

[Books] Deep Learning With Python

Yeah, reviewing a ebook **deep learning with python** could accumulate your near contacts listings. This is just one of the solutions for you to be successful. As understood, talent does not suggest that you have extraordinary points.

Comprehending as with ease as arrangement even more than other will give each success. neighboring to, the message as without difficulty as keenness of this deep learning with python can be taken as without difficulty as picked to act.

Deep Learning with Python-François Chollet 2017-10-28 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

Deep Learning with Python-François Chollet 2018

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.

Advanced Deep Learning with Python-Ivan Vasilev 2019-12-12 Gain expertise in advanced deep learning domains such as neural networks, meta-learning, graph neural networks, and memory augmented neural networks using the Python ecosystem Key Features Get to grips with building faster and more robust deep learning architectures Investigate and train convolutional neural network (CNN) models with GPU-accelerated libraries such as TensorFlow and PyTorch Apply deep neural networks (DNNs) to computer vision problems, NLP, and GANs Book Description In order to build robust deep learning systems, you'll need to understand everything from how neural networks work to training CNN models. In this book, you'll discover newly developed deep learning models, methodologies used in the domain, and their implementation based on areas of application. You'll start by understanding the building blocks and the math behind neural networks, and then move on to CNNs and their advanced applications in computer vision. You'll also learn to apply the most popular CNN architectures in object detection and image segmentation. Further on, you'll focus on variational autoencoders and GANs. You'll then use neural networks to extract sophisticated vector representations of words, before going on to cover various types of recurrent networks, such as LSTM and GRU. You'll even explore the attention mechanism to process sequential data without the help of recurrent neural networks (RNNs). Later, you'll use graph neural networks for processing structured data, along with covering meta-learning, which allows you to train neural networks with fewer training samples. Finally, you'll understand how to apply deep learning to autonomous vehicles. By the end of this book, you'll have mastered key deep learning concepts and the different applications of deep learning models in the real world. What you will learn Cover advanced and state-of-the-art neural network architectures Understand the theory and math behind neural networks Train DNNs and apply them to modern deep learning problems Use CNNs for object detection and image segmentation Implement generative adversarial networks (GANs) and variational autoencoders to generate new images Solve natural language processing (NLP) tasks, such as machine translation, using sequence-to-sequence models Understand DL techniques, such as meta-learning and graph neural networks Who this book is for This book is for data scientists, deep learning engineers and researchers, and AI developers who want to further their knowledge of deep learning and build innovative and unique deep learning projects. Anyone looking to get to grips with advanced use cases and methodologies adopted in the deep learning domain using real-world examples will also find this book useful. Basic understanding of deep learning concepts and working knowledge of the Python programming language is assumed.

Deep Learning with Python-Nikhil Ketkar 2017-04-18 Discover the practical aspects of implementing deep-learning solutions using the rich Python ecosystem. This book bridges the gap between the academic state-of-the-art and the industry state-of-the-practice by introducing you to deep learning frameworks such as Keras, Theano, and Caffe. The practicalities of these frameworks is often acquired by practitioners by reading source code, manuals, and posting questions on community forums, which tends to be a slow and a painful process. Deep Learning with Python allows you to ramp up to such practical know-how in a short period of time and focus more on the domain, models, and algorithms. This book briefly covers the mathematical prerequisites and fundamentals of deep learning, making this book a good starting point for software developers who want to get started in deep learning. A brief survey of deep learning architectures is also included. Deep Learning with Python also introduces you to key concepts of automatic differentiation and GPU computation which, while not central to deep learning, are critical when it comes to conducting large scale experiments. What You Will Learn Leverage deep learning frameworks in Python namely, Keras, Theano, and Caffe Gain the fundamentals of deep learning with mathematical prerequisites Discover the practical considerations of large scale experiments Take deep learning models to production Who This Book Is For Software developers who want to try out deep learning as a practical solution to a particular problem. Software developers in a data science team who want to take deep learning models developed by data scientists to production.

Python Deep Learning-Valentino Zocca 2017-04-28 Take your machine learning skills to the next level by mastering Deep Learning concepts and algorithms using Python. About This Book Explore and create intelligent systems using cutting-edge deep learning techniques Implement deep learning algorithms and work with revolutionary libraries in Python Get real-world examples and easy-to-follow tutorials on Theano, TensorFlow, H2O and more Who This Book Is For This book is for Data Science practitioners as well as aspirants who have a basic foundational understanding of Machine Learning concepts and some programming experience with Python. A mathematical background with a conceptual understanding of calculus and statistics is also desired. What You Will Learn Get a practical deep dive into deep learning algorithms Explore deep learning further with Theano, Caffe, Keras, and TensorFlow Learn about two of the most powerful techniques at the core of many practical deep learning implementations: Auto-Encoders and Restricted Boltzmann Machines Dive into Deep Belief Nets and Deep Neural Networks Discover more deep learning algorithms with Dropout and Convolutional Neural Networks Get to know device strategies so you can use deep learning algorithms and libraries in the real world In Detail With an increasing interest in AI around the world, deep learning has attracted a great deal of public attention. Every day, deep learning algorithms are used broadly across different industries. The book will give you all the practical information available on the subject, including the best practices, using real-world use cases. You will learn to recognize and extract information to increase predictive accuracy and optimize results. Starting with a quick recap of important machine learning concepts, the book will delve straight into deep learning principles using Sci-kit learn. Moving ahead, you will learn to use the latest open source libraries such as Theano, Keras, Google's TensorFlow, and H2O. Use this guide to uncover the difficulties of pattern recognition, scaling data with greater accuracy and discussing deep learning algorithms and techniques. Whether you want to dive deeper into Deep Learning, or want to investigate how to get more out of this powerful technology, you'll find everything inside. Style and approach Python Machine Learning by example follows practical hands on approach. It walks you through the key elements of Python and its powerful machine learning libraries with the help of real world projects.

Deep Learning with Python-Nikhil Ketkar 2021-04-10 Master the practical aspects of implementing deep learning solutions with PyTorch, using a hands-on approach to understanding both theory and practice. This updated edition will prepare you for applying deep learning to real world problems with a sound theoretical foundation and practical know-how with PyTorch, a platform developed by Facebook's Artificial Intelligence Research Group. You'll start with a perspective on how and why deep learning with PyTorch has emerged as an path-breaking framework with a set of tools and techniques to solve real-world problems. Next, the book will ground you with the mathematical fundamentals of linear algebra, vector calculus, probability and optimization. Having established this foundation, you'll move on to key components and functionality of PyTorch including layers, loss functions and optimization algorithms. You'll also gain an understanding of Graphical Processing Unit (GPU) based computation, which is essential for training deep learning models. All the key architectures in deep learning are covered, including feedforward networks, convolution neural networks, recurrent neural networks, long short-term memory networks, autoencoders and generative adversarial networks. Backed by a number of tricks of the trade for training and optimizing deep learning models, this edition of Deep Learning with Python explains the best practices in taking these models to production with PyTorch. What You'll Learn Review machine learning fundamentals such as overfitting, underfitting, and regularization. Understand deep learning fundamentals such as feed-forward networks, convolution neural networks, recurrent neural networks, automatic differentiation, and stochastic gradient descent. Apply in-depth linear algebra with PyTorch Explore PyTorch fundamentals and its building blocks Work with tuning and optimizing models Who This Book Is For Beginners with a working knowledge of Python who want to understand Deep Learning in a practical, hands-on manner.

Deep Learning with Python-Nikhil Ketkar 2017-04-19 Discover the practical aspects of implementing deep-learning solutions using the rich Python ecosystem. This book bridges the gap between the academic state-of-the-art and the industry state-of-the-practice by introducing you to deep learning frameworks such as Keras, Theano, and Caffe. The practicalities of these frameworks is often acquired by practitioners by reading source code, manuals, and posting questions on community forums, which tends to be a slow and a painful process. Deep Learning with Python allows you to ramp up to such practical know-how in a short period of time and focus more on the domain, models, and algorithms. This book briefly covers the mathematical prerequisites and fundamentals of deep learning, making this book a good starting point for software developers who want to get started in deep learning. A brief survey of deep learning architectures is also included. Deep Learning with Python also introduces you to key concepts of automatic differentiation and GPU computation which, while not central to deep learning, are critical when it comes to conducting large scale experiments. What You Will Learn Leverage deep learning frameworks in Python namely, Keras, Theano, and Caffe Gain the fundamentals of deep learning with mathematical prerequisites Discover the practical considerations of large scale experiments Take deep learning models to production Who This Book Is For Software developers who want to try out deep learning as a practical solution to a particular problem. Software developers in a data science team who want to take deep learning models developed by data scientists to production.

Deep Learning with Applications Using Python-Navin Kumar Manaswi 2018-04-04 Explore deep learning applications, such as computer vision, speech recognition, and chatbots, using frameworks such as TensorFlow and Keras. This book helps you to ramp up your practical know-how in a short period of time and focuses you on the domain, models, and algorithms required for deep learning applications. Deep Learning with Applications Using Python covers topics such as chatbots, natural language processing, and face and object recognition. The goal is to equip you with the concepts, techniques, and algorithm implementations needed to create programs capable of performing deep learning. This book covers convolutional neural networks, recurrent neural networks, and multilayer perceptrons. It also discusses popular APIs such as IBM Watson, Microsoft Azure, and scikit-learn. What You Will Learn Work with various deep learning frameworks such as TensorFlow, Keras, and scikit-learn. Use face recognition and face detection capabilities Create speech-to-text and text-to-speech functionality Engage with chatbots using deep learning Who This Book Is For Data scientists and developers who want to adapt and build deep learning applications.

Introduction to Machine Learning with Python-Andreas C. Müller 2016-09-26 Machine learning has become an integral part of many commercial applications and research projects, but this field is not exclusive to large companies with extensive research teams. If you use Python, even as a beginner, this book will teach you practical ways to build your own machine learning solutions. With all the data available today, machine learning applications are limited only by your imagination. You'll learn the steps necessary to create a successful machine-learning application with Python and the scikit-learn library. Authors Andreas Müller and Sarah Guido focus on the practical aspects of using machine learning algorithms, rather than the math behind them. Familiarity with the NumPy and matplotlib libraries will help you get even more from this book. With this book, you'll learn: Fundamental concepts and applications of machine learning Advantages and shortcomings of widely used machine learning algorithms How to represent data processed by machine learning, including which data aspects to focus on Advanced methods for model evaluation and parameter tuning The concept of pipelines for chaining models and encapsulating your workflow Methods for working with text data, including text-specific processing techniques Suggestions for improving your machine learning and data science skills

Learn Keras for Deep Neural Networks-Jojo Moolayil 2018-12-07 Learn, understand, and implement deep neural networks in a math- and programming-friendly approach using Keras and Python. The book focuses on an end-to-end approach to developing supervised learning algorithms in regression and classification with practical business-centric use-cases implemented in Keras. The overall book comprises three sections with two chapters in each section. The first section prepares you with all the necessary basics to get started in deep learning. Chapter 1 introduces you to the world of deep learning and its difference from machine learning, the choices of frameworks for deep learning, and the Keras ecosystem. You will cover a real-life business problem that can be solved by supervised learning algorithms with deep neural networks. You'll tackle one use case for regression and another for classification leveraging popular Kaggle datasets. Later, you will see an interesting and challenging part of deep learning: hyperparameter tuning; helping you further improve your models when building robust deep learning applications. Finally, you'll further hone your skills in deep learning and cover areas of active development and research in deep learning. At the end of Learn Keras for Deep Neural Networks, you will have a thorough understanding of deep learning principles and have practical hands-on experience in developing enterprise-grade deep learning solutions in Keras. What You'll Learn Master fast-paced practical deep learning concepts with math- and programming-friendly abstractions. Design, develop, train, validate, and deploy deep neural networks using the Keras framework Use best practices for debugging and validating deep learning models Deploy and integrate deep learning as a service into a larger software service or product Extend deep learning principles into other popular frameworks Who This Book Is For Software engineers and data engineers with basic programming skills in any language and who are keen on exploring deep learning for a career move or an enterprise project.

Python Deep Learning Projects-Matthew Lamons 2018-10-31 Insightful projects to master deep learning and neural network architectures using Python and Keras Key Features Explore deep learning across computer vision, natural language processing (NLP), and image processing Discover best practices for the training of deep neural networks and their deployment Access popular deep learning models as well as widely used neural network architectures Book Description Deep learning has been gradually revolutionizing every field of artificial intelligence, making application development easier. Python Deep Learning Projects imparts all the knowledge needed to implement complex deep learning projects in the field of computational linguistics and computer vision. Each of these projects is unique, helping you progressively master the subject. You'll learn how to implement a text classifier system using a recurrent neural network (RNN) model and optimize it to understand the shortcomings you might experience while implementing a simple deep learning system. Similarly, you'll discover how to develop various projects, including word vector representation, open domain question answering, and building chatbots using seq-to-seq models and language modeling. In addition to this, you'll cover advanced concepts, such as regularization, gradient clipping, gradient normalization, and bidirectional RNNs, through a series of engaging projects. By the end of this book, you will have gained knowledge to develop your own deep learning systems in a straightforward way and in an efficient way What you will learn Set up a deep learning development environment on Amazon Web Services (AWS) Apply GPU-powered instances as well as the deep learning AMI Implement seq-to-seq models for modeling natural language processing (NLP) Develop an end-to-end speech recognition system Build a system for pixel-wise semantic labeling of an image Create a system that generates images and their regions Who this book is for Python Deep Learning Projects is for you if you want to get insights into deep learning, data science, and artificial intelligence. This book is also for those who want to break into deep learning and develop their own AI projects. It is assumed that you have sound knowledge of Python programming

Python Deep Learning-Ivan Vasilev 2019-01-16 Learn advanced state-of-the-art deep learning techniques and their applications using popular Python libraries Key Features Build a strong foundation in neural networks and deep learning with Python libraries Explore advanced deep learning techniques and their applications across computer vision and NLP Learn how a computer can navigate in complex environments with reinforcement learning Book Description With the surge in artificial intelligence in applications catering to both business and consumer needs, deep learning is more important than ever for meeting current and future market demands. With this book, you'll explore deep learning, and learn how to put machine learning to use in your projects. This second edition of Python Deep Learning will get you up to speed with deep learning, deep neural networks, and how to train them with high-performance algorithms and popular Python frameworks. You'll uncover different neural network architectures, such as convolutional networks, recurrent neural networks, long short-term memory (LSTM) networks, and capsule networks. You'll also learn how to solve problems in the fields of computer vision, natural language processing (NLP), and speech recognition. You'll study generative model approaches such as variational autoencoders and Generative Adversarial Networks (GANs) to generate images. As you delve into newly evolved areas of reinforcement learning, you'll gain an understanding of state-of-the-art algorithms that are the main components behind popular games Go, Atari, and Dota. By the end of the book, you will be well-versed with the theory of deep learning along with its real-world applications. What you will learn Grasp the mathematical theory behind neural networks and deep learning processes Investigate and resolve computer vision challenges using convolutional networks and capsule networks Solve generative tasks using variational autoencoders and Generative Adversarial Networks Implement complex NLP tasks using recurrent networks (LSTM and GRU) and attention models Explore reinforcement learning and understand how agents behave in a complex environment Get up to date with applications of deep learning in autonomous vehicles Who this book is for This book is for data science practitioners, machine learning engineers, and those interested in deep learning who have a basic foundation in machine learning and some Python programming experience. A background in mathematics and conceptual understanding of calculus and statistics will help you gain maximum benefit from this book.

Deep Learning with PyTorch-Eli Stevens 2020-08-04 Every other day we hear about new ways to put deep learning to good use: improved medical imaging, accurate credit card fraud detection, long range weather forecasting, and more. PyTorch puts these superpowers in your hands, providing a comfortable PyTorch experience that gets you started quickly and then grows with you as you—and your deep learning skills—become more sophisticated. Deep Learning with PyTorch will make that journey engaging and fun. Summary Every other day we hear about new ways to put deep learning to good use: improved medical imaging, accurate credit card fraud detection, long range weather forecasting, and more. PyTorch puts these superpowers in your hands, providing a comfortable PyTorch experience that gets you started quickly and then grows with you as you—and your deep learning skills—become more sophisticated. Deep Learning with PyTorch will make that journey engaging and fun. Foreword by Soumith Chintala, Co-creator of PyTorch. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Although many deep learning tools use Python, the PyTorch library is truly Pythonic. Instantly familiar to anyone who knows PyData tools like NumPy and scikit-learn, PyTorch simplifies deep learning without sacrificing advanced features. It's excellent for building quick models, and it scales smoothly from laptop to enterprise. Because companies like Apple, Facebook, and JPMorgan Chase rely on PyTorch, it's a great skill to have as you expand your career options. It's easy to get started with PyTorch. It minimizes cognitive overhead without sacrificing the access to advanced features, meaning you can focus on what matters the most - building and training the latest and greatest deep learning models and contribute to making a dent in the world. PyTorch is also a snap to scale and extend, and it partners well with other Python tooling. PyTorch has been adopted by hundreds of deep learning practitioners and several first-class projects like FAIR, OpenAI, FastAI and Purdue. About the book Deep Learning with PyTorch teaches you to create neural networks and deep learning systems with PyTorch. This practical book quickly gets you to work building a real-world example from scratch: a tumor image classifier. Along the way, it covers best practices for the entire DL pipeline, including the PyTorch Tensor API, loading data in Python, monitoring training, and visualizing results. After covering the basics, the book will take you on a journey through larger projects. The centerpiece of the book is a neural network designed for cancer detection. You'll discover ways for training networks with limited inputs and start processing data to get some results. You'll sift through the unreliable initial results and focus on how to diagnose and fix the problems in your neural network. Finally, you'll look at ways to improve your results by training with augmented data, make improvements to the model architecture, and perform other fine tuning. What's inside Training deep neural networks Implementing modules and loss functions Utilizing pretrained models from PyTorch Hub Exploring code samples in Jupyter Notebooks About the reader For Python programmers with an interest in machine learning. About the author Eli Stevens had roles from software engineer to CTO, and is currently working on machine learning in the self-driving-car industry. Luca Antiga is cofounder of an AI engineering company and an AI tech startup, as well as a former PyTorch contributor. Thomas Viehmann is a PyTorch core developer and machine learning trainer and consultant. consultant based in Munich, Germany and a PyTorch core developer. Table of Contents PART 1 - CORE PYTORCH 1 Introducing deep learning and the PyTorch Library 2 Pretrained networks 3 It starts with a tensor 4 Real-world data representation using tensors 5 The mechanics of learning 6 Using a neural network to fit the data 7 Telling birds from airplanes: Learning from images 8 Using convolutions to generalize PART 2 - LEARNING FROM IMAGES IN THE REAL WORLD: EARLY DETECTION OF LUNG CANCER 9 Using PyTorch to fight cancer 10 Combining data sources into a unified dataset 11 Training a classification model to detect suspected tumors 12 Improving training with metrics and augmentation 13 Using segmentation to find suspected nodules 14 End-to-end nodule analysis, and where to go next PART 3 - DEPLOYMENT 15 Deploying to production

Practical Machine Learning with Python-Dipanjan Sarkar 2017-12-20 Master the essential skills needed to recognize and solve complex problems with machine learning and deep learning. Using real-world examples that leverage the popular Python machine learning ecosystem, this book is your perfect companion for learning the art and science of machine learning to become a successful practitioner. The concepts, techniques, tools, frameworks, and methodologies used in this book will teach you how to think, design, build, and execute machine learning systems and projects successfully. Practical Machine Learning with Python follows a structured and comprehensive three-tiered approach packed with hands-on examples and code. Part 1 focuses on understanding machine learning concepts and tools. This includes machine learning basics with a broad overview of algorithms, techniques, concepts and applications, followed by a tour of the entire Python machine learning ecosystem. Brief guides for useful machine learning tools, libraries and frameworks are also covered. Part 2 details standard machine learning pipelines, with an emphasis on data processing analysis, feature engineering, and modeling. You will learn how to process, wrangle, summarize and visualize data in its various forms. Feature engineering and selection methodologies will be covered in detail with real-world datasets followed by model building, tuning, interpretation and deployment. Part 3 explores multiple real-world case studies spanning diverse domains and industries like retail, transportation, movies, music, marketing, computer vision and finance. For each case study, you will learn the application of various machine learning techniques and methods. The hands-on examples will help you become familiar with state-of-the-art machine learning tools and techniques and understand what algorithms are best suited for any problem. Practical Machine Learning with Python will empower you to start solving your own problems with machine learning today! What You'll Learn Execute end-to-end machine learning projects and systems Implement hands-on examples with industry standard, open source, robust machine learning tools and frameworks Review case studies depicting applications of machine learning and deep learning on diverse domains and industries Apply a wide range of machine learning models including regression, classification, and clustering. Understand and apply the latest models and methodologies from deep learning including CNNs, RNNs, LSTMs and transfer learning. Who This Book Is For IT professionals, analysts, developers, data scientists, engineers, graduate students

Machine Learning with Python Cookbook-Chris Albon 2018-03-09 This practical guide provides nearly 200 self-contained recipes to help you solve machine learning challenges you may encounter in your daily work. If you're comfortable with Python and its libraries, including pandas and scikit-learn, you'll be able to address specific problems such as loading data, handling text or numerical data, model selection, and dimensionality reduction and many other topics. Each recipe includes code that you can copy and paste into a toy dataset to ensure that it actually works. From there, you can insert, combine, or adapt the code to help construct your application. Recipes also include a discussion that explains the solution and provides meaningful context. This cookbook takes you beyond theory and concepts by providing the nuts and bolts you need to construct working machine learning applications. You'll find recipes for: Vectors, matrices, and arrays Handling numerical and categorical data, text, images, and dates and times Dimensionality reduction using feature extraction or feature selection Model evaluation and selection Linear and logical regression, trees and forests, and k-nearest neighbors Support vector machines (SVM), naïve Bayes, clustering, and neural networks Saving and loading trained models

Deep Learning with Python: A Fundamentals Guide to Understanding Machine Learning and Artificial Intelligence with Scikit-Learn, Tensorflow, and Sebastian Dark 2018-11 Curious to discover the revolutionary technology that is shaping our future and changing the world? Deep learning is a part of the field of computer science and a subset of machine learning that involves computer systems being able to "learn" unsupervised with data that is unlabeled or unstructured. In 2017, AlphaGo, which is AI developed by Google DeepMind and started off by only knowing the rules of the game, was eventually able to train itself and beat Ke Jie, the world No. 1 ranked player at the time. Although this was not may seem that impressive at first, it is important to understand that Go is a very complex game that many programmers were not able to trump with AI in the past. Although Go is an interesting example, the possibilities of using machine learning are limitless. From retail to medicine to finance, machine learning has the ability to change each industry it comes into contact with. In fact, this revolution has already begun and will only continue to get bigger. According to statista.com, the artificial intelligence industry is set to grow exponentially in the next few years from \$7 Billion in 2018 to \$90 Billion in 2025! This isn't something you can afford to miss. Without a doubt it is the future. However, it is as complex as it is revolutionary. If you do not have a background or any experience in the field, it is easy to get bogged down by all the complicated concepts and term. And if you are at a more advanced level, the information you find won't be thorough enough. In this book, you will find the perfect balance between the information being very thorough and being able to understand it. Although tailored for beginners, it won't contain simple and easily accessible information. You will dive deep into the field but will be carefully led through it in a way that will make everything easy to understand even if you do not have a technical background in computer programming. In this Guide, you will discover... What Machine Learning and Deep Learning Is And How You Can Use It To Change The World How The Field Can Be Broken Down And Learned In A Manageable Way Various Applications and Potential of Deep Learning That You Can Utilize - That You May Never Have Even Imagined Supervised And Unsupervised Learning - And Breaking It Down Step By Step How You Can Create And Train Deep Learning Models Where and How To Install the Best Programs So You Can Get Started Today Sample Codes And Datasets To Practice Along With And much more! If you are finally prepared to begin grasping this revolutionary technology at a high level despite what your technical background may be, Click "Add to Cart" Now! **Get the Kindle eBook version for FREE when you buy the Paperback version of this book!*

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.

Applied Deep Learning with Python-Alex Galea 2018-08-31 A hands-on guide to deep learning that's filled with intuitive explanations and engaging practical examples Key Features Designed to iteratively develop the skills of Python users who don't have a data science background Covers the key foundational concepts you'll need to know when building deep learning systems Full of step-by-step exercises and activities to help build the skills that you need for the real-world Book Description Taking an approach that uses the latest developments in the Python ecosystem, you'll first be guided through the Jupyter ecosystem, key visualization libraries and powerful data sanitization techniques before we train our first predictive model. We'll explore a variety of approaches to classification like support vector networks, random decision forests and k-nearest neighbours to build out your understanding before we move into more complex territory. It's okay if these terms seem overwhelming; we'll show you how to put them to work. We'll build upon our classification coverage by taking a quick look at ethical web scraping and interactive visualizations to help you professionally gather and present your analysis. It's after this that we start building out our keystone deep learning application, one that aims to predict the future price of Bitcoin based on historical public data. By guiding you through a trained neural network, we'll explore common deep learning network architectures (convolutional, recurrent, generative adversarial) and branch out into deep reinforcement learning before we dive into model optimization and evaluation. We'll do all of this whilst working on a production-ready web application that combines TensorFlow and Keras to produce a meaningful user-friendly result, leaving you with all the skills you need to tackle and develop your own real-world deep learning projects confidently and effectively. What you will learn Discover how you can assemble and clean your very own datasets Develop a tailored machine learning classification strategy Build, train and enhance your own models to solve unique problems Work with production-ready frameworks like TensorFlow and Keras Explain how neural networks operate in clear and simple terms Understand how to deploy your predictions to the web Who this book is for If you're a Python programmer stepping into the world of data science, this is the ideal way to get started.

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.

Practical Deep Learning-Ron Kneusel 2021-03-16 Practical Deep Learning teaches total beginners how to build the datasets and models needed to train neural networks for your own DL projects. If you've been curious about machine learning but didn't know where to start, this is the book you've been waiting for. Focusing on the subfield of machine learning known as deep learning, it explains core concepts and gives you the foundation you need to start building your own models. Rather than simply outlining recipes for using existing toolkits, Practical Deep Learning teaches you the why of deep learning and will inspire you to explore further. All you need is basic familiarity with computer programming and high school math—the book will cover the rest. After an introduction to Python, you'll move through key topics like how to build a good training dataset, work with the scikit-learn and Keras libraries, and evaluate your models' performance. You'll also learn: • How to use classic machine learning models like k-Nearest Neighbors, Random Forests, and Support Vector Machines • How neural networks work and how they're trained • How to use convolutional neural networks • How to develop a successful deep learning model from scratch You'll conduct experiments along the way, building to a final case study that incorporates everything you've learned. The perfect introduction to this dynamic, ever-expanding field, Practical Deep Learning will give you the skills and confidence to dive into your own machine learning projects.

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 with Python-Mark Graph 2019-10-15 This book doesn't have any superpowers or magic formula to help you master the art of neural networks and deep learning. We believe that such learning is all in your heart. You need to learn a concept by heart and then brainstorm its different possibilities. I don't claim that after reading this book you will become an expert in Python and Deep Learning Neural Networks. Instead, you will, for sure, have a basic understanding of deep learning and its implications and real-life applications. Most of the time, what confuses us is the application of a certain thing in our lives. Once we know that, we can relate the subject to that particular thing and learn. An interesting thing is that neural networks also learn the same way. This makes it easier to learn about them when we know the basics. Let's take a look at what this book has to offer: ● The basics of Python including data types, operators and numbers. ● Advanced programming in Python with Python expressions, types and much more. ● A comprehensive overview of deep learning and its link to the smart systems that we are now building. ● An overview of how artificial neural networks work in real life. ● An overview of PyTorch. ● An overview of TensorFlow. ● An overview of Keras. ● How to create a convolutional neural network. ● A comprehensive understanding of deep learning applications and its ethical implications, including in the present and future. This book offers you the basic knowledge about Python and Deep Learning Neural Networks that you will need to lay the foundation for future studies. This book will start you on the road to mastering the art of deep learning neural networks. When I say that I don't have the magic formula to make you learn, I mean it. My point is that you should learn Python coding and Python libraries to build neural networks by practicing hard. The more you practice, the better it is for your skills. It is only after thorough and in depth practice that you will be able to create your own programs. Unlike other books, I don't claim that this book will make you a master of deep learning after a single read. That's not realistic, in fact, it's even a bit absurd. What I claim is that you will definitely learn about the basics. The rest is practice. The more you practice the better you code.

Python Machine Learning-Wei-Meng Lee 2019-04-04 Python makes machine learning easy for beginners and experienced developers With computing power increasing exponentially and costs decreasing at the same time, there is no better time to learn machine learning using Python. Machine learning tasks that once required enormous processing power are now possible on desktop machines. However, machine learning is not for the faint of heart—it requires a good foundation in statistics, as well as programming knowledge. Python Machine Learning will help coders of all levels master one of the most in-demand programming skillsets in use today. Readers will get started by following fundamental topics such as an introduction to Machine Learning and Data Science. For each learning algorithm, readers will use a real-life scenario to show how Python is used to solve the problem at hand. • Python data science—manipulating data and data visualization • Data cleansing • Understanding Machine learning algorithms • Supervised learning algorithms • Unsupervised learning algorithms • Deploying machine learning models Python Machine Learning is essential reading for students, developers, or anyone with a keen interest in taking their coding skills to the next level.

Python Machine Learning-Sebastian Raschka 2019-12-12 Applied machine learning with a solid foundation in theory. Revised and expanded for TensorFlow 2, GANs, and reinforcement learning. Key Features Third edition of the bestselling, widely acclaimed Python machine learning book Clear and intuitive explanations take you deep into the theory and practice of Python machine learning Fully updated and expanded to cover TensorFlow 2, Generative Adversarial Network models, reinforcement learning, and best practices Book Description Python Machine Learning, Third Edition is a comprehensive guide to machine learning and deep learning with Python. It acts as both a step-by-step tutorial, and a reference you'll keep coming back to as you build your machine learning systems. Packed with clear explanations, visualizations, and working examples, the book covers all the essential machine learning techniques in depth. While some books teach you only to follow instructions, with this machine learning book, Raschka and Mirjalili teach the principles behind machine learning, allowing you to build models and applications for yourself. Updated for TensorFlow 2.0, this new third edition introduces readers to its new Keras API features, as well as the latest additions to scikit-learn. It's also expanded to cover cutting-edge reinforcement learning techniques based on deep learning, as well as an introduction to GANs. Finally, this book also explores a subfield of natural language processing (NLP) called sentiment analysis, helping you learn how to use machine learning algorithms to classify documents. This book is your companion to machine learning with Python, whether you're a Python developer new to machine learning or want to deepen your knowledge of the latest developments. What you will learn Master the frameworks, models, and techniques that enable machines to 'learn' from data Use scikit-learn for machine learning and TensorFlow for deep learning Apply machine learning to image classification, sentiment analysis, intelligent web applications, and more Build and train neural networks, GANs, and other models Discover best practices for evaluating and tuning models Predict continuous target outcomes using regression analysis Dig deeper into textual and social media data using sentiment analysis 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 resource. Written for developers and data scientists who want to create practical machine learning and deep learning code, this book is ideal for anyone who wants to teach computers how to learn from data.

Deep Learning for Natural Language Processing-Palash Goyal 2018-06-26 Discover the concepts of deep learning used for natural language processing (NLP), with full-fledged examples of neural network models such as recurrent neural networks, long short-term memory networks, and sequence-2-sequence models. You'll start by covering the mathematical prerequisites and the fundamentals of deep learning and NLP with practical examples. The first three chapters of the book cover the basics of NLP, starting with word-vector representation before moving onto advanced algorithms. The final chapters focus entirely on implementation, and deal with sophisticated architectures such as RNN, LSTM, and Seq2Seq, using Python tools: TensorFlow, and Keras. Deep Learning for Natural Language Processing follows a progressive approach and combines all the knowledge you have gained to build a question-answer chatbot system. This book is a good starting point for people who want to get started in deep learning for NLP. All the code presented in the book will be available in the form of iPython notebooks and scripts, which allow you to try out the examples and extend them in interesting ways. What You Will Learn Gain the fundamentals of deep learning and its mathematical prerequisites Discover deep learning frameworks in Python Develop a chatbot Implement a research paper on sentiment classification Who This Book Is For Software developers who are curious to try out deep learning with NLP.

Deep Learning-Ian Goodfellow 2016-11-10 An introduction to a broad range of topics in deep learning, covering mathematical and conceptual background, deep learning techniques used in industry, and research perspectives. "Written by three experts in the field, Deep Learning is the only comprehensive book on the subject." —Elon Musk, cochair of OpenAI; cofounder and CEO of Tesla and SpaceX Deep learning is a form of machine learning that enables computers to learn from experience and understand the world in terms of a hierarchy of concepts. Because the computer gathers knowledge from experience, there is no need for a human computer operator to formally specify all the knowledge that the computer needs. The hierarchy of concepts allows the computer to learn complicated concepts by building them out of simpler ones; a graph of these hierarchies would be many layers deep. This book introduces a broad range of topics in deep learning. The text offers mathematical and conceptual background, covering relevant concepts in linear algebra, probability theory and information theory, numerical computation, and machine learning. It describes deep learning techniques used by practitioners in industry, including deep feedforward networks, regularization, optimization algorithms, convolutional networks, sequence modeling, and practical methodology; and it surveys such applications as natural language processing, speech recognition, computer vision, online recommendation systems, bioinformatics, and videogames. Finally, the book offers research perspectives, covering such theoretical topics as linear factor models, autoencoders, representation learning, structured probabilistic models, Monte Carlo methods, the partition function, approximate inference, and deep generative models. Deep Learning can be used by undergraduate or graduate students planning careers in either industry or research, and by software engineers who want to begin using deep learning in their products or platforms. A website offers supplementary material for both readers and instructors.

Advanced Machine Learning with Python-John Hearty 2016-07-28 Solve challenging data science problems by mastering cutting-edge machine learning techniques in Python About This Book Resolve complex machine learning problems and explore deep learning Learn to use Python code for implementing a range of machine learning algorithms and techniques A practical tutorial that tackles real-world computing problems through a rigorous and effective approach Who This Book Is For This title is for Python developers and analysts or data scientists who are looking to add to their existing skills by accessing some of the most powerful recent trends in data science. If you've ever considered building your own image or text-tagging solution, or of entering a Kaggle contest for instance, this book is for you! Prior experience of Python and grounding in some of the core concepts of machine learning would be helpful. What You Will Learn Compete with top data scientists by gaining a practical and theoretical understanding of cutting-edge deep learning algorithms Apply your new found skills to solve real problems, through clearly-explained code for every technique and test Automate large sets of complex data and overcome time-consuming practical challenges Improve the accuracy of models and your existing input data using powerful feature engineering techniques Use multiple learning techniques together to improve the consistency of results Understand the hidden structure of datasets using a range of unsupervised techniques Gain insight into how the experts solve challenging data problems with an effective, iterative, and validation-focused approach Improve the effectiveness of your deep learning models further by using powerful ensemble techniques to strap multiple models together In Detail Designed to take you on a guided tour of the most relevant and powerful machine learning techniques in use today by top data scientists, this book is just what you need to push your Python algorithms to maximum potential. Clear examples and detailed code samples demonstrate deep learning techniques, semi-supervised learning, and more - all whilst working with real-world applications that include image, music, text, and financial data. The machine learning techniques covered in this book are at the forefront of commercial practice. They are applicable now for the first time in contexts such as image recognition, NLP and web search, computational creativity, and commercial/financial data modeling. Deep Learning algorithms and ensembles of models are in use by data scientists at top tech and digital companies, but the skills needed to apply them successfully, while in high demand, are still scarce. This book is designed to take the reader on a guided tour of the most relevant and powerful machine learning techniques. Clear descriptions of how techniques work and detailed code examples demonstrate deep learning techniques, semi-supervised learning and more, in real world applications. We will also learn about NumPy and Theano. By the end of this book, you will learn a set of advanced Machine Learning techniques and acquire a broad set of powerful skills in the area of feature selection & feature engineering. Style and approach This book focuses on clarifying the theory and code behind complex algorithms to make them practical, useable, and well-understood. Each topic is described with real-world applications, providing both broad contextual coverage and detailed guidance.

Practical Deep Learning for Cloud, Mobile, and Edge-Anirudh Koul 2019-10-14 Whether you're a software engineer aspiring to enter the world of deep learning, a veteran data scientist, or a hobbyist with a simple dream of making the next viral AI app, you might have wondered where to begin. This step-by-step guide teaches you how to build practical deep learning applications for the cloud, mobile, browsers, and edge devices using a hands-on approach. Relying on years of industry experience transforming deep learning research into award-winning applications, Anirudh Koul, Siddha Ganju, and Meher Kasam guide you through the process of converting an idea into something that people in the real world can use. Train, tune, and deploy computer vision models with Keras, TensorFlow, Core ML, and TensorFlow Lite Develop AI for a range of devices including Raspberry Pi, Jetson Nano, and Google Coral Explore fun projects, from Silicon Valley's Not Hotdog app to 40+ industry case studies Simulate an autonomous car in a video game environment and build a miniature version with reinforcement learning Use transfer learning to train models in minutes Discover 50+ practical tips for maximizing model accuracy and speed, debugging, and scaling to millions of users

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

Advanced Deep Learning with Python-Ivan Vasilev 2019-12-12 Gain expertise in advanced deep learning domains such as neural networks, meta-learning, graph neural networks, and memory augmented neural networks using the Python ecosystem Key Features Get to grips with building faster and more robust deep learning architectures Investigate and train convolutional neural network (CNN) models with GPU-accelerated libraries such as TensorFlow and PyTorch Apply deep neural networks (DNNs) to computer vision problems, NLP, and GANs Book Description In order to build robust deep learning systems, you'll need to understand everything from how neural networks work to training CNN models. In this book, you'll discover newly developed deep learning models, methodologies used in the domain, and their implementation based on areas of application. You'll start by understanding the building blocks and the math behind neural networks, and then move on to CNNs and their advanced applications in computer vision. You'll also learn to apply the most popular CNN architectures in object detection and image segmentation. Further on, you'll focus on variational autoencoders and GANs. You'll then use neural networks to extract sophisticated vector representations of words, before going on to cover various types of recurrent networks, such as LSTM and GRU. You'll even explore the attention mechanism to process sequential data without the help of recurrent neural networks (RNNs). Later, you'll use graph neural networks for processing structured data, along with covering meta-learning, which allows you to train neural networks with fewer training samples. Finally, you'll understand how to apply deep learning to autonomous vehicles. By the end of this book, you'll have mastered key deep learning concepts and the different applications of deep learning models in the real world. What you will learn Cover advanced and state-of-the-art neural network architectures Understand the theory and math behind neural networks Train DNNs and apply them to modern deep learning problems Use CNNs for object detection and image segmentation Implement generative adversarial networks (GANs) and variational autoencoders to generate new images Solve natural language processing (NLP) tasks, such as machine translation, using sequence-to-sequence models Understand DL techniques, such as meta-learning and graph neural networks Who this book is for This book is for data scientists, deep learning engineers and researchers, and AI developers who want to further their knowledge of deep learning and build innovative and unique deep learning projects. Anyone looking to get to grips with advanced use cases and methodologies adopted in the deep learning domain using real-world examples will also find this book useful. Basic understanding of deep learning concepts and working knowledge of the Python programming language is assumed.

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

Introduction to Deep Learning-Eugene Charniak 2019-02-19 A project-based guide to the basics of deep learning. This concise, project-driven guide to deep learning takes readers through a series of program-writing tasks that introduce them to the use of deep learning in such areas of artificial intelligence as computer vision, natural-language processing, and reinforcement learning. The author, a longtime artificial intelligence researcher specializing in natural-language processing, covers feed-forward neural nets, convolutional neural nets, word embeddings, recurrent neural nets, sequence-to-sequence learning, deep reinforcement learning, unsupervised models, and other fundamental concepts and techniques. Students and practitioners learn the basics of deep learning by working through programs in Tensorflow, an open-source machine learning framework. "I find I learn computer science material best by sitting down and writing programs," the author writes, and the book reflects this approach. Each chapter includes a programming project, exercises, and references for further reading. An early chapter is devoted to TensorFlow and its interface with Python, the widely used programming language. Familiarity with linear algebra, multivariate calculus, and probability and statistics is required, as is a rudimentary knowledge of programming in Python. The book can be used in both undergraduate and graduate courses; practitioners will find it an essential reference.

Building Machine Learning Systems with Python - Second Edition-Luis Pedro Coelho 2015-03-26 This book primarily targets Python developers who want to learn and use Python's machine learning capabilities and gain valuable insights from data to develop effective solutions for business problems.

Practical Machine Learning for Data Analysis Using Python-Abdulhamit Subasi 2020-06-05 Practical Machine Learning for Data Analysis Using Python is a problem solver's guide for creating real-world intelligent systems. It provides a comprehensive approach with concepts, practices, hands-on examples, and sample code. The book teaches readers the vital skills required to understand and solve different problems with machine learning. It teaches machine learning techniques necessary to become a successful practitioner, through the presentation of real-world case studies in Python machine learning ecosystems. The book also focuses on building a foundation of machine learning knowledge to solve different real-world case studies across various fields, including biomedical signal analysis, healthcare, security, economics, and finance. Moreover, it covers a wide range of machine learning models, including regression, classification, and forecasting. The goal of the book is to help a broad range of readers, including IT professionals, analysts, developers, data scientists, engineers, and graduate students, to solve their own real-world problems. Offers a comprehensive overview of the application of machine learning tools in data analysis across a wide range of subject areas Teaches readers how to apply machine learning techniques to biomedical signals, financial data, and healthcare data Explores important classification and regression algorithms as well as other machine learning techniques Explains how to use Python to handle data extraction, manipulation, and exploration techniques, as well as how to visualize data spread across multiple dimensions and extract useful features

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 book 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.

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 Bevelveld and beautiful illustrations by Aglæe 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.

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 and Deep Learning Using Python and TensorFlow-Venkata Reddy Konasani 2021-04-29 Understand the principles and practices of machine learning and deep learning This hands-on guide lays out machine learning and deep learning techniques and technologies in a style that is approachable, using just the basic math required. Written by a pair of experts in the field, Machine Learning and Deep Learning Using Python and TensorFlow contains case studies in several industries, including banking, insurance, e-commerce, retail, and healthcare. The book shows how to utilize machine learning and deep learning functions in today's smart devices and apps. You get download links for datasets, code, and sample projects referred to in the text. Coverage includes: Machine learning and deep learning concepts Python programming and statistics fundamentals Regression and logistic regression Decision trees Model selection and cross-validation Cluster analysis Random forests and boosting Artificial neural networks TensorFlow and Keras Deep learning hyperparameters Convolutional neural networks Recurrent neural networks and long short-term memory

Programming Machine Learning-Paolo Perrotta 2020-03-31 You've decided to tackle machine learning - because you're job hunting, embarking on a new project, or just think self-driving cars are cool. But where to start? It's easy to be intimidated, even as a software developer. The good news is that it doesn't have to be that hard. Master machine learning by writing code one line at a time, from simple learning programs all the way to a true deep learning system. Tackle the hard topics by breaking them down so they're easier to understand, and build your confidence by getting your hands dirty. Peel away the obscurities of machine learning, starting from scratch and going all the way to deep learning. Machine learning can be intimidating, with its reliance on math and algorithms that most programmers don't encounter in their regular work. Take a hands-on approach, writing the Python code yourself, without any libraries to obscure what's really going on. Iterate on your design, and add layers of complexity as you go. Build an image recognition application from scratch with supervised learning. Predict the future with linear regression. Dive into gradient descent, a fundamental algorithm that drives most of machine learning. Create perceptrons to classify data. Build neural networks to tackle more complex and sophisticated data sets. Train and refine those networks with backpropagation and batching. Layer the neural networks, eliminate overfitting, and add convolution to transform your neural network into a true deep learning system. Start from the beginning and code your way to machine learning mastery. What You Need: The examples in this book are written in Python, but don't worry if you don't know this language: you'll pick up all the Python you need very quickly. Apart from that, you'll only need your computer, and your code-adept brain.