

Hands On Machine Learning With Scikit Learn And Tensorflow Concepts Tools And Techniques For Building Intelligent Systems

Recognizing the way ways to acquire this books **Hands On Machine Learning With Scikit Learn And Tensorflow Concepts Tools And Techniques For Building Intelligent Systems** is additionally useful. You have remained in right site to begin getting this info. get the Hands On Machine Learning With Scikit Learn And Tensorflow Concepts Tools And Techniques For Building Intelligent Systems partner that we meet the expense of here and check out the link.

You could purchase lead Hands On Machine Learning With Scikit Learn And Tensorflow Concepts Tools And Techniques For Building Intelligent Systems or get it as soon as feasible. You could quickly download this Hands On Machine Learning With Scikit Learn And Tensorflow Concepts Tools And Techniques For Building Intelligent Systems after getting deal. So, subsequent to you require the books swiftly, you can straight acquire it. Its as a result categorically easy and therefore fats, isnt it? You have to favor to in this way of being

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

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

unsupervised, and reinforcement learning Set up and manage machine learning projects end-to-end Build an anomaly detection system to catch credit card fraud 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

[Python Reinforcement Learning](#) - Sudharsan Ravichandiran 2019-04-18
Apply modern reinforcement learning and deep reinforcement learning methods using Python and its powerful libraries Key FeaturesYour entry point into the world of artificial intelligence using the power of PythonAn example-rich guide to master various RL and DRL algorithmsExplore the power of modern Python libraries to gain confidence in building self-trained applicationsBook Description Reinforcement Learning (RL) is the trending and most promising branch of artificial intelligence. This Learning Path will help you master not only the basic reinforcement learning algorithms but also the advanced deep reinforcement learning algorithms. The Learning Path starts with an introduction to RL followed

by OpenAI Gym, and TensorFlow. You will then explore various RL algorithms, such as Markov Decision Process, Monte Carlo methods, and dynamic programming, including value and policy iteration. You'll also work on various datasets including image, text, and video. This example-rich guide will introduce you to deep RL algorithms, such as Dueling DQN, DRQN, A3C, PPO, and TRPO. You will gain experience in several domains, including gaming, image processing, and physical simulations. You'll explore TensorFlow and OpenAI Gym to implement algorithms that also predict stock prices, generate natural language, and even build other neural networks. You will also learn about imagination-augmented agents, learning from human preference, DQfD, HER, and many of the recent advancements in RL. By the end of the Learning Path, you will have all the knowledge and experience needed to implement RL and deep RL in your projects, and you enter the world of artificial intelligence to solve various real-life problems. This Learning Path includes content from the following Packt products: Hands-On Reinforcement Learning with Python by Sudharsan Ravichandiran Python Reinforcement Learning Projects by Sean Saito, Yang Wenzhuo, and Rajalingappaa Shanmugamani What you will learn Train an agent to walk using OpenAI Gym and TensorFlow Solve multi-armed-bandit problems using various algorithms Build intelligent agents using the DRQN algorithm to play the Doom game Teach your agent to play Connect4 using AlphaGo Zero Defeat Atari arcade games using the value iteration method Discover how to deal with discrete and continuous action spaces in various environments Who this book is for If you're an ML/DL enthusiast interested in AI and want to explore RL and deep RL from scratch, this Learning Path is for you. Prior knowledge of linear algebra is expected. *Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow* - Aurélien Géron 2022-10-04 Through a recent series of breakthroughs, deep learning has boosted the entire field of machine learning. Now, even programmers who know close to nothing about this technology can use simple, efficient tools to implement programs capable of learning from data. This bestselling book uses concrete examples, minimal theory, and production-ready Python

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

Python Machine Learning By Example - Yuxi (Hayden) Liu
2020-10-30

Equipped with the latest updates, this third edition of Python Machine Learning By Example provides a comprehensive course for ML enthusiasts to strengthen their command of ML concepts, techniques, and algorithms.

[The Machine Learning Workshop](#) - Hyatt Saleh 2020-07-22

Take a comprehensive and step-by-step approach to understanding machine learning Key Features Discover how to apply the scikit-learn uniform API in all types of machine learning models Understand the difference between supervised and unsupervised learning models Reinforce your understanding of machine learning concepts by working on real-world examples Book Description Machine learning algorithms are an integral part of almost all modern applications. To make the learning process faster and more accurate, you need a tool flexible and powerful enough to help you build machine learning algorithms quickly and easily. With The Machine Learning Workshop, you'll master the scikit-learn library and become proficient in developing

clever machine learning algorithms. The Machine Learning Workshop begins by demonstrating how unsupervised and supervised learning algorithms work by analyzing a real-world dataset of wholesale customers. Once you've got to grips with the basics, you'll develop an artificial neural network using scikit-learn and then improve its performance by fine-tuning hyperparameters. Towards the end of the workshop, you'll study the dataset of a bank's marketing activities and build machine learning models that can list clients who are likely to subscribe to a term deposit. You'll also learn how to compare these models and select the optimal one. By the end of The Machine Learning Workshop, you'll not only have learned the difference between supervised and unsupervised models and their applications in the real world, but you'll also have developed the skills required to get started with programming your very own machine learning algorithms. What you will learn

Understand how to select an algorithm that best fits your dataset and desired outcome

Explore popular real-world algorithms such as K-means, Mean-Shift, and DBSCAN

Discover different approaches to solve machine learning classification problems

Develop neural network structures using the scikit-learn package

Use the NN algorithm to create models for predicting future outcomes

Perform error analysis to improve your model's performance

Who this book is for

The Machine Learning Workshop is perfect for machine learning beginners. You will need Python programming experience, though no prior knowledge of scikit-learn and machine learning is necessary.

Learning scikit-learn: Machine Learning in Python - Raul Garreta
2013-11-25

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 for Cybersecurity - Soma Halder
2018-12-31

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

Deep Learning - Ian Goodfellow 2016-11-10

An introduction to a broad range of topics in deep learning, covering

mathematical and conceptual background, deep learning techniques used in industry, and research perspectives. “Written by three experts in the field, Deep Learning is the only comprehensive book on the subject.” —Elon Musk, cochair of OpenAI; cofounder and CEO of Tesla and SpaceX

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

Hands-On Machine Learning for Algorithmic Trading - Stefan Jansen 2018-12-31

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.

Machine Learning with TensorFlow, Second Edition - Mattmann A. Chris
2021-02-02

Updated with new code, new projects, and new chapters, Machine Learning with TensorFlow, Second Edition gives readers a solid foundation in machine-learning concepts and the TensorFlow library. Summary Updated with new code, new projects, and new chapters, Machine Learning with TensorFlow, Second Edition gives readers a solid foundation in machine-learning concepts and the TensorFlow library. Written by NASA JPL Deputy CTO and Principal Data Scientist Chris Mattmann, all examples are accompanied by downloadable Jupyter Notebooks for a hands-on experience coding TensorFlow with Python. New and revised content expands coverage of core machine learning algorithms, and advancements in neural networks such as VGG-Face facial identification classifiers and deep speech classifiers. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Supercharge your data analysis with machine learning! ML algorithms automatically improve as they process data, so results get better over time. You don't have to be a mathematician to use ML: Tools like Google's TensorFlow library help with complex calculations so you can focus on getting the answers you need. About the book Machine Learning with TensorFlow, Second Edition is a fully revised guide to building machine learning models using Python and TensorFlow. You'll apply core ML concepts to real-world challenges, such as sentiment analysis, text classification, and image recognition. Hands-on examples illustrate neural network techniques for deep speech processing, facial identification, and auto-encoding with CIFAR-10. What's inside Machine Learning with TensorFlow Choosing the best ML approaches Visualizing algorithms with TensorBoard

Sharing results with collaborators Running models in Docker About the reader Requires intermediate Python skills and knowledge of general algebraic concepts like vectors and matrices. Examples use the super-stable 1.15.x branch of TensorFlow and TensorFlow 2.x. About the author Chris Mattmann is the Division Manager of the Artificial Intelligence, Analytics, and Innovation Organization at NASA Jet Propulsion Lab. The first edition of this book was written by Nishant Shukla with Kenneth Fricklas. Table of Contents PART 1 - YOUR MACHINE-LEARNING RIG 1 A machine-learning odyssey 2 TensorFlow essentials PART 2 - CORE LEARNING ALGORITHMS 3 Linear regression and beyond 4 Using regression for call-center volume prediction 5 A gentle introduction to classification 6 Sentiment classification: Large movie-review dataset 7 Automatically clustering data 8 Inferring user activity from Android accelerometer data 9 Hidden Markov models 10 Part-of-speech tagging and word-sense disambiguation PART 3 - THE NEURAL NETWORK PARADIGM 11 A peek into autoencoders 12 Applying autoencoders: The CIFAR-10 image dataset 13 Reinforcement learning 14 Convolutional neural networks 15 Building a real-world CNN: VGG-Face ad VGG-Face Lite 16 Recurrent neural networks 17 LSTMs and automatic speech recognition 18 Sequence-to-sequence models for chatbots 19 Utility landscape

Introduction to Machine Learning with Python - Andreas C. Müller
2016-09-26

Machine learning has become an integral part of many commercial applications and research projects, but this field is not exclusive to large companies with extensive research teams. If you use Python, even as a beginner, this book will teach you practical ways to build your own machine learning solutions. With all the data available today, machine learning applications are limited only by your imagination. You'll learn the steps necessary to create a successful machine-learning application with Python and the scikit-learn library. Authors Andreas Müller and Sarah Guido focus on the practical aspects of using machine learning algorithms, rather than the math behind them. Familiarity with the NumPy and matplotlib libraries will help you get even more from this

book. With this book, you'll learn: Fundamental concepts and applications of machine learning Advantages and shortcomings of widely used machine learning algorithms How to represent data processed by machine learning, including which data aspects to focus on Advanced methods for model evaluation and parameter tuning The concept of pipelines for chaining models and encapsulating your workflow Methods for working with text data, including text-specific processing techniques Suggestions for improving your machine learning and data science skills [Machine Learning](#) - Stephen Marsland 2011-03-23

Traditional books on machine learning can be divided into two groups—those aimed at advanced undergraduates or early postgraduates with reasonable mathematical knowledge and those that are primers on how to code algorithms. The field is ready for a text that not only demonstrates how to use the algorithms that make up machine learning methods, but

Grokking Deep Learning - Andrew W. Trask 2019-01-23

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

Python Data Science Handbook - Jake VanderPlas 2016-11-21

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

[The Self-Assembling Brain](#) - Peter Robin Hiesinger 2021-05-04

"In this book, Peter Robin Hiesinger explores historical and contemporary attempts to understand the information needed to make biological and artificial neural networks. Developmental neurobiologists and computer scientists with an interest in artificial intelligence - driven by the promise and resources of biomedical research on the one hand, and by the promise and advances of computer technology on the other - are trying to understand the fundamental principles that guide the generation of an intelligent system. Yet, though researchers in these disciplines share a common interest, their perspectives and approaches are often quite different. The book makes the case that "the information problem" underlies both fields, driving the questions that are driving forward the frontiers, and aims to encourage cross-disciplinary communication and understanding, to help both fields make progress. The questions that challenge researchers in these fields include the following. How does genetic information unfold during the years-long process of human brain development, and can this be a short-cut to create human-level artificial intelligence? Is the biological brain just messy hardware that can be improved upon by running learning algorithms in computers? Can artificial intelligence bypass evolutionary programming of "grown" networks? These questions are tightly linked, and answering them requires an understanding of how information unfolds algorithmically to generate functional neural networks. Via a

series of closely linked "discussions" (fictional dialogues between researchers in different disciplines) and pedagogical "seminars," the author explores the different challenges facing researchers working on neural networks, their different perspectives and approaches, as well as the common ground and understanding to be found amongst those sharing an interest in the development of biological brains and artificial intelligent systems"--

Thoughtful Machine Learning with Python - Matthew Kirk
2017-01-16

Gain the confidence you need to apply machine learning in your daily work. With this practical guide, author Matthew Kirk shows you how to integrate and test machine learning algorithms in your code, without the academic subtext. Featuring graphs and highlighted code examples throughout, the book features tests with Python's Numpy, Pandas, Scikit-Learn, and SciPy data science libraries. If you're a software engineer or business analyst interested in data science, this book will help you: Reference real-world examples to test each algorithm through engaging, hands-on exercises Apply test-driven development (TDD) to write and run tests before you start coding Explore techniques for improving your machine-learning models with data extraction and feature development Watch out for the risks of machine learning, such as underfitting or overfitting data Work with K-Nearest Neighbors, neural networks, clustering, and other algorithms

Mastering Predictive Analytics with scikit-learn and TensorFlow - Alvaro Fuentes 2018-09-29

Learn advanced techniques to improve the performance and quality of your predictive models Key Features Use ensemble methods to improve the performance of predictive analytics models Implement feature selection, dimensionality reduction, and cross-validation techniques Develop neural network models and master the basics of deep learning Book Description Python is a programming language that provides a wide range of features that can be used in the field of data science. Mastering Predictive Analytics with scikit-learn and TensorFlow covers various implementations of ensemble methods, how they are used

with real-world datasets, and how they improve prediction accuracy in classification and regression problems. This book starts with ensemble methods and their features. You will see that scikit-learn provides tools for choosing hyperparameters for models. As you make your way through the book, you will cover the nitty-gritty of predictive analytics and explore its features and characteristics. You will also be introduced to artificial neural networks and TensorFlow, and how it is used to create neural networks. In the final chapter, you will explore factors such as computational power, along with improvement methods and software enhancements for efficient predictive analytics. By the end of this book, you will be well-versed in using deep neural networks to solve common problems in big data analysis. What you will learn

- Use ensemble algorithms to obtain accurate predictions
- Apply dimensionality reduction techniques to combine features and build better models
- Choose the optimal hyperparameters using cross-validation
- Implement different techniques to solve current challenges in the predictive analytics domain
- Understand various elements of deep neural network (DNN) models
- Implement neural networks to solve both classification and regression problems

Who this book is for Mastering Predictive Analytics with scikit-learn and TensorFlow is for data analysts, software engineers, and machine learning developers who are interested in implementing advanced predictive analytics using Python. Business intelligence experts will also find this book indispensable as it will teach them how to progress from basic predictive models to building advanced models and producing more accurate predictions. Prior knowledge of Python and familiarity with predictive analytics concepts are assumed.

Deep Learning for Coders with fastai and PyTorch - Jeremy Howard
2020-06-29

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

Python Machine Learning - Sebastian Raschka 2015-09-23

Unlock deeper insights into Machine Learning with this vital guide to cutting-edge predictive analytics About This Book Leverage Python's most powerful open-source libraries for deep learning, data wrangling, and data visualization Learn effective strategies and best practices to improve and optimize machine learning systems and algorithms Ask - and answer - tough questions of your data with robust statistical models, built for a range of datasets Who This Book Is For If you want to find out how to use Python to start answering critical questions of your data, pick up Python Machine Learning - whether you want to get started from scratch or want to extend your data science knowledge, this is an essential and unmissable resource. What You Will Learn Explore how to use different machine learning models to ask different questions of your data Learn how to build neural networks using Keras and Theano Find out how to write clean and elegant Python code that will optimize the strength of your algorithms Discover how to embed your machine learning model in a web application for increased accessibility Predict continuous target outcomes using regression analysis Uncover hidden patterns and structures in data with clustering Organize data using effective pre-processing techniques Get to grips with sentiment analysis to delve deeper into textual and social media data In Detail Machine learning and predictive analytics are transforming the way businesses and other organizations operate. Being able to understand trends and

patterns in complex data is critical to success, becoming one of the key strategies for unlocking growth in a challenging contemporary marketplace. Python can help you deliver key insights into your data – its unique capabilities as a language let you build sophisticated algorithms and statistical models that can reveal new perspectives and answer key questions that are vital for success. Python Machine Learning gives you access to the world of predictive analytics and demonstrates why Python is one of the world's leading data science languages. If you want to ask better questions of data, or need to improve and extend the capabilities of your machine learning systems, this practical data science book is invaluable. Covering a wide range of powerful Python libraries, including scikit-learn, Theano, and Keras, and featuring guidance and tips on everything from sentiment analysis to neural networks, you'll soon be able to answer some of the most important questions facing you and your organization. Style and approach Python Machine Learning connects the fundamental theoretical principles behind machine learning to their practical application in a way that focuses you on asking and answering the right questions. It walks you through the key elements of Python and its powerful machine learning libraries, while demonstrating how to get to grips with a range of statistical models.

Analytical Skills for AI and Data Science - Daniel Vaughan
2020-05-21

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

complete with examples you can apply when working to solve your own issues. Break business decisions into stages that can be tackled using different skills from the analytical toolbox Identify and embrace uncertainty in decision making and protect against common human biases Customize optimal decisions to different customers using predictive and prescriptive methods and technologies Ask business questions that create high value through AI- and data-driven technologies

Machine Learning avec Scikit-Learn - Aurélien Géron 2019-11-20

Hands-On Ensemble Learning with Python - George Kyriakides
2019-07-19

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.

Mathematics for Machine Learning - Marc Peter Deisenroth 2020-04-23

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.

Deep Learning with Applications Using Python - Navin Kumar Manaswi 2018-04-04

Explore deep learning applications, such as computer vision, speech recognition, and chatbots, using frameworks such as TensorFlow and

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

AI and Machine Learning for Coders - Laurence Moroney 2020-10-01

If you're looking to make a career move from programmer to AI specialist, this is the ideal place to start. Based on Laurence Moroney's extremely successful AI courses, this introductory book provides a hands-on, code-first approach to help you build confidence while you learn key topics. You'll understand how to implement the most common scenarios in machine learning, such as computer vision, natural language processing (NLP), and sequence modeling for web, mobile, cloud, and embedded runtimes. Most books on machine learning begin with a daunting amount of advanced math. This guide is built on practical lessons that let you work directly with the code. You'll learn: How to build models with TensorFlow using skills that employers desire
The basics of machine learning by working with code samples
How to implement computer vision, including feature detection in images
How to use NLP to tokenize and sequence words and sentences
Methods for embedding models in Android and iOS
How to serve models over the web and in the cloud with TensorFlow
Serving

Machine Learning with Pytorch and Scikit-Learn - Sebastian Raschka 2022-02-25

PyTorch book of the bestselling and widely acclaimed Python Machine Learning series expanded to include transformers, XGBoost, and graph neural networks

Key Features: Learn applied machine learning with a solid foundation in theory Clear, intuitive explanations take you deep into the theory and practice of Python machine learning Fully updated and expanded to cover PyTorch, transformers, XGBoost, graph neural networks, and best practices

Book Description: Machine Learning with PyTorch and Scikit-Learn is a comprehensive guide to machine learning and deep learning with PyTorch. It acts as both a step-by-step tutorial, and a reference you'll keep coming back to as you build your machine learning systems. Packed with clear explanations, visualizations, and examples, this book covers all the essential machine learning techniques in depth. While some books teach you only to follow instructions, with this machine learning book, we teach you the principles to build models and applications for yourself. Updated to cover deep learning using PyTorch, this book also introduces readers to the latest additions to scikit-learn. Moreover, this book covers various machine learning and deep learning techniques for text and image classification. You will also learn about generative adversarial networks (GANs) for generating new data and training intelligent agents with reinforcement learning. Finally, this new edition is also expanded to cover the latest trends in deep learning, including introductions to graph neural networks and large-scale transformers used for natural language processing (NLP). This PyTorch book is your companion to machine learning with Python, whether you're a Python developer new to machine learning or want to deepen your knowledge of the latest developments.

What You Will Learn: Explore frameworks, models, and techniques for machines to 'learn' from data Use scikit-learn for machine learning and PyTorch for deep learning Train machine learning classifiers on images, text, and more Build and train neural networks, transformers, and graph neural networks Discover best practices for evaluating and tuning models Predict continuous target outcomes using regression analysis Dig deeper into textual and social media data using sentiment analysis

Who this book is for: If you know some Python and you want to use machine learning and deep

learning, pick up this book. Whether you want to start from scratch or extend your machine learning knowledge, this is an essential resource. Written for developers and data scientists who want to create practical machine learning with Python and PyTorch deep learning code. This Python book is ideal for anyone who wants to teach computers how to learn from data. Working knowledge of the Python programming language, along with a good understanding of calculus and linear algebra is a must.

Hands-On Machine Learning with scikit-learn and Scientific Python Toolkits - Tarek Amr 2020-07-24

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.

Machine Learning with scikit-learn Quick Start Guide - Kevin Jolly
2018-10-30

Deploy supervised and unsupervised machine learning algorithms using scikit-learn to perform classification, regression, and clustering. Key Features Build your first machine learning model using scikit-learn Train supervised and unsupervised models using popular techniques such as classification, regression and clustering Understand how scikit-learn can be applied to different types of machine learning problems Book Description 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 various supervised

and unsupervised machine learning models. You will learn classification, regression, and clustering techniques to work with different types of datasets and train your models. Finally, you will learn about an effective pipeline to help you build a machine learning project from scratch. By the end of this book, you will be confident in building your own machine learning models for accurate predictions. What you will learn Learn how to work with all scikit-learn's machine learning algorithms Install and set up scikit-learn to build your first machine learning model Employ Unsupervised Machine Learning Algorithms to cluster unlabelled data into groups Perform classification and regression machine learning Use an effective pipeline to build a machine learning project from scratch Who this book is for This book is for aspiring machine learning developers who want to get started with scikit-learn. Intermediate knowledge of Python programming and some fundamental knowledge of linear algebra and probability will help.

Mastering Machine Learning with scikit-learn - Gavin Hackeling
2017-07-24

Use scikit-learn to apply machine learning to real-world problems About This Book Master popular machine learning models including k-nearest neighbors, random forests, logistic regression, k-means, naive Bayes, and artificial neural networks Learn how to build and evaluate performance of efficient models using scikit-learn Practical guide to master your basics and learn from real life applications of machine learning Who This Book Is For This book is intended for software engineers who want to understand how common machine learning algorithms work and develop an intuition for how to use them, and for data scientists who want to learn about the scikit-learn API. Familiarity with machine learning fundamentals and Python are helpful, but not required. What You Will Learn Review fundamental concepts such as bias and variance Extract features from categorical variables, text, and images Predict the values of continuous variables using linear regression and K Nearest Neighbors Classify documents and images using logistic regression and support vector machines Create ensembles of estimators using bagging and boosting techniques Discover hidden structures in data using K-Means

clustering Evaluate the performance of machine learning systems in common tasks In Detail Machine learning is the buzzword bringing computer science and statistics together to build smart and efficient models. Using powerful algorithms and techniques offered by machine learning you can automate any analytical model. This book examines a variety of machine learning models including popular machine learning algorithms such as k-nearest neighbors, logistic regression, naive Bayes, k-means, decision trees, and artificial neural networks. It discusses data preprocessing, hyperparameter optimization, and ensemble methods. You will build systems that classify documents, recognize images, detect ads, and more. You will learn to use scikit-learn's API to extract features from categorical variables, text and images; evaluate model performance, and develop an intuition for how to improve your model's performance. By the end of this book, you will master all required concepts of scikit-learn to build efficient models at work to carry out advanced tasks with the practical approach. Style and approach This book is motivated by the belief that you do not understand something until you can describe it simply. Work through toy problems to develop your understanding of the learning algorithms and models, then apply your learnings to real-life problems.

Python Machine Learning for Beginners - Ai Publishing 2020-10-23
Python Machine Learning for Beginners Machine Learning (ML) and Artificial Intelligence (AI) are here to stay. Yes, that's right. Based on a significant amount of data and evidence, it's obvious that ML and AI are here to stay. Consider any industry today. The practical applications of ML are really driving business results. Whether it's healthcare, e-commerce, government, transportation, social media sites, financial services, manufacturing, oil and gas, marketing and sales You name it. The list goes on. There's no doubt that ML is going to play a decisive role in every domain in the future. But what does a Machine Learning professional do? A Machine Learning specialist develops intelligent algorithms that learn from data and also adapt to the data quickly. Then, these high-end algorithms make accurate predictions. *Python Machine Learning for Beginners* presents you with a hands-on approach to learn

ML fast. How Is This Book Different? AI Publishing strongly believes in learning by doing methodology. With this in mind, we have crafted this book with care. You will find that the emphasis on the theoretical aspects of machine learning is equal to the emphasis on the practical aspects of the subject matter. You'll learn about data analysis and visualization in great detail in the first half of the book. Then, in the second half, you'll learn about machine learning and statistical models for data science. Each chapter presents you with the theoretical framework behind the different data science and machine learning techniques, and practical examples illustrate the working of these techniques. When you buy this book, your learning journey becomes so much easier. The reason is you get instant access to all the related learning material presented with this book--references, PDFs, Python codes, and exercises--on the publisher's website. All this material is available to you at no extra cost. You can download the ML datasets used in this book at runtime, or you can access them via the Resources/Datasets folder. You'll also find the short course on Python programming in the second chapter immensely useful, especially if you are new to Python. Since this book gives you access to all the Python codes and datasets, you only need access to a computer with the internet to get started. The topics covered include: Introduction and Environment Setup Python Crash Course Python NumPy Library for Data Analysis Introduction to Pandas Library for Data Analysis Data Visualization via Matplotlib, Seaborn, and Pandas Libraries Solving Regression Problems in ML Using Sklearn Library Solving Classification Problems in ML Using Sklearn Library Data Clustering with ML Using Sklearn Library Deep Learning with Python TensorFlow 2.0 Dimensionality Reduction with PCA and LDA Using Sklearn Click the BUY NOW button to start your Machine Learning journey.

scikit-learn : Machine Learning Simplified - Raul Garreta 2017-11-10
Implement scikit-learn into every step of the data science pipeline About This Book Use Python and scikit-learn to create intelligent applications Discover how to apply algorithms in a variety of situations to tackle common and not-so common challenges in the machine learning domain A practical, example-based guide to help you gain expertise in

implementing and evaluating machine learning systems using scikit-learn
Who This Book Is For If you are a programmer and want to explore machine learning and data-based methods to build intelligent applications and enhance your programming skills, this is the course for you. No previous experience with machine-learning algorithms is required. What You Will Learn Review fundamental concepts including supervised and unsupervised experiences, common tasks, and performance metrics Classify objects (from documents to human faces and flower species) based on some of their features, using a variety of methods from Support Vector Machines to Naive Bayes Use Decision Trees to explain the main causes of certain phenomena such as passenger survival on the Titanic Evaluate the performance of machine learning systems in common tasks Master algorithms of various levels of complexity and learn how to analyze data at the same time Learn just enough math to think about the connections between various algorithms Customize machine learning algorithms to fit your problem, and learn how to modify them when the situation calls for it Incorporate other packages from the Python ecosystem to munge and visualize your dataset Improve the way you build your models using parallelization techniques In Detail Machine learning, the art of creating applications that learn from experience and data, has been around for many years. Python is quickly becoming the go-to language for analysts and data scientists due to its simplicity and flexibility; moreover, within the Python data space, scikit-learn is the unequivocal choice for machine learning. The course combines an introduction to some of the main concepts and methods in machine learning with practical, hands-on examples of real-world problems. The course starts by walking through different methods to prepare your data—be it a dataset with missing values or text columns that require the categories to be turned into indicator variables. After the data is ready, you'll learn different techniques aligned with different objectives—be it a dataset with known outcomes such as sales by state, or more complicated problems such as clustering similar customers. Finally, you'll learn how to polish your algorithm to ensure that it's both accurate and resilient to new datasets. You will learn to incorporate

machine learning in your applications. Ranging from handwritten digit recognition to document classification, examples are solved step-by-step using scikit-learn and Python. By the end of this course you will have learned how to build applications that learn from experience, by applying the main concepts and techniques of machine learning. Style and Approach Implement scikit-learn using engaging examples and fun exercises, and with a gentle and friendly but comprehensive "learn-by-doing" approach. This is a practical course, which analyzes compelling data about life, health, and death with the help of tutorials. It offers you a useful way of interpreting the data that's specific to this course, but that can also be applied to any other data. This course is designed to be both a guide and a reference for moving beyond the basics of scikit-learn.

Hands-On Machine Learning with R - Brad Boehmke 2019-11-07
Hands-on Machine Learning with R provides a practical and applied approach to learning and developing intuition into today's most popular machine learning methods. This book serves as a practitioner's guide to the machine learning process and is meant to help the reader learn to apply the machine learning stack within R, which includes using various R packages such as glmnet, h2o, ranger, xgboost, keras, and others to effectively model and gain insight from their data. The book favors a hands-on approach, providing an intuitive understanding of machine learning concepts through concrete examples and just a little bit of theory. Throughout this book, the reader will be exposed to the entire machine learning process including feature engineering, resampling, hyperparameter tuning, model evaluation, and interpretation. The reader will be exposed to powerful algorithms such as regularized regression, random forests, gradient boosting machines, deep learning, generalized low rank models, and more! By favoring a hands-on approach and using real word data, the reader will gain an intuitive understanding of the architectures and engines that drive these algorithms and packages, understand when and how to tune the various hyperparameters, and be able to interpret model results. By the end of this book, the reader should have a firm grasp of R's machine learning stack and be able to implement a systematic approach for producing high quality modeling

results. Features:

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

Python Scikit-Learn for Beginners - Ai Publishing 2021-03-28

Python for Data Scientists -- Scikit-Learn Specialization Scikit-Learn, also known as Sklearn, is a free, open-source machine learning (ML) library used for the Python language. In February 2010, this library was first made public. And in less than three years, it became one of the most popular machine learning libraries on Github. Scikit-learn is the best place to start for access to easy-to-use, top-notch implementations of popular algorithms. This library speeds up the development of ML models. The main features of the Scikit-learn library are regression, classification, and clustering algorithms (random forests, K-means, gradient boosting, DBSCAN, AND support vector machines). The Scikit-learn library also integrates well with other Python libraries, such as NumPy, Pandas, IPython, SciPy, Sympy, and Matplotlib, to fulfill different tasks. Python for Data Scientists: Scikit-Learn Specialization presents you with a hands-on, simple approach to learn Scikit-learn fast. How Is This Book Different? Most Python books assume you know how to code using Pandas, NumPy, and Matplotlib. But this book does not. The author spends a lot of time teaching you how actually write the simplest codes in Python to achieve machine learning models. In-depth coverage of the Scikit-learn library starts from the third chapter itself. Jumping straight to Scikit-learn makes it easy for you to follow along. The other advantage is Jupyter Notebook is used to write and explain the code right through this book. You can access the datasets used in this book easily by downloading them at runtime. You can also access them through the Datasets folder in the SharePoint and GitHub repositories. You also get to work on three hands-on mini-projects: Spam Email Detection with Scikit-Learn IMDB Movies Sentimental Analysis Image Classification with Scikit-Learn The scripts, graphs, and images in the book are clear and provide easy-to-understand visuals to the text description. If you're new to data science, you will find this book a great option for self-study.

Overall, you can count on this learning by doing book to help you accomplish your data science career goals faster. The topics covered include: Introduction to Scikit-Learn and Other Machine Learning Libraries Environment Setup and Python Crash Course Data Preprocessing with Scikit-Learn Feature Selection with Python Scikit-Learn Library Solving Regression Problems in Machine Learning Using Sklearn Library Solving Classification Problems in Machine Learning Using Sklearn Library Clustering Data with Scikit-Learn Library Dimensionality Reduction with PCA and LDA Using Sklearn Selecting Best Models with Scikit-Learn Natural Language Processing with Scikit-Learn Image Classification with Scikit-Learn Hit the BUY NOW button and start your Data Science Learning journey.

Machine Learning in Python: Hands on Machine Learning with Python Tools, Concepts and Techniques - Bob Mather 2019-07-15

The ability to crunch data effectively can propel your career or business to great heights. Machine Learning is the most effective data analysis tool. While it is a complex topic, it can be broken down into simpler steps, as show in this book. We are using Python, which is a great programming language for beginners.

Tensorflow Machine Learning - Benjamin Smith 2020-08-24

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

- This book includes a comprehensive and detailed explanation of the concepts. No chapter has idle talk. Every line in this book has been written while keeping the convenience and interest of the reader in mind.
- This book features some really cool tips and tricks that build upon some very basic and fundamental practices of machine learning. Using these tips and tricks will help increase the productivity of your models.
- Each topic addresses some of the most important issues that users experience when

working with machine learning. For instance, in the later parts of this book, after discussing deep learning, we shift our focus towards the main challenges that arise when creating and implementing a complex and large deep neural network. This book aims to give readers a productive reading session. In order to accomplish this, each chapter has fragmented sections that highlight interesting topics. Furthermore, the chapter layout guides the reader through the many concepts of machine learning very easily. If you're interested in tips and tricks to machine learning with the use of scikit-learn, keras and Tensorflow, then click the BUY NOW button to get started today!

Hands-On Automated Machine Learning - Sibanja Das 2018-04-26

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 Unsupervised Learning with Python - Giuseppe Bonaccorso 2019-02-28

Discover the skill-sets required to implement various approaches to Machine Learning with Python
Key Features
Explore unsupervised learning with clustering, autoencoders, restricted Boltzmann machines, and more
Build your own neural network models using modern Python libraries
Practical examples show you how to implement different machine learning and deep learning techniques
Book Description
Unsupervised learning is about making use of raw, untagged data and applying learning algorithms to it to help a machine predict its outcome. With this book, you will explore the concept of unsupervised learning to cluster large sets of data and analyze them repeatedly until the desired outcome is found using Python. This book starts with the key differences between supervised, unsupervised, and semi-supervised learning. You will be introduced to the best-used libraries and frameworks from the Python ecosystem and address unsupervised learning in both the machine learning and deep learning domains. You will explore various algorithms, techniques that are used to implement unsupervised learning in real-world use cases. You will learn a variety of unsupervised learning approaches, including randomized optimization, clustering, feature selection and transformation, and information theory. You will get hands-

on experience with how neural networks can be employed in unsupervised scenarios. You will also explore the steps involved in building and training a GAN in order to process images. By the end of this book, you will have learned the art of unsupervised learning for different real-world challenges. What you will learn Use cluster algorithms to identify and optimize natural groups of data Explore advanced non-linear and hierarchical clustering in action Soft label assignments for fuzzy c-means and Gaussian mixture models Detect anomalies through density estimation Perform principal component analysis using neural network models Create unsupervised models using GANs Who this book is for This book is intended for statisticians, data scientists, machine learning developers, and deep learning practitioners who want to build smart applications by implementing key building block unsupervised learning, and master all the new techniques and algorithms offered in machine learning and deep learning using real-world examples. Some prior knowledge of machine learning concepts and statistics is desirable.

Hands-On Gradient Boosting with XGBoost and scikit-learn - Corey Wade
2020-10-16

Get to grips with building robust XGBoost models using Python and scikit-learn for deployment Key Features Get up and running with machine learning and understand how to boost models with XGBoost in no time Build real-world machine learning pipelines and fine-tune hyperparameters to achieve optimal results Discover tips and tricks and gain innovative insights from XGBoost Kaggle winners Book Description XGBoost is an industry-proven, open-source software library that provides a gradient boosting framework for scaling billions of data points quickly and efficiently. The book introduces machine learning and XGBoost in scikit-learn before building up to the theory behind gradient boosting. You'll cover decision trees and analyze bagging in the machine learning context, learning hyperparameters that extend to XGBoost along the way. You'll build gradient boosting models from scratch and extend gradient boosting to big data while recognizing speed limitations using timers. Details in XGBoost are explored with a focus on speed

enhancements and deriving parameters mathematically. With the help of detailed case studies, you'll practice building and fine-tuning XGBoost classifiers and regressors using scikit-learn and the original Python API. You'll leverage XGBoost hyperparameters to improve scores, correct missing values, scale imbalanced datasets, and fine-tune alternative base learners. Finally, you'll apply advanced XGBoost techniques like building non-correlated ensembles, stacking models, and preparing models for industry deployment using sparse matrices, customized transformers, and pipelines. By the end of the book, you'll be able to build high-performing machine learning models using XGBoost with minimal errors and maximum speed. What you will learn Build gradient boosting models from scratch Develop XGBoost regressors and classifiers with accuracy and speed Analyze variance and bias in terms of fine-tuning XGBoost hyperparameters Automatically correct missing values and scale imbalanced data Apply alternative base learners like dart, linear models, and XGBoost random forests Customize transformers and pipelines to deploy XGBoost models Build non-correlated ensembles and stack XGBoost models to increase accuracy Who this book is for This book is for data science professionals and enthusiasts, data analysts, and developers who want to build fast and accurate machine learning models that scale with big data. Proficiency in Python, along with a basic understanding of linear algebra, will help you to get the most out of this book.

Hands-on Scikit-Learn for Machine Learning Applications - David Paper
2019-11-16

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.

Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow - Aurélien Géron 2019-09-05

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