

Machine Learning With Python Cookbook

Yeah, reviewing a ebook **Machine Learning With Python Cookbook** could add your near connections listings. This is just one of the solutions for you to be successful. As understood, success does not suggest that you have astounding points.

Comprehending as capably as contract even more than supplementary will allow each success. neighboring to, the proclamation as with ease as insight of this Machine Learning With Python Cookbook can be taken as without difficulty as picked to act.

Ensemble Machine Learning Cookbook - Dipayan Sarkar 2019-01-31
Implement machine learning algorithms to build ensemble models using Keras, H2O, Scikit-Learn, Pandas and more Key FeaturesApply popular machine learning algorithms using a recipe-based approachImplement boosting, bagging, and stacking ensemble methods to improve machine learning modelsDiscover real-world ensemble applications and encounter complex challenges in Kaggle competitionsBook Description Ensemble modeling is an approach used to improve the performance of machine learning models. It combines two or more similar or dissimilar machine learning algorithms to deliver superior intellectual powers. This book will help you to implement popular machine learning algorithms to cover different paradigms of ensemble machine learning such as boosting, bagging, and stacking. The Ensemble Machine Learning Cookbook will start by getting you acquainted with the basics of ensemble techniques and exploratory data analysis. You'll then learn to implement tasks related to statistical and machine learning algorithms to understand the ensemble of multiple heterogeneous algorithms. It will also ensure that you don't miss out on key topics, such as like resampling methods. As you progress, you'll get a better understanding of bagging, boosting, stacking, and working with the Random Forest algorithm using real-world examples. The book will highlight how these ensemble methods use multiple models to improve machine learning results, as compared to a single model. In the concluding chapters, you'll delve into advanced ensemble models using neural networks, natural language processing, and more. You'll also be able to implement models such as fraud detection, text categorization, and sentiment analysis. By the end of this book, you'll be able to harness ensemble techniques and the working mechanisms of machine learning algorithms to build intelligent models using individual recipes. What you will learnUnderstand how to use machine learning algorithms for regression and classification problemsImplement ensemble techniques such as averaging, weighted averaging, and max-votingGet to grips with advanced ensemble methods, such as bootstrapping, bagging, and stackingUse Random Forest for tasks such as classification and regressionImplement an ensemble of homogeneous and heterogeneous machine learning algorithmsLearn and implement various boosting techniques, such as AdaBoost, Gradient Boosting Machine, and XGBoostWho this book is for This book is designed for data scientists, machine learning developers, and deep learning enthusiasts who want to delve into machine learning algorithms to build powerful ensemble models. Working knowledge of Python programming and basic statistics is a must to help you grasp the concepts in the book.

Python Machine Learning Cookbook - Chris Albon 2018-04-06
This practical guide provides nearly 200 self-contained recipes to help you solve machine learning challenges you may encounter in your daily work. If you're comfortable with Python and its libraries, including pandas and scikit-learn, you'll be able to address specific problems such as loading data, handling text or numerical data, model selection, and dimensionality reduction and many other topics. Each recipe includes code that you can copy and paste into a toy dataset to ensure that it actually works. From there, you can insert, combine, or adapt the code to help construct your application. Recipes also include a discussion that explains the solution and provides meaningful context. This cookbook takes you beyond theory and concepts by providing the nuts and bolts you need to construct working machine learning applications. You'll find recipes for: Vectors, matrices, and arrays Handling numerical and categorical data, text, images, and dates and times Dimensionality reduction using feature extraction or feature selection Model evaluation and selection Linear and logical regression, trees and forests, and k-nearest neighbors Support vector machines (SVM), naïve Bayes, clustering, and neural networks Saving and loading trained models
Machine Learning for Cybersecurity Cookbook - Emmanuel Tsukerman 2019-11-25
Learn how to apply modern AI to create powerful cybersecurity solutions

for malware, pentesting, social engineering, data privacy, and intrusion detection Key FeaturesManage data of varying complexity to protect your system using the Python ecosystemApply ML to pentesting, malware, data privacy, intrusion detection system(IDS) and social engineeringAutomate your daily workflow by addressing various security challenges using the recipes covered in the bookBook 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 learnLearn how to build malware classifiers to detect suspicious activitiesApply ML to generate custom malware to pentest your securityUse ML algorithms with complex datasets to implement cybersecurity conceptsCreate neural networks to identify fake videos and imagesSecure your organization from one of the most popular threats - insider threatsDefend against zero-day threats by constructing an anomaly detection systemDetect web vulnerabilities effectively by combining Metasploit and MLUnderstand how to train a model without exposing the training dataWho 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.

Machine Learning with Python Cookbook - Chris Albon 2018-03-09
This practical guide provides nearly 200 self-contained recipes to help you solve machine learning challenges you may encounter in your daily work. If you're comfortable with Python and its libraries, including pandas and scikit-learn, you'll be able to address specific problems such as loading data, handling text or numerical data, model selection, and dimensionality reduction and many other topics. Each recipe includes code that you can copy and paste into a toy dataset to ensure that it actually works. From there, you can insert, combine, or adapt the code to help construct your application. Recipes also include a discussion that explains the solution and provides meaningful context. This cookbook takes you beyond theory and concepts by providing the nuts and bolts you need to construct working machine learning applications. You'll find recipes for: Vectors, matrices, and arrays Handling numerical and categorical data, text, images, and dates and times Dimensionality reduction using feature extraction or feature selection Model evaluation and selection Linear and logical regression, trees and forests, and k-nearest neighbors Support vector machines (SVM), naïve Bayes, clustering, and neural networks Saving and loading trained models
PyTorch 1.x Reinforcement Learning Cookbook - Yuxi (Hayden) Liu 2019-10-31
Implement reinforcement learning techniques and algorithms with the

help of real-world examples and recipes

Key Features

- Use PyTorch 1.x to design and build self-learning artificial intelligence (AI) models
- Implement RL algorithms to solve control and optimization challenges faced by data scientists today
- Apply modern RL libraries to simulate a controlled environment for your projects

Book Description

Reinforcement learning (RL) is a branch of machine learning that has gained popularity in recent times. It allows you to train AI models that learn from their own actions and optimize their behavior. PyTorch has also emerged as the preferred tool for training RL models because of its efficiency and ease of use. With this book, you'll explore the important RL concepts and the implementation of algorithms in PyTorch 1.x. The recipes in the book, along with real-world examples, will help you master various RL techniques, such as dynamic programming, Monte Carlo simulations, temporal difference, and Q-learning. You'll also gain insights into industry-specific applications of these techniques. Later chapters will guide you through solving problems such as the multi-armed bandit problem and the cartpole problem using the multi-armed bandit algorithm and function approximation. You'll also learn how to use Deep Q-Networks to complete Atari games, along with how to effectively implement policy gradients. Finally, you'll discover how RL techniques are applied to Blackjack, Gridworld environments, internet advertising, and the Flappy Bird game. By the end of this book, you'll have developed the skills you need to implement popular RL algorithms and use RL techniques to solve real-world problems. What you will learn

Use Q-learning and the state-action-reward-state-action (SARSA) algorithm to solve various Gridworld problems

- Develop a multi-armed bandit algorithm to optimize display advertising
- Scale up learning and control processes using Deep Q-Networks
- Simulate Markov Decision Processes, OpenAI Gym environments, and other common control problems
- Select and build RL models, evaluate their performance, and optimize and deploy them
- Use policy gradient methods to solve continuous RL problems

Who this book is for Machine learning engineers, data scientists and AI researchers looking for quick solutions to different reinforcement learning problems will find this book useful. Although prior knowledge of machine learning concepts is required, experience with PyTorch will be useful but not necessary.

Deep Learning Cookbook - Douwe Osinga 2018-06-05

Deep learning doesn't have to be intimidating. Until recently, this machine-learning method required years of study, but with frameworks such as Keras and Tensorflow, software engineers without a background in machine learning can quickly enter the field. With the recipes in this cookbook, you'll learn how to solve deep-learning problems for classifying and generating text, images, and music. Each chapter consists of several recipes needed to complete a single project, such as training a music recommending system. Author Douwe Osinga also provides a chapter with half a dozen techniques to help you if you're stuck. Examples are written in Python with code available on GitHub as a set of Python notebooks. You'll learn how to:

- Create applications that will serve real users
- Use word embeddings to calculate text similarity
- Build a movie recommender system based on Wikipedia links
- Learn how AIs see the world by visualizing their internal state
- Build a model to suggest emojis for pieces of text
- Reuse pretrained networks to build an inverse image search service
- Compare how GANs, autoencoders and LSTMs generate icons
- Detect music styles and index song collections

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

Natural Language Processing with Python Cookbook - Krishna Bhavsar

2017-11-24

Learn the tricks and tips that will help you design Text Analytics solutions

About This Book

- Independent recipes that will teach you how to efficiently perform Natural Language Processing in Python
- Use dictionaries to create your own named entities using this easy-to-follow guide
- Learn how to implement NLTK for various scenarios with the help of example-rich recipes to take you beyond basic Natural Language Processing

Who This Book Is For

This book is intended for data scientists, data analysts, and data science professionals who want to upgrade their existing skills to implement advanced text analytics using NLP. Some basic knowledge of Natural Language Processing is recommended.

What You Will Learn

- Explore corpus management using internal and external corpora
- Learn WordNet usage and a couple of simple application assignments using WordNet
- Operate on raw text
- Learn to perform tokenization, stemming, lemmatization, and spelling corrections, stop words removals, and more
- Understand regular expressions for pattern matching
- Learn to use and write your own POS taggers and grammars
- Learn to evaluate your own trained models
- Explore Deep Learning techniques in NLP
- Generate Text from Nietzsche's writing using LSTM
- Utilize the BABI dataset and LSTM to model episodes

In Detail

Natural Language Processing (NLP) is a field of computer science, artificial intelligence, and computational linguistics concerned with the interactions between computers and human (natural) languages; in particular, it's about programming computers to fruitfully process large natural language corpora. This book includes unique recipes that will teach you various aspects of performing Natural Language Processing with NLTK—the leading Python platform for the task. You will come across various recipes during the course, covering (among other topics) natural language understanding, Natural Language Processing, and syntactic analysis. You will learn how to understand language, plan sentences, and work around various ambiguities. You will learn how to efficiently use NLTK and implement text classification, identify parts of speech, tag words, and more. You will also learn how to analyze sentence structures and master lexical analysis, syntactic and semantic analysis, pragmatic analysis, and the application of deep learning techniques. By the end of this book, you will have all the knowledge you need to implement Natural Language Processing with Python.

Style and Approach

This book's rich collection of recipes will come in handy when you are working with Natural Language Processing with Python. Addressing your common and not-so-common pain points, this is a book that you must have on the shelf.

TensorFlow 2 Reinforcement Learning Cookbook - Praveen Palanisamy 2021-01-15

This cookbook will help you to gain a solid understanding of deep reinforcement learning (RL) algorithms with the help of concise, easy-to-follow implementations from scratch. You'll learn how to implement these algorithms with minimal code and develop AI applications to solve real-world and business problems using RL.

Machine Learning Cookbook with Python - Rehan Guha 2020-11-12

A Cookbook that will help you implement Machine Learning algorithms and techniques by building real-world projects

KEY FEATURES

- Learn how to handle an entire Machine Learning Pipeline supported with adequate mathematics. Create Predictive Models and choose the right model for various types of Datasets. Learn the art of tuning a model to improve accuracy as per Business requirements. Get familiar with concepts related to Data Analytics with Visualization, Data Science and Machine Learning.

DESCRIPTION

Machine Learning does not have to be intimidating at all. This book focuses on the concepts of Machine Learning and Data Analytics with mathematical explanations and programming examples. All the codes are written in Python as it is one of the most popular programming languages used for Data Science and Machine Learning. Here I have leveraged multiple libraries like NumPy, Pandas, scikit-learn, etc. to ease our task and not reinvent the wheel. There are five projects in total, each addressing a unique problem. With the recipes in this cookbook, one will learn how to solve Machine Learning problems for real-time data and perform Data Analysis and Analytics, Classification, and beyond. The datasets used are also unique and will help one to think, understand the problem and proceed towards the goal. The book is not saturated with Mathematics, but mostly all the Mathematical concepts are covered for the important topics. Every chapter typically starts with some theory and prerequisites, and then it gradually dives into the implementation of the same concept using Python, keeping a project in the background.

WHAT WILL YOU LEARN

- Understand the working of the O.S.E.M.N. framework in Data Science.
- Get familiar with the end-to-end implementation of Machine

Learning Pipeline. Learn how to implement Machine Learning algorithms and concepts using Python. Learn how to build a Predictive Model for a Business case. WHO THIS BOOK IS FORÉ This cookbook is meant for anybody who is passionate enough to get into the World of Machine Learning and has a preliminary understanding of the Basics of Linear Algebra, Calculus, Probability, and Statistics. This book also serves as a reference guidebook for intermediate Machine Learning practitioners. Ê TABLE OF CONTENTS 1. Boston Crime 2. World Happiness Report 3. Iris Species 4. Credit Card Fraud Detection 5. Heart Disease UCI [Machine Learning with R Cookbook](#) - Yu-Wei, Chiu (David Chiu) 2015-03-26

The R language is a powerful open source functional programming language. At its core, R is a statistical programming language that provides impressive tools to analyze data and create high-level graphics. This book covers the basics of R by setting up a user-friendly programming environment and performing data ETL in R. Data exploration examples are provided that demonstrate how powerful data visualization and machine learning is in discovering hidden relationships. You will then dive into important machine learning topics, including data classification, regression, clustering, association rule mining, and dimension reduction.

[Python Machine Learning Cookbook](#) - Giuseppe Ciaburro 2019-03-29
Discover powerful ways to effectively solve real-world machine learning problems using key libraries including scikit-learn, TensorFlow, and PyTorch Key Features Learn and implement machine learning algorithms in a variety of real-life scenarios Cover a range of tasks catering to supervised, unsupervised and reinforcement learning techniques Find easy-to-follow code solutions for tackling common and not-so-common challenges Book Description This eagerly anticipated second edition of the popular Python Machine Learning Cookbook will enable you to adopt a fresh approach to dealing with real-world machine learning and deep learning tasks. With the help of over 100 recipes, you will learn to build powerful machine learning applications using modern libraries from the Python ecosystem. The book will also guide you on how to implement various machine learning algorithms for classification, clustering, and recommendation engines, using a recipe-based approach. With emphasis on practical solutions, dedicated sections in the book will help you to apply supervised and unsupervised learning techniques to real-world problems. Toward the concluding chapters, you will get to grips with recipes that teach you advanced techniques including reinforcement learning, deep neural networks, and automated machine learning. By the end of this book, you will be equipped with the skills you need to apply machine learning techniques and leverage the full capabilities of the Python ecosystem through real-world examples. What you will learn Use predictive modeling and apply it to real-world problems Explore data visualization techniques to interact with your data Learn how to build a recommendation engine Understand how to interact with text data and build models to analyze it Work with speech data and recognize spoken words using Hidden Markov Models Get well versed with reinforcement learning, automated ML, and transfer learning Work with image data and build systems for image recognition and biometric face recognition Use deep neural networks to build an optical character recognition system Who this book is for This book is for data scientists, machine learning developers, deep learning enthusiasts and Python programmers who want to solve real-world challenges using machine-learning techniques and algorithms. If you are facing challenges at work and want ready-to-use code solutions to cover key tasks in machine learning and the deep learning domain, then this book is what you need. Familiarity with Python programming and machine learning concepts will be useful. *Practical Machine Learning with Python* - Dipanjan Sarkar 2017-12-20 Master the essential skills needed to recognize and solve complex problems with machine learning and deep learning. Using real-world examples that leverage the popular Python machine learning ecosystem, this book is your perfect companion for learning the art and science of machine learning to become a successful practitioner. The concepts, techniques, tools, frameworks, and methodologies used in this book will teach you how to think, design, build, and execute machine learning systems and projects successfully. *Practical Machine Learning with Python* follows a structured and comprehensive three-tiered approach packed with hands-on examples and code. Part 1 focuses on understanding machine learning concepts and tools. This includes machine learning basics with a broad overview of algorithms, techniques, concepts and applications, followed by a tour of the entire Python machine learning ecosystem. Brief guides for useful machine learning tools, libraries and frameworks are also covered. Part 2 details standard

machine learning pipelines, with an emphasis on data processing analysis, feature engineering, and modeling. You will learn how to process, wrangle, summarize and visualize data in its various forms. Feature engineering and selection methodologies will be covered in detail with real-world datasets followed by model building, tuning, interpretation and deployment. Part 3 explores multiple real-world case studies spanning diverse domains and industries like retail, transportation, movies, music, marketing, computer vision and finance. For each case study, you will learn the application of various machine learning techniques and methods. The hands-on examples will help you become familiar with state-of-the-art machine learning tools and techniques and understand what algorithms are best suited for any problem. *Practical Machine Learning with Python* will empower you to start solving your own problems with machine learning today! What You'll Learn Execute end-to-end machine learning projects and systems Implement hands-on examples with industry standard, open source, robust machine learning tools and frameworks Review case studies depicting applications of machine learning and deep learning on diverse domains and industries Apply a wide range of machine learning models including regression, classification, and clustering. Understand and apply the latest models and methodologies from deep learning including CNNs, RNNs, LSTMs and transfer learning. Who This Book Is For IT professionals, analysts, developers, data scientists, engineers, graduate students

Artificial Intelligence with Python Cookbook - Ben Auffarth 2020-10-30

Work through practical recipes to learn how to solve complex machine learning and deep learning problems using Python Key FeaturesGet up and running with artificial intelligence in no time using hands-on problem-solving recipesExplore popular Python libraries and tools to build AI solutions for images, text, sounds, and imagesImplement NLP, reinforcement learning, deep learning, GANs, Monte-Carlo tree search, and much moreBook Description Artificial intelligence (AI) plays an integral role in automating problem-solving. This involves predicting and classifying data and training agents to execute tasks successfully. This book will teach you how to solve complex problems with the help of independent and insightful recipes ranging from the essentials to advanced methods that have just come out of research. *Artificial Intelligence with Python Cookbook* starts by showing you how to set up your Python environment and taking you through the fundamentals of data exploration. Moving ahead, you'll be able to implement heuristic search techniques and genetic algorithms. In addition to this, you'll apply probabilistic models, constraint optimization, and reinforcement learning. As you advance through the book, you'll build deep learning models for text, images, video, and audio, and then delve into algorithmic bias, style transfer, music generation, and AI use cases in the healthcare and insurance industries. Throughout the book, you'll learn about a variety of tools for problem-solving and gain the knowledge needed to effectively approach complex problems. By the end of this book on AI, you will have the skills you need to write AI and machine learning algorithms, test them, and deploy them for production. What you will learnImplement data preprocessing steps and optimize model hyperparametersDelve into representational learning with adversarial autoencodersUse active learning, recommenders, knowledge embedding, and SAT solversGet to grips with probabilistic modeling with TensorFlow probabilityRun object detection, text-to-speech conversion, and text and music generationApply swarm algorithms, multi-agent systems, and graph networksGo from proof of concept to production by deploying models as microservicesUnderstand how to use modern AI in practiceWho this book is for This AI machine learning book is for Python developers, data scientists, machine learning engineers, and deep learning practitioners who want to learn how to build artificial intelligence solutions with easy-to-follow recipes. You'll also find this book useful if you're looking for state-of-the-art solutions to perform different machine learning tasks in various use cases. Basic working knowledge of the Python programming language and machine learning concepts will help you to work with code effectively in this book.

IPython Interactive Computing and Visualization Cookbook - Cyrille Rossant 2014-09-25

Intended to anyone interested in numerical computing and data science: students, researchers, teachers, engineers, analysts, hobbyists... Basic knowledge of Python/NumPy is recommended. Some skills in mathematics will help you understand the theory behind the computational methods.

Machine Learning with Python - Abhishek Vijayvargia 2018-03-01

Providing code examples in python, this book introduces the concepts of machine learning with mathematical explanations and programming fundamentals. --

Python for Finance Cookbook - Eryk Lewinson 2020-01-31

Solve common and not-so-common financial problems using Python libraries such as NumPy, SciPy, and pandas Key Features Use powerful Python libraries such as pandas, NumPy, and SciPy to analyze your financial data Explore unique recipes for financial data analysis and processing with Python Estimate popular financial models such as CAPM and GARCH using a problem-solution approach Book Description Python is one of the most popular programming languages used in the financial industry, with a huge set of accompanying libraries. In this book, you'll cover different ways of downloading financial data and preparing it for modeling. You'll calculate popular indicators used in technical analysis, such as Bollinger Bands, MACD, RSI, and backtest automatic trading strategies. Next, you'll cover time series analysis and models, such as exponential smoothing, ARIMA, and GARCH (including multivariate specifications), before exploring the popular CAPM and the Fama-French three-factor model. You'll then discover how to optimize asset allocation and use Monte Carlo simulations for tasks such as calculating the price of American options and estimating the Value at Risk (VaR). In later chapters, you'll work through an entire data science project in the financial domain. You'll also learn how to solve the credit card fraud and default problems using advanced classifiers such as random forest, XGBoost, LightGBM, and stacked models. You'll then be able to tune the hyperparameters of the models and handle class imbalance. Finally, you'll focus on learning how to use deep learning (PyTorch) for approaching financial tasks. By the end of this book, you'll have learned how to effectively analyze financial data using a recipe-based approach. What you will learn Download and preprocess financial data from different sources Backtest the performance of automatic trading strategies in a real-world setting Estimate financial econometrics models in Python and interpret their results Use Monte Carlo simulations for a variety of tasks such as derivatives valuation and risk assessment Improve the performance of financial models with the latest Python libraries Apply machine learning and deep learning techniques to solve different financial problems Understand the different approaches used to model financial time series data Who this book is for This book is for financial analysts, data analysts, and Python developers who want to learn how to implement a broad range of tasks in the finance domain. Data scientists looking to devise intelligent financial strategies to perform efficient financial analysis will also find this book useful. Working knowledge of the Python programming language is mandatory to grasp the concepts covered in the book effectively.

Deep Learning with Python - Francois Chollet 2017-11-30

Summary Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Machine learning has made remarkable progress in recent years. We went from near-unusable speech and image recognition, to near-human accuracy. We went from machines that couldn't beat a serious Go player, to defeating a world champion. Behind this progress is deep learning—a combination of engineering advances, best practices, and theory that enables a wealth of previously impossible smart applications. About the Book Deep Learning with Python introduces the field of deep learning using the Python language and the powerful Keras library. Written by Keras creator and Google AI researcher François Chollet, this book builds your understanding through intuitive explanations and practical examples. You'll explore challenging concepts and practice with applications in computer vision, natural-language processing, and generative models. By the time you finish, you'll have the knowledge and hands-on skills to apply deep learning in your own projects. What's Inside Deep learning from first principles Setting up your own deep-learning environment Image-classification models Deep learning for text and sequences Neural style transfer, text generation, and image generation About the Reader Readers need intermediate Python skills. No previous experience with Keras, TensorFlow, or machine learning is required. About the Author François Chollet works on deep learning at Google in Mountain View, CA. He is the creator of the Keras deep-learning library, as well as a contributor to the TensorFlow machine-learning framework. He also does deep-learning research, with a focus on computer vision and the application of machine

learning to formal reasoning. His papers have been published at major conferences in the field, including the Conference on Computer Vision and Pattern Recognition (CVPR), the Conference and Workshop on Neural Information Processing Systems (NIPS), the International Conference on Learning Representations (ICLR), and others. Table of Contents PART 1 - FUNDAMENTALS OF DEEP LEARNING What is deep learning? Before we begin: the mathematical building blocks of neural networks Getting started with neural networks Fundamentals of machine learning PART 2 - DEEP LEARNING IN PRACTICE Deep learning for computer vision Deep learning for text and sequences Advanced deep-learning best practices Generative deep learning Conclusions appendix A - Installing Keras and its dependencies on Ubuntu appendix B - Running Jupyter notebooks on an EC2 GPU instance

Python Machine Learning Cookbook - Giuseppe Ciaburro 2019-03-30

Discover powerful ways to effectively solve real-world machine learning problems using key libraries including scikit-learn, TensorFlow, and PyTorch Key Features Learn and implement machine learning algorithms in a variety of real-life scenarios Cover a range of tasks catering to supervised, unsupervised and reinforcement learning techniques Find easy-to-follow code solutions for tackling common and not-so-common challenges Book Description This eagerly anticipated second edition of the popular Python Machine Learning Cookbook will enable you to adopt a fresh approach to dealing with real-world machine learning and deep learning tasks. With the help of over 100 recipes, you will learn to build powerful machine learning applications using modern libraries from the Python ecosystem. The book will also guide you on how to implement various machine learning algorithms for classification, clustering, and recommendation engines, using a recipe-based approach. With emphasis on practical solutions, dedicated sections in the book will help you to apply supervised and unsupervised learning techniques to real-world problems. Toward the concluding chapters, you will get to grips with recipes that teach you advanced techniques including reinforcement learning, deep neural networks, and automated machine learning. By the end of this book, you will be equipped with the skills you need to apply machine learning techniques and leverage the full capabilities of the Python ecosystem through real-world examples. What you will learn Use predictive modeling and apply it to real-world problems Explore data visualization techniques to interact with your data Learn how to build a recommendation engine Understand how to interact with text data and build models to analyze it Work with speech data and recognize spoken words using Hidden Markov Models Get well versed with reinforcement learning, automated ML, and transfer learning Work with image data and build systems for image recognition and biometric face recognition Use deep neural networks to build an optical character recognition system Who this book is for This book is for data scientists, machine learning developers, deep learning enthusiasts and Python programmers who want to solve real-world challenges using machine-learning techniques and algorithms. If you are facing challenges at work and want ready-to-use code solutions to cover key tasks in machine learning and the deep learning domain, then this book is what you need. Familiarity with Python programming and machine learning concepts will be useful.

Applied Text Analysis with Python - Benjamin Bengfort 2018-06-11

From news and speeches to informal chatter on social media, natural language is one of the richest and most underutilized sources of data. Not only does it come in a constant stream, always changing and adapting in context; it also contains information that is not conveyed by traditional data sources. The key to unlocking natural language is through the creative application of text analytics. This practical book presents a data scientist's approach to building language-aware products with applied machine learning. You'll learn robust, repeatable, and scalable techniques for text analysis with Python, including contextual and linguistic feature engineering, vectorization, classification, topic modeling, entity resolution, graph analysis, and visual steering. By the end of the book, you'll be equipped with practical methods to solve any number of complex real-world problems. Preprocess and vectorize text into high-dimensional feature representations Perform document classification and topic modeling Steer the model selection process with visual diagnostics Extract key phrases, named entities, and graph structures to reason about data in text Build a dialog framework to enable chatbots and language-driven interaction Use Spark to scale processing power and neural networks to scale model complexity

Pragmatic Machine Learning with Python - Avishek Nag 2020-04-30

An easy-to-understand guide to learn practical Machine Learning techniques with Mathematical foundations KEY FEATURES - A balanced combination of underlying mathematical theories & practical examples

with Python code - Coverage of latest topics like multi-label classification, Text Mining, Doc2Vec, Word2Vec, XMeans clustering, unsupervised outlier detection, techniques to deploy ML models in production-grade systems with PMML, etc - Coverage of sufficient & relevant visualization techniques specific to any topic

DESCRIPTION
This book will be ideal for working professionals who want to learn Machine Learning from scratch. The first chapter will be an introductory chapter to make readers comfortable with the idea of Machine Learning and the required mathematical theories. There will be a balanced combination of underlying mathematical theories corresponding to any Machine Learning topic and its implementation using Python. Most of the implementations will be based on 'scikit-learn,' but other Python libraries like 'Gensim' or 'PyTorch' will also be used for some topics like text analytics or deep learning. The book will be divided into chapters based on primary Machine Learning topics like Classification, Regression, Clustering, Deep Learning, Text Mining, etc. The book will also explain different techniques of putting Machine Learning models into production-grade systems using Big Data or Non-Big Data flavors and standards for exporting models.

WHAT WILL YOU LEARN - Get familiar with practical concepts of Machine Learning from ground zero - Learn how to deploy Machine Learning models in production - Understand how to do "Data Science Storytelling" - Explore the latest topics in the current industry about Machine Learning

WHO THIS BOOK IS FOR This book would be ideal for experienced Software Professionals who are trying to get into the field of Machine Learning. Anyone who wishes to Learn Machine Learning concepts and models in the production lifecycle.

TABLE OF CONTENTS
1. Introduction to Machine Learning & Mathematical preliminaries
2. Classification
3. Regression
4. Clustering
5. Deep Learning & Neural Networks
6. Miscellaneous Unsupervised Learning
7. Text Mining
8. Machine Learning models in production
9. Case Studies & Data Science Storytelling

TensorFlow Machine Learning Cookbook - Nick McClure 2017-02-14
Explore machine learning concepts using the latest numerical computing library — TensorFlow — with the help of this comprehensive cookbook

About This Book Your quick guide to implementing TensorFlow in your day-to-day machine learning activities Learn advanced techniques that bring more accuracy and speed to machine learning Upgrade your knowledge to the second generation of machine learning with this guide on TensorFlow

Who This Book Is For This book is ideal for data scientists who are familiar with C++ or Python and perform machine learning activities on a day-to-day basis. Intermediate and advanced machine learning implementers who need a quick guide they can easily navigate will find it useful.

What You Will Learn Become familiar with the basics of the TensorFlow machine learning library Get to know Linear Regression techniques with TensorFlow Learn SVMs with hands-on recipes Implement neural networks and improve predictions Apply NLP and sentiment analysis to your data Master CNN and RNN through practical recipes Take TensorFlow into production

In Detail TensorFlow is an open source software library for Machine Intelligence. The independent recipes in this book will teach you how to use TensorFlow for complex data computations and will let you dig deeper and gain more insights into your data than ever before. You'll work through recipes on training models, model evaluation, sentiment analysis, regression analysis, clustering analysis, artificial neural networks, and deep learning - each using Google's machine learning library TensorFlow. This guide starts with the fundamentals of the TensorFlow library which includes variables, matrices, and various data sources. Moving ahead, you will get hands-on experience with Linear Regression techniques with TensorFlow. The next chapters cover important high-level concepts such as neural networks, CNN, RNN, and NLP. Once you are familiar and comfortable with the TensorFlow ecosystem, the last chapter will show you how to take it to production.

Style and approach This book takes a recipe-based approach where every topic is explicated with the help of a real-world example.

Python Deep Learning Cookbook - Indra den Bakker 2017-10-27
Solve different problems in modelling deep neural networks using Python, Tensorflow, and Keras with this practical guide

About This Book Practical recipes on training different neural network models and tuning them for optimal performance Use Python frameworks like TensorFlow, Caffe, Keras, Theano for Natural Language Processing, Computer Vision, and more A hands-on guide covering the common as well as the not so common problems in deep learning using Python

Who This Book Is For This book is intended for machine learning professionals who are looking to use deep learning algorithms to create real-world applications using Python. Thorough understanding of the machine learning concepts and

Python libraries such as NumPy, SciPy and scikit-learn is expected. Additionally, basic knowledge in linear algebra and calculus is desired.

What You Will Learn Implement different neural network models in Python Select the best Python framework for deep learning such as PyTorch, Tensorflow, MXNet and Keras Apply tips and tricks related to neural networks internals, to boost learning performances Consolidate machine learning principles and apply them in the deep learning field Reuse and adapt Python code snippets to everyday problems Evaluate the cost/benefits and performance implication of each discussed solution

In Detail Deep Learning is revolutionizing a wide range of industries. For many applications, deep learning has proven to outperform humans by making faster and more accurate predictions. This book provides a top-down and bottom-up approach to demonstrate deep learning solutions to real-world problems in different areas. These applications include Computer Vision, Natural Language Processing, Time Series, and Robotics. The Python Deep Learning Cookbook presents technical solutions to the issues presented, along with a detailed explanation of the solutions. Furthermore, a discussion on corresponding pros and cons of implementing the proposed solution using one of the popular frameworks like TensorFlow, PyTorch, Keras and CNTK is provided. The book includes recipes that are related to the basic concepts of neural networks. All techniques s, as well as classical networks topologies. The main purpose of this book is to provide Python programmers a detailed list of recipes to apply deep learning to common and not-so-common scenarios. Style and approach Unique blend of independent recipes arranged in the most logical manner

Machine Learning for Time-Series with Python - Ben Auffarth 2021-10-29

Get better insights from time-series data and become proficient in model performance analysis

Key Features Explore popular and modern machine learning methods including the latest online and deep learning algorithms Learn to increase the accuracy of your predictions by matching the right model with the right problem Master time series via real-world case studies on operations management, digital marketing, finance, and healthcare

Book Description The Python time-series ecosystem is huge and often quite hard to get a good grasp on, especially for time-series since there are so many new libraries and new models. This book aims to deepen your understanding of time series by providing a comprehensive overview of popular Python time-series packages and help you build better predictive systems. Machine Learning for Time-Series with Python starts by re-introducing the basics of time series and then builds your understanding of traditional autoregressive models as well as modern non-parametric models. By observing practical examples and the theory behind them, you will become confident with loading time-series datasets from any source, deep learning models like recurrent neural networks and causal convolutional network models, and gradient boosting with feature engineering. This book will also guide you in matching the right model to the right problem by explaining the theory behind several useful models. You'll also have a look at real-world case studies covering weather, traffic, biking, and stock market data. By the end of this book, you should feel at home with effectively analyzing and applying machine learning methods to time-series.

What you will learn Understand the main classes of time series and learn how to detect outliers and patterns Choose the right method to solve time-series problems Characterize seasonal and correlation patterns through autocorrelation and statistical techniques Get to grips with time-series data visualization Understand classical time-series models like ARMA and ARIMA Implement deep learning models, like Gaussian processes, transformers, and state-of-the-art machine learning models Become familiar with many libraries like Prophet, XGboost, and TensorFlow

Who this book is for This book is ideal for data analysts, data scientists, and Python developers who want instantly useful and practical recipes to implement today, and a comprehensive reference book for tomorrow. Basic knowledge of the Python Programming language is a must, while familiarity with statistics will help you get the most out of this book.

Hands-On Data Science and Python Machine Learning - Frank Kane 2017-07-31

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.

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.

Python Cookbook - David Beazley 2013-05-10

If you need help writing programs in Python 3, or want to update older

Python 2 code, this book is just the ticket. Packed with practical recipes written and tested with Python 3.3, this unique cookbook is for experienced Python programmers who want to focus on modern tools and idioms. Inside, you'll find complete recipes for more than a dozen topics, covering the core Python language as well as tasks common to a wide variety of application domains. Each recipe contains code samples you can use in your projects right away, along with a discussion about how and why the solution works. Topics include: Data Structures and Algorithms Strings and Text Numbers, Dates, and Times Iterators and Generators Files and I/O Data Encoding and Processing Functions Classes and Objects Metaprogramming Modules and Packages Network and Web Programming Concurrency Utility Scripting and System Administration Testing, Debugging, and Exceptions C Extensions

Python Data Analysis Cookbook - Ivan Idris 2016-07-22

Over 140 practical recipes to help you make sense of your data with ease and build production-ready data apps About This Book Analyze Big Data sets, create attractive visualizations, and manipulate and process various data types Packed with rich recipes to help you learn and explore amazing algorithms for statistics and machine learning Authored by Ivan Idris, expert in python programming and proud author of eight highly reviewed books Who This Book Is For This book teaches Python data analysis at an intermediate level with the goal of transforming you from journeyman to master. Basic Python and data analysis skills and affinity are assumed. What You Will Learn Set up reproducible data analysis Clean and transform data Apply advanced statistical analysis Create attractive data visualizations Web scrape and work with databases, Hadoop, and Spark Analyze images and time series data Mine text and analyze social networks Use machine learning and evaluate the results Take advantage of parallelism and concurrency In Detail Data analysis is a rapidly evolving field and Python is a multi-paradigm programming language suitable for object-oriented application development and functional design patterns. As Python offers a range of tools and libraries for all purposes, it has slowly evolved as the primary language for data science, including topics on: data analysis, visualization, and machine learning. Python Data Analysis Cookbook focuses on reproducibility and creating production-ready systems. You will start with recipes that set the foundation for data analysis with libraries such as matplotlib, NumPy, and pandas. You will learn to create visualizations by choosing color maps and palettes then dive into statistical data analysis using distribution algorithms and correlations. You'll then help you find your way around different data and numerical problems, get to grips with Spark and HDFS, and then set up migration scripts for web mining. In this book, you will dive deeper into recipes on spectral analysis, smoothing, and bootstrapping methods. Moving on, you will learn to rank stocks and check market efficiency, then work with metrics and clusters. You will achieve parallelism to improve system performance by using multiple threads and speeding up your code. By the end of the book, you will be capable of handling various data analysis techniques in Python and devising solutions for problem scenarios. Style and Approach The book is written in "cookbook" style striving for high realism in data analysis. Through the recipe-based format, you can read each recipe separately as required and immediately apply the knowledge gained.

Keras Deep Learning Cookbook - Rajdeep Dua 2018-10-31

Leverage the power of deep learning and Keras to develop smarter and more efficient data models Key Features Understand different neural networks and their implementation using Keras Explore recipes for training and fine-tuning your neural network models Put your deep learning knowledge to practice with real-world use-cases, tips, and tricks Book Description Keras has quickly emerged as a popular deep learning library. Written in Python, it allows you to train convolutional as well as recurrent neural networks with speed and accuracy. The Keras Deep Learning Cookbook shows you how to tackle different problems encountered while training efficient deep learning models, with the help of the popular Keras library. Starting with installing and setting up Keras, the book demonstrates how you can perform deep learning with Keras in the TensorFlow. From loading data to fitting and evaluating your model for optimal performance, you will work through a step-by-step process to tackle every possible problem faced while training deep models. You will implement convolutional and recurrent neural networks, adversarial networks, and more with the help of this handy guide. In addition to this, you will learn how to train these models for real-world image and language processing tasks. By the end of this book, you will have a practical, hands-on understanding of how you can leverage the power of Python and Keras to perform effective deep learning What you will learn Install and configure Keras in TensorFlow Master neural

network programming using the Keras library Understand the different Keras layers Use Keras to implement simple feed-forward neural networks, CNNs and RNNs Work with various datasets and models used for image and text classification Develop text summarization and reinforcement learning models using Keras Who this book is for Keras Deep Learning Cookbook is for you if you are a data scientist or machine learning expert who wants to find practical solutions to common problems encountered while training deep learning models. A basic understanding of Python and some experience in machine learning and neural networks is required for this book.

TensorFlow 1.x Deep Learning Cookbook - Antonio Gulli 2017-12-12
Take the next step in implementing various common and not-so-common neural networks with Tensorflow 1.x About This Book Skill up and implement tricky neural networks using Google's TensorFlow 1.x An easy-to-follow guide that lets you explore reinforcement learning, GANs, autoencoders, multilayer perceptrons and more. Hands-on recipes to work with Tensorflow on desktop, mobile, and cloud environment Who This Book Is For This book is intended for data analysts, data scientists, machine learning practitioners and deep learning enthusiasts who want to perform deep learning tasks on a regular basis and are looking for a handy guide they can refer to. People who are slightly familiar with neural networks, and now want to gain expertise in working with different types of neural networks and datasets, will find this book quite useful. What You Will Learn Install TensorFlow and use it for CPU and GPU operations Implement DNNs and apply them to solve different AI-driven problems. Leverage different data sets such as MNIST, CIFAR-10, and Youtube8m with TensorFlow and learn how to access and use them in your code. Use TensorBoard to understand neural network architectures, optimize the learning process, and peek inside the neural network black box. Use different regression techniques for prediction and classification problems Build single and multilayer perceptrons in TensorFlow Implement CNN and RNN in TensorFlow, and use it to solve real-world use cases. Learn how restricted Boltzmann Machines can be used to recommend movies. Understand the implementation of Autoencoders and deep belief networks, and use them for emotion detection. Master the different reinforcement learning methods to implement game playing agents. GANs and their implementation using TensorFlow. In Detail Deep neural networks (DNNs) have achieved a lot of success in the field of computer vision, speech recognition, and natural language processing. The entire world is filled with excitement about how deep networks are revolutionizing artificial intelligence. This exciting recipe-based guide will take you from the realm of DNN theory to implementing them practically to solve the real-life problems in artificial intelligence domain. In this book, you will learn how to efficiently use TensorFlow, Google's open source framework for deep learning. You will implement different deep learning networks such as Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), Deep Q-learning Networks (DQNs), and Generative Adversarial Networks (GANs) with easy to follow independent recipes. You will learn how to make Keras as backend with TensorFlow. With a problem-solution approach, you will understand how to implement different deep neural architectures to carry out complex tasks at work. You will learn the performance of different DNNs on some popularly used data sets such as MNIST, CIFAR-10, Youtube8m, and more. You will not only learn about the different mobile and embedded platforms supported by TensorFlow but also how to set up cloud platforms for deep learning applications. Get a sneak peek of TPU architecture and how they will affect DNN future. By using crisp, no-nonsense recipes, you will become an expert in implementing deep learning techniques in growing real-world applications and research areas such as reinforcement learning, GANs, autoencoders and more. Style and approach This book consists of hands-on recipes where you'll deal with real-world problems. You'll execute a series of tasks as you walk through data mining challenges using TensorFlow 1.x. Your one-stop solution for common and not-so-common pain points, this is a book that you must have on the shelf.

Machine Learning Using TensorFlow Cookbook - Alexia Audevert 2021-02-08

Comprehensive recipes to give you valuable insights on Transformers, Reinforcement Learning, and more Key Features Deep Learning solutions from Kaggle Masters and Google Developer Experts Get to grips with the fundamentals including variables, matrices, and data sources Learn advanced techniques to make your algorithms faster and more accurate Book Description The independent recipes in Machine Learning Using TensorFlow Cookbook will teach you how to perform complex data computations and gain valuable insights into your data. Dive into recipes

on training models, model evaluation, sentiment analysis, regression analysis, artificial neural networks, and deep learning - each using Google's machine learning library, TensorFlow. This cookbook covers the fundamentals of the TensorFlow library, including variables, matrices, and various data sources. You'll discover real-world implementations of Keras and TensorFlow and learn how to use estimators to train linear models and boosted trees, both for classification and regression. Explore the practical applications of a variety of deep learning architectures, such as recurrent neural networks and Transformers, and see how they can be used to solve computer vision and natural language processing (NLP) problems. With the help of this book, you will be proficient in using TensorFlow, understand deep learning from the basics, and be able to implement machine learning algorithms in real-world scenarios. What you will learn Take TensorFlow into production Implement and fine-tune Transformer models for various NLP tasks Apply reinforcement learning algorithms using the TF-Agents framework Understand linear regression techniques and use Estimators to train linear models Execute neural networks and improve predictions on tabular data Master convolutional neural networks and recurrent neural networks through practical recipes Who this book is for If you are a data scientist or a machine learning engineer, and you want to skip detailed theoretical explanations in favor of building production-ready machine learning models using TensorFlow, this book is for you. Basic familiarity with Python, linear algebra, statistics, and machine learning is necessary to make the most out of this book.

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 **Python Feature Engineering Cookbook** - Soledad Galli 2020-01-22
Extract accurate information from data to train and improve machine learning models using NumPy, SciPy, pandas, and scikit-learn libraries Key Features Discover solutions for feature generation, feature extraction, and feature selection Uncover the end-to-end feature engineering process across continuous, discrete, and unstructured datasets Implement modern feature extraction techniques using Python's pandas, scikit-learn, SciPy and NumPy libraries Book Description Feature engineering is invaluable for developing and enriching your machine learning models. In this cookbook, you will work with the best tools to streamline your feature engineering pipelines and techniques and simplify and improve the quality of your code. Using Python libraries such as pandas, scikit-learn, Featuretools, and Feature-engine, you'll learn how to work with both continuous and discrete datasets and be able to transform features from unstructured datasets. You will develop the skills necessary to select the best features as well as the most suitable extraction techniques. This book will cover Python recipes that will help you automate feature engineering to simplify complex processes. You'll also get to grips with different feature engineering strategies, such as the box-cox transform, power transform, and log transform across machine learning, reinforcement learning, and natural language processing (NLP) domains. By the end of this book, you'll have discovered tips and practical solutions to all of your feature engineering problems. What you will learn Simplify your feature engineering pipelines with powerful Python packages Get to grips with imputing missing values Encode categorical variables with a wide set of techniques Extract insights from text quickly and effortlessly Develop features from transactional data and time series data Derive new features by combining existing variables Understand how to transform, discretize, and scale

your variables Create informative variables from date and time Who this book is for This book is for machine learning professionals, AI engineers, data scientists, and NLP and reinforcement learning engineers who want to optimize and enrich their machine learning models with the best features. Knowledge of machine learning and Python coding will assist you with understanding the concepts covered in this book.

OpenCV 3 Computer Vision with Python Cookbook - Aleksei Spizhevoi 2018-03-23

Recipe-based approach to tackle the most common problems in Computer Vision by leveraging the functionality of OpenCV using Python APIs Key Features ● Build computer vision applications with OpenCV functionality via Python API ● Get to grips with image processing, multiple view geometry, and machine learning ● Learn to use deep learning models for image classification, object detection, and face recognition Book Description OpenCV 3 is a native cross-platform library for computer vision, machine learning, and image processing. OpenCV's convenient high-level APIs hide very powerful internals designed for computational efficiency that can take advantage of multicore and GPU processing. This book will help you tackle increasingly challenging computer vision problems by providing a number of recipes that you can use to improve your applications. In this book, you will learn how to process an image by manipulating pixels and analyze an image using histograms. Then, we'll show you how to apply image filters to enhance image content and exploit the image geometry in order to relay different views of a pictured scene. We'll explore techniques to achieve camera calibration and perform a multiple-view analysis. Later, you'll work on reconstructing a 3D scene from images, converting low-level pixel information to high-level concepts for applications such as object detection and recognition. You'll also discover how to process video from files or cameras and how to detect and track moving objects. Finally, you'll get acquainted with recent approaches in deep learning and neural networks. By the end of the book, you'll be able to apply your skills in OpenCV to create computer vision applications in various domains. What you will learn ● Get familiar with low-level image processing methods ● See the common linear algebra tools needed in computer vision ● Work with different camera models and epipolar geometry ● Find out how to detect interesting points in images and compare them ● Binarize images and mask out regions of interest ● Detect objects and track them in videos Who this book is for This book is for developers who have a basic knowledge of Python. If you are aware of the basics of OpenCV and are ready to build computer vision systems that are smarter, faster, more complex, and more practical than the competition, then this book is for you.

Statistics for Machine Learning - Pratap Dangeti 2017-07-21

Build Machine Learning models with a sound statistical understanding. About This Book Learn about the statistics behind powerful predictive models with p-value, ANOVA, and F- statistics. Implement statistical computations programmatically for supervised and unsupervised learning through K-means clustering. Master the statistical aspect of Machine Learning with the help of this example-rich guide to R and Python. Who This Book Is For This book is intended for developers with little to no background in statistics, who want to implement Machine Learning in their systems. Some programming knowledge in R or Python will be useful. What You Will Learn Understand the Statistical and Machine Learning fundamentals necessary to build models Understand the major differences and parallels between the statistical way and the Machine Learning way to solve problems Learn how to prepare data and feed models by using the appropriate Machine Learning algorithms from the more-than-adequate R and Python packages Analyze the results and tune the model appropriately to your own predictive goals Understand the concepts of required statistics for Machine Learning Introduce yourself to necessary fundamentals required for building supervised & unsupervised deep learning models Learn reinforcement learning and its application in the field of artificial intelligence domain In Detail Complex statistics in Machine Learning worry a lot of developers. Knowing statistics helps you build strong Machine Learning models that are optimized for a given problem statement. This book will teach you all it takes to perform complex statistical computations required for Machine Learning. You will gain information on statistics behind supervised learning, unsupervised learning, reinforcement learning, and more. Understand the real-world examples that discuss the statistical side of Machine Learning and familiarize yourself with it. You will also design programs for performing tasks such as model, parameter fitting, regression, classification, density collection, and more. By the end of the book, you will have mastered the required statistics for Machine

Learning and will be able to apply your new skills to any sort of industry problem. Style and approach This practical, step-by-step guide will give you an understanding of the Statistical and Machine Learning fundamentals you'll need to build models.

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

Machine Learning with Amazon SageMaker Cookbook - Joshua Arvin Lat 2021-10-29

A step-by-step solution-based guide to preparing building, training, and deploying high-quality machine learning models with Amazon SageMaker Key Features Perform ML experiments with built-in and custom algorithms in SageMaker Explore proven solutions when working with TensorFlow, PyTorch, Hugging Face Transformers, and scikit-learn Use the different features and capabilities of SageMaker to automate relevant ML processes Book Description Amazon SageMaker is a fully managed machine learning (ML) service that helps data scientists and ML practitioners manage ML experiments. In this book, you'll use the different capabilities and features of Amazon SageMaker to solve relevant data science and ML problems. This step-by-step guide features 80 proven recipes designed to give you the hands-on machine learning experience needed to contribute to real-world experiments and projects. You'll cover the algorithms and techniques that are commonly used when training and deploying NLP, time series forecasting, and computer vision models to solve ML problems. You'll explore various solutions for working with deep learning libraries and frameworks such as TensorFlow, PyTorch, and Hugging Face Transformers in Amazon SageMaker. You'll also learn how to use SageMaker Clarify, SageMaker Model Monitor, SageMaker Debugger, and SageMaker Experiments to debug, manage, and monitor multiple ML experiments and deployments. Moreover, you'll have a better understanding of how SageMaker Feature Store, Autopilot, and Pipelines can meet the specific needs of data science teams. By the end of this book, you'll be able to combine the different solutions you've learned as building blocks to solve real-world ML problems. What you will learn Train and deploy NLP, time series forecasting, and computer vision models to solve different business problems Push the limits of customization in SageMaker using custom container images Use AutoML capabilities with SageMaker Autopilot to create high-quality models Work with effective data analysis and preparation techniques Explore solutions for debugging and managing ML experiments and deployments Deal with bias detection and ML explainability requirements using SageMaker Clarify Automate intermediate and complex deployments and workflows using a variety of solutions Who this book is for This book is for developers, data scientists, and machine learning practitioners interested in using Amazon SageMaker to build, analyze, and deploy machine learning models with 80 step-by-step recipes. All you need is an AWS account to get things running. Prior knowledge of AWS, machine learning, and the Python programming language will help you to grasp the concepts covered in this book more effectively.

Machine Learning with Python for Everyone - Mark Fenner 2019-07-30

The Complete Beginner's Guide to Understanding and Building Machine Learning Systems with Python Machine Learning with Python for Everyone will help you master the processes, patterns, and strategies

you need to build effective learning systems, even if you're an absolute beginner. If you can write some Python code, this book is for you, no matter how little college-level math you know. Principal instructor Mark E. Fenner relies on plain-English stories, pictures, and Python examples to communicate the ideas of machine learning. Mark begins by discussing machine learning and what it can do; introducing key mathematical and computational topics in an approachable manner; and walking you through the first steps in building, training, and evaluating learning systems. Step by step, you'll fill out the components of a practical learning system, broaden your toolbox, and explore some of the field's most sophisticated and exciting techniques. Whether you're a student, analyst, scientist, or hobbyist, this guide's insights will be applicable to every learning system you ever build or use. Understand machine learning algorithms, models, and core machine learning concepts Classify examples with classifiers, and quantify examples with regressors Realistically assess performance of machine learning systems Use feature engineering to smooth rough data into useful forms Chain multiple components into one system and tune its performance Apply machine learning techniques to images and text Connect the core concepts to neural networks and graphical models Leverage the Python scikit-learn library and other powerful tools Register your book for convenient access to downloads, updates, and/or corrections as they become available. See inside book for details.

[Deep Learning with fastai Cookbook](#) - Mark Ryan 2021-09-24

Harness the power of the easy-to-use, high-performance fastai framework to rapidly create complete deep learning solutions with few lines of code Key Features Discover how to apply state-of-the-art deep learning techniques to real-world problems Build and train neural networks using the power and flexibility of the fastai framework Use deep learning to tackle problems such as image classification and text classification Book Description fastai is an easy-to-use deep learning framework built on top of PyTorch that lets you rapidly create complete deep learning solutions with as few as 10 lines of code. Both predominant low-level deep learning frameworks, TensorFlow and PyTorch, require a lot of code, even for straightforward applications. In contrast, fastai handles the messy details for you and lets you focus on applying deep learning to actually solve problems. The book begins by summarizing the value of fastai and showing you how to create a simple 'hello world' deep learning application with fastai. You'll then learn how to use fastai for all four application areas that the framework explicitly supports: tabular data, text data (NLP), recommender systems, and vision data. As you advance, you'll work through a series of practical examples that illustrate how to create real-world applications of each type. Next, you'll learn how to deploy fastai models, including creating a simple web application that predicts what object is depicted in an image. The book wraps up with an overview of the advanced features of fastai. By the end of this fastai book, you'll be able to create your own deep learning applications using fastai. You'll also have learned how to use fastai to prepare raw datasets, explore datasets, train deep learning models, and deploy trained models. What you will learn Prepare real-

world raw datasets to train fastai deep learning models Train fastai deep learning models using text and tabular data Create recommender systems with fastai Find out how to assess whether fastai is a good fit for a given problem Deploy fastai deep learning models in web applications Train fastai deep learning models for image classification Who this book is for This book is for data scientists, machine learning developers, and deep learning enthusiasts looking to explore the fastai framework using a recipe-based approach. Working knowledge of the Python programming language and machine learning basics is strongly recommended to get the most out of this deep learning book.

[Python Machine Learning Cookbook](#) - Prateek Joshi 2016-06-23

100 recipes that teach you how to perform various machine learning tasks in the real world About This Book Understand which algorithms to use in a given context with the help of this exciting recipe-based guide Learn about perceptrons and see how they are used to build neural networks Stuck while making sense of images, text, speech, and real estate? This guide will come to your rescue, showing you how to perform machine learning for each one of these using various techniques Who This Book Is For This book is for Python programmers who are looking to use machine-learning algorithms to create real-world applications. This book is friendly to Python beginners, but familiarity with Python programming would certainly be useful to play around with the code. What You Will Learn Explore classification algorithms and apply them to the income bracket estimation problem Use predictive modeling and apply it to real-world problems Understand how to perform market segmentation using unsupervised learning Explore data visualization techniques to interact with your data in diverse ways Find out how to build a recommendation engine Understand how to interact with text data and build models to analyze it Work with speech data and recognize spoken words using Hidden Markov Models Analyze stock market data using Conditional Random Fields Work with image data and build systems for image recognition and biometric face recognition Grasp how to use deep neural networks to build an optical character recognition system In Detail Machine learning is becoming increasingly pervasive in the modern data-driven world. It is used extensively across many fields such as search engines, robotics, self-driving cars, and more. With this book, you will learn how to perform various machine learning tasks in different environments. We'll start by exploring a range of real-life scenarios where machine learning can be used, and look at various building blocks. Throughout the book, you'll use a wide variety of machine learning algorithms to solve real-world problems and use Python to implement these algorithms. You'll discover how to deal with various types of data and explore the differences between machine learning paradigms such as supervised and unsupervised learning. We also cover a range of regression techniques, classification algorithms, predictive modeling, data visualization techniques, recommendation engines, and more with the help of real-world examples. Style and approach You will explore various real-life scenarios in this book where machine learning can be used, and learn about different building blocks of machine learning using independent recipes in the book.