

# Tensorflow For Machine Intelligence A Hands On Introduction To Learning Algorithms

Right here, we have countless books **Tensorflow For Machine Intelligence A Hands On Introduction To Learning Algorithms** and collections to check out. We additionally pay for variant types and afterward type of the books to browse. The welcome book, fiction, history, novel, scientific research, as with ease as various further sorts of books are readily genial here.

As this Tensorflow For Machine Intelligence A Hands On Introduction To Learning Algorithms , it ends in the works creature one of the favored book Tensorflow For Machine Intelligence A Hands On Introduction To Learning Algorithms collections that we have. This is why you remain in the best website to see the unbelievable books to have.

*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.

**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 Convolutional Neural Networks with TensorFlow** - Iffat Zafar 2018-08-28

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.

**Tensorflow Machine Learning** - Benjamin Smith 2020-08-24  
Machine Learning is an emerging field in the discipline of computer science. The possibilities are virtually endless and the things we can achieve with machine learning bridge the gap between reality and science fiction. If you are one of those people who developed an interest and learned the basics of machine learning and want to improve your foundation, then this is the right book for you. Here's a list of some of the distinct features of this

book that set it apart from others: - This book includes a comprehensive and detailed explanation of the concepts. No chapter has idle talk. Every line in this book has been written while keeping the convenience and interest of the reader in mind.- This book features some really cool tips and tricks that build upon some very basic and fundamental practices of machine learning. Using these tips and tricks will help increase the productivity of your models.-Each topic addresses some of the most important issues that users experience when working with machine learning. For instance, in the later parts of this book, after discussing deep learning, we shift our focus towards the main challenges that arise when creating and implementing a complex and large deep neural network.-This book aims to give readers a productive reading session. In order to accomplish this, each chapter has fragmented sections that highlight interesting topics. Furthermore, the chapter layout guides the reader through the many concepts of machine learning very easily.If you're interested in tips and tricks to machine learning with the use of scikit-learn, keras and Tensorflow, then click the BUY NOW button to get started today!

*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 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 *Deep Learning For Dummies* - John Paul Mueller 2019-05-14 Take a deep dive into deep learning

Deep learning provides the means for discerning patterns in the data that drive online business and social media outlets. Deep Learning for Dummies gives you the information you need to take the mystery out of the topic—and all of the underlying technologies associated with it. In no time, you'll make sense of those increasingly confusing algorithms, and find a simple and safe environment to experiment with deep learning. The book develops a sense of precisely what deep learning can do at a high level and then provides examples of the major deep learning application types. Includes sample code Provides real-world examples within the approachable text Offers hands-on activities to make learning easier Shows you how to use Deep Learning more effectively with the right tools This book is perfect for those who want to better understand the basis of the underlying technologies that we use each and every day.

**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.

Hands-On Unsupervised Learning Using Python - Ankur A. Patel 2019-02-21

Many industry experts consider unsupervised learning the next frontier in artificial intelligence, one that may hold the key to general artificial intelligence. Since the majority of the world's data is unlabeled, conventional supervised learning cannot be applied. Unsupervised learning, on the other hand, can be applied to unlabeled datasets to discover meaningful patterns buried deep in the data, patterns that may be near impossible for humans to uncover. Author Ankur Patel shows you how to apply unsupervised learning using two simple, production-ready Python frameworks: Scikit-learn and TensorFlow using Keras. With code and hands-on examples, data scientists will identify difficult-to-find patterns in data and gain deeper business insight, detect anomalies, perform automatic feature engineering and selection, and generate synthetic datasets. All you need is programming and some machine learning experience to get started. Compare the strengths and weaknesses of the different machine learning approaches: supervised, unsupervised, and reinforcement learning

Set up and manage machine learning projects end-to-end  
Build an anomaly detection system to catch credit card fraud  
Cluster users into distinct and homogeneous groups  
Perform semisupervised learning  
Develop movie recommender systems using restricted Boltzmann machines  
Generate synthetic images using generative adversarial networks

*Deep Learning Illustrated* - Jon Krohn 2019-08-05

"The authors' clear visual style

provides a comprehensive look at what's currently possible with artificial neural networks as well as a glimpse of the magic that's to come." –Tim Urban, author of Wait But Why Fully Practical, Insightful Guide to Modern Deep Learning Deep learning is transforming software, facilitating powerful new artificial intelligence capabilities, and driving unprecedented algorithm performance. Deep Learning Illustrated is uniquely intuitive and offers a complete introduction to the discipline's techniques. Packed with full-color figures and easy-to-follow code, it sweeps away the complexity of building deep learning models, making the subject approachable and fun to learn. World-class instructor and practitioner Jon Krohn—with visionary content from Grant Beyleveld and beautiful illustrations by Aglaé Bassens—presents straightforward analogies to explain what deep learning is, why it has become so popular, and how it relates to other machine learning approaches. Krohn has created a practical reference and tutorial for developers, data scientists, researchers, analysts, and students who want to start applying it. He illuminates theory with hands-on Python code in accompanying Jupyter notebooks. To help you progress quickly, he focuses on the versatile deep learning library Keras to nimbly construct efficient TensorFlow models; PyTorch, the leading alternative library, is also covered. You'll gain a pragmatic understanding of all major deep learning approaches and their uses in applications ranging from machine vision and natural language processing to image generation and game-playing algorithms. Discover what makes deep learning systems unique, and the implications for practitioners Explore new tools that make deep

learning models easier to build, use, and improve Master essential theory: artificial neurons, training, optimization, convolutional nets, recurrent nets, generative adversarial networks (GANs), deep reinforcement learning, and more Walk through building interactive deep learning applications, and move forward with your own artificial intelligence projects Register your book for convenient access to downloads, updates, and/or corrections as they become available. See inside book for details. Analytical Skills for AI and Data Science - Daniel Vaughan 2020-05-21 While several market-leading companies have successfully transformed their business models by following data- and AI-driven paths, the vast majority have yet to reap the benefits. How can your business and analytics units gain a competitive advantage by capturing the full potential of this predictive revolution? This practical guide presents a battle-tested end-to-end method to help you translate business decisions into tractable prescriptive solutions using data and AI as fundamental inputs. Author Daniel Vaughan shows data scientists, analytics practitioners, and others interested in using AI to transform their businesses not only how to ask the right questions but also how to generate value using modern AI technologies and decision-making principles. You'll explore several use cases common to many enterprises, complete with examples you can apply when working to solve your own issues. Break business decisions into stages that can be tackled using different skills from the analytical toolbox Identify and embrace uncertainty in decision making and protect against common human biases Customize optimal decisions to different customers using predictive

and prescriptive methods and technologies Ask business questions that create high value through AI- and data-driven technologies

Hands-on Computer Vision with TensorFlow 2 - Benjamin Planche 2019

Computer vision is achieving a new frontier of capabilities in fields like health, automobile or robotics. This book explores TensorFlow 2, Google's open-source AI framework, and teaches how to leverage deep neural networks for visual tasks. It will help you acquire the insight and skills to be a part of the exciting advances in computer vision.

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 projectsGet to grips with TensorFlow's impressive range of module offeringsImplement 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 recommendationsBuild projects using CNNs, NLP, and Bayesian neural networksPlay Pac-Man using deep reinforcement learningDeploy scalable TensorFlow-based machine learning systemsGenerate 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

**Hands-On Artificial Intelligence for Beginners** - Patrick D. Smith 2018-10-31

Grasp the fundamentals of Artificial Intelligence and build your own intelligent systems with ease

**Key Features**Enter the world of AI with the help of solid concepts and real-world use casesExplore AI components

to build real-world automated intelligence Become well versed with machine learning and deep learning concepts Book Description Virtual Assistants, such as Alexa and Siri, process our requests, Google's cars have started to read addresses, and Amazon's prices and Netflix's recommended videos are decided by AI. Artificial Intelligence is one of the most exciting technologies and is becoming increasingly significant in the modern world. Hands-On Artificial Intelligence for Beginners will teach you what Artificial Intelligence is and how to design and build intelligent applications. This book will teach you to harness packages such as TensorFlow in order to create powerful AI systems. You will begin with reviewing the recent changes in AI and learning how artificial neural networks (ANNs) have enabled more intelligent AI. You'll explore feedforward, recurrent, convolutional, and generative neural networks (FFNNs, RNNs, CNNs, and GNNs), as well as reinforcement learning methods. In the concluding chapters, you'll learn how to implement these methods for a variety of tasks, such as generating text for chatbots, and playing board and video games. By the end of this book, you will be able to understand exactly what you need to consider when optimizing ANNs and how to deploy and maintain AI applications. What you will learn Use TensorFlow packages to create AI systems Build feedforward, convolutional, and recurrent neural networks Implement generative models for text generation Build reinforcement learning algorithms to play games Assemble RNNs, CNNs, and decoders to create an intelligent assistant Utilize RNNs to predict stock market behavior Create and scale training pipelines and deployment architectures for AI systems Who this book is for This book is designed for

beginners in AI, aspiring AI developers, as well as machine learning enthusiasts with an interest in leveraging various algorithms to build powerful AI applications. *Python Machine Learning* - Sebastian Raschka 2017-09-20 Unlock modern machine learning and deep learning techniques with Python by using the latest cutting-edge open source Python libraries. About This Book Second edition of the bestselling book on Machine Learning A practical approach to key frameworks in data science, machine learning, and deep learning Use the most powerful Python libraries to implement machine learning and deep learning Get to know the best practices to improve and optimize your machine learning systems and algorithms Who This Book Is For If you know some Python and you want to use machine learning and deep learning, pick up this book. Whether you want to start from scratch or extend your machine learning knowledge, this is an essential and unmissable resource. Written for developers and data scientists who want to create practical machine learning and deep learning code, this book is ideal for developers and data scientists who want to teach computers how to learn from data. What You Will Learn Understand the key frameworks in data science, machine learning, and deep learning Harness the power of the latest Python open source libraries in machine learning Explore machine learning techniques using challenging real-world data Master deep neural network implementation using the TensorFlow library Learn the mechanics of classification algorithms to implement the best tool for the job Predict continuous target outcomes using regression analysis Uncover hidden patterns and structures in data with clustering



Delve deeper into textual and social media data using sentiment analysis In Detail Machine learning is eating the software world, and now deep learning is extending machine learning. Understand and work at the cutting edge of machine learning, neural networks, and deep learning with this second edition of Sebastian Raschka's bestselling book, Python Machine Learning. Thoroughly updated using the latest Python open source libraries, this book offers the practical knowledge and techniques you need to create and contribute to machine learning, deep learning, and modern data analysis. Fully extended and modernized, Python Machine Learning Second Edition now includes the popular TensorFlow deep learning library. The scikit-learn code has also been fully updated to include recent improvements and additions to this versatile machine learning library. Sebastian Raschka and Vahid Mirjalili's unique insight and expertise introduce you to machine learning and deep learning algorithms from scratch, and show you how to apply them to practical industry challenges using realistic and interesting examples. By the end of the book, you'll be ready to meet the new data analysis opportunities in today's world. If you've read the first edition of this book, you'll be delighted to find a new balance of classical ideas and modern insights into machine learning. Every chapter has been critically updated, and there are new chapters on key technologies. You'll be able to learn and work with TensorFlow more deeply than ever before, and get essential coverage of the Keras neural network library, along with the most recent updates to scikit-learn. Style and Approach Python Machine Learning Second Edition takes a practical, hands-on coding approach so you can learn about machine learning by

coding with Python. This book moves fluently between the theoretical principles of machine learning and the practical details of implementation with Python. *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 ad VGG-Face Lite 16 Recurrent neural networks 17 LSTMs and automatic speech recognition 18 Sequence-to-sequence models for chatbots 19 Utility landscape *Machine Learning* - Michael White 2018-07-20 The rate that AI is growing is perhaps too fast to understand. The

last 40 years have seen a few false starts but, in the last decade, we've seen huge advances in the processing power of computers and in data storage. All of a sudden, the game has changed in a big way. These days, much of the technology we depend on and rely on daily is powered by AI. Take Google Translate, for example; next time you see something in another language, point your phone camera at it and Google Translate will translate it for you - instantly. AI is also being used to design treatment plans (evidence-based) for cancer sufferers, analyzing medical results instantly so a patient is sent to the right specialist quickly, and helping in research for drug recovery. In this guide, we will be exploring machine learning, the concepts that run these technologies and by the time you get to the end you will have more knowledge than many and will be equipped to start building your own applications.

[Deep Learning with Python](#) - Francois Chollet 2017-11-30

Summary Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Machine learning has made remarkable progress in recent years. We went from near-unusable speech and image recognition, to near-human accuracy. We went from machines that couldn't beat a serious Go player, to defeating a world champion. Behind this progress is deep learning—a combination of engineering advances, best practices, and theory that

enables a wealth of previously impossible smart applications. About the Book Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. You'll explore challenging concepts and practice with applications in computer vision, natural-language processing, and generative models. By the time you finish, you'll have the knowledge and hands-on skills to apply deep learning in your own projects. What's Inside Deep learning from first principles Setting up your own deep-learning environment Image-classification models Deep learning for text and sequences Neural style transfer, text generation, and image generation About the Reader Readers need intermediate Python skills. No previous experience with Keras, TensorFlow, or machine learning is required. About the Author François Chollet works on deep learning at Google in Mountain View, CA. He is the creator of the Keras deep-learning library, as well as a contributor to the TensorFlow machine-learning framework. He also does deep-learning research, with a focus on computer vision and the application of machine learning to formal reasoning. His papers have been published at major conferences in the field, including the Conference on Computer Vision and Pattern Recognition (CVPR), the Conference and Workshop on Neural Information Processing Systems (NIPS), the International Conference on Learning Representations (ICLR), and others. Table of Contents PART 1 - FUNDAMENTALS OF DEEP LEARNING What is deep learning? Before we begin: the mathematical building blocks of

neural networks Getting started with neural networks Fundamentals of machine learning PART 2 - DEEP LEARNING IN PRACTICE Deep learning for computer vision Deep learning for text and sequences Advanced deep-learning best practices Generative deep learning Conclusions appendix A - Installing Keras and its dependencies on Ubuntu appendix B - Running Jupyter notebooks on an EC2 GPU instance 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.

*Hands-On Artificial Intelligence with TensorFlow* - Amir Ziai 2018-10-31

Book Description

Artificial

Intelligence (AI) is a popular area with an emphasis on creating intelligent machines that can reason, evaluate, and understand the same way as humans. It is used extensively across many fields, such as image recognition, robotics, language processing, healthcare, finance, and more. Hands-On Artificial Intelligence with TensorFlow gives you a rundown of essential AI concepts and their implementation with TensorFlow, also highlighting different approaches to solving AI problems using machine learning and deep learning techniques. In addition to this, the book covers advanced concepts, such as reinforcement learning, generative adversarial networks (GANs), and multimodal learning. Once you have grasped all this, you'll move on to exploring GPU computing and neuromorphic computing, along with the latest trends in quantum computing. You'll work through case studies that will help you examine AI applications in the important areas of computer vision, healthcare, and FinTech, and analyze their datasets. In the concluding chapters, you'll briefly investigate possible developments in AI that we can expect to see in the future. By the end of this book, you will be well-versed with the essential concepts of AI and their implementation using TensorFlow. What you will learn

Explore the core concepts of AI and its different approaches

Use the TensorFlow framework for smart applications

Implement various machine and deep learning algorithms with TensorFlow

Design self-learning RL systems and implement generative models

Perform GPU computing efficiently using best practices

Build enterprise-grade apps for computer vision, NLP, and healthcare

Who this book is for

Hands-On Artificial Intelligence with TensorFlow is for you if you are a

machine learning developer, data scientist, AI researcher, or anyone who wants to build artificial intelligence applications using TensorFlow. You need to have some working knowledge of machine learning to get the most out of this book.

### **Hands-On Machine Learning with TensorFlow.js** - Kai Sasaki 2019-11-27

Get hands-on with the browser-based JavaScript library for training and deploying machine learning models effectively Key Features Build, train and run machine learning models in the browser using TensorFlow.js Create smart web applications from scratch with the help of useful examples Use flexible and intuitive APIs from TensorFlow.js to understand how machine learning algorithms function Book Description

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

equipped with the skills you need to create your own web-based ML applications and fine-tune models to achieve high performance. What you will learn Use the t-SNE algorithm in TensorFlow.js to reduce dimensions in an input dataset Deploy tfjs-converter to convert Keras models and load them into TensorFlow.js Apply the Bellman equation to solve MDP problems Use the k-means algorithm in TensorFlow.js to visualize prediction results Create tf.js packages with Parcel, Webpack, and Rollup to deploy web apps Implement tf.js backend frameworks to tune and accelerate app performance Who this book is for This book is for web developers who want to learn how to integrate machine learning techniques with web-based applications from scratch. This book will also appeal to data scientists, machine learning practitioners, and deep learning enthusiasts who are looking to perform accelerated, browser-based machine learning on Web using TensorFlow.js. Working knowledge of JavaScript programming language is all you need to get started.

### **Hands-On Machine Learning with ML.NET** - Jarred Capellman 2020-03-27

Create, train, and evaluate various machine learning models such as regression, classification, and clustering using ML.NET, Entity Framework, and ASP.NET Core Key Features Get well-versed with the ML.NET framework and its components and APIs using practical examples Learn how to build, train, and evaluate popular machine learning algorithms with ML.NET offerings Extend your existing machine learning models by integrating with TensorFlow and other libraries Book Description Machine learning (ML) is widely used in many industries such as science, healthcare, and research and its popularity is only growing. In March 2018, Microsoft introduced ML.NET to

help .NET enthusiasts in working with ML. With this book, you'll explore how to build ML.NET applications with the various ML models available using C# code. The book starts by giving you an overview of ML and the types of ML algorithms used, along with covering what ML.NET is and why you need it to build ML apps. You'll then explore the ML.NET framework, its components, and APIs. The book will serve as a practical guide to helping you build smart apps using the ML.NET library. You'll gradually become well versed in how to implement ML algorithms such as regression, classification, and clustering with real-world examples and datasets. Each chapter will cover the practical implementation, showing you how to implement ML within .NET applications. You'll also learn to integrate TensorFlow in ML.NET applications. Later you'll discover how to store the regression model housing price prediction result to the database and display the real-time predicted results from the database on your web application using ASP.NET Core Blazor and SignalR. By the end of this book, you'll have learned how to confidently perform basic to advanced-level machine learning tasks in ML.NET. What you will learn

- Understand the framework, components, and APIs of ML.NET using C#
- Develop regression models using ML.NET for employee attrition and file classification
- Evaluate classification models for sentiment prediction of restaurant reviews
- Work with clustering models for file type classifications
- Use anomaly detection to find anomalies in both network traffic and login history
- Work with ASP.NET Core Blazor to create an ML.NET enabled web application
- Integrate pre-trained TensorFlow and ONNX models in a WPF ML.NET application for image classification

and object detection Who this book is for If you are a .NET developer who wants to implement machine learning models using ML.NET, then this book is for you. This book will also be beneficial for data scientists and machine learning developers who are looking for effective tools to implement various machine learning algorithms. A basic understanding of C# or .NET is mandatory to grasp the concepts covered in this book effectively.

**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

### **Hands-On Machine Learning with Scikit-Learn and TensorFlow** -

Aurélien Géron 2017-04

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 Apply practical code examples without acquiring excessive machine learning theory or algorithm details

**Machine Learning** - Stephen Marsland 2011-03-23

Traditional books on machine learning can be divided into two groups- those

aimed at advanced undergraduates or early postgraduates with reasonable mathematical knowledge and those that are primers on how to code

algorithms. The field is ready for a text that not only demonstrates how to use the algorithms that make up machine learning methods, but

### **Hands-On Deep Learning Algorithms with Python** - Sudharsan Ravichandiran 2019-07-25

Understand basic to advanced deep learning algorithms, the mathematical principles behind them, and their practical applications. Key Features Get up-to-speed with building your own neural networks from scratch Gain insights into the mathematical principles behind deep learning algorithms Implement popular deep learning algorithms such as CNNs, RNNs, and more using TensorFlow Book Description Deep learning is one of the most popular domains in the AI space, allowing you to develop multi-layered models of varying complexities. This book introduces you to popular deep learning algorithms--from basic to advanced--and shows you how to implement them from scratch using TensorFlow.

Throughout the book, you will gain insights into each algorithm, the mathematical principles behind it, and how to implement it in the best possible manner. The book starts by explaining how you can build your own neural networks, followed by introducing you to TensorFlow, the powerful Python-based library for machine learning and deep learning. Moving on, you will get up to speed with gradient descent variants, such as NAG, AMSGrad, AdaDelta, Adam, and Nadam. The book will then provide you with insights into RNNs and LSTM and how to generate song lyrics with RNN. Next, you will master the math for convolutional and capsule networks, widely used for image recognition tasks. Then you learn how machines

understand the semantics of words and documents using CBOW, skip-gram, and PV-DM. Afterward, you will explore various GANs, including InfoGAN and LSGAN, and autoencoders, such as contractive autoencoders and VAE. By the end of this book, you will be equipped with all the skills you need to implement deep learning in your own projects. What you will learn

- Implement basic-to-advanced deep learning algorithms
- Master the mathematics behind deep learning algorithms
- Become familiar with gradient descent and its variants, such as AMSGrad, AdaDelta, Adam, and Nadam
- Implement recurrent networks, such as RNN, LSTM, GRU, and seq2seq models
- Understand how machines interpret images using CNN and capsule networks
- Implement different types of generative adversarial network, such as CGAN, CycleGAN, and StackGAN
- Explore various types of autoencoder, such as Sparse autoencoders, DAE, CAE, and VAE

Who this book is for If you are a machine learning engineer, data scientist, AI developer, or simply want to focus on neural networks and deep learning, this book is for you. Those who are completely new to deep learning, but have some experience in machine learning and Python programming, will also find the book very helpful.

*TensorFlow for Machine Intelligence* - Sam Abrahams 2016

*Hands-On Deep Learning for Images with TensorFlow* - Will Ballard  
2018-07-31

Explore TensorFlow's capabilities to perform efficient deep learning on images

- Key Features
- Discover image processing for machine vision
- Build an effective image classification system using the power of CNNs
- Leverage TensorFlow's capabilities to perform efficient deep learning

Book Description TensorFlow is Google's popular offering for machine learning

and deep learning, quickly becoming a favorite tool for performing fast, efficient, and accurate deep learning tasks. Hands-On Deep Learning for Images with TensorFlow shows you the practical implementations of real-world projects, teaching you how to leverage TensorFlow's capabilities to perform efficient image processing using the power of deep learning. With the help of this book, you will get to grips with the different paradigms of performing deep learning such as deep neural nets and convolutional neural networks, followed by understanding how they can be implemented using TensorFlow. By the end of this book, you will have mastered all the concepts of deep learning and their implementation with TensorFlow and Keras. What you will learn

- Build machine learning models particularly focused on the MNIST digits
- Work with Docker and Keras to build an image classifier
- Understand natural language models to process text and images
- Prepare your dataset for machine learning
- Create classical, convolutional, and deep neural networks
- Create a RESTful image classification server

Who this book is for Hands-On Deep Learning for Images with TensorFlow is for you if you are an application developer, data scientist, or machine learning practitioner looking to integrate machine learning into application software and master deep learning by implementing practical projects in TensorFlow. Knowledge of Python programming and basics of deep learning are required to get the best out of this book.

**Python Machine Learning** - Sebastian Raschka 2015-09-23

Unlock deeper insights into Machine Learning with this vital guide to cutting-edge predictive analytics

About This Book Leverage Python's most powerful open-source libraries



for deep learning, data wrangling, and data visualization Learn effective strategies and best practices to improve and optimize machine learning systems and algorithms Ask – and answer – tough questions of your data with robust statistical models, built for a range of datasets Who This Book Is For If you want to find out how to use Python to start answering critical questions of your data, pick up Python Machine Learning – whether you want to get started from scratch or want to extend your data science knowledge, this is an essential and unmissable resource. What You Will Learn Explore how to use different machine learning models to ask different questions of your data Learn how to build neural networks using Keras and Theano Find out how to write clean and elegant Python code that will optimize the strength of your algorithms Discover how to embed your machine learning model in a web application for increased accessibility Predict continuous target outcomes using regression analysis Uncover hidden patterns and structures in data with clustering Organize data using effective pre-processing techniques Get to grips with sentiment analysis to delve deeper into textual and social media data In Detail Machine learning and predictive analytics are transforming the way businesses and other organizations operate. Being able to understand trends and patterns in complex data is critical to success, becoming one of the key strategies for unlocking growth in a challenging contemporary marketplace. Python can help you deliver key insights into your data – its unique capabilities as a language let you build sophisticated algorithms and statistical models that can reveal new perspectives and answer key questions that are vital for success.

Python Machine Learning gives you access to the world of predictive analytics and demonstrates why Python is one of the world's leading data science languages. If you want to ask better questions of data, or need to improve and extend the capabilities of your machine learning systems, this practical data science book is invaluable. Covering a wide range of powerful Python libraries, including scikit-learn, Theano, and Keras, and featuring guidance and tips on everything from sentiment analysis to neural networks, you'll soon be able to answer some of the most important questions facing you and your organization. Style and approach Python Machine Learning connects the fundamental theoretical principles behind machine learning to their practical application in a way that focuses you on asking and answering the right questions. It walks you through the key elements of Python and its powerful machine learning libraries, while demonstrating how to get to grips with a range of statistical models.

[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.

*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.

*Machine Learning* - Gabriel Rhys  
2017-10-18

Can Machines Really Learn? Machine learning (ML) is a type of artificial intelligence (AI) that provides computers with the ability to learn without being explicitly programmed. Machine learning has become an essential pillar of IT in all aspects, even though it has been hidden in the recent past. We are increasingly being surrounded by several machine learning-based apps across a broad spectrum of industries. From search engines to anti-spam filters to credit card fraud detection systems, list of machine learning applications is ever-expanding in scope and applications. The goal of this book is to provide you with a hands-on, project-based overview of machine learning systems and how they are applied over a vast spectrum of applications that underpins AI technology from Absolute Beginners to

Experts. This book is a fast-paced, thorough introduction to Machine Learning that will have you writing programs, solving problems, and making things that work in no time. This book presents algorithms and approaches in such a way that grounds them in larger systems as you learn about a variety of topics, including: Supervised and Unsupervised learning methods Artificial Neural Networks Hands-on projects based on Real-world applications Bayesian learning method Reinforcement learning And much more By the end of this book, you should have a strong understanding of machine learning so that you can pursue any further and more advanced learning. Learning Outcomes: By the end of this book, you will be able to: Identify potential applications of machine learning in practice Describe the core differences in analyses enabled by regression, classification, and clustering Select the appropriate machine learning task for a potential application Apply regression, classification, and clustering Represent your data as features to serve as input to machine learning models Utilize a dataset to fit a model to analyze new data Build an end-to-end application that uses machine learning at its core Implement these techniques in Python If you've been thinking seriously about digging into ML, this book will get you up to speed. Why wait any longer?

**Deep Learning Essentials** - Anurag Bhardwaj  
2018-01-30

Get to grips with the essentials of deep learning by leveraging the power of Python Key Features Your one-stop solution to get started with the essentials of deep learning and neural network modeling Train different kinds of neural networks to tackle various problems in Natural Language Processing, computer vision,

speech recognition, and more Covers popular Python libraries such as Tensorflow, Keras, and more, along with tips on training, deploying and optimizing your deep learning models in the best possible manner Book Description Deep Learning a trending topic in the field of Artificial Intelligence today and can be considered to be an advanced form of machine learning, which is quite tricky to master. This book will help you take your first steps in training efficient deep learning models and applying them in various practical scenarios. You will model, train, and deploy different kinds of neural networks such as Convolutional Neural Network, Recurrent Neural Network, and will see some of their applications in real-world domains including computer vision, natural language processing, speech recognition, and so on. You will build practical projects such as chatbots, implement reinforcement learning to build smart games, and develop expert systems for image captioning and processing. Popular Python library such as TensorFlow is used in this book to build the models. This book also covers solutions for different problems you might come across while training models, such as noisy datasets, small datasets, and more. This book does not assume any prior knowledge of deep learning. By the end of this book, you will have a firm understanding of the basics of deep learning and neural network modeling, along with their practical applications. What you will learn Get to grips with the core concepts of deep learning and neural networks Set up deep learning library such as TensorFlow Fine-tune your deep learning models for NLP and Computer Vision applications Unify different information sources, such as images, text, and speech through deep

learning Optimize and fine-tune your deep learning models for better performance Train a deep reinforcement learning model that plays a game better than humans Learn how to make your models get the best out of your GPU or CPU Who this book is for Aspiring data scientists and machine learning experts who have limited or no exposure to deep learning will find this book to be very useful. If you are looking for a resource that gets you up and running with the fundamentals of deep learning and neural networks, this book is for you. As the models in the book are trained using the popular Python-based libraries such as Tensorflow and Keras, it would be useful to have sound programming knowledge of Python.

*Deep Learning By Example* - Ahmed Menshawy 2018-02-28

Grasp the fundamental concepts of deep learning using Tensorflow in a hands-on manner Key Features Get a first-hand experience of the deep learning concepts and techniques with this easy-to-follow guide Train different types of neural networks using Tensorflow for real-world problems in language processing, computer vision, transfer learning, and more Designed for those who believe in the concept of 'learn by doing', this book is a perfect blend of theory and code examples Book Description Deep learning is a popular subset of machine learning, and it allows you to build complex models that are faster and give more accurate predictions. This book is your companion to take your first steps into the world of deep learning, with hands-on examples to boost your understanding of the topic. This book starts with a quick overview of the essential concepts of data science and machine learning which are required to get started with deep learning. It introduces you

to Tensorflow, the most widely used machine learning library for training deep learning models. You will then work on your first deep learning problem by training a deep feed-forward neural network for digit classification, and move on to tackle other real-world problems in computer vision, language processing, sentiment analysis, and more. Advanced deep learning models such as generative adversarial networks and their applications are also covered in this book. By the end of this book, you will have a solid understanding of all the essential concepts in deep learning. With the help of the examples and code provided in this book, you will be equipped to train your own deep learning models with more confidence. What you will learn Understand the fundamentals of deep learning and how it is different from machine learning Get familiarized with Tensorflow, one of the most popular libraries for advanced machine learning Increase the predictive power of your model using feature engineering Understand the basics of deep learning by solving a digit classification problem of MNIST Demonstrate face generation based on the CelebA database, a promising application of generative models Apply deep learning to other domains like language modeling, sentiment analysis, and machine translation Who this book is for This book targets data scientists and machine learning developers who wish to get started with deep learning. If you know what deep learning is but are not quite sure of how to use it, this book will help you as well. An understanding of statistics and data science concepts is required. Some familiarity with Python programming will also be beneficial.

Hands-On Machine Learning with R - Brad Boehmke 2019-11-07

Hands-on Machine Learning with R provides a practical and applied approach to learning and developing intuition into today's most popular machine learning methods. This book serves as a practitioner's guide to the machine learning process and is meant to help the reader learn to apply the machine learning stack within R, which includes using various R packages such as glmnet, h2o, ranger, xgboost, keras, and others to effectively model and gain insight from their data. The book favors a hands-on approach, providing an intuitive understanding of machine learning concepts through concrete examples and just a little bit of theory. Throughout this book, the reader will be exposed to the entire machine learning process including feature engineering, resampling, hyperparameter tuning, model evaluation, and interpretation. The reader will be exposed to powerful algorithms such as regularized regression, random forests, gradient boosting machines, deep learning, generalized low rank models, and more! By favoring a hands-on approach and using real word data, the reader will gain an intuitive understanding of the architectures and engines that drive these algorithms and packages, understand when and how to tune the various hyperparameters, and be able to interpret model results. By the end of this book, the reader should have a firm grasp of R's machine learning stack and be able to implement a systematic approach for producing high quality modeling results. Features:

- Offers a practical and applied introduction to the most popular machine learning methods.
- Topics covered include feature engineering, resampling, deep learning and more.
- Uses a hands-on approach and real world data.

**Hands-On Machine Learning with Scikit-Learn, Keras, and Tensorflow** -

Aurélien Géron 2022-11-15

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

Aurélien Géron 2022-10-04

Through a recent series of 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 bestselling book uses concrete examples, minimal theory, and production-ready Python frameworks (Scikit-Learn, Keras, and TensorFlow) to help you gain an intuitive understanding of the concepts and tools for building intelligent systems. With this updated third edition, author Aurélien Géron explores a range of techniques, starting with simple linear regression and progressing to deep neural networks. Numerous code examples and exercises throughout the book help you apply what you've learned. Programming experience is all you need to get started. Use Scikit-learn to track an example ML project end to end Explore several models, including support vector machines, decision trees, random forests, and ensemble methods Exploit unsupervised learning techniques such as dimensionality reduction, clustering, and anomaly detection Dive into neural net architectures, including convolutional nets, recurrent nets, generative adversarial networks, autoencoders, diffusion models, and transformers Use TensorFlow and Keras to build and train neural nets for computer vision, natural language processing, generative models, and deep reinforcement learning

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

### **Hands-On Unsupervised Learning with Python** - Giuseppe Bonaccorso

2019-02-28

Discover the skill-sets required to implement various approaches to Machine Learning with Python Key Features Explore unsupervised learning with clustering, autoencoders, restricted Boltzmann machines, and more Build your own neural network models using modern Python libraries Practical examples show you how to implement different machine learning and deep learning techniques Book Description Unsupervised learning is about making use of raw, untagged data and applying learning algorithms to it to

help a machine predict its outcome. With this book, you will explore the concept of unsupervised learning to cluster large sets of data and analyze them repeatedly until the desired outcome is found using Python. This book starts with the key differences between supervised, unsupervised, and semi-supervised learning. You will be introduced to the best-used libraries and frameworks from the Python ecosystem and address unsupervised learning in both the machine learning and deep learning domains. You will explore various algorithms, techniques that are used to implement unsupervised learning in real-world use cases. You will learn a variety of unsupervised learning approaches, including randomized optimization, clustering, feature selection and transformation, and information theory. You will get hands-on experience with how neural networks can be employed in unsupervised scenarios. You will also explore the steps involved in building and training a GAN in order to process images. By the end of this book, you will have learned the art of unsupervised learning for different real-world challenges. What you will learn

Use cluster algorithms to identify and optimize natural groups of data

Explore advanced non-linear and hierarchical clustering in action

Soft label assignments for fuzzy c-means and Gaussian mixture models

Detect anomalies through density estimation

Perform principal component analysis using neural network models

Create unsupervised models using GANs

Who this book is for

This book is intended for statisticians, data scientists, machine learning developers, and deep learning practitioners who want to build smart applications by implementing key building block unsupervised learning, and master all the new techniques and algorithms

offered in machine learning and deep learning using real-world examples. Some prior knowledge of machine learning concepts and statistics is desirable.

*Hands-On Machine Learning with TensorFlow.js* - Kai Sasaki 2019-11-27

Get hands-on with the browser-based JavaScript library for training and deploying machine learning models effectively

Key Features

- Build, train and run machine learning models in the browser using TensorFlow.js
- Create smart web applications from scratch with the help of useful examples
- Use flexible and intuitive APIs from TensorFlow.js to understand how machine learning algorithms function

Book Description

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

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

performance Who this book is for This book is for web developers who want to learn how to integrate machine learning techniques with web-based applications from scratch. This book will also appeal to data scientists, machine learning practitioners, and deep learning enthusiasts who are looking to perform accelerated, browser-based machine learning on Web using TensorFlow.js. Working knowledge of JavaScript programming language is all you need to get started.