Handson Machine Learning With Scikitlearn And Tensorflow

Convex OptimizationHands-On Machine Learning with Scikit-LearnDeep Learning for Coders with fastai and PyTorchHands-On Machine Learning with C++Hands-on Scikit-learn for Machine LearningHands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow, 2nd EditionHands-On Machine Learning with Scikit-Learn and TensorFlowIntroduction to Machine Learning with PythonMachine Learning for Cybersecurity CookbookHands-On Automated Machine LearningMastering Machine Learning with scikit-learnHands-On Image Processing with PythonBuilding Machine Learning Powered ApplicationsApplied Machine LearningHands-on Scikit-Learn for Machine Learning ApplicationsA Modern Approach to Regression with RDeep LearningHands-On Machine Learning for CybersecurityLearning Scikit-LearnPython for Data AnalysisHands-On Machine Learning with Scikit-Learn and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent SystemsHands-On Machine Learning for Algorithmic TradingApplied Deep Learning with PythonHands-On Unsupervised Learning Using PythonPython Data Science HandbookHands-On Machine Learning with scikit-learn and Scientific Python ToolkitsHands-On Deep Learning Algorithms with PythonHands-On Machine Learning with Scikit-Learn and TensorFlowHands-On Machine Learning with RHands-On Ensemble Learning with PythonMATLAB for Machine LearningHands on Machine Learning with PythonIntroduction to Data ScienceHands-On Machine Learning with TensorFlow.jsBuilding Machine Learning Systems with Python - Second EditionHands-On Machine Learning with Scikit-Learn and TensorFlowHands-On Machine Learning with Scikit-Learn, Keras, and TensorFlowHands-On Data Science and Python Machine LearningHands-on Machine Learning with Python and Scikit-LearnPython Machine Learning

Convex Optimization

Learn the skills necessary to design, build, and deploy applications powered by machine learning (ML). Through the course of this hands-on book, you'll build an example ML-driven application from initial idea to deployed product. Data scientists, software engineers, and product managers—including experienced practitioners and novices alike—will learn the tools, best practices, and challenges involved in building a real-world ML application step by step. Author Emmanuel Ameisen, an experienced data scientist who led an AI education program, demonstrates practical ML concepts using code snippets, illustrations, screenshots, and interviews with industry leaders. Part I teaches you how to plan an ML application and measure success. Part II explains how to build a working ML model. Part III demonstrates ways to improve the model until it fulfills your original vision. Part IV covers deployment and monitoring strategies. This book will help you: Define your product goal and set up a machine learning problem Build your first end-to-end pipeline quickly and acquire an initial dataset Train and evaluate your ML models and address performance bottlenecks Deploy and monitor your models in a production environment

Hands-On Machine Learning with Scikit-Learn

Aspiring data science professionals can learn the Scikit-Learn library along with the fundamentals of machine learning with this book. The book combines the Anaconda Python distribution with the popular Scikit-Learn library to demonstrate a wide range of supervised and unsupervised machine learning algorithms. Care is taken to walk you through the principles of machine learning through clear examples written in Python that you can try out and experiment with at home on your own machine. All applied math and programming skills required to master the content are covered in this book. In-depth knowledge of object-oriented programming is not required as working and complete examples are provided and explained. Coding examples are in-depth and complex when necessary. They are also concise, accurate, and complete, and complement the machine learning concepts introduced. Working the examples helps to build the skills necessary to understand and apply complex machine learning algorithms. Hands-on Scikit-Learn for Machine Learning Applications is an excellent starting point for those pursuing a career in machine learning. Students of this book will learn the fundamentals that are a prerequisite to competency. Readers will be exposed to the Anaconda distribution of Python that is designed specifically for data science professionals, and will build skills in the popular Scikit-Learn library that underlies many machine learning applications in the world of Python. What You'll Learn Work with simple and complex datasets common to Scikit-Learn Manipulate data into vectors and matrices for algorithmic processing Become familiar with the Anaconda distribution used in data science Apply machine learning with Classifiers, Regressors, and Dimensionality Reduction Tune algorithms and find the best algorithms for each dataset Load data from and save to CSV, JSON, Numpy, and Pandas formats Who This Book Is For The aspiring data scientist yearning to break into machine learning through mastering the underlying fundamentals that are sometimes skipped over in the rush to be productive. Some knowledge of object-oriented programming and very basic applied linear algebra will make learning easier, although anyone can benefit from this book.

Deep Learning for Coders with fastai and PyTorch

If you are a software developer who wants to learn how machine learning models work and how to apply them effectively, this book is for you. Familiarity with machine learning fundamentals and Python will be helpful, but is not essential.

Hands-On Machine Learning with C++

Hands-on Scikit-learn for Machine Learning

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.

Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow, 2nd Edition

Get hands-on with the browser-based JavaScript library for training and deploying machine learning models effectively Key Features Build, train and run machine learning models in the browser using TensorFlow.js Create smart web applications from scratch with the help of useful examples Use flexible and intuitive APIs from TensorFlow.js to understand how machine learning algorithms function Book Description TensorFlow.js is a framework that enables you to create performant machine learning (ML) applications that run smoothly in a web browser. With this book, you will learn how to use TensorFlow.js to implement various ML models through an example-based approach. Starting with the basics, you'll understand how ML models can be built on the web. Moving on, you will get to grips with the TensorFlow.js ecosystem to develop applications more efficiently. The book will then guide you through implementing ML techniques and algorithms such as regression, clustering, fast Fourier transform (FFT), and dimensionality reduction. You will later cover the Bellman equation to solve Markov decision process (MDP) problems and understand how it is related to reinforcement learning. Finally, you will explore techniques for deploying ML-based web applications and training models with TensorFlow Core. Throughout this ML book, you'll discover useful tips and tricks that will build on your knowledge. By the end of this book, you will be equipped with the skills you need to create your own web-based ML applications and fine-tune models to achieve high performance. What you will learn Use the t-SNE algorithm in TensorFlow.js to reduce dimensions in an input dataset Deploy tfjs-converter to convert Keras models and load them into TensorFlow.js Apply the Bellman equation to solve MDP problems Use the k-means algorithm in TensorFlow.js to visualize prediction results Create tf.js packages with Parcel, Webpack, and Rollup to deploy web apps Implement tf.js backend frameworks to tune and accelerate app performance Who this book is for This book is for web developers who want to learn how to integrate machine learning techniques with web-based applications from scratch. This book will also appeal to data scientists, machine learning practitioners, and deep learning enthusiasts who are looking to perform accelerated, browser-based machine learning on Web using TensorFlow.js. Working knowledge of JavaScript programming language is all you need to get started.

Hands-On Machine Learning with Scikit-Learn and TensorFlow

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 Features Become familiar with data processing, performance measuring, and model selection using various C++ libraries Implement practical machine learning and deep learning techniques to build smart models Deploy machine learning models to work on mobile and

embedded devices Book 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 learn Explore how to load and preprocess various data types to suitable C++ data structures Employ key machine learning algorithms with various C++ libraries Understand the grid-search approach to find the best parameters for a machine learning model Implement an algorithm for filtering anomalies in user data using Gaussian distribution Improve collaborative filtering to deal with dynamic user preferences Use C++ libraries and APIs to manage model structures and parameters Implement a C++ program to solve image classification tasks with LeNet architecture Who this book is for You will find this C++ machine learning book useful if 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 this book.

Introduction to Machine Learning with Python

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.

Machine Learning for Cybersecurity Cookbook

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 Automated Machine Learning

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.

Mastering Machine Learning with scikit-learn

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.

Hands-On Image Processing with Python

For many researchers, Python is a first-class tool mainly because of its libraries for storing, manipulating, and gaining insight from data. Several resources exist for individual pieces of this data science stack, but only with the Python Data Science Handbook do you get them all—IPython, NumPy, Pandas, Matplotlib, Scikit-Learn, and other related tools. Working scientists and data crunchers familiar with reading and writing Python code will find this comprehensive desk reference ideal for tackling day-to-day issues: manipulating, transforming, and cleaning data; visualizing different types of data; and using data to build statistical or machine learning models. Quite simply, this is the must-have reference for scientific computing in Python. With this handbook, you'll learn how to use: IPython and Jupyter: provide computational environments for data scientists using Python NumPy: includes the ndarray for efficient storage and manipulation of dense data arrays in Python Pandas: features the DataFrame for efficient storage and manipulation of labeled/columnar data in Python Matplotlib: includes capabilities for a flexible range of data visualizations in Python Scikit-Learn: for efficient and clean Python implementations of the most important and established machine learning algorithms

Building Machine Learning Powered Applications

Explore the mathematical computations and algorithms for image processing using popular Python tools and frameworks. Key Features Practical coverage of every image processing task with popular Python libraries Includes topics such as pseudo-coloring, noise smoothing, computing image descriptors Covers popular machine learning and deep learning techniques for complex image processing tasks Book Description Image processing plays an important role in our daily lives with various applications such as in social media (face detection), medical imaging (X-ray, CT-scan), security (fingerprint recognition) to robotics & space. This book will touch the core of image processing, from concepts to code using Python. The book will start from the classical image processing techniques and explore the evolution of image processing algorithms up to the recent advances in image processing or computer vision with deep learning. We will learn how to use

image processing libraries such as PIL, scikit-mage, and scipy ndimage in Python. This book will enable us to write code snippets in Python 3 and quickly implement complex image processing algorithms such as image enhancement, filtering, segmentation, object detection, and classification. We will be able to use machine learning models using the scikit-learn library and later explore deep CNN, such as VGG-19 with Keras, and we will also use an end-to-end deep learning model called YOLO for object detection. We will also cover a few advanced problems, such as image inpainting, gradient blending, variational denoising, seam carving, quilting, and morphing. By the end of this book, we will have learned to implement various algorithms for efficient image processing. What you will learn Perform basic data pre-processing tasks such as image denoising and spatial filtering in Python Implement Fast Fourier Transform (FFT) and Frequency domain filters (e.g., Weiner) in Python Do morphological image processing and segment images with different algorithms Learn techniques to extract features from images and match images Write Python code to implement supervised / unsupervised machine learning algorithms for image processing Use deep learning models for image classification, segmentation, object detection and style transfer Who this book is for This book is for Computer Vision Engineers, and machine learning developers who are good with Python programming and want to explore details and complexities of image processing. No prior knowledge of the image processing techniques is expected.

Applied Machine Learning

Hands-On Machine Learning with Scikit-Learn Book Description In this Book Hands-On Machine Learning with Scikit Learn. The author covered both Supervised and Unsupervised Machine Learning Algorithms. The authors explain both Theoretical and Practical Implementation in depth and Explain Each Algorithm from Scratch. For Practical Implementation uses the Scikit-learn Library in this book. Scikit-Learn is a robust machine learning library for the Python programming language. It provides a set of supervised and unsupervised learning algorithms. This book is the easiest way to learn how to deploy, optimize, and evaluate all of the important machine learning algorithms that scikit-learn provides. This book teaches you how to use scikit-learn for machine learning. You will start by setting up and configuring your machine learning environment with scikit-learn. To put scikit-learn to use, you will learn how to implement variously supervised and unsupervised machine learning models. You will learn classification, regression, Association Rule, clustering techniques and Dimensionality Reduction Techniques to work with different types of datasets and train your models. Key Features ● Learn Supervised & Unsupervised Machine Learning Algorithms in Depth. ●Build your first machine learning model using scikit-learn ●Train supervised and unsupervised models using popular techniques such as classification, regression, clustering and Dimensionality Reduction. ●Understand how scikit-learn can be applied to different types of machine learning problems What you will learn ●Perform classification and regression machine learning ●Employ Unsupervised Machine Learning Algorithms to cluster unlabeled data into groups ●Apply the Dimensionality Reduction Technique for reducing the Dimensionality of the dataset Who this book is for ●Anyone who interesting in Machine Learning. ●Fundamental knowledge

of linear algebra and probability will help. ●Intermediate knowledge of Python programming Table of Contents 1. Introduction to Machine Learning 2. Linear Regression 3. Naïve Bayes 4. Decision Tree (classification & Regression) 5. Random Forrest(classification & Regression) 6. K-Nearest Neighbors 7. Logistic Regression 8. Support Vector Machine 9. Association Rule (Apriori & Eclat) 10. Clustering (K-Mean & Hierarchical) 11. Dimensionality Reduction (PCA & LDA)

Hands-on Scikit-Learn for Machine Learning Applications

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.

A Modern Approach to Regression with R

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.

Deep Learning

Hands-On Machine Learning for Cybersecurity

Convex optimization problems arise frequently in many different fields. This book provides a comprehensive introduction to the subject, and shows in detail how such problems can be solved numerically with great efficiency. The book begins with the basic elements of convex sets and functions, and then describes various classes of convex optimization problems. Duality and approximation techniques are then covered, as are statistical estimation techniques. Various geometrical problems are then presented, and there is detailed discussion of unconstrained and constrained minimization problems,

and interior-point methods. The focus of the book is on recognizing convex optimization problems and then finding the most appropriate technique for solving them. It contains many worked examples and homework exercises and will appeal to students, researchers and practitioners in fields such as engineering, computer science, mathematics, statistics, finance and economics.

Learning Scikit-Learn

***** BUY NOW (will soon return to 24.77 $***** MONEY BACK GUARANTEE BY AMAZON (See Below FAQ) *****Are you thinking of learning more about Machine Learning using Python? (For Beginners)This book is for you. It would seek to explain you all need to know about machine learning and its application using Python in an intuitive way. From AI Sciences Publisher Our books may be the best one for beginners; it's a step-by-step guide for any person who wants to start learning Artificial Intelligence and Data Science from scratch. It will help you in preparing a solid foundation and learn any other high-level courses.To get the most out of the concepts that would be covered, readers are advised to adopt a hands on approach which would lead to better mental representations. Target UsersThe book designed for a variety of target audiences. The most suitable users would include: Anyone who is intrigued by how algorithms arrive at predictions but has no previous knowledge of the field. Software developers and engineers with a strong programming background but seeking to break into the field of machine learning. Seasoned professionals in the field of artificial intelligence and machine learning who desire a bird's eye view of current techniques and approaches. What's Inside This Book? Overview of Python Programming Language Statistics Probability The Data Science Process Machine Learning Supervised Learning Algorithms Unsupervised Learning Algorithms Semi-supervised Learning Algorithms Reinforcement Learning Algorithms Overfitting and Underfitting Python Data Science Tools Jupyter Notebook Numerical Python (Numpy) Pandas Scientific Python (Scipy) Matplotlib Scikit-Learn K-Nearest Neighbors Naive Bayes Simple and Multiple Linear Regression Logistic Regression Generalized Linear Models Decision Trees and Random Forest Neural Networks Perceptrons Backpropagation Clustering K-means with Scikit-Learn Bottom-up Hierarchical Clustering K-means Clustering Network Analysis Betweenness centrality Eigenvector Centrality Recommender Systems Multi-Class Classification Popular Classification Algorithms Support Vector Machine Deep Learning using TensorFlow Deep Learning Case Studies Frequently Asked Questions Q: Is this book for me and do I need programming experience?A: If you want to smash Machine Learning from scratch, this book is for you. If you already wrote a few lines of code and recognize basic programming statements, you'll be OK. Q: Does this book include everything I need to become a Machine Learning expert?A: Unfortunately, no. This book is designed for readers taking their first steps in Machine Learning and further learning will be required beyond this book to master all aspects of Machine Learning. Q: Can I have a refund if this book doesn't fit for me?A: Yes, Amazon refund you if you aren't satisfied, for more information about the amazon refund service please go to the amazon help platform. We will also be happy to help you if you send us an email (email address inside the book).***** MONEY BACK GUARANTEE BY AMAZON ***** Editorial Reviews"This book

succeeds in covering most important techniques in a clear, intuitive way that is perfect for newbies and those seeking to improve their practice in the Machine LearningFields VERY QUICKLY ." --Adrian B. Machine Learning Researcher Consulting AI company

Python for Data Analysis

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

Hands-On Machine Learning with Scikit-Learn and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems

Combine popular machine learning techniques to create ensemble models using Python Key Features Implement ensemble models using algorithms such as random forests and AdaBoost Apply boosting, bagging, and stacking ensemble methods to improve the prediction accuracy of your model Explore real-world data sets and practical examples coded in scikit-learn and Keras Book Description Ensembling is a technique of combining two or more similar or dissimilar machine learning algorithms to create a model that delivers superior predictive power. This book will demonstrate how you can use a variety of weak algorithms to make a strong predictive model. With its hands-on approach, you'll not only get up to speed on the basic theory but also the application of various ensemble learning techniques. Using examples and real-world datasets, you'll be able to produce better machine learning models to solve supervised learning problems such as classification and regression. Furthermore, you'll go on to leverage ensemble learning techniques such as clustering to produce unsupervised machine learning models. As you progress, the chapters will cover different machine learning algorithms that are widely used in the practical world to make predictions and classifications. You'll even get to grips with the use of Python libraries such as scikit-learn and Keras for implementing different ensemble models. By the end of this book, you will be well-versed

in ensemble learning, and have the skills you need to understand which ensemble method is required for which problem, and successfully implement them in real-world scenarios. What you will learn Implement ensemble methods to generate models with high accuracy Overcome challenges such as bias and variance Explore machine learning algorithms to evaluate model performance Understand how to construct, evaluate, and apply ensemble models Analyze tweets in real time using Twitter's streaming API Use Keras to build an ensemble of neural networks for the MovieLens dataset Who this book is for This book is for data analysts, data scientists, machine learning engineers and other professionals who are looking to generate advanced models using ensemble techniques. An understanding of Python code and basic knowledge of statistics is required to make the most out of this book.

Hands-On Machine Learning for Algorithmic Trading

This book focuses on tools and techniques for building regression models using real-world data and assessing their validity. A key theme throughout the book is that it makes sense to base inferences or conclusions only on valid models. Plots are shown to be an important tool for both building regression models and assessing their validity. We shall see that deciding what to plot and how each plot should be interpreted will be a major challenge. In order to overcome this challenge we shall need to understand the mathematical properties of the fitted regression models and associated diagnostic procedures. As such this will be an area of focus throughout the book. In particular, we shall carefully study the properties of resi- als in order to understand when patterns in residual plots provide direct information about model misspecification and when they do not. The regression output and plots that appear throughout the book have been gen- ated using R. The output from R that appears in this book has been edited in minor ways. On the book web site you will find the R code used in each example in the text.

Applied Deep Learning with Python

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 Unsupervised Learning Using Python

Get into the world of smart data security using machine learning algorithms and Python libraries Key Features Learn machine learning algorithms and cybersecurity fundamentals Automate your daily workflow by applying use cases to many facets of security Implement smart machine learning solutions to detect various cybersecurity problems Book 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 learn Use machine learning algorithms with complex datasets to implement cybersecurity concepts Implement machine learning algorithms such as clustering, k-means, and Naive Bayes to solve real-world problems Learn to speed up a system using Python libraries with NumPy, Scikit-learn, and CUDA Understand how to combat malware, detect spam, and fight financial fraud to mitigate cyber crimes Use TensorFlow in the cybersecurity domain and implement real-world examples Learn how machine learning and Python can be used in complex cyber issues Who 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

Python Data Science Handbook

Extract patterns and knowledge from your data in easy way using MATLAB About This Book Get your first steps into machine learning with the help of this easy-to-follow guide Learn regression, clustering, classification, predictive analytics, artificial neural networks and more with MATLAB Understand how your data works and identify hidden layers in the data with the power of machine learning. Who This Book Is For This book is for data analysts, data scientists, students, or anyone who is looking to get started with machine learning and want to build efficient data processing and predicting applications.

A mathematical and statistical background will really help in following this book well. What You Will Learn Learn the introductory concepts of machine learning. Discover different ways to transform data using SAS XPORT, import and export tools, Explore the different types of regression techniques such as simple & multiple linear regression, ordinary least squares estimation, correlations and how to apply them to your data. Discover the basics of classification methods and how to implement Naive Bayes algorithm and Decision Trees in the Matlab environment. Uncover how to use clustering methods like hierarchical clustering to grouping data using the similarity measures. Know how to perform data fitting, pattern recognition, and clustering analysis with the help of MATLAB Neural Network Toolbox. Learn feature selection and extraction for dimensionality reduction leading to improved performance. In Detail MATLAB is the language of choice for many researchers and mathematics experts for machine learning. This book will help you build a foundation in machine learning using MATLAB for beginners. You'll start by getting your system ready with t he MATLAB environment for machine learning and you'll see how to easily interact with the Matlab workspace. We'll then move on to data cleansing, mining and analyzing various data types in machine learning and you'll see how to display data values on a plot. Next, you'll get to know about the different types of regression techniques and how to apply them to your data using the MATLAB functions. You'll understand the basic concepts of neural networks and perform data fitting, pattern recognition, and clustering analysis. Finally, you'll explore feature selection and extraction techniques for dimensionality reduction for performance improvement. At the end of the book, you will learn to put it all together into real-world cases covering major machine learning algorithms and be comfortable in performing machine learning with MATLAB. Style and approach The book takes a very comprehensive approach to enhance your understanding of machine learning using MATLAB. Sufficient real-world examples and use cases are included in the book to help you grasp the concepts quickly and apply them easily in your day-to-day work.

Hands-On Machine Learning with scikit-learn and Scientific Python Toolkits

Introduction to Data Science: Data Analysis and Prediction Algorithms with R introduces concepts and skills that can help you tackle real-world data analysis challenges. It covers concepts from probability, statistical inference, linear regression, and machine learning. It also helps you develop skills such as R programming, data wrangling, data visualization, predictive algorithm building, file organization with UNIX/Linux shell, version control with Git and GitHub, and reproducible document preparation. This book is a textbook for a first course in data science. No previous knowledge of R is necessary, although some experience with programming may be helpful. The book is divided into six parts: R, data visualization, statistics with R, data wrangling, machine learning, and productivity tools. Each part has several chapters meant to be presented as one lecture. The author uses motivating case studies that realistically mimic a data scientist's experience. He starts by asking specific questions and answers these through data analysis so concepts are learned as a means to answering the questions. Examples of the case studies included are: US murder rates by state, self-reported student heights, trends in world health

and economics, the impact of vaccines on infectious disease rates, the financial crisis of 2007-2008, election forecasting, building a baseball team, image processing of hand-written digits, and movie recommendation systems. The statistical concepts used to answer the case study questions are only briefly introduced, so complementing with a probability and statistics textbook is highly recommended for in-depth understanding of these concepts. If you read and understand the chapters and complete the exercises, you will be prepared to learn the more advanced concepts and skills needed to become an expert.

Hands-On Deep Learning Algorithms with Python

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 Machine Learning with Scikit-Learn and TensorFlow

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 Features Delve into machine learning with this comprehensive guide to scikit-learn and scientific Python Master the art of data-driven problem-solving with hands-on examples Foster your theoretical and practical knowledge of supervised and unsupervised machine learning algorithms Book 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 learn Understand when to use supervised, unsupervised, or reinforcement learning algorithms Find out how to collect and prepare your data for machine learning tasks Tackle imbalanced data and optimize your algorithm for a bias or variance tradeoff Apply supervised and unsupervised algorithms to overcome various machine learning challenges Employ best practices for tuning your algorithm's hyper parameters Discover how to use neural networks for classification and regression Build, evaluate, and deploy your machine learning solutions to production Who 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 Machine Learning with R

"Machine learning and artificial intelligence are the new big data--at least as far as buzzwords in the workplace go. The scikit-learn library is one of the most popular platforms for everyday Machine Learning and data science because it is built upon Python, a fully featured programming language. This course will help you discover the magical black box that is Machine Learning by teaching a practical approach to modeling using Python along with the Scikit-Learn library. We begin our journey by observing the end result of a Machine Learning deployment before moving back to the fundamentals and into exploratory data analysis. Moving on, we learn to develop complex pipelines and techniques for building custom transformer objects for feature extraction, manipulation, and other effective data cleansing techniques. Finally, we discover

how to select a model, apply optimal hyper-parameters, and deploy it."--Resource description page.

Hands-On Ensemble Learning with Python

Many industry experts consider unsupervised learning the next frontier in artificial intelligence, one that may hold the key to general artificial intelligence. Since the majority of the world's data is unlabeled, conventional supervised learning cannot be applied. Unsupervised learning, on the other hand, can be applied to unlabeled datasets to discover meaningful patterns buried deep in the data, patterns that may be near impossible for humans to uncover. Author Ankur Patel shows you how to apply unsupervised learning using two simple, production-ready Python frameworks: Scikit-learn and TensorFlow using Keras. With code and hands-on examples, data scientists will identify difficult-to-find patterns in data and gain deeper business insight, detect anomalies, perform automatic feature engineering and selection, and generate synthetic datasets. All you need is programming and some machine learning experience to get started. Compare the strengths and weaknesses of the different machine learning approaches: supervised, unsupervised, and reinforcement learning Set up and manage machine learning projects end-to-end Build an anomaly detection system to catch credit card fraud Clusters users into distinct and homogeneous groups Perform semisupervised learning Develop movie recommender systems using restricted Boltzmann machines Generate synthetic images using generative adversarial networks

MATLAB for Machine Learning

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 Machine Learning with Python

Explore effective trading strategies in real-world markets using NumPy, spaCy, pandas, scikit-learn, and Keras Key Features Implement machine learning algorithms to build, train, and validate algorithmic models Create your own algorithmic design process to apply probabilistic machine learning approaches to trading decisions Develop neural networks for algorithmic trading to perform time series forecasting and smart analytics Book Description The explosive growth of digital data has boosted the demand for expertise in trading strategies that use machine learning (ML). This book enables you to use a broad range of supervised and unsupervised algorithms to extract signals from a wide variety of data sources and create powerful investment strategies. This book shows how to access market, fundamental, and alternative data via API or web scraping and offers a framework to evaluate alternative data. You'll practice the ML workflow from model design, loss metric definition, and parameter tuning to performance evaluation in a time series context. You will understand ML algorithms such as Bayesian and ensemble methods and manifold learning, and will know how to train and tune these models using pandas, statsmodels, sklearn, PyMC3, xgboost, lightgbm, and catboost. This book also teaches you how to extract features from text data using spaCy, classify news and assign sentiment scores, and to use gensim to model topics and learn word embeddings from financial reports. You will also build and evaluate neural networks, including RNNs and CNNs, using Keras and PyTorch to exploit unstructured data for sophisticated strategies. Finally, you will apply transfer learning to satellite images to predict economic activity and use reinforcement learning to build agents that learn to trade in the OpenAI Gym. What you will learn Implement machine learning techniques to solve investment and trading problems Leverage market, fundamental, and alternative data to research alpha factors Design and fine-tune supervised, unsupervised, and reinforcement learning models Optimize portfolio risk and performance using pandas, NumPy, and scikit-learn Integrate machine learning models into a live trading strategy on Quantopian Evaluate strategies using reliable

backtesting methodologies for time series Design and evaluate deep neural networks using Keras, PyTorch, and TensorFlow Work with reinforcement learning for trading strategies in the OpenAI Gym Who this book is for Hands-On Machine Learning for Algorithmic Trading is for data analysts, data scientists, and Python developers, as well as investment analysts and portfolio managers working within the finance and investment industry. If you want to perform efficient algorithmic trading by developing smart investigating strategies using machine learning algorithms, this is the book for you. Some understanding of Python and machine learning techniques is mandatory.

Introduction to Data Science

Learn how to apply modern AI to create powerful cybersecurity solutions for malware, pentesting, social engineering, data privacy, and intrusion detection Key Features Manage data of varying complexity to protect your system using the Python ecosystem Apply ML to pentesting, malware, data privacy, intrusion detection system(IDS) and social engineering Automate your daily workflow by addressing various security challenges using the recipes covered in the book Book Description Organizations today face a major threat in terms of cybersecurity, from malicious URLs to credential reuse, and having robust security systems can make all the difference. With this book, you'll learn how to use Python libraries such as TensorFlow and scikit-learn to implement the latest artificial intelligence (AI) techniques and handle challenges faced by cybersecurity researchers. You'll begin by exploring various machine learning (ML) techniques and tips for setting up a secure lab environment. Next, you'll implement key ML algorithms such as clustering, gradient boosting, random forest, and XGBoost. The book will guide you through constructing classifiers and features for malware, which you'll train and test on real samples. As you progress, you'll build self-learning, reliant systems to handle cybersecurity tasks such as identifying malicious URLs, spam email detection, intrusion detection, network protection, and tracking user and process behavior. Later, you'll apply generative adversarial networks (GANs) and autoencoders to advanced security tasks. Finally, you'll delve into secure and private AI to protect the privacy rights of consumers using your ML models. By the end of this book, you'll have the skills you need to tackle real-world problems faced in the cybersecurity domain using a recipe-based approach. What you will learn Learn how to build malware classifiers to detect suspicious activities Apply ML to generate custom malware to pentest your security Use ML algorithms with complex datasets to implement cybersecurity concepts Create neural networks to identify fake videos and images Secure your organization from one of the most popular threats – insider threats Defend against zero-day threats by constructing an anomaly detection system Detect web vulnerabilities effectively by combining Metasploit and ML Understand how to train a model without exposing the training data Who this book is for This book is for cybersecurity professionals and security researchers who are looking to implement the latest machine learning techniques to boost computer security, and gain insights into securing an organization using red and blue team ML. This recipe-based book will also be useful for data scientists and machine learning developers who want to experiment with smart techniques in the cybersecurity domain. Working knowledge of Python programming and familiarity

with cybersecurity fundamentals will help you get the most out of this book.

Hands-On Machine Learning with TensorFlow.js

This text covers all the fundamentals and presents basic theoretical concepts and a wide range of techniques (algorithms) applicable to challenges in our day-to-day lives. The book recognizes that most of the ideas behind machine learning are simple and straightforward. It provides a platform for hands-on experience through self-study machine learning projects. Datasets for some benchmark applications have been explained to encourage the use of algorithms covered in this book. This is a comprehensive text book on machine learning for undergraduates in computer science and all engineering degree programs. Post graduates and research scholars will find it a useful initial exposure to the subject, before they go for highly theoretical depth in the specific areas of their research. For engineers, scientists, business managers and other practitioners, the book will help build the foundations of machine learning.

Building Machine Learning Systems with Python - Second Edition

Presents case studies and instructions on how to solve data analysis problems using Python.

Hands-On Machine Learning with Scikit-Learn and TensorFlow

The book adopts a tutorial-based approach to introduce the user to Scikit-learn.If you are a programmer who wants to explore machine learning and data-based methods to build intelligent applications and enhance your programming skills, this the book for you. No previous experience with machine-learning algorithms is required.

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

Automate data and model pipelines for faster machine learning applications Key Features Build automated modules for different machine learning components Understand each component of a machine learning pipeline in depth Learn to use different open source AutoML and feature engineering platforms Book Description AutoML is designed to automate parts of Machine Learning. Readily available AutoML tools are making data science practitioners' work easy and are received well in the advanced analytics community. Automated Machine Learning covers the necessary foundation needed to create automated machine learning modules and helps you get up to speed with them in the most practical way possible. In this book, you'll learn how to automate different tasks in the machine learning pipeline such as data preprocessing, feature selection, model training, model optimization, and much more. In addition to this, it demonstrates how you can use the

available automation libraries, such as auto-sklearn and MLBox, and create and extend your own custom AutoML components for Machine Learning. By the end of this book, you will have a clearer understanding of the different aspects of automated Machine Learning, and you'll be able to incorporate automation tasks using practical datasets. You can leverage your learning from this book to implement Machine Learning in your projects and get a step closer to winning various machine learning competitions. What you will learn Understand the fundamentals of Automated Machine Learning systems Explore auto-sklearn and MLBox for AutoML tasks Automate your preprocessing methods along with feature transformation Enhance feature selection and generation using the Python stack Assemble individual components of ML into a complete AutoML framework Demystify hyperparameter tuning to optimize your ML models Dive into Machine Learning concepts such as neural networks and autoencoders Understand the information costs and trade-offs associated with AutoML Who this book is for If you're a budding data scientist, data analyst, or Machine Learning enthusiast and are new to the concept of automated machine learning, this book is ideal for you. You'll also find this book useful if you're an ML engineer or data professional interested in developing quick machine learning pipelines for your projects. Prior exposure to Python programming will help you get the best out of this book.

Hands-On Data Science and Python Machine Learning

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 Python and Scikit-Learn

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

Python Machine Learning

"Scikit-learn is arguably the most popular Python library for Machine Learning today. Thousands of Data Scientists and Machine Learning practitioners use it for day to day tasks throughout a Machine Learning project's life cycle. Due to its popularity and coverage of a wide variety of ML models and built-in utilities, jobs for Scikit-learn are in high demand, both in industry and academia. If you're an aspiring machine learning engineer ready to take real-world projects head-on, Hands-on Scikit-Learn for Machine Learning will walk you through the most commonly used models, libraries, and utilities offered by Scikit-learn. By the end of the course, you will have a set of ML problem-solving tools in the form of code modules and utility functions based on Scikit-learn in one place, instead of spread over several books and courses, which you can easily use on real-world projects and data sets."--Resource description page.

ROMANCE ACTION & ADVENTURE MYSTERY & THRILLER BIOGRAPHIES & HISTORY CHILDREN'S YOUNG ADULT FANTASY HISTORICAL FICTION HORROR LITERARY FICTION NON-FICTION SCIENCE FICTION