

# Tensorflow

This is likewise one of the factors by obtaining the soft documents of this **Tensorflow** by online. You might not require more period to spend to go to the ebook initiation as well as search for them. In some cases, you likewise pull off not discover the revelation Tensorflow that you are looking for. It will completely squander the time.

However below, next you visit this web page, it will be consequently totally easy to get as competently as download guide Tensorflow

It will not admit many mature as we accustom before. You can realize it while piece of legislation something else at home and even in your workplace. for that reason easy! So, are you question? Just exercise just what we come up with the money for below as competently as review **Tensorflow** what you later than to read!

**Machine Learning Using TensorFlow Cookbook** - Alexia Audevert 2021-02-08

Comprehensive recipes to give you valuable insights on Transformers, Reinforcement Learning, and more Key Features Deep Learning solutions from Kaggle Masters and Google Developer Experts Get to grips with the fundamentals including variables, matrices, and data sources Learn advanced techniques to make your algorithms faster and more accurate Book Description The independent recipes in Machine Learning Using TensorFlow Cookbook will teach you how to perform complex data computations and gain valuable insights into your data. Dive into recipes on training models, model evaluation, sentiment analysis, regression analysis, artificial neural networks, and deep learning - each using Google's machine learning library, TensorFlow. This cookbook covers the fundamentals of the TensorFlow library, including variables, matrices, and various data sources.

You'll discover real-world implementations of Keras and TensorFlow and learn how to use estimators to train linear models and boosted trees, both for classification and regression. Explore the practical applications of a variety of deep learning architectures, such as recurrent neural networks and Transformers, and see how they can be used to solve computer vision and natural language processing (NLP) problems. With the help of this book, you will be proficient in using TensorFlow, understand deep learning from the basics, and be able to implement machine learning algorithms in real-world scenarios. What you will learn Take TensorFlow into production Implement and fine-tune Transformer models for various NLP tasks Apply reinforcement learning algorithms using the TF-Agents framework Understand linear regression techniques and use Estimators to train linear models Execute neural networks and improve predictions on tabular data Master convolutional neural networks and

recurrent neural networks through practical recipes Who this book is for If you are a data scientist or a machine learning engineer, and you want to skip detailed theoretical explanations in favor of building production-ready machine learning models using TensorFlow, this book is for you. Basic familiarity with Python, linear algebra, statistics, and machine learning is necessary to make the most out of this book.

*Deep Learning with TensorFlow* - Giancarlo Zaccone  
2018-03-30

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.

**Pro Deep Learning with TensorFlow** - Santanu Pattanayak  
2017-12-06

Deploy deep learning solutions in production with ease using TensorFlow. You'll also develop the mathematical understanding and intuition required to invent new deep learning architectures and solutions on your own. Pro Deep Learning with TensorFlow provides practical, hands-on expertise so you can learn deep learning from scratch and deploy meaningful deep learning solutions. This book will allow you to get up to speed quickly using TensorFlow and to optimize different deep learning architectures. All of the practical aspects of deep learning that are relevant in any industry are

emphasized in this book. You will be able to use the prototypes demonstrated to build new deep learning applications. The code presented in the book is available in the form of iPython notebooks and scripts which allow you to try out examples and extend them in interesting ways. You will be equipped with the mathematical foundation and scientific knowledge to pursue research in this field and give back to the community. What You'll Learn Understand full stack deep learning using TensorFlow and gain a solid mathematical foundation for deep learning Deploy complex deep learning solutions in production using TensorFlow Carry out research on deep learning and perform experiments using TensorFlow Who This Book Is For Data scientists and machine learning professionals, software developers, graduate students, and open source enthusiasts

*TensorFlow For Dummies* - Matthew Scarpino 2018-04-03

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.

*Applied Neural Networks with TensorFlow 2* - Orhan Gazi Yalçın 2020-11-30

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.

**TensorFlow in Action** - Thushan Ganegedara 2022-10-18  
Unlock the TensorFlow design secrets behind successful deep learning applications! Deep learning StackOverflow contributor Thushan Ganegedara teaches you the new features of TensorFlow 2 in this hands-on guide. In TensorFlow in Action you will learn: Fundamentals of TensorFlow Implementing deep learning networks Picking a high-level Keras API for model building with confidence Writing comprehensive end-to-end data pipelines Building models for computer vision and natural language processing Utilizing pretrained NLP models Recent algorithms including transformers, attention models, and ELMo In TensorFlow in Action, you'll dig into the newest version of Google's amazing TensorFlow framework as you learn to create incredible deep learning applications. Author Thushan Ganegedara uses quirky stories, practical examples, and behind-the-scenes explanations to demystify concepts otherwise trapped in dense academic papers. As you dive into modern deep learning techniques like transformer and attention models, you'll benefit from the unique insights of a top StackOverflow contributor for deep learning and NLP. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Google's TensorFlow framework sits at the

heart of modern deep learning. Boasting practical features like multi-GPU support, network data visualization, and easy production pipelines using TensorFlow Extended (TFX), TensorFlow provides the most efficient path to professional AI applications. And the Keras library, fully integrated into TensorFlow 2, makes it a snap to build and train even complex models for vision, language, and more. About the book TensorFlow in Action teaches you to construct, train, and deploy deep learning models using TensorFlow 2. In this practical tutorial, you'll build reusable skill hands-on as you create production-ready applications such as a French-to-English translator and a neural network that can write fiction. You'll appreciate the in-depth explanations that go from DL basics to advanced applications in NLP, image processing, and MLOps, complete with important details that you'll return to reference over and over. What's inside Covers TensorFlow 2.9 Recent algorithms including transformers, attention models, and ELMo Build on pretrained models Writing end-to-end data pipelines with TFX About the reader For Python programmers with basic deep learning skills. About the author Thushan Ganegedara is a senior ML engineer at Canva and TensorFlow expert. He holds a PhD in machine learning from the University of Sydney. Table of Contents PART 1 FOUNDATIONS OF TENSORFLOW 2 AND DEEP LEARNING 1 The amazing world of TensorFlow 2 TensorFlow 2 3 Keras and data retrieval in TensorFlow 2 4 Dipping toes in deep learning 5 State-of-the-art in deep learning: Transformers PART 2 LOOK MA, NO HANDS! DEEP NETWORKS IN THE REAL WORLD 6 Teaching machines to see: Image classification with CNNs 7 Teaching machines to see better: Improving CNNs and making them confess 8 Telling things apart: Image segmentation 9 Natural

language processing with TensorFlow: Sentiment analysis  
10 Natural language processing with TensorFlow: Language  
modeling PART 3 ADVANCED DEEP NETWORKS FOR COMPLEX  
PROBLEMS 11 Sequence-to-sequence learning: Part 1 12  
Sequence-to-sequence learning: Part 2 13 Transformers 14  
TensorBoard: Big brother of TensorFlow 15 TFX: MLOps and  
deploying models with TensorFlow

## **Machine Learning for Economics and Finance in TensorFlow 2 - Isaiah Hull 2020-11-26**

Work on economic problems and solutions with tools from machine learning. ML has taken time to move into the space of academic economics. This is because empirical work in economics is concentrated on the identification of causal relationships in parsimonious statistical models; whereas machine learning is oriented towards prediction and is generally uninterested in either causality or parsimony. That leaves a gap for both students and professionals in the economics industry without a standard reference. This book focuses on economic problems with an empirical dimension, where machine learning methods may offer something of value. This includes coverage of a variety of discriminative deep learning models (DNNs, CNNs, RNNs, LSTMs, the Transformer Model, etc.), generative machine learning models, random forests, gradient boosting, clustering, and feature extraction. You'll also learn about the intersection of empirical methods in economics and machine learning, including regression analysis, text analysis, and dimensionality reduction methods, such as principal components analysis. TensorFlow offers a toolset that can be used to setup and solve any mathematical model, including those commonly used in economics. This book is structured to teach through a sequence of complete examples, each framed in terms of a

specific economic problem of interest or topic. Otherwise complicated content is then distilled into accessible examples, so you can use TensorFlow to solve workhorse models in economics and finance. What You'll Learn Define, train, and evaluate machine learning models in TensorFlow 2 Apply fundamental concepts in machine learning, such as deep learning and natural language processing, to economic and financial problems Solve workhorse models in economics and finance Who This Book Is For Students and data scientists working in the economics industry. Academic economists and social scientists who have an interest in machine learning are also likely to find this book useful.

## Hands-On Image Generation with TensorFlow - Soon Yau Cheong 2020-12-24

Implement various state-of-the-art architectures, such as GANs and autoencoders, for image generation using TensorFlow 2.x from scratch Key Features Understand the different architectures for image generation, including autoencoders and GANs Build models that can edit an image of your face, turn photos into paintings, and generate photorealistic images Discover how you can build deep neural networks with advanced TensorFlow 2.x features Book Description The emerging field of Generative Adversarial Networks (GANs) has made it possible to generate indistinguishable images from existing datasets. With this hands-on book, you'll not only develop image generation skills but also gain a solid understanding of the underlying principles. Starting with an introduction to the fundamentals of image generation using TensorFlow, this book covers Variational Autoencoders (VAEs) and GANs. You'll discover how to build models for different applications as you get to grips with performing face swaps using

deepfakes, neural style transfer, image-to-image translation, turning simple images into photorealistic images, and much more. You'll also understand how and why to construct state-of-the-art deep neural networks using advanced techniques such as spectral normalization and self-attention layer before working with advanced models for face generation and editing. You'll also be introduced to photo restoration, text-to-image synthesis, video retargeting, and neural rendering. Throughout the book, you'll learn to implement models from scratch in TensorFlow 2.x, including PixelCNN, VAE, DCGAN, WGAN, pix2pix, CycleGAN, StyleGAN, GauGAN, and BigGAN. By the end of this book, you'll be well versed in TensorFlow and be able to implement image generative technologies confidently. What you will learn

Train on face datasets and use them to explore latent spaces for editing new faces  
Get to grips with swapping faces with deepfakes  
Perform style transfer to convert a photo into a painting  
Build and train pix2pix, CycleGAN, and BicycleGAN for image-to-image translation  
Use iGAN to understand manifold interpolation and GauGAN to turn simple images into photorealistic images  
Become well versed in attention generative models such as SAGAN and BigGAN  
Generate high-resolution photos with Progressive GAN and StyleGAN  
Who this book is for  
The Hands-On Image Generation with TensorFlow book is for deep learning engineers, practitioners, and researchers who have basic knowledge of convolutional neural networks and want to learn various image generation techniques using TensorFlow 2.x. You'll also find this book useful if you are an image processing professional or computer vision engineer looking to explore state-of-the-art architectures to improve and enhance images and videos. Knowledge of Python and TensorFlow will help you to get

the best out of this book.

*TensorFlow Pocket Primer* - Oswald Campesato 2019-05-09  
As part of the best-selling Pocket Primer series, this book is designed to introduce beginners to TensorFlow 1.x fundamentals for basic machine learning algorithms in TensorFlow. It is intended to be a fast-paced introduction to various "core" features of TensorFlow, with code samples that cover deep learning and TensorFlow basics. The material in the chapters illustrates how to solve a variety of tasks after which you can do further reading to deepen your knowledge. Companion files with all of the code samples are available for downloading from the publisher by writing to [info@merclearning.com](mailto:info@merclearning.com). Features: Uses Python for code samples  
Covers TensorFlow APIs and Datasets  
Assumes the reader has very limited experience  
Companion files with all of the source code examples (download from the publisher)

**Learning TensorFlow.js** - Gant Laborde 2021-05-10  
Given the demand for AI and the ubiquity of JavaScript, TensorFlow.js was inevitable. With this Google framework, seasoned AI veterans and web developers alike can help propel the future of AI-driven websites. In this guide, author Gant Laborde--Google Developer Expert in machine learning and the web--provides a hands-on end-to-end approach to TensorFlow.js fundamentals for a broad technical audience that includes data scientists, engineers, web developers, students, and researchers. You'll begin by working through some basic examples in TensorFlow.js before diving deeper into neural network architectures, DataFrames, TensorFlow Hub, model conversion, transfer learning, and more. Once you finish this book, you'll know how to build and deploy production-ready deep learning systems with

TensorFlow.js. Explore tensors, the most fundamental structure of machine learning Convert data into tensors and back with a real-world example Combine AI with the web using TensorFlow.js Use resources to convert, train, and manage machine learning data Build and train your own training models from scratch

*Learn TensorFlow in 24 Hours* - Alex Nordeen 2020-10-31  
Tensorflow is the most popular Deep Learning Library out there. It has fantastic graph computations feature which helps data scientist to visualize his designed neural network using TensorBoard. This Machine learning library supports both Convolution as well as Recurrent Neural network. It supports parallel processing on CPU as well as GPU. Prominent machine learning algorithms supported by TensorFlow are Deep Learning Classification, wibe & deep, Boston Tree amongst others. The book is very hands-on and gives you industry ready deep learnings practices. Here is what is covered in the book – Table Of Content Chapter 1: What is Deep learning? Chapter 2: Machine Learning vs Deep Learning Chapter 3: What is TensorFlow? Chapter 4: Comparison of Deep Learning Libraries Chapter 5: How to Download and Install TensorFlow Windows and Mac Chapter 6: Jupyter Notebook Tutorial Chapter 7: Tensorflow on AWS Chapter 8: TensorFlow Basics: Tensor, Shape, Type, Graph, Sessions & Operators Chapter 9: Tensorboard: Graph Visualization with Example Chapter 10: NumPy Chapter 11: Pandas Chapter 12: Scikit-Learn Chapter 13: Linear Regression Chapter 14: Linear Regression Case Study Chapter 15: Linear Classifier in TensorFlow Chapter 16: Kernel Methods Chapter 17: TensorFlow ANN (Artificial Neural Network) Chapter 18: ConvNet(Convolutional Neural Network): TensorFlow Image Classification Chapter 19: Autoencoder with TensorFlow Chapter 20: RNN(Recurrent

Neural Network) TensorFlow  
Programming with TensorFlow - Kolla Bhanu Prakash  
2021-01-22

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 deep learning, Natural Language Processing (NLP), speech recognition, and general predictive analytics. The book provides a hands-on approach to TensorFlow fundamentals for a broad technical audience—from data scientists and engineers to students and researchers. The authors begin by working through some basic examples in TensorFlow before diving deeper into topics such as CNN, RNN, LSTM, and GNN. The book is written for those who want to build powerful, robust, and accurate predictive models with the power of TensorFlow, combined with other open source Python libraries. The authors demonstrate TensorFlow projects on Single Board Computers (SBCs).

Machine Learning with TensorFlow - Nishant Shukla  
2018-02-12

Summary Machine Learning with TensorFlow gives readers a solid foundation in machine-learning concepts plus hands-on experience coding TensorFlow with Python. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology TensorFlow, Google's library for large-scale machine learning, simplifies often-complex computations by representing them as graphs and efficiently mapping parts of the graphs to machines in a cluster or to the processors of a single machine. About the Book Machine Learning with TensorFlow gives readers a solid foundation in machine-learning concepts plus hands-on experience coding TensorFlow with Python. You'll learn the basics by working with classic

prediction, classification, and clustering algorithms. Then, you'll move on to the money chapters: exploration of deep-learning concepts like autoencoders, recurrent neural networks, and reinforcement learning. Digest this book and you will be ready to use TensorFlow for machine-learning and deep-learning applications of your own. What's Inside Matching your tasks to the right machine-learning and deep-learning approaches Visualizing algorithms with TensorBoard Understanding and using neural networks About the Reader Written for developers experienced with Python and algebraic concepts like vectors and matrices. About the Author Author Nishant Shukla is a computer vision researcher focused on applying machine-learning techniques in robotics. Senior technical editor, Kenneth Fricklas, is a seasoned developer, author, and machine-learning practitioner. Table of Contents PART 1 - YOUR MACHINE-LEARNING RIG A machine-learning odyssey TensorFlow essentials PART 2 - CORE LEARNING ALGORITHMS Linear regression and beyond A gentle introduction to classification Automatically clustering data Hidden Markov models PART 3 - THE NEURAL NETWORK PARADIGM A peek into autoencoders Reinforcement learning Convolutional neural networks Recurrent neural networks Sequence-to-sequence models for chatbots Utility landscape

**TensorFlow Machine Learning Projects** - Ankit Jain  
2018-11-30

Implement TensorFlow's offerings such as TensorBoard, TensorFlow.js, TensorFlow Probability, and TensorFlow Lite to build smart automation projects Key Features Use machine learning and deep learning principles to build real-world projects Get to grips with TensorFlow's impressive range of module offerings Implement projects

on GANs, reinforcement learning, and capsule network Book Description TensorFlow has transformed the way machine learning is perceived. TensorFlow Machine Learning Projects teaches you how to exploit the benefits--simplicity, efficiency, and flexibility--of using TensorFlow in various real-world projects. With the help of this book, you'll not only learn how to build advanced projects using different datasets but also be able to tackle common challenges using a range of libraries from the TensorFlow ecosystem. To start with, you'll get to grips with using TensorFlow for machine learning projects; you'll explore a wide range of projects using TensorForest and TensorBoard for detecting exoplanets, TensorFlow.js for sentiment analysis, and TensorFlow Lite for digit classification. As you make your way through the book, you'll build projects in various real-world domains, incorporating natural language processing (NLP), the Gaussian process, autoencoders, recommender systems, and Bayesian neural networks, along with trending areas such as Generative Adversarial Networks (GANs), capsule networks, and reinforcement learning. You'll learn how to use the TensorFlow on Spark API and GPU-accelerated computing with TensorFlow to detect objects, followed by how to train and develop a recurrent neural network (RNN) model to generate book scripts. By the end of this book, you'll have gained the required expertise to build full-fledged machine learning projects at work. What you will learn Understand the TensorFlow ecosystem using various datasets and techniques Create recommendation systems for quality product recommendations Build projects using CNNs, NLP, and Bayesian neural networks Play Pac-Man using deep reinforcement learning Deploy scalable TensorFlow-based machine learning systems Generate your



own book script using RNNs Who this book is for TensorFlow Machine Learning Projects is for you if you are a data analyst, data scientist, machine learning professional, or deep learning enthusiast with basic knowledge of TensorFlow. This book is also for you if you want to build end-to-end projects in the machine learning domain using supervised, unsupervised, and reinforcement learning techniques

[Learning TensorFlow.js](#) - Gant Laborde 2021-05-10

Given the demand for AI and the ubiquity of JavaScript, TensorFlow.js was inevitable. With this Google framework, seasoned AI veterans and web developers alike can help propel the future of AI-driven websites. In this guide, author Gant Laborde--Google Developer Expert in machine learning and the web--provides a hands-on end-to-end approach to TensorFlow.js fundamentals for a broad technical audience that includes data scientists, engineers, web developers, students, and researchers. You'll begin by working through some basic examples in TensorFlow.js before diving deeper into neural network architectures, DataFrames, TensorFlow Hub, model conversion, transfer learning, and more. Once you finish this book, you'll know how to build and deploy production-ready deep learning systems with TensorFlow.js. Explore tensors, the most fundamental structure of machine learning Convert data into tensors and back with a real-world example Combine AI with the web using TensorFlow.js Use resources to convert, train, and manage machine learning data Build and train your own training models from scratch

TensorFlow.js. Explore tensors, the most fundamental structure of machine learning Convert data into tensors and back with a real-world example Combine AI with the web using TensorFlow.js Use resources to convert, train, and manage machine learning data Build and train your own training models from scratch

**Mastering TensorFlow 1.x** - Armando Fandango 2018-01-22

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 TensorFlow 1.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 Machine Learning with Scikit-Learn, Keras, and TensorFlow** - Aurélien Géron 2019-09-05

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

**The TensorFlow Workshop** - Matthew Moocarme 2021-12-15 Get started with TensorFlow fundamentals to build and train deep learning models with real-world data, practical exercises, and challenging activities Key Features Understand the fundamentals of tensors, neural networks, and deep learning Discover how to implement and fine-tune deep learning models for real-world datasets Build your experience and confidence with hands-on exercises and activities Book Description Getting to grips with tensors, deep learning, and neural networks can be intimidating and confusing for anyone, no matter their experience level. The breadth of information out there, often written at a very high level and aimed at advanced practitioners, can make getting started even more challenging. If this sounds familiar to you, The TensorFlow Workshop is here to help. Combining clear explanations, realistic examples, and plenty of hands-on practice, it'll quickly get you up and running. You'll

start off with the basics – learning how to load data into TensorFlow, perform tensor operations, and utilize common optimizers and activation functions. As you progress, you'll experiment with different TensorFlow development tools, including TensorBoard, TensorFlow Hub, and Google Colab, before moving on to solve regression and classification problems with sequential models. Building on this solid foundation, you'll learn how to tune models and work with different types of neural network, getting hands-on with real-world deep learning applications such as text encoding, temperature forecasting, image augmentation, and audio processing. By the end of this deep learning book, you'll have the skills, knowledge, and confidence to tackle your own ambitious deep learning projects with TensorFlow. What you will learn

- Get to grips with TensorFlow's mathematical operations
- Pre-process a wide variety of tabular, sequential, and image data
- Understand the purpose and usage of different deep learning layers
- Perform hyperparameter-tuning to prevent overfitting of training data
- Use pre-trained models to speed up the development of learning models
- Generate new data based on existing patterns using generative models

Who this book is for This TensorFlow book is for anyone who wants to develop their understanding of deep learning and get started building neural networks with TensorFlow. Basic knowledge of Python programming and its libraries, as well as a general understanding of the fundamentals of data science and machine learning, will help you grasp the topics covered in this book more easily.

**TinyML** - Pete Warden 2019-12-16

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

**Intelligent Mobile Projects with TensorFlow** - Jeff Tang 2018-05-22

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.

**TensorFlow 2.0 Quick Start Guide** - Tony Holdroyd  
2019-03-29

TensorFlow is one of the most popular machine learning frameworks in Python. With this book, you will improve your knowledge of some of the latest TensorFlow features and will be able to perform supervised and unsupervised machine learning and also train neural networks.

*TensorFlow in Action* - Thushan Ganegedara 2022-11-01

Unlock the TensorFlow design secrets behind successful deep learning applications! Deep learning StackOverflow contributor Thushan Ganegedara teaches you the new features of TensorFlow 2 in this hands-on guide. In *TensorFlow in Action* you will learn:

- Fundamentals of TensorFlow
- Implementing deep learning networks
- Picking a high-level Keras API for model building with confidence
- Writing comprehensive end-to-end data pipelines
- Building models for computer vision and natural language processing
- Utilizing pretrained NLP models
- Recent algorithms including transformers, attention models, and ElMo

In *TensorFlow in Action*, you'll dig into the newest version of Google's amazing TensorFlow framework as you learn to create incredible deep learning applications. Author Thushan Ganegedara uses quirky stories, practical examples, and behind-the-scenes explanations to demystify concepts otherwise trapped in dense academic papers. As you dive into modern deep learning techniques like transformer and attention models, you'll benefit from the unique insights of a top StackOverflow contributor for deep learning and NLP. About the technology Google's TensorFlow framework sits at the heart of modern deep learning. Boasting practical features like multi-GPU support, network data visualization, and easy production pipelines using

TensorFlow Extended (TFX), TensorFlow provides the most efficient path to professional AI applications. And the Keras library, fully integrated into TensorFlow 2, makes it a snap to build and train even complex models for vision, language, and more. About the book TensorFlow in Action teaches you to construct, train, and deploy deep learning models using TensorFlow 2. In this practical tutorial, you'll build reusable skill hands-on as you create production-ready applications such as a French-to-English translator and a neural network that can write fiction. You'll appreciate the in-depth explanations that go from DL basics to advanced applications in NLP, image processing, and MLOps, complete with important details that you'll return to reference over and over. What's inside Covers TensorFlow 2.9 Recent algorithms including transformers, attention models, and ElMo Build on pretrained models Writing end-to-end data pipelines with TFX About the reader For Python programmers with basic deep learning skills. About the author Thushan Ganegedara is a senior ML engineer at Canva and TensorFlow expert. He holds a PhD in machine learning from the University of Sydney. Table of Contents PART 1 FOUNDATIONS OF TENSORFLOW 2 AND DEEP LEARNING 1 The amazing world of TensorFlow 2 TensorFlow 2 3 Keras and data retrieval in TensorFlow 2 4 Dipping toes in deep learning 5 State-of-the-art in deep learning: Transformers PART 2 LOOK MA, NO HANDS! DEEP NETWORKS IN THE REAL WORLD 6 Teaching machines to see: Image classification with CNNs 7 Teaching machines to see better: Improving CNNs and making them confess 8 Telling things apart: Image segmentation 9 Natural language processing with TensorFlow: Sentiment analysis 10 Natural language processing with TensorFlow: Language modeling PART 3 ADVANCED DEEP NETWORKS FOR COMPLEX

PROBLEMS 11 Sequence-to-sequence learning: Part 1 12 Sequence-to-sequence learning: Part 2 13 Transformers 14 TensorBoard: Big brother of TensorFlow 15 TFX: MLOps and deploying models with TensorFlow

**Learning TensorFlow** - Tom Hope 2017-08-09

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

**Hands-On Computer Vision with TensorFlow 2** - Benjamin Planche 2019-05-30

A practical guide to building high performance systems

for object detection, segmentation, video processing, smartphone applications, and more Key Features Discover how to build, train, and serve your own deep neural networks with TensorFlow 2 and Keras Apply modern solutions to a wide range of applications such as object detection and video analysis Learn how to run your models on mobile devices and web pages and improve their performance Book 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 learn Create your own neural networks from scratch Classify images with modern architectures including Inception and ResNet Detect and segment objects in images with YOLO, Mask R-CNN, and U-Net Tackle problems faced when developing self-driving cars and facial emotion recognition systems Boost your application's performance with transfer learning, GANs, and domain adaptation Use recurrent neural networks (RNNs) for video analysis Optimize and deploy your networks on mobile devices and in the browser Who 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.

**Learn TensorFlow 2.0** - Pramod Singh 2019-12-17  
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.

*Deep Learning with TensorFlow* - Giancarlo Zaccone  
2017-04-24

Delve into neural networks, implement deep learning algorithms, and explore layers of data abstraction with the help of this comprehensive TensorFlow guide About This Book 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 Real-world contextualization through some deep learning problems concerning research and application Who This Book Is For The book is intended for a general audience of 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 You Will Learn Learn about machine learning landscapes along with the historical development and progress of deep learning Learn about deep machine intelligence and GPU computing with the latest TensorFlow 1.x Access public datasets

and utilize them using TensorFlow to load, process, and transform data Use TensorFlow on real-world datasets, including images, text, and more Learn how to evaluate the performance of your deep learning models Using 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 In Detail Deep learning is the step that comes after machine learning, and has more advanced implementations. Machine learning is not just for academics anymore, but is becoming a mainstream practice through wide adoption, and deep learning has taken the front seat. As a data scientist, if you want to explore data abstraction layers, this book will be your guide. This book shows how this can be exploited in the real world with complex raw data using TensorFlow 1.x. Throughout the book, you'll learn how to implement deep learning algorithms for machine learning systems and integrate them into your product offerings, including search, image recognition, and language processing. Additionally, you'll learn how to analyze and improve the performance of deep learning models. This can be done by comparing algorithms against benchmarks, along with machine intelligence, to learn from the information and determine ideal behaviors within a specific context. After finishing the book, you will be familiar with machine learning techniques, in particular the use of TensorFlow for deep learning, and will be ready to apply your knowledge to research or commercial projects. Style and approach This step-by-step guide will explore common, and not so common, deep neural networks and show how these can be exploited in the real world with complex raw data. With the help of practical examples,

you will learn how to implement different types of neural nets to build smart applications related to text, speech, and image data processing.

*Machine Learning with TensorFlow, Second Edition* - Mattmann A. Chris 2021-02-02

Updated with new code, new projects, and new chapters, *Machine Learning with TensorFlow, Second Edition* gives readers a solid foundation in machine-learning concepts and the TensorFlow library. Summary Updated with new code, new projects, and new chapters, *Machine Learning with TensorFlow, Second Edition* gives readers a solid foundation in machine-learning concepts and the TensorFlow library. Written by NASA JPL Deputy CTO and Principal Data Scientist Chris Mattmann, all examples are accompanied by downloadable Jupyter Notebooks for a hands-on experience coding TensorFlow with Python. New and revised content expands coverage of core machine learning algorithms, and advancements in neural networks such as VGG-Face facial identification classifiers and deep speech classifiers. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Supercharge your data analysis with machine learning! ML algorithms automatically improve as they process data, so results get better over time. You don't have to be a mathematician to use ML: Tools like Google's TensorFlow library help with complex calculations so you can focus on getting the answers you need. About the book *Machine Learning with TensorFlow, Second Edition* is a fully revised guide to building machine learning models using Python and TensorFlow. You'll apply core ML concepts to real-world challenges, such as sentiment analysis, text classification, and image recognition. Hands-on examples illustrate neural network techniques for deep speech

processing, facial identification, and auto-encoding with CIFAR-10. What's inside *Machine Learning with TensorFlow* Choosing the best ML approaches Visualizing algorithms with TensorBoard Sharing results with collaborators Running models in Docker About the reader Requires intermediate Python skills and knowledge of general algebraic concepts like vectors and matrices. Examples use the super-stable 1.15.x branch of TensorFlow and TensorFlow 2.x. About the author Chris Mattmann is the Division Manager of the Artificial Intelligence, Analytics, and Innovation Organization at NASA Jet Propulsion Lab. The first edition of this book was written by Nishant Shukla with Kenneth Fricklas. Table of Contents PART 1 - YOUR MACHINE-LEARNING RIG 1 A machine-learning odyssey 2 TensorFlow essentials PART 2 - CORE LEARNING ALGORITHMS 3 Linear regression and beyond 4 Using regression for call-center volume prediction 5 A gentle introduction to classification 6 Sentiment classification: Large movie-review dataset 7 Automatically clustering data 8 Inferring user activity from Android accelerometer data 9 Hidden Markov models 10 Part-of-speech tagging and word-sense disambiguation PART 3 - THE NEURAL NETWORK PARADIGM 11 A peek into autoencoders 12 Applying autoencoders: The CIFAR-10 image dataset 13 Reinforcement learning 14 Convolutional neural networks 15 Building a real-world CNN: VGG-Face and VGG-Face Lite 16 Recurrent neural networks 17 LSTMs and automatic speech recognition 18 Sequence-to-sequence models for chatbots 19 Utility landscape

**Natural Language Processing with TensorFlow** - Thushan Ganegedara 2018-05-31

Write modern natural language processing applications using deep learning algorithms and TensorFlow Key Features Focuses on more efficient natural language



processing using TensorFlow Covers NLP as a field in its own right to improve understanding for choosing TensorFlow tools and other deep learning approaches Provides choices for how to process and evaluate large unstructured text datasets Learn to apply the TensorFlow toolbox to specific tasks in the most interesting field in artificial intelligence Book Description Natural language processing (NLP) supplies the majority of data available to deep learning applications, while TensorFlow is the most important deep learning framework currently available. Natural Language Processing with TensorFlow brings TensorFlow and NLP together to give you invaluable tools to work with the immense volume of unstructured data in today's data streams, and apply these tools to specific NLP tasks. Thushan Ganegedara starts by giving you a grounding in NLP and TensorFlow basics. You'll then learn how to use Word2vec, including advanced extensions, to create word embeddings that turn sequences of words into vectors accessible to deep learning algorithms. Chapters on classical deep learning algorithms, like convolutional neural networks (CNN) and recurrent neural networks (RNN), demonstrate important NLP tasks as sentence classification and language generation. You will learn how to apply high-performance RNN models, like long short-term memory (LSTM) cells, to NLP tasks. You will also explore neural machine translation and implement a neural machine translator. After reading this book, you will gain an understanding of NLP and you'll have the skills to apply TensorFlow in deep learning NLP applications, and how to perform specific NLP tasks. What you will learn Core concepts of NLP and various approaches to natural language processing How to solve NLP tasks by applying TensorFlow functions to create neural networks Strategies to

process large amounts of data into word representations that can be used by deep learning applications Techniques for performing sentence classification and language generation using CNNs and RNNs About employing state-of-the art advanced RNNs, like long short-term memory, to solve complex text generation tasks How to write automatic translation programs and implement an actual neural machine translator from scratch The trends and innovations that are paving the future in NLP Who this book is for This book is for Python developers with a strong interest in deep learning, who want to learn how to leverage TensorFlow to simplify NLP tasks. Fundamental Python skills are assumed, as well as some knowledge of machine learning and undergraduate-level calculus and linear algebra. No previous natural language processing experience required, although some background in NLP or computational linguistics will be helpful.

[Deep Learning with TensorFlow 2 and Keras - Second Edition](#) - Antonio Gulli 2019-12-20

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.

**Mastering Computer Vision with TensorFlow 2.x** -

Krishnendu Kar 2020-05-15

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 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.

*TensorFlow Machine Learning* - Benjamin Smith 2021-01-04 Have you ever wondered how machine learning works? These days, machine learning, deep learning and neural nets are common terms and they are here to stay as a part of our everyday language. Machine learning is not the easiest of topics to teach, purely because there is so much to it. Machine learning, deep learning and artificial intelligence are used in more applications than most humans even think about - email, Amazon, Netflix, Spotify, and other popular online marketplaces use machine learning to weed out spam emails and bring you recommendations based on your shopping or streaming preferences. Machine learning is used in healthcare, in finance, in just about every industry you can think of - it's here to stay, whether we like it or not. One of the most important parts of learning machine learning is

knowing which algorithm to choose and which library. Python is the most popular machine learning programming language and it has a huge advantage over other languages - the large amount of built-in libraries; three of the most important are TensorFlow, Keras and Scikit-Learn. And that's what this book is about - machine learning with TensorFlow, Keras and Scikit-learn. Here's what you will learn: -What machine learning is -How it applies in the real world -Different models and learning types -Different machine learning algorithms -Deep learning vs. machine learning -What TensorFlow is and how to use it -What TensorFlow comprises -Operators, variables, placeholders, and more -What Keras is and how to use it -Keras vs. TensorFlow -How to use Keras for linear regression -How to use Keras to build a neural net -What Scikit-Learn is and how to use it -Using Scikit-Learn to build regression and classification trees -How to build a random forest model -How to install Keras, TensorFlow and Scikit-Learn And much more! All the practical examples in the book use Python, so you are expected to need some knowledge of the language before you start. If you're looking to advanced your skills in machine learning, then this is the book for you! Grab your copy of this book today!

*TensorFlow Deep Learning Projects* - Alexey Grigorev 2018-03-28

Leverage the power of Tensorflow to design deep learning systems for a variety of real-world scenarios Key Features Build efficient deep learning pipelines using the popular Tensorflow framework Train neural networks such as ConvNets, generative models, and LSTMs Includes projects related to Computer Vision, stock prediction, chatbots and more Book Description TensorFlow is one of the most popular frameworks used for machine learning

and, more recently, deep learning. It provides a fast and efficient framework for training different kinds of deep learning models, with very high accuracy. This book is your guide to master deep learning with TensorFlow with the help of 10 real-world projects. TensorFlow Deep Learning Projects starts with setting up the right TensorFlow environment for deep learning. Learn to train different types of deep learning models using TensorFlow, including Convolutional Neural Networks, Recurrent Neural Networks, LSTMs, and Generative Adversarial Networks. While doing so, you will build end-to-end deep learning solutions to tackle different real-world problems in image processing, recommendation systems, stock prediction, and building chatbots, to name a few. You will also develop systems that perform machine translation, and use reinforcement learning techniques to play games. By the end of this book, you will have mastered all the concepts of deep learning and their implementation with TensorFlow, and will be able to build and train your own deep learning models with TensorFlow confidently. What you will learn

- Set up the TensorFlow environment for deep learning
- Construct your own ConvNets for effective image processing
- Use LSTMs for image caption generation
- Forecast stock prediction accurately with an LSTM architecture
- Learn what semantic matching is by detecting duplicate Quora questions
- Set up an AWS instance with TensorFlow to train GANs
- Train and set up a chatbot to understand and interpret human input
- Build an AI capable of playing a video game by itself –and win it!

Who this book is for This book is for data scientists, machine learning developers as well as deep learning practitioners, who want to build interesting deep learning projects that leverage the power of Tensorflow. Some understanding of machine

learning and deep learning, and familiarity with the TensorFlow framework is all you need to get started with this book.

**Machine Learning With Tensorflow** - Frank Millstein  
2020-07-03

Machine Learning with TensorFlow TensorFlow is a powerful open source software library for performing various numerical data flow graphs. With its powerful resources, TensorFlow is perfect for machine learning enthusiasts offering plenty of workspace where you can improve your machine learning techniques and build your own machine learning algorithms. Thanks to its capability, in recent times TensorFlow definitely has made its way into the software mainstream, so everyone who is interested in machine learnings definitely should consider getting hands on TensorFlow practices. With this book as your guide, you will get your hands on TensorFlow machine learning techniques, learn how to perform different neural network operations, learn how to deal with massive datasets and finally build your first machine learning model for data classification. Here Is a Preview of What You'll Learn Here... What is machine learning Main uses and benefits of machine learning How to get started with TensorFlow, installing and loading data Data flow graphs and basic TensorFlow expressions How to define your data flow graphs and how to use TensorBoard for data visualization Main TensorFlow operations and building tensors How to perform data transformation using different techniques How to build high performance data pipelines using TensorFlow Dataset framework How to create TensorFlow iterators Creating MNIST classifiers with one-hot transformation Get this book NOW and learn how to do various machine learning tasks using TensorFlow!

*Mastering TensorFlow 2.x* - Rajdeep 2022-03-24

Work with TensorFlow and Keras for real performance of deep learning

**KEY FEATURES**

- Combines theory and implementation with in-detail use-cases.
- Coverage on both, TensorFlow 1.x and 2.x with elaborated concepts.
- Exposure to Distributed Training, GANs and Reinforcement Learning.

**DESCRIPTION** Mastering TensorFlow 2.x is a must to read and practice if you are interested in building various kinds of neural networks with high level TensorFlow and Keras APIs. The book begins with the basics of TensorFlow and neural network concepts, and goes into specific topics like image classification, object detection, time series forecasting and Generative Adversarial Networks. While we are practicing TensorFlow 2.6 in this book, the version of Tensorflow will change with time; however you can still use this book to witness how Tensorflow outperforms. This book includes the use of a local Jupyter notebook and the use of Google Colab in various use cases including GAN and Image classification tasks. While you explore the performance of TensorFlow, the book also covers various concepts and in-detail explanations around reinforcement learning, model optimization and time series models.

**WHAT YOU WILL LEARN**

- Getting started with Tensorflow 2.x and basic building blocks.
- Get well versed in functional programming with TensorFlow.
- Practice Time Series analysis along with strong understanding of concepts.
- Get introduced to use of TensorFlow in Reinforcement learning and Generative Adversarial Networks.
- Train distributed models and how to optimize them.

**WHO THIS BOOK IS FOR** This book is designed for machine learning engineers, NLP engineers and deep learning practitioners who want to utilize the performance of TensorFlow in their ML and AI projects.

Readers are expected to have some familiarity with Tensorflow and the basics of machine learning would be helpful.

**TABLE OF CONTENTS**

1. Getting started with TensorFlow 2.x
2. Machine Learning with TensorFlow 2.x
3. Keras based APIs
4. Convolutional Neural Networks in Tensorflow
5. Text Processing with TensorFlow 2.x
6. Time Series Forecasting with TensorFlow 2.x
7. Distributed Training and DataInput pipelines
8. Reinforcement Learning
9. Model Optimization
10. Generative Adversarial Networks

*TensorFlow for Deep Learning* - Bharath Ramsundar  
2018-03-01

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  
**Hands-On Neural Networks with TensorFlow 2.0** - Paolo Galeone 2019-09-18

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.

**Neural Network Programming with TensorFlow** - Manpreet Singh Ghotra 2017-11-10

Neural Networks and their implementation decoded with TensorFlow About This Book Develop a strong background in neural network programming from scratch, using the popular Tensorflow library. Use Tensorflow to implement different kinds of neural networks – from simple feedforward neural networks to multilayered perceptrons, CNNs, RNNs and more. A highly practical guide including real-world datasets and use-cases to simplify your understanding of neural networks and their implementation. Who This Book Is For This book is meant for developers with a statistical background who want to work with neural networks. Though we will be using TensorFlow as the underlying library for neural networks, book can be used as a generic resource to bridge the gap between the math and the implementation of deep learning. If you have some understanding of

Tensorflow and Python and want to learn what happens at a level lower than the plain API syntax, this book is for you. What You Will Learn Learn Linear Algebra and mathematics behind neural network. Dive deep into Neural networks from the basic to advanced concepts like CNN, RNN Deep Belief Networks, Deep Feedforward Networks. Explore Optimization techniques for solving problems like Local minima, Global minima, Saddle points Learn through real world examples like Sentiment Analysis. Train different types of generative models and explore autoencoders. Explore TensorFlow as an example of deep learning implementation. In Detail If you're aware of the buzz surrounding the terms such as "machine learning," "artificial intelligence," or "deep learning," you might know what neural networks are. Ever wondered how they help in solving complex computational problem efficiently, or how to train efficient neural networks? This book will teach you just that. You will start by getting a quick overview of the popular TensorFlow library and how it is used to train different neural networks. You will get a thorough understanding of the fundamentals and basic math for neural networks and why TensorFlow is a popular choice Then, you will proceed to implement a simple feed forward neural network. Next you will master optimization techniques and algorithms for neural networks using TensorFlow. Further, you will learn to implement some more complex types of neural networks such as convolutional neural networks, recurrent neural networks, and Deep Belief Networks. In the course of the book, you will be working on real-world datasets to get a hands-on understanding of neural network programming. You will also get to train generative models and will learn the applications of autoencoders. By the end of this book, you will have

a fair understanding of how you can leverage the power of TensorFlow to train neural networks of varying complexities, without any hassle. While you are learning about various neural network implementations you will learn the underlying mathematics and linear algebra and how they map to the appropriate TensorFlow constructs. Style and Approach This book is designed to give you just the right number of concepts to back up the examples. With real-world use cases and problems solved, this book is a handy guide for you. Each concept is backed by a generic and real-world problem, followed by a variation, making you independent and able to solve any problem with neural networks. All of the content is demystified by a simple and straightforward approach.

### **TensorFlow: Powerful Predictive Analytics with**

**TensorFlow** - Md. Rezaul Karim 2018-03-14

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. Probabilistic Deep Learning - Oliver Duerr 2020-11-10 Probabilistic Deep Learning is a hands-on guide to the

principles that support neural networks. Learn to improve network performance with the right distribution for different data types, and discover Bayesian variants that can state their own uncertainty to increase accuracy. This book provides easy-to-apply code and uses popular frameworks to keep you focused on practical applications. Summary Probabilistic Deep Learning: With Python, Keras and TensorFlow Probability teaches the increasingly popular probabilistic approach to deep learning that allows you to refine your results more quickly and accurately without much trial-and-error testing. Emphasizing practical techniques that use the Python-based Tensorflow Probability Framework, you'll learn to build highly-performant deep learning applications that can reliably handle the noise and uncertainty of real-world data. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology The world is a noisy and uncertain place. Probabilistic deep learning models capture that noise and uncertainty, pulling it into real-world scenarios. Crucial for self-driving cars and scientific testing, these techniques help deep learning engineers assess the accuracy of their results, spot errors, and improve their understanding of how algorithms work. About the book Probabilistic Deep Learning is a hands-on guide to the principles that support neural networks. Learn to improve network performance with the right distribution for different data types, and discover Bayesian variants that can state their own uncertainty to increase accuracy. This book provides easy-to-apply code and uses popular frameworks to keep you focused on practical applications. What's inside Explore maximum likelihood and the statistical basis of deep learning Discover



probabilistic models that can indicate possible outcomes  
Learn to use normalizing flows for modeling and  
generating complex distributions Use Bayesian neural  
networks to access the uncertainty in the model About  
the reader For experienced machine learning developers.  
About the author Oliver Dürr is a professor at the  
University of Applied Sciences in Konstanz, Germany.  
Beate Sick holds a chair for applied statistics at ZHAW  
and works as a researcher and lecturer at the University  
of Zurich. Elvis Murina is a data scientist. Table of  
Contents PART 1 - BASICS OF DEEP LEARNING 1 Introduction

to probabilistic deep learning 2 Neural network  
architectures 3 Principles of curve fitting PART 2 -  
MAXIMUM LIKELIHOOD APPROACHES FOR PROBABILISTIC DL  
MODELS 4 Building loss functions with the likelihood  
approach 5 Probabilistic deep learning models with  
TensorFlow Probability 6 Probabilistic deep learning  
models in the wild PART 3 - BAYESIAN APPROACHES FOR  
PROBABILISTIC DL MODELS 7 Bayesian learning 8 Bayesian  
neural networks  
**Deep Learning avec Keras et TensorFlow** - Aurélien Géron  
2020-05-20