Hands On Machine Learning With Scikit Learn And Tensorflow

Getting the books **Hands On Machine Learning With Scikit Learn And Tensorflow** now is not type of inspiring means. You could not isolated going in imitation of books accrual or library or borrowing from your friends to gain access to them. This is an unconditionally simple means to specifically acquire lead by on-line. This online declaration Hands On Machine Learning With Scikit Learn And Tensorflow can be one of the options to accompany you once having supplementary time.

It will not waste your time. take on me, the e-book will extremely melody you supplementary business to read. Just invest tiny get older to entrance this on-line message **Hands On Machine Learning With Scikit Learn And Tensorflow** as skillfully as evaluation them wherever you are now.

Hands-on Scikit-Learn for

Machine Learning Applications
Packt Publishing Ltd
Get into the world of smart data security using machine learning algorithms and Python libraries Key FeaturesLearn machine learning algorithms and cybersecurity fundamentalsAutomate your daily workflow by applying use cases to many facets of securityImplement smart machine learning solutions to detect various cybersecurity problemsBook Description Cyber threats today are one of the costliest losses that an organization can face. In this book, we use the most efficient tool to solve the big problems that exist in the cybersecurity domain. The book begins by giving you the basics of ML in cybersecurity using Python and its libraries. You will explore various ML domains (such as time series analysis and ensemble modeling) to get your foundations right. You will implement various examples such as building system to identify malicious URLs, and building a program to detect fraudulent emails and spam. Later, you will learn how to make effective use of K-means algorithm to develop a solution to detect and alert you to any malicious activity in the network. Also learn how to implement biometrics and fingerprint to validate whether the user is a legitimate user or not. Finally, you will see how we change the game with TensorFlow and learn how deep learning is effective for creating models and training systems What you will learnUse machine learning algorithms with complex datasets to implement cybersecurity conceptsImplement machine learning algorithms such as clustering, k-means, and Naive Bayes to solve real-world

problemsLearn to speed up a system using Python libraries with NumPy, Scikit-learn, and CUDAUnderstand how to combat malware, detect spam, and fight financial fraud to mitigate cyber crimesUse TensorFlow in the cybersecurity domain and implement real-world examplesLearn how machine learning and Python can be used in complex cyber issuesWho this book is for This book is for the data scientists, machine learning developers, security researchers, and anyone keen to apply machine learning to up-skill computer security. Having some working knowledge of Python and being familiar with the basics of machine learning and cybersecurity fundamentals will help to get the most out of the book

Hands-On Transfer Learning with Python Packt Publishing Ltd

Concepts, tools, and techniques to explore deep learning architectures and methodologies Key FeaturesExplore advanced deep learning architectures using various datasets and frameworksImplement deep architectures for neural network models such as CNN, RNN, GAN, and many moreDiscover design patterns and different challenges for various deep learning architecturesBook Description Deep

learning architectures are composed of multilevel nonlinear operations that represent high-level abstractions; this allows you to learn useful feature representations from the data. This book will help you learn and implement deep learning architectures to resolve various deep learning research problems. Hands-On Deep Learning Architectures with Python explains the essential learning algorithms used for deep and shallow architectures. Packed with practical implementations and ideas to help you build efficient artificial intelligence systems (AI), this book will help you learn how neural networks play a major role in building deep architectures. You will understand various deep learning architectures (such as AlexNet, VGG Net, GoogleNet) with easy-to-follow code and diagrams. In addition to this, the book will also guide you in building and

training various deep architectures such as the Boltzmann mechanism, autoencoders, convolutional neural networks (CNNs), recurrent neural networks (RNNs), natural language processing (NLP), GAN, and more—all with practical implementations. By the end of this book, you will be able to construct deep models using popular frameworks and datasets with the required design patterns for each architecture. You will be ready to explore the potential of deep architectures in today's world. What you will learnImplement CNNs, RNNs, and other commonly used architectures with PythonExplore architectures such as VGGNet, AlexNet, and GoogLeNetBuild deep learning architectures for AI applications such as face and image recognition, fraud detection, and many moreUnderstand the architectures and applications of Boltzmann machines and autoencoders with concrete examples Master artificial

intelligence and neural network concepts and apply them to your architectureUnderstand deep learning architectures for mobile and embedded systemsWho this book is for If you're a data scientist, machine learning developer/engineer, or deep learning practitioner, or are curious about AI and want to upgrade your knowledge of various deep learning architectures, this book will appeal to you. You are expected to have some knowledge of statistics and machine learning algorithms to get the best out of this book

Machine Learning with PyTorch and Scikit-Learn CRC Press

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

Hands-On Machine Learning for Algorithmic Trading "O'Reilly Media, Inc."

A definitive guide to creating an intelligent web

application with the best of machine learning and JavaScript Key Features Solve complex computational problems in browser with JavaScript Teach your browser how to learn from rules using the power of machine learning Understand discoveries on web interface and API in machine learning Book Description In over 20 years of existence, JavaScript has been pushing beyond the boundaries of web evolution with proven existence on servers, embedded devices, Smart TVs, IoT, Smart Cars, and more. Today, with the added advantage of machine learning research and support for JS libraries, JavaScript makes your browsers smarter than ever with the ability to learn patterns and reproduce them to become a part of innovative products and applications. Hands-on Machine Learning with JavaScript presents various avenues of machine learning in a practical and objective way, and helps implement them using the JavaScript language. Predicting behaviors, analyzing feelings, grouping data, and building neural models are some of the skills you will build from this book. You will learn how to train your machine learning models and work with different kinds of data. During this journey, you will come across use cases such as face

detection, spam filtering, recommendation systems, character recognition, and more. Moreover, you will learn how to work with deep neural networks and guide your applications to gain insights from data. By the end of this book, you'll have gained hands-on knowledge on evaluating and implementing the right model, along with choosing from different JS libraries, such as NaturalNode, brain, harthur, classifier, and many more to design smarter applications. What you will learn Get an overview of state-of-the-art machine learning Understand the pre-processing of data handling, cleaning, and preparation Learn Mining and Pattern Extraction with JavaScript Build your own model for classification, clustering, and prediction Identify the most appropriate model for each type of problem Apply machine learning techniques to real-world applications Learn how JavaScript can be a powerful language for machine learning Who this book is for This book is for you if you are a JavaScript developer who wants to implement machine learning to make applications smarter, gain insightful information from the data, and enter the field of machine learning without switching to another language. Working knowledge of JavaScript language is expected to get the most out of the book.

Machine Learning Packt Publishing Ltd

Dig deep into the data with a hands-on guide to machine learning with updated examples and more! Machine Learning: Hands-On for Developers and Technical Professionals provides hands-on instruction and fully-coded working examples for the most common machine learning techniques used by developers and technical professionals. The book contains a breakdown of each ML variant, explaining how it works and how it is used within certain industries, allowing readers to incorporate the presented techniques into their own work as they follow along. A core tenant of machine learning is a strong focus on data preparation, and a full exploration of the various types of learning algorithms illustrates how the proper tools can help any developer extract information and insights from existing data. The book includes a full complement of Instructor's Materials to facilitate use in the classroom, making this resource useful for students and as a professional reference. At its core, machine learning is a mathematical, algorithm-based technology that forms the basis of historical data mining and modern big data science. Scientific analysis of big data requires a working knowledge of machine learning, which forms predictions based on known properties learned from training data. Machine Learning is an accessible, comprehensive guide for the non-mathematician, providing clear guidance that allows readers to: Learn the

languages of machine learning including Hadoop, Mahout, and Weka Understand decision trees, Bayesian networks, and artificial neural networks Implement Association Rule, Real Time, and Batch learning Develop a strategic plan for safe, effective, and efficient machine learning By learning to construct a system that can learn from data, readers can increase their utility across industries. Machine learning sits at the core of deep dive data analysis and visualization, which is increasingly in demand as companies discover the goldmine hiding in their existing data. For the tech professional involved in data science, Machine Learning: Hands-On for Developers and Technical Professionals provides the skills and techniques required to dig deeper.

Hands-On Machine Learning with R Packt Publishing Ltd Speed up the design and implementation of deep learning solutions using Apache Spark Key FeaturesExplore the world of distributed deep learning with Apache SparkTrain neural networks with deep learning libraries such as BigDL and TensorFlowDevelop Spark deep learning applications to intelligently handle large and complex datasetsBook Description Deep learning is a subset of machine learning where datasets with several layers of complexity can be processed. Hands-On Deep Learning with Apache Spark addresses the sheer complexity of technical and analytical parts and the speed at which deep learning solutions can be implemented on Apache

Spark. The book starts with the fundamentals of Apache Spark and deep learning. You will set up Spark for deep learning, learn principles of distributed modeling, and understand different types of neural nets. You will then implement deep learning models, such as convolutional neural networks (CNNs), recurrent neural networks (RNNs), and long short-term memory (LSTM) on Spark. As you progress through the book, you will gain hands-on experience of what it takes to understand the complex datasets you are dealing with. During the course of this book, you will use popular deep learning frameworks, such as TensorFlow, Deeplearning4j, and Keras to train your distributed models. By the end of this book, you'll have gained experience with the implementation of your models on a variety of use cases. What you will learnUnderstand the basics of deep learningSet up Apache Spark for deep learningUnderstand the principles of distribution modeling and different types of neural networksObtain an understanding of deep learning algorithmsDiscover textual analysis and deep learning with SparkUse popular deep learning frameworks, such as Deeplearning4j, TensorFlow, and KerasExplore popular deep learning algorithms Who this book is for If you are a Scala developer, data scientist, or data analyst who wants to learn how to use Spark for implementing efficient deep learning models, Hands-On Deep Learning with Apache Spark is for you. Knowledge of

the core machine learning concepts and some exposure to Spark will be helpful.

Deep Learning for Coders with fastai and PyTorch Packt Publishing Ltd

Graphics in this book are printed in black and white. Through a series of recent breakthroughs, deep learning has boosted the entire field of machine learning. Now, even programmers who know close to nothing about this technology can use simple, efficient tools to implement programs capable of learning from data. This practical book shows you how. By using concrete examples, minimal theory, and two production-ready Python frameworks—scikit-learn and TensorFlow—author Aurélien Géron helps you gain an intuitive understanding of the concepts and tools for building intelligent systems. You'll learn a range of techniques, starting with simple linear regression and progressing to deep neural networks. With exercises in each chapter to help you apply what you've learned, all you need is programming experience to get started. Explore the machine learning landscape, particularly neural nets Use scikit-learn to track an example machine-learning project end-to-end Explore several training models, including support vector machines, decision trees, random forests, and ensemble methods Use the TensorFlow library to build and train neural nets Dive into neural net architectures, including convolutional nets, recurrent nets, and deep reinforcement learning Learn techniques for

training and scaling deep neural nets Apply practical code examples without acquiring excessive machine learning theory or algorithm details

Hands-On Machine Learning with Scikit-Learn and TensorFlow Packt Publishing Ltd

A practical guide to getting the most out of Excel, using it for data preparation, applying machine learning models (including cloud services) and understanding the outcome of the data analysis. Key FeaturesUse Microsoft's product Excel to build advanced forecasting models using varied examples Cover range of machine learning tasks such as data mining, data analytics, smart visualization, and more Derive data-driven techniques using Excel plugins and APIs without much code required Book Description We have made huge progress in teaching computers to perform difficult tasks, especially those that are repetitive and time-consuming for humans. Excel users, of all levels, can feel left behind by this innovation wave. The truth is that a large amount of the work needed to develop and use a machine learning model can be done in Excel. The book starts by giving a general introduction to machine learning, making every concept clear and understandable. Then, it shows every step of a machine learning project, from data collection, reading from different data sources, developing models, and visualizing the results using Excel features and offerings. In every chapter, there are several examples and hands-on exercises that will show the reader how to combine Excel functions, add-ins, and connections to databases and to cloud services to reach the desired goal: building a full data analysis flow. Different machine learning models are shown, tailored to the type of data to be analyzed. At the end of the book, the reader is presented with some advanced use cases using Automated Machine Learning, and artificial neural network, which simplifies the analysis task and represents the future of machine learning. What you will learnUse

Excel to preview and cleanse datasetsUnderstand correlations between variables and optimize the input to machine learning modelsUse and evaluate different machine learning models from ExcelUnderstand the use of different visualizationsLearn the basic concepts and calculations to understand how artificial neural networks workLearn how to connect Excel to the Microsoft Azure cloudGet beyond proof of concepts and build fully functional data analysis flowsWho this book is for This book is for data analysis, machine learning enthusiasts, project managers, and someone who doesn't want to code much for performing core tasks of machine learning. Each example will help you perform end-to-end smart analytics. Working knowledge of Excel is required.

The Hundred-page Machine Learning Book Packt Publishing Ltd

Integrate scikit-learn with various tools such as NumPy, pandas, imbalanced-learn, and scikit-surprise and use it to solve real-world machine learning problems Key FeaturesDelve into machine learning with this comprehensive guide to scikit-learn and scientific PythonMaster the art of data-driven problem-solving with hands-on examplesFoster your theoretical and practical knowledge of supervised and unsupervised machine learning algorithmsBook Description Machine learning is applied everywhere, from business to research and academia, while scikit-learn is a versatile library that is popular among machine learning practitioners. This book serves as a practical guide for anyone looking to provide hands-on machine learning solutions with scikit-learn and Python toolkits. The book begins with an

explanation of machine learning concepts and fundamentals, and strikes a balance between theoretical concepts and their applications. Each chapter covers a different set of algorithms, and shows you how to use them to solve real-life problems. You ' ll also learn about various key supervised and unsupervised machine learning algorithms using practical examples. Whether it is an instance-based learning algorithm, Bayesian estimation, a deep neural network, a tree-based ensemble, or a recommendation system, you ' ll gain a thorough understanding of its theory and learn when to apply it. As you advance, you ' ll learn how to deal with unlabeled data and when to use different clustering and anomaly detection algorithms. By the end of this machine learning book, you ' ll have learned how to take a data-driven approach to provide end-to-end machine learning solutions. You ' ll also have discovered how to formulate the problem at hand, prepare required data, and evaluate and deploy models in production. What you will learnUnderstand when to use supervised, unsupervised, or reinforcement learning algorithmsFind out how to collect and prepare your data for machine learning tasksTackle imbalanced data and optimize your algorithm for a bias or variance tradeoffApply supervised and unsupervised algorithms to overcome various machine learning challengesEmploy best practices for tuning your algorithm ' s hyper

parametersDiscover how to use neural networks for classification and regressionBuild, evaluate, and deploy your machine learning solutions to productionWho this book is for This book is for data scientists, machine learning practitioners, and anyone who wants to learn how machine learning algorithms work and to build different machine learning models using the Python ecosystem. The book will help you take your knowledge of machine learning to the next level by grasping its ins and outs and tailoring it to your needs. Working knowledge of Python and a basic understanding of underlying mathematical and statistical concepts is required.

Hands-On Unsupervised Learning Using Python Packt Publishing Ltd

Deep learning is often viewed as the exclusive domain of math PhDs and big tech companies. But as this hands-on guide demonstrates, programmers comfortable with Python can achieve impressive results in deep learning with little math background, small amounts of data, and minimal code. How? With fastai, the first library to provide a consistent interface to the most frequently used deep learning applications. Authors Jeremy Howard and Sylvain Gugger, the creators of fastai, show you how to train a model on a wide range of tasks using fastai and PyTorch. You ' ll also dive progressively further into deep learning theory to gain a complete understanding of the algorithms behind the scenes. Train models in computer vision, natural language

processing, tabular data, and collaborative filtering Learn the latest deep learning techniques that matter most in practice Improve accuracy, speed, and reliability by understanding how deep learning models work Discover how to turn your models into web applications Implement deep learning algorithms from scratch Consider the ethical implications of your work Gain insight from the foreword by PyTorch cofounder, Soumith Chintala

Hands-On Machine Learning with C++ Packt Publishing Ltd

The fundamental mathematical tools needed to understand machine learning include linear algebra, analytic geometry, matrix decompositions, vector calculus, optimization, probability and statistics. These topics are traditionally taught in disparate courses, making it hard for data science or computer science students, or professionals, to efficiently learn the mathematics. This self-contained textbook bridges the gap between mathematical and machine learning texts, introducing the mathematical concepts with a minimum of prerequisites. It uses these concepts to derive four central machine learning methods: linear regression, principal component analysis, Gaussian mixture models and support vector machines. For students and others with a mathematical background, these derivations provide a starting point to machine learning texts. For those learning the mathematics for the first time, the methods help build intuition and practical experience with applying mathematical concepts. Every chapter includes worked examples and exercises to test understanding. Programming

tutorials are offered on the book's web site.

Grokking Deep Learning O'Reilly Media

This book covers the fundamentals of machine learning with Python in a concise and dynamic manner. It covers data mining and large-scale machine learning using Apache Spark. About This Book Take your first steps in the world of data science by understanding the tools and techniques of data analysis Train efficient Machine Learning models in Python using the supervised and unsupervised learning methods Learn how to use Apache Spark for processing Big Data efficiently Who This Book Is For If you are a budding data scientist or a data analyst who wants to analyze and gain actionable insights from data using Python, this book is for you. Programmers with some experience in Python who want to enter the lucrative world of Data Science will also find this book to be very useful, but you don't need to be an expert Python coder or mathematician to get the most from this book. What You Will Learn Learn how to clean your data and ready it for analysis Implement the popular clustering and regression methods in Python Train efficient machine learning models using decision trees and random forests Visualize the results of your analysis using Python's Matplotlib library Use Apache Spark's MLlib package to perform machine learning on large datasets In Detail Join Frank Kane, who worked on Amazon and IMDb's machine learning algorithms, as he guides you on your first steps into the world of data science. Hands-On Data Science and Python Machine Learning gives you the tools that you need to understand and explore the core topics in the field, and the confidence and practice to build and analyze your own machine learning models. With the help of interesting and easy-to-follow practical examples, Frank Kane explains potentially complex topics such as Bayesian methods and K-means clustering in a way that anybody can understand them. Based on Frank's successful data science course, Hands-On Data Science and Python Machine

Learning empowers you to conduct data analysis and perform efficient machine learning using Python. Let Frank help you unearth the value in your data using the various data mining and data analysis techniques available in Python, and to develop efficient predictive models to predict future results. You will also learn how to perform large-scale machine learning on Big Data using Apache Spark. The book covers preparing your data for analysis, training machine learning models, and visualizing the final data analysis. Style and approach This comprehensive book is a perfect blend of theory and hands-on code examples in Python which can be used for your reference at any time.

Machine Learning "O'Reilly Media, Inc."

Implement supervised and unsupervised machine learning algorithms using C++ libraries such as PyTorch C++ API, Caffe2, Shogun, Shark-ML, mlpack, and dlib with the help of real-world examples and datasets Key FeaturesBecome familiar with data processing, performance measuring, and model selection using various C++ librariesImplement practical machine learning and deep learning techniques to build smart modelsDeploy machine learning models to work on mobile and embedded devicesBook Description C++ can make your machine learning models run faster and more efficiently. This handy guide will help you learn the fundamentals of machine learning (ML), showing you how to use C++ libraries to get the most out of your data. This book makes machine learning with C++ for beginners easy with its example-based approach, demonstrating how to implement supervised and unsupervised ML algorithms through real-world examples.

This book will get you hands-on with tuning and optimizing a model for different use cases, assisting you with model selection and the measurement of performance. You ' ll cover techniques such as product recommendations, ensemble learning, and anomaly detection using modern C++ libraries such as PyTorch C++ API, Caffe2, Shogun, Shark-ML, mlpack, and dlib. Next, you ' ll explore neural networks and deep learning using examples such as image classification and sentiment analysis, which will help you solve various problems. Later, you ' ll learn how to handle production and deployment challenges on mobile and cloud platforms, before discovering how to export and import models using the ONNX format. By the end of this C++ book, you will have real-world machine learning and C++ knowledge, as well as the skills to use C++ to build powerful ML systems. What you will learnExplore how to load and preprocess various data types to suitable C++ data structuresEmploy key machine learning algorithms with various C++ librariesUnderstand the grid-search approach to find the best parameters for a machine learning modelImplement an algorithm for filtering anomalies in user data using Gaussian distributionImprove collaborative filtering to deal with dynamic user preferencesUse C++ libraries and APIs to manage model structures and parametersImplement a C++ program to solve image classification tasks with LeNet architectureWho this book is for You will find this C++

machine learning book useful ifthis book.
you want to get started with machine learning algorithms and techniques using the popular C++ language. As well as being a useful first course in machine learning with C++, this book will also appeal to data analysts, data scientists, and machine learning developers who are looking to implement different machine learning models in production using varied datasets and examples. Working knowledge of the C++ programming language is mandatory to get started with

Hands-On Deep Learning for Games Packt Publishing Ltd
A hands-on guide enriched with examples to master deep reinforcement learning algorithms with Python Key Features Your entry point into the world of artificial intelligence using the power of Python An example-rich guide to master various RL and DRL algorithms Explore various state-of-the-art architectures along with math Book Description Reinforcement Learning (RL) is the trending and most promising branch of artificial intelligence. Hands-On Reinforcement learning with Python will help you master not only the basic reinforcement learning algorithms but also the advanced deep reinforcement learning algorithms. The book starts with an introduction to Reinforcement Learning followed by OpenAI Gym, and TensorFlow. You will then explore various RL algorithms and concepts, such as Markov Decision Process, Monte Carlo methods, and dynamic programming, including value and policy iteration. This example-rich guide will introduce you to deep reinforcement learning

algorithms, such as Dueling DQN, DRQN, A3C, PPO, and TRPO. You will also learn about imagination-augmented agents, learning from human preference, DQfD, HER, and many more of the recent advancements in reinforcement learning. By the end of the book, you will have all the knowledge and experience needed to implement reinforcement learning and deep reinforcement learning in your projects, and you will be all set to enter the world of artificial intelligence. What you will learn Understand the basics of reinforcement learning methods, algorithms, and elements Train an agent to walk using OpenAI Gym and Tensorflow Understand the Markov Decision Process, Bellman ' s optimality, and TD learning Solve multi-armed-bandit problems using various algorithms Master deep learning algorithms, such as RNN, LSTM, and CNN with applications Build intelligent agents using the DRQN algorithm to play the Doom game Teach agents to play the Lunar Lander game using DDPG Train an agent to win a car racing game using dueling DQN Who this book is for If you ' re a machine learning developer or deep learning enthusiast interested in artificial intelligence and want to learn about reinforcement learning from scratch, this book is for you. Some knowledge of linear algebra, calculus, and the Python programming language will help you understand the concepts covered in this book.

Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow O'Reilly Media

Hands-on Machine Learning with R provides a practical and applied approach to learning and developing intuition into today's most popular machine learning

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

learning and more. · Uses a hands-on approach and real world data.

Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow, 3rd Edition

Cambridge University Press

A comprehensive guide to getting well-versed with the mathematical techniques for building modern deep learning architectures Key FeaturesUnderstand linear algebra, calculus, gradient algorithms, and other concepts essential for training deep neural networksLearn the mathematical concepts needed to understand how deep learning models functionUse deep learning for solving problems related to vision, image, text, and sequence applicationsBook Description Most programmers and data scientists struggle with mathematics, having either overlooked or forgotten core mathematical concepts. This book uses Python libraries to help you understand the math required to build deep learning (DL) models. You'll begin by learning about core mathematical and modern computational techniques used to design and implement DL algorithms. This book will cover essential topics, such as linear algebra, eigenvalues and eigenvectors, the singular value decomposition concept, and gradient algorithms, to help you understand how to train deep neural networks. Later chapters focus on important neural networks, such as the linear neural network and multilayer perceptrons, with a primary focus on helping you learn how each model works. As you advance, you will delve into the math used for regularization, multi-layered DL, forward propagation, optimization, and backpropagation techniques to understand what it takes to build full-fledged DL models. Finally, you ' ll explore CNN, recurrent neural network (RNN), and GAN models and their application. By the end of this book, you'll have built a strong foundation in neural networks and DL mathematical concepts, which will help you to confidently research and build custom models in DL. What you will learnUnderstand the key

mathematical concepts for building neural network modelsDiscover core multivariable calculus conceptsImprove the performance of deep learning models using optimization techniquesCover optimization algorithms, from basic stochastic gradient descent (SGD) to the advanced Adam optimizerUnderstand computational graphs and their importance in DLExplore the backpropagation algorithm to reduce output errorCover DL algorithms such as convolutional neural networks (CNNs), sequence models, and generative adversarial networks (GANs)Who this book is for This book is for data scientists, machine learning developers, aspiring deep learning developers, or anyone who wants to understand the foundation of deep learning by learning the math behind it. Working knowledge of the Python programming language and machine learning basics is required.

Hands-On Machine Learning on Google Cloud Platform Packt Publishing Ltd

While several market-leading companies have successfully transformed their business models by following data- and AI-driven paths, the vast majority have yet to reap the benefits. How can your business and analytics units gain a competitive advantage by capturing the full potential of this predictive revolution? This practical guide presents a battle-tested end-to-end method to help you translate business decisions into tractable prescriptive solutions using data and AI as fundamental inputs. Author Daniel Vaughan shows data scientists, analytics practitioners, and others interested in using AI to transform their businesses not only how to ask the right questions but also how to generate value using modern AI technologies and decision-making principles. You ' ll explore several use cases common to many enterprises, complete with examples you can apply when working to solve your own issues. Break business decisions into stages that can be tackled using different skills from the analytical toolbox Identify and embrace uncertainty in decision making and protect against

common human biases Customize optimal decisions to different customers using predictive and prescriptive methods and technologies Ask business questions that create high value through AI- and data-driven technologies

Hands-On Reinforcement Learning for Games Packt Publishing Ltd

Summary Grokking Deep Learning teaches you to build deep learning neural networks from scratch! In his engaging style, seasoned deep learning expert Andrew Trask shows you the science under the hood, so you grok for yourself every detail of training neural networks. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Deep learning, a branch of artificial intelligence, teaches computers to learn by using neural networks, technology inspired by the human brain. Online text translation, self-driving cars, personalized product recommendations, and virtual voice assistants are just a few of the exciting modern advancements possible thanks to deep learning. About the Book Grokking Deep Learning teaches you to build deep learning neural networks from scratch! In his engaging style, seasoned deep learning expert Andrew Trask shows you the science under the hood, so you grok for yourself every detail of training neural networks. Using only Python and its math-supporting library, NumPy, you'll train your own neural networks to see and understand images, translate text into different languages, and even write like Shakespeare! When you're done, you'll be fully prepared to move on to mastering deep learning frameworks. What's inside The science behind deep learning Building and training your own neural networks Privacy concepts, including federated learning Tips for continuing your pursuit of deep learning About the Reader For readers with high school-level math and intermediate programming skills. About the Author Andrew Trask is a PhD student at Oxford University and a research scientist at

DeepMind. Previously, Andrew was a researcher and analytics product manager at Digital Reasoning, where he trained the world's largest artificial neural network and helped guide the analytics roadmap for the Synthesys cognitive computing platform. Table of Contents Introducing deep learning: why you should learn it Fundamental concepts: how do machines learn? Introduction to neural prediction: forward propagation Introduction to neural learning: gradient descent Learning multiple weights at a time: generalizing gradient descent Building your first deep neural network: introduction to backpropagation How to picture neural networks: in your head and on paper Learning signal and ignoring noise:introduction to regularization and batching Modeling probabilities and nonlinearities: activation functions Neural learning about edges and corners: intro to convolutional neural networks Neural networks that understand language: king - man + woman == ? Neural networks that write like Shakespeare: recurrent layers for variable-length data Introducing automatic optimization: let's build a deep learning framework Learning to write like Shakespeare: long short-term memory Deep learning on unseen data: introducing federated learning Where to go from here: a brief guide

Hands-On Machine Learning for Cybersecurity CRC Press

Through a recent series of breakthroughs, deep learning has boosted the entire field of machine learning. Now, even programmers who know close to nothing about this technology can use simple, efficient tools to implement programs capable of learning from data. This best-selling book uses concrete examples, minimal theory, and production-ready Python frameworks--scikit-learn, Keras, and TensorFlow--to help you gain an intuitive understanding of the concepts and tools for building intelligent systems. With this updated third edition, author Aurelien Geron explores

a range of techniques, starting with simple linear regression and progressing to deep neural networks. Numerous code examples and exercises throughout the book help you apply what you've learned. Programming experience is all you need to get started. Use scikit-learn to track an example machine learning project end to end Explore several models, including support vector machines, decision trees, random forests, and ensemble methods Exploit unsupervised learning techniques such as dimensionality reduction, clustering, and anomaly detection Dive into neural net architectures, including convolutional nets, recurrent nets, generative adversarial networks, and transformers Use TensorFlow and Keras to build and train neural nets for computer vision, natural language processing, generative models, and deep reinforcement learning Train neural nets using multiple GPUs and deploy them at scale using Google's Vertex AI.

Hands-On Machine Learning with Microsoft Excel 2019 Packt Publishing Ltd

Understand basic to advanced deep learning algorithms, the mathematical principles behind them, and their practical applications. Key FeaturesGet up-to-speed with building your own neural networks from scratch Gain insights into the mathematical principles behind deep learning algorithmsImplement popular deep learning algorithms such as CNNs, RNNs, and more using TensorFlowBook 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 learnImplement basic-to-advanced deep learning algorithmsMaster the mathematics behind deep learning algorithmsBecome familiar with gradient descent and its variants, such as AMSGrad, AdaDelta, Adam, and NadamImplement recurrent networks, such as RNN, LSTM, GRU, and seq2seq modelsUnderstand how machines interpret images using CNN and capsule networksImplement different types of generative adversarial network, such as CGAN, CycleGAN, and

StackGANExplore various types of autoencoder, such as Sparse autoencoders, DAE, CAE, and VAEWho 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.