hidden patterns from structured and

unstructured data for automated decision making in business intelligence. Predictive decisions are becoming a huge trend worldwide, catering to wide industry sectors by predicting which decisions are more likely to give maximum results. TensorFlow, Google's brainchild, is immensely popular and extensively used for predictive analysis. This book is a quick learning guide on all the three types of machine learning, that is, supervised, unsupervised, and reinforcement learning with TensorFlow. This book will teach you predictive analytics for high-dimensional and sequence data. In particular, you will learn the linear regression model for regression analysis. You will also learn how to use regression for predicting continuous values. You will learn

supervised learning algorithms for predictive analytics. You will explore unsupervised learning and clustering using K-means. You will then learn how to predict neighborhoods using K-means, and then, see another example of clustering audio clips based on their audio features. This book is ideal for developers, data analysts, machine learning practitioners, and deep learning enthusiasts who want to build powerful, robust, and accurate predictive models with the power of TensorFlow. This book is embedded with useful assessments that will help you revise the concepts you have learned in this book. What you will learn: Learn TensorFlow features in a real-life problem, followed by detailed TensorFlow installation and configuration. Explore computation graphs, data, and

programming models also get an insight into an example of implementing linear regression model for predictive analytics. Solve the Titanic survival problem using logistic regression, random forests, and SVMs for predictive analytics. Dig deeper into predictive analytics and find out how to take advantage of it to cluster records belonging to the certain group or class for a dataset of unsupervised observations. Learn several examples of how to apply reinforcement learning algorithms for developing predictive models on real-life datasets. Who this book is for: This book is aimed at developers, data analysts, machine learning practitioners, and deep learning enthusiasts who want to build powerful, robust, and accurate predictive models with the power of TensorFlow.

Learning TensorFlow "O'Reilly Media, Inc."

Get hands-on with the browser-based JavaScript library for training and deploying machine learning models effectively Key Features Build, train and run machine learning models in the browser using TensorFlow.js Create smart web applications from scratch with the help of useful examples Use flexible and intuitive APIs from TensorFlow.js to understand how machine learning algorithms function Book Description TensorFlow.js is a framework that enables you to create performant machine learning (ML) applications that run smoothly in a web browser. With this book, you will learn how to use TensorFlow.js to implement various ML models through an example-based

approach. Starting with the basics, you'll understand how ML models can be built on the web. Moving on, you will get to grips with the TensorFlow.js ecosystem to develop applications more efficiently. The book will then guide you through implementing ML techniques and algorithms such as regression, clustering, fast Fourier transform (FFT), and dimensionality reduction. You will later cover the Bellman equation to solve Markov decision process (MDP) problems and understand how it is related to reinforcement learning. Finally, you will explore techniques for deploying ML-based web applications and training models with TensorFlow Core. Throughout this ML book, you'll discover useful tips and tricks that will build on your knowledge. By the end of this book,

you will be equipped with the skills you need to create your own web-based ML applications and fine-tune models to achieve high performance. What you will learn

- Use the t-SNE algorithm in TensorFlow.js to reduce dimensions in an input dataset
- Deploy tfjs-converter to convert Keras models and load them into TensorFlow.js
- Apply the Bellman equation to solve MDP problems
- Use the k-means algorithm in TensorFlow.js to visualize prediction results
- Create tf.js packages with Parcel, Webpack, and Rollup to deploy web apps
- Implement tf.js backend frameworks to tune and accelerate app performance

Who this book is for
This book is for web developers who want to learn how to integrate machine learning techniques with web-based applications from

scratch. This book will also appeal to data scientists, machine learning practitioners, and deep learning enthusiasts who are looking to perform accelerated, browser-based machine learning on Web using TensorFlow.js. Working knowledge of JavaScript programming language is all you need to get started.

[TensorFlow 2 Pocket Primer](#) Packt Publishing Ltd

Learn how to solve challenging machine learning problems with TensorFlow, Google's revolutionary new software library for deep learning. If you have some background in basic linear algebra and calculus, this practical book introduces machine-learning fundamentals by showing you how to design systems capable of detecting

objects in images, understanding text, analyzing video, and predicting the properties of potential medicines. TensorFlow for Deep Learning teaches concepts through practical examples and helps you build knowledge of deep learning foundations from the ground up. It's ideal for practicing developers with experience designing software systems, and useful for scientists and other professionals familiar with scripting but not necessarily with designing learning algorithms. Learn TensorFlow fundamentals, including how to perform basic computation Build simple learning systems to understand their mathematical foundations Dive into fully connected deep networks used in thousands of applications Turn prototypes into high-quality models with

hyperparameter optimization Process images with convolutional neural networks Handle natural language datasets with recurrent neural networks Use reinforcement learning to solve games such as tic-tac-toe Train deep networks with hardware including GPUs and tensor processing units [Intelligent Mobile Projects with TensorFlow](#) Packt Publishing Ltd This book is a comprehensive introduction for those who are new to scalable and optimized TensorFlow for production. You will learn how to deliver enterprise-grade support for your existing and newly built AI applications. You will address the various needs of AI-enabled organizations to manage and scale machine learning workloads in production.

Teach language to machines using Python's deep learning library Springer Nature

Build machine and deep learning systems with the newly released TensorFlow 2 and Keras for the lab, production, and mobile devices Key Features Introduces and then uses TensorFlow 2 and Keras right from the start Teaches key machine and deep learning techniques Understand the fundamentals of deep learning and machine learning through clear explanations and extensive code samples Book Description Deep Learning with TensorFlow 2 and Keras, Second Edition teaches neural networks and deep learning techniques alongside TensorFlow (TF) and Keras. You'll learn how to write deep learning applications

in the most powerful, popular, and scalable machine learning stack available. TensorFlow is the machine learning library of choice for professional applications, while Keras offers a simple and powerful Python API for accessing TensorFlow. TensorFlow 2 provides full Keras integration, making advanced machine learning easier and more convenient than ever before. This book also introduces neural networks with TensorFlow, runs through the main applications (regression, ConvNets (CNNs), GANs, RNNs, NLP), covers two working example apps, and then dives into TF in production, TF mobile, and using TensorFlow with AutoML. What you will learn Build machine learning and deep learning systems with TensorFlow 2 and the Keras API Use Regression

analysis, the most popular approach to machine learning Understand ConvNets (convolutional neural networks) and how they are essential for deep learning systems such as image classifiers Use GANs (generative adversarial networks) to create new data that fits with existing patterns Discover RNNs (recurrent neural networks) that can process sequences of input intelligently, using one part of a sequence to correctly interpret another Apply deep learning to natural human language and interpret natural language texts to produce an appropriate response Train your models on the cloud and put TF to work in real environments Explore how Google tools can automate simple ML workflows without the need for complex modeling Who this book is for This book is for

Python developers and data scientists who want to build machine learning and deep learning systems with TensorFlow. Whether or not you have done machine learning before, this book gives you the theory and practice required to use Keras, TensorFlow 2, and AutoML to build machine learning systems.

Solve computer vision problems with modeling in TensorFlow and Python Manning Publications

Tackle common commercial machine learning problems with Google's TensorFlow 1.x library and build deployable solutions. About This Book Enter the new era of second-generation machine learning with Python with this practical and insightful guide Set up TensorFlow 1.x for actual industrial use, including high-performance setup

aspects such as multi-GPU support
Create pipelines for training and using
applying classifiers using raw real-world
data Who This Book Is For This book is
for data scientists and researchers who
are looking to either migrate from an
existing machine learning library or jump
into a machine learning platform
headfirst. The book is also for software
developers who wish to learn deep
learning by example. Particular focus is
placed on solving commercial deep
learning problems from several
industries using TensorFlow's unique
features. No commercial domain
knowledge is required, but familiarity
with Python and matrix math is
expected. What You Will Learn Explore
how to use different machine learning
models to ask different questions of your

data Learn how to build deep neural
networks using TensorFlow 1.x Cover
key tasks such as clustering, sentiment
analysis, and regression analysis using
TensorFlow 1.x Find out how to write
clean and elegant Python code that will
optimize the strength of your algorithms
Discover how to embed your machine
learning model in a web application for
increased accessibility Learn how to use
multiple GPUs for faster training using
AWS In Detail Google's TensorFlow is a
game changer in the world of machine
learning. It has made machine learning
faster, simpler, and more accessible
than ever before. This book will teach
you how to easily get started with
machine learning using the power of
Python and TensorFlow 1.x. Firstly, you'll
cover the basic installation procedure

and explore the capabilities of TensorFlow 1.x. This is followed by training and running the first classifier, and coverage of the unique features of the library including data flow graphs, training, and the visualization of performance with TensorBoard—all within an example-rich context using problems from multiple industries. You'll be able to further explore text and image analysis, and be introduced to CNN models and their setup in TensorFlow 1.x. Next, you'll implement a complete real-life production system from training to serving a deep learning model. As you advance you'll learn about Amazon Web Services (AWS) and create a deep neural network to solve a video action recognition problem. Lastly, you'll convert the Caffe model to TensorFlow

and be introduced to the high-level TensorFlow library, TensorFlow-Slim. By the end of this book, you will be geared up to take on any challenges of implementing TensorFlow 1.x in your machine learning environment. Style and approach This comprehensive guide will enable you to understand the latest advances in machine learning and will empower you to implement this knowledge in your machine learning environment.

Machine Learning with TensorFlow

"O'Reilly Media, Inc."

Build, scale, and deploy deep neural network models using the star libraries in Python Key Features Delve into advanced machine learning and deep learning use cases using Tensorflow and Keras Build, deploy, and scale end-to-

end deep neural network models in a production environment Learn to deploy TensorFlow on mobile, and distributed TensorFlow on GPU, Clusters, and Kubernetes Book Description TensorFlow is the most popular numerical computation library built from the ground up for distributed, cloud, and mobile environments. TensorFlow represents the data as tensors and the computation as graphs. This book is a comprehensive guide that lets you explore the advanced features of TensorFlow 1.x. Gain insight into TensorFlow Core, Keras, TF Estimators, TFLearn, TF Slim, Pretty Tensor, and Sonnet. Leverage the power of TensorFlow and Keras to build deep learning models, using concepts such as transfer learning, generative adversarial

networks, and deep reinforcement learning. Throughout the book, you will obtain hands-on experience with varied datasets, such as MNIST, CIFAR-10, PTB, text8, and COCO-Images. You will learn the advanced features of TensorFlow1.x, such as distributed TensorFlow with TF Clusters, deploy production models with TensorFlow Serving, and build and deploy TensorFlow models for mobile and embedded devices on Android and iOS platforms. You will see how to call TensorFlow and Keras API within the R statistical software, and learn the required techniques for debugging when the TensorFlow API-based code does not work as expected. The book helps you obtain in-depth knowledge of TensorFlow, making you the go-to person for solving artificial intelligence

problems. By the end of this guide, you will have mastered the offerings of TensorFlow and Keras, and gained the skills you need to build smarter, faster, and efficient machine learning and deep learning systems. What you will learn

- Master advanced concepts of deep learning such as transfer learning, reinforcement learning, generative models and more, using TensorFlow and Keras
- Perform supervised (classification and regression) and unsupervised (clustering) learning to solve machine learning tasks
- Build end-to-end deep learning (CNN, RNN, and Autoencoders) models with TensorFlow
- Scale and deploy production models with distributed and high-performance computing on GPU and clusters
- Build TensorFlow models to work with

- multilayer perceptrons using Keras, TFLearn, and R
- Learn the functionalities of smart apps by building and deploying TensorFlow models on iOS and Android devices
- Supercharge TensorFlow with distributed training and deployment on Kubernetes and TensorFlow Clusters

Who this book is for This book is for data scientists, machine learning engineers, artificial intelligence engineers, and for all TensorFlow users who wish to upgrade their TensorFlow knowledge and work on various machine learning and deep learning problems. If you are looking for an easy-to-follow guide that underlines the intricacies and complex use cases of machine learning, you will find this book extremely useful. Some basic understanding of TensorFlow is required to get the most out of the book.

Hands-On Computer Vision with TensorFlow 2 Packt Publishing Ltd
Through a series of recent breakthroughs, deep learning has boosted the entire field of machine learning. Now, even programmers who know close to nothing about this technology can use simple, efficient tools to implement programs capable of learning from data. This practical book shows you how. By using concrete examples, minimal theory, and two production-ready Python frameworks—Scikit-Learn and TensorFlow—author Aurélien Géron helps you gain an intuitive understanding of the concepts and tools for building intelligent systems. You'll learn a range of techniques, starting with simple linear regression and progressing

to deep neural networks. With exercises in each chapter to help you apply what you've learned, all you need is programming experience to get started. Explore the machine learning landscape, particularly neural nets Use Scikit-Learn to track an example machine-learning project end-to-end Explore several training models, including support vector machines, decision trees, random forests, and ensemble methods Use the TensorFlow library to build and train neural nets Dive into neural net architectures, including convolutional nets, recurrent nets, and deep reinforcement learning Learn techniques for training and scaling deep neural nets **A hands-on guide to building deep learning models from scratch using real-world datasets** Mercury Learning

and Information

Implement deep learning applications using TensorFlow while learning the “why” through in-depth conceptual explanations. You’ll start by learning what deep learning offers over other machine learning models. Then familiarize yourself with several technologies used to create deep learning models. While some of these technologies are complementary, such as Pandas, Scikit-Learn, and Numpy—others are competitors, such as PyTorch, Caffe, and Theano. This book clarifies the positions of deep learning and TensorFlow among their peers. You’ll then work on supervised deep learning models to gain applied experience with the technology. A single-layer of multiple perceptrons will be used to build a

shallow neural network before turning it into a deep neural network. After showing the structure of the ANNs, a real-life application will be created with TensorFlow 2.0 Keras API. Next, you’ll work on data augmentation and batch normalization methods. Then, the Fashion MNIST dataset will be used to train a CNN. CIFAR10 and Imagenet pre-trained models will be loaded to create already advanced CNNs. Finally, move into theoretical applications and unsupervised learning with auto-encoders and reinforcement learning with tf-agent models. With this book, you’ll delve into applied deep learning practical functions and build a wealth of knowledge about how to use TensorFlow effectively. What You’ll Learn Compare competing technologies and see why

TensorFlow is more popular Generate text, image, or sound with GANs Predict the rating or preference a user will give to an item Sequence data with recurrent neural networks Who This Book Is For Data scientists and programmers new to the fields of deep learning and machine learning APIs.

Mastering Computer Vision with

TensorFlow 2.x Packt Publishing Ltd

A comprehensive guide to developing neural network-based solutions using TensorFlow 2.0 Key Features Understand the basics of machine learning and discover the power of neural networks and deep learning Explore the structure of the TensorFlow framework and understand how to transition to TF 2.0 Solve any deep learning problem by developing neural network-based

solutions using TF 2.0 Book Description TensorFlow, the most popular and widely used machine learning framework, has made it possible for almost anyone to develop machine learning solutions with ease. With TensorFlow (TF) 2.0, you'll explore a revamped framework structure, offering a wide variety of new features aimed at improving productivity and ease of use for developers. This book covers machine learning with a focus on developing neural network-based solutions. You'll start by getting familiar with the concepts and techniques required to build solutions to deep learning problems. As you advance, you'll learn how to create classifiers, build object detection and semantic segmentation networks, train generative models, and speed up the

development process using TF 2.0 tools such as TensorFlow Datasets and TensorFlow Hub. By the end of this TensorFlow book, you'll be ready to solve any machine learning problem by developing solutions using TF 2.0 and putting them into production. What you will learn Grasp machine learning and neural network techniques to solve challenging tasks Apply the new features of TF 2.0 to speed up development Use TensorFlow Datasets (tfds) and the tf.data API to build high-efficiency data input pipelines Perform transfer learning and fine-tuning with TensorFlow Hub Define and train networks to solve object detection and semantic segmentation problems Train Generative Adversarial Networks (GANs) to generate images and data distributions Use the

SavedModel file format to put a model, or a generic computational graph, into production Who this book is for If you're a developer who wants to get started with machine learning and TensorFlow, or a data scientist interested in developing neural network solutions in TF 2.0, this book is for you. Experienced machine learning engineers who want to master the new features of the TensorFlow framework will also find this book useful. Basic knowledge of calculus and a strong understanding of Python programming will help you grasp the topics covered in this book.

Natural Language Processing with TensorFlow Packt Publishing Ltd

Computer vision is achieving a new frontier of capabilities in fields like health, automobile or robotics. This book

explores TensorFlow 2, Google's open-source AI framework, and teaches how to leverage deep neural networks for visual tasks. It will help you acquire the insight and skills to be a part of the exciting advances in computer vision.

Pro Deep Learning with TensorFlow Apress

Roughly inspired by the human brain, deep neural networks trained with large amounts of data can solve complex tasks with unprecedented accuracy. This practical book provides an end-to-end guide to TensorFlow, the leading open source software library that helps you build and train neural networks for computer vision, natural language processing (NLP), speech recognition, and general predictive analytics. Authors Tom Hope, Yehezkel Resheff, and Itay

Lieder provide a hands-on approach to TensorFlow fundamentals for a broad technical audience—from data scientists and engineers to students and researchers. You'll begin by working through some basic examples in TensorFlow before diving deeper into topics such as neural network architectures, TensorBoard visualization, TensorFlow abstraction libraries, and multithreaded input pipelines. Once you finish this book, you'll know how to build and deploy production-ready deep learning systems in TensorFlow. Get up and running with TensorFlow, rapidly and painlessly Learn how to use TensorFlow to build deep learning models from the ground up Train popular deep learning models for computer vision and NLP Use extensive abstraction

libraries to make development easier and faster Learn how to scale TensorFlow, and use clusters to distribute model training Deploy TensorFlow in a production setting

TensorFlow 2.0 Quick Start Guide
Mercury Learning and Information
A comprehensive guide to developing neural network-based solutions using TensorFlow 2.0 Key Features Understand the basics of machine learning and discover the power of neural networks and deep learning Explore the structure of the TensorFlow framework and understand how to transition to TF 2.0 Solve any deep learning problem by developing neural network-based solutions using TF 2.0 Book Description TensorFlow, the most popular and widely used machine learning framework, has

made it possible for almost anyone to develop machine learning solutions with ease. With TensorFlow (TF) 2.0, you'll explore a revamped framework structure, offering a wide variety of new features aimed at improving productivity and ease of use for developers. This book covers machine learning with a focus on developing neural network-based solutions. You'll start by getting familiar with the concepts and techniques required to build solutions to deep learning problems. As you advance, you'll learn how to create classifiers, build object detection and semantic segmentation networks, train generative models, and speed up the development process using TF 2.0 tools such as TensorFlow Datasets and TensorFlow Hub. By the end of this

TensorFlow book, you'll be ready to solve any machine learning problem by developing solutions using TF 2.0 and putting them into production. What you will learn Grasp machine learning and neural network techniques to solve challenging tasks Apply the new features of TF 2.0 to speed up development Use TensorFlow Datasets (tfds) and the tf.data API to build high-efficiency data input pipelines Perform transfer learning and fine-tuning with TensorFlow Hub Define and train networks to solve object detection and semantic segmentation problems Train Generative Adversarial Networks (GANs) to generate images and data distributions Use the SavedModel file format to put a model, or a generic computational graph, into production Who this book is for If you're

a developer who wants to get started with machine learning and TensorFlow, or a data scientist interested in developing neural network solutions in TF 2.0, this book is for you. Experienced machine learning engineers who want to master the new features of the TensorFlow framework will also find this book useful. Basic knowledge of calculus and a strong understanding of Python programming will help you grasp the topics covered in this book.

Get up to speed with the newly introduced features of TensorFlow 2.0

Packt Publishing Ltd

Learn how to apply TensorFlow to a wide range of deep learning and Machine Learning problems with this practical guide on training CNNs for image classification, image recognition, object

detection and many computer vision challenges. Key Features Learn the fundamentals of Convolutional Neural Networks Harness Python and TensorFlow to train CNNs Build scalable deep learning models that can process millions of items Book Description Convolutional Neural Networks (CNN) are one of the most popular architectures used in computer vision apps. This book is an introduction to CNNs through solving real-world problems in deep learning while teaching you their implementation in popular Python library - TensorFlow. By the end of the book, you will be training CNNs in no time! We start with an overview of popular machine learning and deep learning models, and then get you set up with a TensorFlow development environment.

This environment is the basis for implementing and training deep learning models in later chapters. Then, you will use Convolutional Neural Networks to work on problems such as image classification, object detection, and semantic segmentation. After that, you will use transfer learning to see how these models can solve other deep learning problems. You will also get a taste of implementing generative models such as autoencoders and generative adversarial networks. Later on, you will see useful tips on machine learning best practices and troubleshooting. Finally, you will learn how to apply your models on large datasets of millions of images. What you will learn Train machine learning models with TensorFlow Create systems that can evolve and scale

during their life cycle Use CNNs in image recognition and classification Use TensorFlow for building deep learning models Train popular deep learning models Fine-tune a neural network to improve the quality of results with transfer learning Build TensorFlow models that can scale to large datasets and systems Who this book is for This book is for Software Engineers, Data Scientists, or Machine Learning practitioners who want to use CNNs for solving real-world problems. Knowledge of basic machine learning concepts, linear algebra and Python will help. [Build, manage, and scale machine learning workloads seamlessly using Google's TensorFlow Enterprise](#) Manning Publications
Create Deep Learning and

Reinforcement Learning apps for multiple platforms with TensorFlow Key Features Build TensorFlow-powered AI applications for mobile and embedded devices Learn modern AI topics such as computer vision, NLP, and deep reinforcement learning Get practical insights and exclusive working code not available in the TensorFlow documentation Book Description As a developer, you always need to keep an eye out and be ready for what will be trending soon, while also focusing on what's trending currently. So, what's better than learning about the integration of the best of both worlds, the present and the future? Artificial Intelligence (AI) is widely regarded as the next big thing after mobile, and Google's TensorFlow is the leading open

source machine learning framework, the hottest branch of AI. This book covers more than 10 complete iOS, Android, and Raspberry Pi apps powered by TensorFlow and built from scratch, running all kinds of cool TensorFlow models offline on-device: from computer vision, speech and language processing to generative adversarial networks and AlphaZero-like deep reinforcement learning. You'll learn how to use or retrain existing TensorFlow models, build your own models, and develop intelligent mobile apps running those TensorFlow models. You'll learn how to quickly build such apps with step-by-step tutorials and how to avoid many pitfalls in the process with lots of hard-earned troubleshooting tips. What you will learn Classify images with transfer learning Detect objects and

their locations Transform pictures with amazing art styles Understand simple speech commands Describe images in natural language Recognize drawing with Convolutional Neural Network and Long Short-Term Memory Predict stock price with Recurrent Neural Network in TensorFlow and Keras Generate and enhance images with generative adversarial networks Build AlphaZero-like mobile game app in TensorFlow and Keras Use TensorFlow Lite and Core ML on mobile Develop TensorFlow apps on Raspberry Pi that can move, see, listen, speak, and learn Who this book is for If you're an iOS/Android developer interested in building and retraining others' TensorFlow models and running them in your mobile apps, or if you're a TensorFlow developer and want to run

your new and amazing TensorFlow models on mobile devices, this book is for you. You'll also benefit from this book if you're interested in TensorFlow Lite, Core ML, or TensorFlow on Raspberry Pi. [Hands-On Neural Networks with TensorFlow 2.0](#) [Guru99](#)

Learn how to use TensorFlow 2.0 to build machine learning and deep learning models with complete examples. The book begins with introducing TensorFlow 2.0 framework and the major changes from its last release. Next, it focuses on building Supervised Machine Learning models using TensorFlow 2.0. It also demonstrates how to build models using customer estimators. Further, it explains how to use TensorFlow 2.0 API to build machine learning and deep learning models for image classification using the

standard as well as custom parameters. You'll review sequence predictions, saving, serving, deploying, and standardized datasets, and then deploy these models to production. All the code presented in the book will be available in the form of executable scripts at Github which allows you to try out the examples and extend them in interesting ways.

What You'll Learn Review the new features of TensorFlow 2.0 Use TensorFlow 2.0 to build machine learning and deep learning models Perform sequence predictions using TensorFlow 2.0 Deploy TensorFlow 2.0 models with practical examples **Who This Book Is For** Data scientists, machine and deep learning engineers.

Create images, text, and music with VAEs, GANs, LSTMs, GPT models and

more John Wiley & Sons

This book is a step-by-step guide to show you how to implement generative models in TensorFlow 2.x from scratch. You'll get to grips with the image

generative technology by covering autoencoders, style transfer, and GANs as well as fundamental and state-of-the-art models.

Related with Tensorflow:

- Plantar Surface Of Foot Anatomy : [click here](#)