Keras Your Simple Guide English Edition

As recognized, adventure as without difficulty as experience practically lesson, amusement, as competently as union can be gotten by just checking out a book **Keras Your Simple Guide English Edition** with it is not directly done, you could endure even more as regards this life, on the subject of the world.

We find the money for you this proper as competently as easy pretension to acquire those all. We pay for Keras Your Simple Guide English Edition and numerous ebook collections from fictions to scientific research in any way. in the middle of them is this Keras Your Simple Guide English Edition that can be your partner.

Keras Your Simple Guide English Edition ***2019-02-08***

CABRERA ASHTYN

Deep Learning with TensorFlow 2 and Keras "O'Reilly Media, Inc."
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 Fran�ois 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 Fran�ois 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 Fran�ois 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 for Natural Language Processing Packt Publishing Ltd
"The manner in which computers are now able to mimic human thinking to process information is rapidly exceeding human capabilities in everything from chess to picking the winner of a song contest. In the modern age of machine learning, computers do not strictly need to receive an 'input command' to perform a task, but rather 'input data'. From the input of data they are able to form their own decisions and take actions virtually as a human world. But given it is a machine, it can consider many more scenarios and execute far more complicated calculations to solve complex problems. This is the element that excites data scientists and machine learning engineers the most. The ability to solve complex problems never before attempted. This book will dive in to introduce machine learning, and is ideal for beginners starting out in machine learning."--page 4 of cover.
Hands-On Transfer Learning with Python Packt Publishing Ltd
Deep learning is the most interesting and powerful machine learning technique right now. Top deep learning libraries are available on the Python ecosystem like Theano and TensorFlow. Tap into their power in a few lines of code using Keras, the best-of-breed applied deep learning library. In this Ebook, learn exactly how to get started and apply deep learning to your own machine learning projects.
Text Analytics with Python No Starch Press
A Practicing Guide to TensorFlow and Deep Learning KEY FEATURES ● Equipped with a necessary introduction to Deep Learning and AI. ● Includes demos and templates to give your projects a good start. ● Find more on the most important facets of AI, image, and speech recognition. DESCRIPTION This book begins with the configuration of an Anaconda development environment, essential for practicing the deep learning process. The basics of machine learning, which are needed for Deep Learning, are explained in this book. TensorFlow is the industry-standard library for Deep Learning, and thereby, it is covered extensively with both versions, 1.x and 2.x. As neural networks are the heart of Deep Learning, the book explains them in great detail and systematic fashion, beginning with a single neuron and progressing through deep multilayer neural networks. The emphasis of this book is on the practical application of key concepts rather than going in-depth with them. After establishing a firm basis in TensorFlow and Neural Networks, the book explains the concepts of image recognition using Convolutional Neural Networks (CNN), followed by speech recognition, and

natural language processing (NLP). Additionally, this book discusses Transformers, the most recent advancement in NLP. WHAT YOU WILL LEARN ● Create machine learning models for classification and regression. ● Utilize TensorFlow 1.x to implement neural networks. ● Work with the Keras API and TensorFlow 2. ● Learn to design and train image categorization models. ● Construct translation and Q & A apps using transformer-based language models. WHO THIS BOOK IS FOR This book is intended for those new to Deep Learning who want to learn about its principles and applications. Before you begin, you'll need to be familiar with Python. TABLE OF CONTENTS 1. Introduction to Artificial Intelligence 2. Machine Learning 3. TensorFlow Programming 4. Neural Networks 5. TensorFlow 2 6. Image Recognition 7. Speech Recognition

Deep Learning with Python Packt Publishing Ltd

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 may not 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!**

Deep Learning with Keras "O'Reilly Media, Inc."

Learn advanced state-of-the-art deep learning techniques and their applications using popular Python libraries Key FeaturesBuild a strong foundation in neural networks and deep learning with Python librariesExplore advanced deep learning techniques and their applications across computer vision and NLPLearn how a computer can navigate in complex environments with reinforcement learningBook 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 learnGrasp the mathematical theory behind neural networks and deep learning processesInvestigate and resolve computer vision challenges using convolutional networks and capsule networksSolve generative tasks using variational autoencoders and Generative Adversarial NetworksImplement complex NLP tasks using recurrent networks (LSTM and GRU) and attention modelsExplore reinforcement learning and understand how agents behave in a complex environmentGet up to date with applications of deep learning in autonomous vehiclesWho 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 Python Data Science

Develop generative models for a variety of real-world use-cases and deploy them to production Key FeaturesDiscover various GAN architectures using Python and Keras libraryUnderstand how GAN models function with the help of theoretical and practical examplesApply your learnings to become an active contributor to open source GAN applicationsBook Description Generative Adversarial Networks (GANs) have revolutionized the fields of machine learning and deep learning. This book will be your first step towards understanding GAN architectures and tackling the challenges involved in training them. This book opens with an introduction to deep learning and generative models, and their applications in artificial intelligence (AI). You will then learn how to build, evaluate, and improve your first GAN with the help of easy-to-follow examples. The next few chapters will guide you through training a GAN model to produce and improve high-resolution images. You will also learn how to implement conditional GANs that give you the ability to control characteristics of GAN outputs. You will build on your knowledge further by exploring a new training methodology for progressive growing of GANs. Moving on, you'll gain insights into state-of-the-art models in image synthesis, speech enhancement, and natural language generation using GANs. In addition to this,

you'll be able to identify GAN samples with TequilaGAN. By the end of this book, you will be well-versed with the latest advancements in the GAN framework using various examples and datasets, and you will have the skills you need to implement GAN architectures for several tasks and domains, including computer vision, natural language processing (NLP), and audio processing. Foreword by Ting-Chun Wang, Senior Research Scientist, NVIDIA What you will learnLearn how GANs work and the advantages and challenges of working with themControl the output of GANs with the help of conditional GANs, using embedding and space manipulationApply GANs to computer vision, NLP, and audio processingUnderstand how to implement progressive growing of GANsUse GANs for image synthesis and speech enhancementExplore the future of GANs in visual and sonic artsImplement pix2pixHD to turn semantic label maps into photorealistic imagesWho this book is for This book is for machine learning practitioners, deep learning researchers, and AI enthusiasts who are looking for a perfect mix of theory and hands-on content in order to implement GANs using Keras. Working knowledge of Python is expected.

Python Deep Learning "O'Reilly Media, Inc."

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

Learn AI with Python Packt Publishing Ltd

Roughly inspired by the human brain, deep neural networks trained with large amounts of data can solve complex tasks with unprecedented accuracy. This practical book provides an end-to-end guide to TensorFlow, the leading open source software library that helps you build and train neural networks for computer vision, natural language processing (NLP), speech recognition, and general predictive analytics. Authors Tom Hope, Yehezkel Resheff, and Itay Lieder provide a hands-on approach to TensorFlow fundamentals for a broad technical audience—from data scientists and engineers to students and researchers. You’ll begin by working through some basic examples in TensorFlow before diving deeper into topics such as neural network architectures, TensorBoard visualization, TensorFlow abstraction libraries, and multithreaded input pipelines. Once you finish this book, you’ll know how to build and deploy production-ready deep learning systems in TensorFlow. Get up and running with TensorFlow, rapidly and painlessly Learn how to use TensorFlow to build deep learning models from the ground up Train popular deep learning models for computer vision and NLP Use extensive abstraction libraries to make development easier and faster Learn how to scale TensorFlow, and use clusters to distribute model training Deploy TensorFlow in a production setting

Advanced Deep Learning with Keras Frank Millstein

A comprehensive guide to advanced deep learning techniques, including Autoencoders, GANs, VAEs, and Deep Reinforcement Learning, that drive today's most impressive AI results Key FeaturesExplore the most advanced deep learning techniques that drive modern AI resultsImplement Deep Neural Networks, Autoencoders, GANs, VAEs, and Deep Reinforcement LearningA wide study of GANs, including Improved GANs, Cross-Domain GANs and Disentangled Representation GANsBook Description Recent developments in deep learning, including GANs, Variational Autoencoders, and Deep Reinforcement Learning, are creating impressive AI results in our news headlines - such as AlphaGo Zero beating world chess champions, and generative AI that can create art paintings that sell for over $400k because they are so human-like. Advanced Deep Learning with Keras is a comprehensive guide to the advanced deep learning techniques available today, so you can create your own cutting-edge AI. Using Keras as an open-source deep learning library, you'll find hands-on projects throughout that show you how to create more effective AI with the latest techniques. The journey begins with an overview of MLPs, CNNs, and RNNs, which are the building blocks for the more advanced techniques in the book. You’ll learn how to implement deep learning models with Keras and Tensorflow, and move forwards to advanced techniques, as you explore deep neural network architectures, including ResNet and DenseNet, and how to create Autoencoders. You then learn all about Generative Adversarial Networks (GANs), and how they can open new levels of AI performance. Variational AutoEncoders (VAEs) are implemented, and you’ll see how GANs and VAEs have the generative power to synthesize data that can be extremely convincing to humans - a major stride forward for modern AI. To complete this set of advanced techniques, you'll learn how to implement Deep Reinforcement Learning (DRL) such as Deep Q-Learning and Policy Gradient Methods, which are critical to many modern results in AI. What you will learnCutting-edge techniques in human-like AI performanceImplement advanced deep learning models using KerasThe building blocks for advanced techniques - MLPs, CNNs, and RNNsDeep neural networks – ResNet and DenseNetAutoencoders and Variational AutoEncoders (VAEs)Generative Adversarial Networks (GANs) and creative AI techniquesDisentangled Representation GANs, and Cross-Domain GANsDeep Reinforcement Learning (DRL) methods and implementationProduce industry-standard applications using OpenAI gymDeep Q-Learning and Policy Gradient MethodsWho this book is for Some fluency with Python is assumed. As an advanced book, you'll be familiar with some machine learning approaches, and some practical experience with DL will be helpful. Knowledge of Keras or TensorFlow is not required but would be helpful.

Deep Learning for Computer Vision Independently Published

Take a hands-on approach to understanding deep learning and build smart applications that can recognize images and interpret text Key Features Understand how to implement deep learning with TensorFlow and Keras Learn the fundamentals of computer vision and image recognition Study the architecture of different neural networks Book Description Are you fascinated by how deep learning powers intelligent applications such as self-driving cars, virtual assistants, facial recognition devices, and chatbots to process data and solve complex problems? Whether you are familiar with machine

learning or are new to this domain, The Deep Learning Workshop will make it easy for you to understand deep learning with the help of interesting examples and exercises throughout. The book starts by highlighting the relationship between deep learning, machine learning, and artificial intelligence and helps you get comfortable with the TensorFlow 2.0 programming structure using hands-on exercises. You'll understand neural networks, the structure of a perceptron, and how to use TensorFlow to create and train models. The book will then let you explore the fundamentals of computer vision by performing image recognition exercises with convolutional neural networks (CNNs) using Keras. As you advance, you'll be able to make your model more powerful by implementing text embedding and sequencing the data using popular deep learning solutions. Finally, you'll get to grips with bidirectional recurrent neural networks (RNNs) and build generative adversarial networks (GANs) for image synthesis. By the end of this deep learning book, you'll have learned the skills essential for building deep learning models with TensorFlow and Keras. What you will learn Understand how deep learning, machine learning, and artificial intelligence are different Develop multilayer deep neural networks with TensorFlow Implement deep neural networks for multiclass classification using Keras Train CNN models for image recognition Handle sequence data and use it in conjunction with RNNs Build a GAN to generate high-quality synthesized images Who this book is for If you are interested in machine learning and want to create and train deep learning models using TensorFlow and Keras, this workshop is for you. A solid understanding of Python and its packages, along with basic machine learning concepts, will help you to learn the topics quickly.

Deep Learning BPB Publications

Implement supervised, unsupervised, and generative deep learning (DL) models using Keras and Dopamine with TensorFlow Key FeaturesUnderstand the fundamental machine learning concepts useful in deep learningLearn the underlying mathematical concepts as you implement deep learning models from scratchExplore easy-to-understand examples and use cases that will help you build a solid foundation in DLBook Description With information on the web exponentially increasing, it has become more difficult than ever to navigate through everything to find reliable content that will help you get started with deep learning. This book is designed to help you if you're a beginner looking to work on deep learning and build deep learning models from scratch, and you already have the basic mathematical and programming knowledge required to get started. The book begins with a basic overview of machine learning, guiding you through setting up popular Python frameworks. You will also understand how to prepare data by cleaning and preprocessing it for deep learning, and gradually go on to explore neural networks. A dedicated section will give you insights into the working of neural networks by helping you get hands-on with training single and multiple layers of neurons. Later, you will cover popular neural network architectures such as CNNs, RNNs, AEs, VAEs, and GANs with the help of simple examples, and learn how to build models from scratch. At the end of each chapter, you will find a question and answer section to help you test what you've learned through the course of the book. By the end of this book, you'll be well-versed with deep learning concepts and have the knowledge you need to use specific algorithms with various tools for different tasks. What you will learnImplement recurrent neural networks (RNNs) and long short-term memory (LSTM) for image classification and natural language processing tasksExplore the role of convolutional neural networks (CNNs) in computer vision and signal processingDiscover the ethical implications of deep learning modelingUnderstand the mathematical terminology associated with deep learningCode a generative adversarial network (GAN) and a variational autoencoder (VAE) to generate images from a learned latent spaceImplement visualization techniques to compare AEs and VAEsWho this book is for This book is for aspiring data scientists and deep learning engineers who want to get started with the fundamentals of deep learning and neural networks. Although no prior knowledge of deep learning or machine learning is required, familiarity with linear algebra and Python programming is necessary to get started.

Deep Learning with Keras Packt Publishing Ltd

Deep learning methods are achieving state-of-the-art results on challenging machine learning problems such as describing photos and translating text from one language to another. In this new laser-focused Ebook, finally cut through the math, research papers and patchwork descriptions about natural language processing. Using clear explanations, standard Python libraries and step-by-step tutorial lessons you will discover what natural language processing is, the promise of deep learning in the field, how to clean and prepare text data for modeling, and how to develop deep learning models for your own natural language processing projects.

Deep Learning with Python: A Fundamentals Guide to Understanding Machine Learning and Artificial Intelligence with Scikit-Learn, Tensorflow, and Packt Publishing Ltd

Learn how to model and train advanced neural networks to implement a variety of Computer Vision tasks Key Features Train different kinds of deep learning model from scratch to solve specific problems in Computer Vision Combine the power of Python, Keras, and TensorFlow to build deep learning models for object detection, image classification, similarity learning, image captioning, and more Includes tips on optimizing and improving the performance of your models under various constraints Book Description Deep learning has shown its power in several application areas of Artificial Intelligence, especially in Computer Vision. Computer Vision is the science of understanding and manipulating images, and finds enormous applications in the areas of robotics, automation, and so on. This book will also show you, with practical examples, how to develop Computer Vision applications by leveraging the power of deep learning. In this book, you will learn different techniques related to object classification, object detection, image segmentation, captioning, image generation, face analysis, and more. You will also explore their applications using popular Python libraries such as TensorFlow and Keras. This book will help you master state-of-the-art, deep learning algorithms and their implementation. What you will learn Set up an environment for deep learning with Python, TensorFlow, and Keras Define and train a model for image and video classification Use features from a pre-trained Convolutional Neural Network model for image retrieval Understand and implement object detection using the real-world Pedestrian Detection scenario Learn about various problems in image captioning and how to overcome them by training images and text together Implement similarity matching and train a model for face recognition Understand the concept of generative models and use them for image generation Deploy your deep learning models and optimize them for high performance Who this book is for This book is targeted at data scientists and Computer Vision practitioners who wish to apply the concepts of Deep Learning to overcome any problem related to Computer Vision. A basic knowledge of programming in Python—and some understanding of machine learning concepts—is required to get the best out of

this book.

Python Deep Learning BPB Publications

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 H20. 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.

fastText Quick Start Guide Packt Publishing Ltd

Unlock deeper insights into Machine Leaning 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 Generative Adversarial Networks with Keras MIT Press

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

Deep Learning Packt Publishing Ltd

Deep Learning with Keras This book will introduce you to various supervised and unsupervised deep learning algorithms like the multilayer perceptron, linear regression and other more advanced deep convolutional and recurrent neural networks. You will also learn about image processing, handwritten recognition, object recognition and much more. Furthermore, you will get familiar with recurrent neural networks like LSTM and GAN as you explore processing sequence data like time series, text, and audio. The book will definitely be your best companion on this great deep learning

journey with Keras introducing you to the basics you need to know in order to take next steps and learn more advanced deep neural networks. Here Is a Preview of What You'll Learn Here... The difference between deep learning and machine learning Deep neural networks Convolutional neural networks Building deep learning models with Keras Multi-layer perceptron network models Activation functions Handwritten recognition using MNIST Solving multi-class classification problems Recurrent neural networks and sequence classification And much more... Get this book NOW and learn more about Deep Learning with Keras!

Beginning with Deep Learning Using TensorFlow Apress

Have you come across the terms machine learning and neural networks in most articles you have recently read? Do you also want to learn how to build a machine learning model that will answer your questions within a blink of your eyes? If you responded yes to any of the above questions, you have come to the right place.

Tensorflow Machine Learning Packt Publishing Ltd

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.