

Mobilenet Tensorflow

[Hands-On Deep Learning with TensorFlow](#)
[TensorFlow 2 Pocket Primer](#)
[The TensorFlow Workshop](#)
[Machine Learning with TensorFlow 1.x](#)
[TensorFlow For Dummies](#)
[TensorFlow for Deep Learning](#)
[TensorFlow 2.0 Quick Start Guide](#)
[What's New in TensorFlow 2.0](#)
[TinyML](#)
[Hands-On Neural Networks with TensorFlow 2.0](#)
[Deep Learning with TensorFlow and Keras](#)
[Mastering Computer Vision with TensorFlow 2.x](#)
[Natural Language Processing with TensorFlow](#)
[TensorFlow in Action](#)
[AI and Machine Learning for Coders](#)
[Deep Learning Patterns and Practices](#)
[Intelligent Mobile Projects with TensorFlow](#)
[Learning TensorFlow](#)
[Programming with TensorFlow](#)
[AI and Machine Learning for Coders](#)
[Learning TensorFlow.js](#)
[Hands-On Computer Vision with TensorFlow 2](#)
[Deep Learning Pipeline](#)
[TensorFlow Reinforcement Learning Quick Start Guide](#)
[Practical Deep Learning for Cloud, Mobile, and Edge](#)
[Deep Learning with TensorFlow](#)
[Mastering TensorFlow 2.x](#)
[Hands-On Deep Learning Architectures with Python](#)
[Machine Learning with TensorFlow, Second Edition](#)
[Generative AI with Python and TensorFlow 2](#)
[Advanced Deep Learning with TensorFlow 2 and Keras](#)
[Neural Network Programming with TensorFlow](#)
[TensorFlow Pocket Primer](#)
[Learn TensorFlow 2.0](#)
[Deep Learning with JavaScript](#)
[Hands-On Machine Learning with TensorFlow.js](#)
[TensorFlow 2 Pocket Reference](#)
[Hands-On Convolutional Neural Networks with TensorFlow](#)
[Visual Inference for IoT Systems: A Practical Approach](#)

Mobilenet Tensorflow

Downloaded from [archive.imba.com](#) by guest

HARRISON BECKER

Hands-On Deep Learning with TensorFlow Packt Publishing Ltd

Hands-On Machine Learning with TensorFlow.js is a comprehensive guide that will help you easily get started with machine learning algorithms and techniques using TensorFlow.js. By the end of this book, you will be able to create and optimize your own web-based machine learning applications using practical examples.

TensorFlow 2 Pocket Primer "O'Reilly Media, Inc."

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.

The TensorFlow Workshop Packt Publishing Ltd 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 *Machine Learning with TensorFlow 1.x* O'Reilly Media Deep learning networks are getting smaller. Much smaller. The Google Assistant team can detect words with a model just 14 kilobytes in size—small enough to run on a microcontroller. With this practical book you'll enter the field of TinyML, where deep learning and embedded systems combine to make astounding things possible with tiny devices. Pete Warden and Daniel Situnayake explain how you can train models small enough to fit into any environment. Ideal for software and hardware developers who want to build embedded systems using machine learning, this guide walks you through creating a series of TinyML projects, step-by-step. No machine learning or microcontroller experience is necessary. Build a speech recognizer, a camera that detects people, and a magic wand that responds to gestures Work with Arduino and ultra-low-power microcontrollers Learn the essentials of ML and how to train your own models Train models to understand audio, image, and accelerometer data Explore TensorFlow Lite for Microcontrollers, Google's toolkit for TinyML Debug applications and provide safeguards for privacy and security Optimize latency, energy usage, and model and binary size

TensorFlow For Dummies BPB Publications

Become a machine learning pro! Google TensorFlow has become

the darling of financial firms and research organizations, but the technology can be intimidating and the learning curve is steep. Luckily, TensorFlow For Dummies is here to offer you a friendly, easy-to-follow book on the subject. Inside, you'll find out how to write applications with TensorFlow, while also grasping the concepts underlying machine learning—all without ever losing your cool! Machine learning has become ubiquitous in modern society, and its applications include language translation, robotics, handwriting analysis, financial prediction, and image recognition. TensorFlow is Google's preeminent toolset for machine learning, and this hands-on guide makes it easy to understand, even for those without a background in artificial intelligence. Install TensorFlow on your computer Learn the fundamentals of statistical regression and neural networks Visualize the machine learning process with TensorBoard Perform image recognition with convolutional neural networks (CNNs) Analyze sequential data with recurrent neural networks (RNNs) Execute TensorFlow on mobile devices and the Google Cloud Platform (GCP) If you're a manager or software developer looking to use TensorFlow for machine learning, this is the book you'll want to have close by.

TensorFlow for Deep Learning Mercury Learning and Information Build your own pipeline based on modern TensorFlow approaches rather than outdated engineering concepts. This book shows you how to build a deep learning pipeline for real-life TensorFlow projects. You'll learn what a pipeline is and how it works so you can build a full application easily and rapidly. Then troubleshoot and overcome basic TensorFlow obstacles to easily create functional apps and deploy well-trained models. Step-by-step and example-oriented instructions help you understand each step of the deep learning pipeline while you apply the most straightforward and effective tools to demonstrative problems and datasets. You'll also develop a deep learning project by preparing data, choosing the model that fits that data, and debugging your model to get the best fit to data all using Tensorflow techniques. Enhance your skills by accessing some of the most powerful recent trends in data science. If you've ever considered building your own image or text-tagging solution or entering a Kaggle contest, Deep Learning Pipeline is for you! What You'll Learn Develop a deep learning project using dataStudy and apply various models to your data Debug and troubleshoot the proper model suited for your data Who This Book Is For Developers, analysts, and data scientists looking to add to or enhance their existing skills by accessing some of the most powerful recent trends in data science. Prior experience in Python or other TensorFlow related languages and mathematics would be helpful.

TensorFlow 2.0 Quick Start Guide Packt Publishing Ltd

Delve into neural networks, implement deep learning algorithms, and explore layers of data abstraction with the help of TensorFlow. Key Features Learn how to implement advanced techniques in deep learning with Google's brainchild, TensorFlow Explore deep neural networks and layers of data abstraction with the help of this comprehensive guide Gain real-world contextualization through some deep learning problems concerning research and application Book Description Deep learning is a branch of machine learning algorithms based on learning multiple levels of abstraction. Neural networks, which are at the core of deep learning, are being used in predictive analytics, computer vision, natural language processing, time series forecasting, and to perform a myriad of other complex tasks. This book is conceived 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, combined with other open source Python libraries. Throughout the book, you'll learn how to develop deep learning applications for machine learning systems using Feedforward Neural Networks, Convolutional Neural Networks, Recurrent Neural Networks, Autoencoders, and Factorization Machines. Discover how to attain deep learning programming on GPU in a distributed way. You'll come away with an in-depth knowledge of machine learning techniques and the skills to apply them to real-world projects. What you will learn Apply deep machine intelligence and GPU computing with TensorFlow Access public datasets and use TensorFlow to load, process, and transform the data Discover how to use the high-level TensorFlow API to build more powerful applications Use deep learning for scalable object detection and mobile computing Train machines quickly to learn from data by exploring reinforcement learning techniques Explore active areas of deep learning research and applications Who this book is for The book is for people interested in machine learning and machine intelligence. A rudimentary level of programming in one language is assumed, as is a basic familiarity with computer science techniques and technologies, including a basic awareness of computer hardware and algorithms. Some competence in mathematics is needed to the level of elementary linear algebra and calculus. **What's New in TensorFlow 2.0** "O'Reilly Media, Inc." Summary Deep learning has transformed the fields of computer vision, image processing, and natural language applications. Thanks to TensorFlow.js, now JavaScript developers can build deep learning apps without relying on Python or R. Deep Learning with JavaScript shows developers how they can bring DL technology to the web. Written by the main authors of the TensorFlow library, this new book provides fascinating use cases and in-depth instruction for deep learning apps in JavaScript in your browser or on Node. Foreword by Nikhil Thorat and Daniel Smilkov. About the technology Running deep learning applications in the browser or on Node-based backends opens up exciting possibilities for smart web applications. With the TensorFlow.js library, you build and train deep learning models with JavaScript. Offering uncompromising production-quality scalability, modularity, and responsiveness, TensorFlow.js really shines for its portability. Its models run anywhere JavaScript runs, pushing ML farther up the application stack. About the book In Deep Learning with JavaScript, you'll learn to use TensorFlow.js to build deep learning models that run directly in the browser. This fast-paced book, written by Google engineers, is practical, engaging, and easy to follow. Through diverse examples featuring text analysis, speech processing, image recognition, and self-learning game AI, you'll master all the basics of deep learning and explore advanced concepts, like retraining existing models for transfer learning and image generation. What's inside - Image and language processing in the browser - Tuning ML models with client-side data - Text and image creation with generative deep learning - Source code samples to test and modify About the reader For JavaScript programmers interested in deep learning. About the author Shangting Cai, Stanley Bileschi and Eric D. Nielsen are software engineers with experience on the Google Brain team, and were crucial to the development of the high-level API of TensorFlow.js. This book is based in part on the classic, Deep Learning with Python by François Chollet. TOC: PART 1 - MOTIVATION AND BASIC CONCEPTS 1 • Deep learning and JavaScript PART 2 - A GENTLE INTRODUCTION TO TENSORFLOW.JS 2 • Getting started: Simple linear regression in TensorFlow.js 3 • Adding nonlinearity: Beyond weighted sums 4 • Recognizing images and sounds using convnets 5 • Transfer learning: Reusing pretrained neural networks PART 3 - ADVANCED DEEP LEARNING WITH TENSORFLOW.JS 6 • Working with data 7 • Visualizing data and models 8 • Underfitting, overfitting, and the universal workflow of machine learning 9 • Deep learning for sequences and text 10 • Generative deep learning 11 • Basics of deep reinforcement learning PART 4 - SUMMARY AND CLOSING WORDS 12 • Testing, optimizing, and deploying models 13 • Summary, conclusions, and beyond **TinyML** Packt Publishing Ltd Build cutting edge machine and deep learning systems for the lab, production, and mobile devices Key Features Understand the fundamentals of deep learning and machine learning through clear explanations and extensive code samples Implement graph

neural networks, transformers using Hugging Face and TensorFlow Hub, and joint and contrastive learning Learn cutting-edge machine and deep learning techniques Book Description Deep Learning with TensorFlow and Keras teaches you neural networks and deep learning techniques using 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 2.x focuses on simplicity and ease of use, with updates like eager execution, intuitive higher-level APIs based on Keras, and flexible model building on any platform. This book uses the latest TF 2.0 features and libraries to present an overview of supervised and unsupervised machine learning models and provides a comprehensive analysis of deep learning and reinforcement learning models using practical examples for the cloud, mobile, and large production environments. This book also shows you how to create neural networks with TensorFlow, runs through popular algorithms (regression, convolutional neural networks (CNNs), transformers, generative adversarial networks (GANs), recurrent neural networks (RNNs), natural language processing (NLP), and graph neural networks (GNNs)), covers working example apps, and then dives into TF in production, TF mobile, and TensorFlow with AutoML. What you will learn Learn how to use the popular GNNs with TensorFlow to carry out graph mining tasks Discover the world of transformers, from pretraining to fine-tuning to evaluating them Apply self-supervised learning to natural language processing, computer vision, and audio signal processing Combine probabilistic and deep learning models using TensorFlow Probability Train your models on the cloud and put TF to work in real environments Build machine learning and deep learning systems with TensorFlow 2.x and the Keras API Who this book is for This hands-on machine learning book is for Python developers and data scientists who want to build machine learning and deep learning systems with TensorFlow. This book gives you the theory and practice required to use Keras, TensorFlow, and AutoML to build machine learning systems. Some machine learning knowledge would be useful. We don't assume TF knowledge. **Hands-On Neural Networks with TensorFlow 2.0** Packt Publishing Ltd Fun and exciting projects to learn what artificial minds can create Key Features Code examples are in TensorFlow 2, which make it easy for PyTorch users to follow along Look inside the most famous deep generative models, from GPT to MuseGAN Learn to build and adapt your own models in TensorFlow 2.x Explore exciting, cutting-edge use cases for deep generative AI Book Description Machines are excelling at creative human skills such as painting, writing, and composing music. Could you be more creative than generative AI? In this book, you'll explore the evolution of generative models, from restricted Boltzmann machines and deep belief networks to VAEs and GANs. You'll learn how to implement models yourself in TensorFlow and get to grips with the latest research on deep neural networks. There's been an explosion in potential use cases for generative models. You'll look at Open AI's news generator, deepfakes, and training deep learning agents to navigate a simulated environment. Recreate the code that's under the hood and uncover surprising links between text, image, and music generation. What you will learn Export the code from GitHub into Google Colab to see how everything works for yourself Compose music using LSTM models, simple GANs, and MuseGAN Create deepfakes using facial landmarks, autoencoders, and pix2pix GAN Learn how attention and transformers have changed NLP Build several text generation pipelines based on LSTMs, BERT, and GPT-2 Implement paired and unpaired style transfer with networks like StyleGAN Discover emerging applications of generative AI like folding proteins and creating videos from images Who this book is for This is a book for Python programmers who are keen to create and have some fun using generative models. To make the most out of this book, you should have a basic familiarity with math and statistics for machine learning. **Deep Learning with TensorFlow and Keras** Apress This easy-to-use reference for TensorFlow 2 design patterns in Python will help you make informed decisions for various use cases. Author KC Tung addresses common topics and tasks in enterprise data science and machine learning practices rather than focusing on TensorFlow itself. When and why would you feed training data as using NumPy or a streaming dataset? How would you set up cross-validations in the training process? How do you leverage a pretrained model using transfer learning? How do you perform hyperparameter tuning? Pick up this pocket reference and reduce the time you spend searching through options for your TensorFlow use cases. Understand best practices in TensorFlow model patterns and ML workflows Use code snippets as templates in building TensorFlow models and workflows Save development time by integrating prebuilt models in TensorFlow Hub Make informed design choices about data ingestion, training paradigms, model saving, and inferencing Address common scenarios such as model design style, data ingestion workflow, model training, and tuning **Mastering Computer Vision with TensorFlow 2.x** Simon and Schuster Discover best practices, reproducible architectures, and design

patterns to help guide deep learning models from the lab into production. In Deep Learning Patterns and Practices you will learn: Internal functioning of modern convolutional neural networks Procedural reuse design pattern for CNN architectures Models for mobile and IoT devices Assembling large-scale model deployments Optimizing hyperparameter tuning Migrating a model to a production environment The big challenge of deep learning lies in taking cutting-edge technologies from R&D labs through to production. Deep Learning Patterns and Practices is here to help. This unique guide lays out the latest deep learning insights from author Andrew Ferlitsch's work with Google Cloud AI. In it, you'll find deep learning models presented in a unique new way: as extendable design patterns you can easily plug-and-play into your software projects. Each valuable technique is presented in a way that's easy to understand and filled with accessible diagrams and code samples. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Discover best practices, design patterns, and reproducible architectures that will guide your deep learning projects from the lab into production. This awesome book collects and illuminates the most relevant insights from a decade of real world deep learning experience. You'll build your skills and confidence with each interesting example. About the book Deep Learning Patterns and Practices is a deep dive into building successful deep learning applications. You'll save hours of trial-and-error by applying proven patterns and practices to your own projects. Tested code samples, real-world examples, and a brilliant narrative style make even complex concepts simple and engaging. Along the way, you'll get tips for deploying, testing, and maintaining your projects. What's inside Modern convolutional neural networks Design pattern for CNN architectures Models for mobile and IoT devices Large-scale model deployments Examples for computer vision About the reader For machine learning engineers familiar with Python and deep learning. About the author Andrew Ferlitsch is an expert on computer vision, deep learning, and operationalizing ML in production at Google Cloud AI Developer Relations. Table of Contents PART 1 DEEP LEARNING FUNDAMENTALS 1 Designing modern machine learning 2 Deep neural networks 3 Convolutional and residual neural networks 4 Training fundamentals PART 2 BASIC DESIGN PATTERN 5 Procedural design pattern 6 Wide convolutional neural networks 7 Alternative connectivity patterns 8 Mobile convolutional neural networks 9 Autoencoders PART 3 WORKING WITH PIPELINES 10 Hyperparameter tuning 11 Transfer learning 12 Data distributions 13 Data pipeline 14 Training and deployment pipeline **Natural Language Processing with TensorFlow** Simon and Schuster 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. **TensorFlow in Action** John Wiley & Sons 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

AI and Machine Learning for Coders O'Reilly Media

This book presents a systematic approach to the implementation of Internet of Things (IoT) devices achieving visual inference through deep neural networks. Practical aspects are covered, with a focus on providing guidelines to optimally select hardware and software components as well as network architectures according to prescribed application requirements. The monograph includes a remarkable set of experimental results and functional procedures supporting the theoretical concepts and methodologies introduced. A case study on animal recognition based on smart camera traps is also presented and thoroughly analyzed. In this case study, different system alternatives are explored and a particular realization is completely developed. Illustrations, numerous plots from simulations and experiments, and supporting information in the form of charts and tables make Visual Inference and IoT Systems: A Practical Approach a clear and detailed guide to the topic. It will be of interest to researchers, industrial practitioners, and graduate students in the fields of computer vision and IoT.

Deep Learning Patterns and Practices Springer Nature

A practical guide to building high performance systems for object detection, segmentation, video processing, smartphone applications, and more Key FeaturesDiscover how to build, train, and serve your own deep neural networks with TensorFlow 2 and KerasApply modern solutions to a wide range of applications such as object detection and video analysisLearn how to run your models on mobile devices and web pages and improve their performanceBook Description Computer vision solutions are becoming increasingly common, making their way into fields such as health, automobile, social media, and robotics. This book will help you explore TensorFlow 2, the brand new version of Google's open source framework for machine learning. You will understand how to benefit from using convolutional neural networks (CNNs) for visual tasks. Hands-On Computer Vision with TensorFlow 2 starts with the fundamentals of computer vision and deep learning, teaching you how to build a neural network from scratch. You will discover the features that have made TensorFlow the most widely used AI library, along with its intuitive Keras interface. You'll then move on to building, training, and deploying CNNs efficiently. Complete with concrete code examples, the book demonstrates how to classify images with modern solutions, such as Inception and ResNet, and extract specific content using You Only Look Once (YOLO), Mask R-CNN, and U-Net. You will also build generative adversarial networks (GANs) and variational autoencoders (VAEs) to create and edit images, and long short-term memory networks (LSTMs) to analyze videos. In the process, you will acquire advanced insights into transfer learning, data augmentation, domain adaptation, and mobile and web deployment, among other key concepts. By the end of the book, you will have both the theoretical understanding and practical skills to solve advanced computer vision problems with TensorFlow 2.0. What you will learnCreate your own neural networks from scratchClassify images with modern architectures including Inception and ResNetDetect and segment objects in images with YOLO, Mask R-CNN, and U-NetTackle problems faced when developing self-driving cars and facial emotion recognition systemsBoost your application's performance with transfer learning, GANs, and domain adaptationUse recurrent neural networks (RNNs) for video analysisOptimize and deploy your networks on mobile devices and in the browserWho this book is for If you're new to deep learning and have some background in

Python programming and image processing, like reading/writing image files and editing pixels, this book is for you. Even if you're an expert curious about the new TensorFlow 2 features, you'll find this book useful. While some theoretical concepts require knowledge of algebra and calculus, the book covers concrete examples focused on practical applications such as visual recognition for self-driving cars and smartphone apps.

Packt Publishing Ltd

If you're looking to make a career move from programmer to AI specialist, this is the ideal place to start. Based on Laurence Moroney's extremely successful AI courses, this introductory book provides a hands-on, code-first approach to help you build confidence while you learn key topics. You'll understand how to implement the most common scenarios in machine learning, such as computer vision, natural language processing (NLP), and sequence modeling for web, mobile, cloud, and embedded runtimes. Most books on machine learning begin with a daunting amount of advanced math. This guide is built on practical lessons that let you work directly with the code. You'll learn: How to build models with TensorFlow using skills that employers desire The basics of machine learning by working with code samples How to implement computer vision, including feature detection in images How to use NLP to tokenize and sequence words and sentences Methods for embedding models in Android and iOS How to serve models over the web and in the cloud with TensorFlow Serving *Intelligent Mobile Projects with TensorFlow* Packt Publishing Ltd Updated and revised second edition of the bestselling guide to advanced deep learning with TensorFlow 2 and Keras Key FeaturesExplore the most advanced deep learning techniques that drive modern AI resultsNew coverage of unsupervised deep learning using mutual information, object detection, and semantic segmentationCompletely updated for TensorFlow 2.xBook Description Advanced Deep Learning with TensorFlow 2 and Keras, Second Edition is a completely updated edition of the bestselling guide to the advanced deep learning techniques available today. Revised for TensorFlow 2.x, this edition introduces you to the practical side of deep learning with new chapters on unsupervised learning using mutual information, object detection (SSD), and semantic segmentation (FCN and PSPNet), further allowing you to create your own cutting-edge AI projects. Using Keras as an open-source deep learning library, the book features hands-on projects that show you how to create more effective AI with the most up-to-date techniques. Starting with an overview of multi-layer perceptrons (MLPs), convolutional neural networks (CNNs), and recurrent neural networks (RNNs), the book then introduces more cutting-edge techniques as you explore deep neural network architectures, including ResNet and DenseNet, and how to create autoencoders. You will then learn about GANs, and how they can unlock new levels of AI performance. Next, you'll discover how a variational autoencoder (VAE) is implemented, and how GANs and VAEs have the generative power to synthesize data that can be extremely convincing to humans. You'll also learn to implement DRL such as Deep Q-Learning and Policy Gradient Methods, which are critical to many modern results in AI. What you will learnUse mutual information maximization techniques to perform unsupervised learningUse segmentation to identify the pixel-wise class of each object in an imageIdentify both the bounding box and class of objects in an image using object detectionLearn the building blocks for advanced techniques - MLPs, CNN, and RNNsUnderstand deep neural networks - including ResNet and DenseNetUnderstand and build autoregressive models - autoencoders, VAEs, and GANsDiscover and implement deep reinforcement learning methodsWho this book is for This is not an introductory book, so fluency with Python is required. The reader should also be familiar with some machine learning approaches, and practical experience with DL will also be helpful. Knowledge of Keras or TensorFlow 2.0 is not required but is recommended.

Learning TensorFlow "O'Reilly Media, Inc."

Concepts, tools, and techniques to explore deep learning architectures and methodologies Key FeaturesExplore advanced deep learning architectures using various datasets and frameworksImplement deep architectures for neural network models such as CNN, RNN, GAN, and many moreDiscover design patterns and different challenges for various deep learning architecturesBook Description Deep learning architectures are composed of multilevel nonlinear operations that represent high-level abstractions; this allows you to learn useful feature representations from the data. This book will help you learn and implement deep learning architectures to resolve various deep

learning research problems. Hands-On Deep Learning Architectures with Python explains the essential learning algorithms used for deep and shallow architectures. Packed with practical implementations and ideas to help you build efficient artificial intelligence systems (AI), this book will help you learn how neural networks play a major role in building deep architectures. You will understand various deep learning architectures (such as AlexNet, VGG Net, GoogleNet) with easy-to-follow code and diagrams. In addition to this, the book will also guide you in building and training various deep architectures such as the Boltzmann mechanism, autoencoders, convolutional neural networks (CNNs), recurrent neural networks (RNNs), natural language processing (NLP), GAN, and more—all with practical implementations. By the end of this book, you will be able to construct deep models using popular frameworks and datasets with the required design patterns for each architecture. You will be ready to explore the potential of deep architectures in today's world. What you will learnImplement CNNs, RNNs, and other commonly used architectures with PythonExplore architectures such as VGGNet, AlexNet, and GoogLeNetBuild deep learning architectures for AI applications such as face and image recognition, fraud detection, and many moreUnderstand the architectures and applications of Boltzmann machines and autoencoders with concrete examples Master artificial intelligence and neural network concepts and apply them to your architectureUnderstand deep learning architectures for mobile and embedded systemsWho this book is for If you're a data scientist, machine learning developer/engineer, or deep learning practitioner, or are curious about AI and want to upgrade your knowledge of various deep learning architectures, this book will appeal to you. You are expected to have some knowledge of statistics and machine learning algorithms to get the best out of this book

Programming with TensorFlow Packt Publishing Ltd

Apply neural network architectures to build state-of-the-art computer vision applications using the Python programming language Key FeaturesGain a fundamental understanding of advanced computer vision and neural network models in use todayCover tasks such as low-level vision, image classification, and object detectionDevelop deep learning models on cloud platforms and optimize them using TensorFlow Lite and the OpenVINO toolkitBook 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 learnExplore methods of feature extraction and image retrieval and visualize different layers of the neural network modelUse TensorFlow for various visual search methods for real-world scenariosBuild neural networks or adjust parameters to optimize the performance of modelsUnderstand TensorFlow DeepLab to perform semantic segmentation on images and DCGAN for image inpaintingEvaluate your model and optimize and integrate it into your application to operate at scaleGet up to speed with techniques for performing manual and automated image annotationWho 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.

Related with Mobilenet Tensorflow:

- Human Organ Anatomy Chart : [click here](#)