

[Books] Machine Learning With Python Cookbook Practical Solutions From Preprocessing To Deep Learning

This is likewise one of the factors by obtaining the soft documents of this **machine learning with python cookbook practical solutions from preprocessing to deep learning** by online. You might not require more time to spend to go to the book instigation as without difficulty as search for them. In some cases, you likewise attain not discover the publication machine learning with python cookbook practical solutions from preprocessing to deep learning that you are looking for. It will no question squander the time.

However below, once you visit this web page, it will be as a result agreed simple to acquire as skillfully as download lead machine learning with python cookbook practical solutions from preprocessing to deep learning

It will not believe many times as we accustom before. You can reach it even if feign something else at house and even in your workplace. fittingly easy! So, are you question? Just exercise just what we pay for under as without difficulty as evaluation **machine learning with python cookbook practical solutions from preprocessing to deep learning** what you as soon as to read!

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

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

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.

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

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.

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

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.

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

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.

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.

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.

Python Machine Learning By Example-Yuxi (Hayden) Liu 2017-05-31 Take tiny steps to enter the big world of data science through this interesting guide About This Book Learn the fundamentals of machine learning and build your own intelligent applications Master the art of building your own machine learning systems with this example-based practical guide Work with important classification and regression algorithms and other machine learning techniques Who This Book Is For This book is for anyone interested in entering the data science stream with machine learning. Basic familiarity with Python is assumed. What You Will Learn Exploit the power of Python to handle data extraction, manipulation, and exploration techniques Use Python to visualize data spread across multiple dimensions and extract useful features Dive deep into the world of analytics to predict situations correctly Implement machine learning classification and regression algorithms from scratch in Python Be amazed to see the algorithms in action Evaluate the performance of a machine learning model and optimize it Solve interesting real-world problems using machine learning and Python as the journey unfolds In Detail Data science and machine learning are some of the top buzzwords in the technical world today. A resurging interest in machine learning is due to the same factors that have made data mining and Bayesian analysis more popular than ever. This book is your entry point to machine learning. This book starts with an introduction to machine learning and the Python language and shows you

how to complete the setup. Moving ahead, you will learn all the important concepts such as, exploratory data analysis, data preprocessing, feature extraction, data visualization and clustering, classification, regression and model performance evaluation. With the help of various projects included, you will find it intriguing to acquire the mechanics of several important machine learning algorithms - they are no more obscure as they thought. Also, you will be guided step by step to build your own models from scratch. Toward the end, you will gather a broad picture of the machine learning ecosystem and best practices of applying machine learning techniques. Through this book, you will learn to tackle data-driven problems and implement your solutions with the powerful yet simple language, Python. Interesting and easy-to-follow examples, to name some, news topic classification, spam email detection, online ad click-through prediction, stock prices forecast, will keep you glued till you reach your goal. Style and approach This book is an enticing journey that starts from the very basics and gradually picks up pace as the story unfolds. Each concept is first succinctly defined in the larger context of things, followed by a detailed explanation of their application. Every concept is explained with the help of a project that solves a real-world problem, and involves hands-on work—giving you a deep insight into the world of machine learning. With simple yet rich language—Python—you will understand and be able to implement the examples with ease.

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.

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.

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

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 Features Apply popular machine learning algorithms using a recipe-based approach Implement boosting, bagging, and stacking ensemble methods to improve machine learning models Discover real-world ensemble applications and encounter complex challenges in Kaggle competitions Book 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 learn Understand how to use machine learning algorithms for regression and classification problems Implement ensemble techniques such as averaging, weighted averaging, and max-voting Get to grips with advanced ensemble methods, such as bootstrapping, bagging, and stacking Use Random Forest for tasks such as classification and regression Implement an ensemble of homogeneous and heterogeneous machine learning algorithms Learn and implement various boosting techniques, such as AdaBoost, Gradient Boosting Machine, and XGBoost Who 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 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

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 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.

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

Deep Learning with Python-Francois Chollet 2017-10-28 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 Francois 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 Francois 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 Francois 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

Scikit-Learn Cookbook - Second Edition-Julian Avila 2017-11-15 Learn to use scikit-learn operations and functions for Machine Learning and deep learning applications.About This Book* Handle a variety of machine learning tasks effortlessly by leveraging the power of scikit-learn* Perform supervised and unsupervised learning with ease, and evaluate the performance of your model* Practical, easy to understand recipes aimed at helping you choose the right machine learning algorithmWho This Book Is ForData Analysts already familiar with Python but not so much with scikit-learn, who want quick solutions to the common machine learning problems will find this book to be very useful. If you are a Python programmer who wants to take a dive into the world of machine learning in a practical manner, this book will help you too.What You Will Learn* Build predictive models in minutes by using scikit-learn* Understand the differences and relationships between Classification and Regression, two types of Supervised Learning.* Use distance metrics to predict in Clustering, a type of Unsupervised Learning* Find points with similar characteristics with Nearest Neighbors.* Use automation and cross-validation to find a best model and focus on it for a data product* Choose among the best algorithm of many or use them together in an ensemble.* Create your own estimator with the simple syntax of sklearn* Explore the feed-forward neural networks available in scikit-learnIn DetailPython is quickly becoming the go-to language for analysts and data scientists due to its simplicity and flexibility, and within the Python data space, scikit-learn is the unequivocal choice for machine learning. This book includes walk throughs and solutions to the common as well as the not-so-common problems in machine learning, and how scikit-learn can be leveraged to perform various machine learning tasks effectively.The second edition begins with taking you through recipes on evaluating the statistical properties of data and generates synthetic data for machine learning modelling. As you progress through the chapters, you will come across recipes that will teach you to implement techniques like data pre-processing, linear regression, logistic regression, K-NN, Naive Bayes, classification, decision trees, Ensembles and much more. Furthermore, you'll learn to optimize your models with multi-class classification, cross validation, model evaluation and dive deeper in to implementing deep learning with scikit-learn. Along with covering the enhanced features on model section, API and new features like classifiers, regressors and estimators the book also contains recipes on evaluating and fine-tuning the performance of your model.By the end of this book, you will have explored plethora of features offered by scikit-learn for Python to solve any machine learning problem you come across.Style and ApproachThis book consists of practical recipes on scikit-learn that target novices as well as intermediate users. It goes deep into the technical issues, covers additional protocols, and many more real-live examples so that you are able to implement it in your daily life scenarios.

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.

Deep Learning with R Cookbook-Swarna Gupta 2020-02-21 Tackle the

complex challenges faced while building end-to-end deep learning models using modern R libraries

Key Features Understand the intricacies of R deep learning packages to perform a range of deep learning tasks Implement deep learning techniques and algorithms for real-world use cases Explore various state-of-the-art techniques for fine-tuning neural network models

Book Description Deep learning (DL) has evolved in recent years with developments such as generative adversarial networks (GANs), variational autoencoders (VAEs), and deep reinforcement learning. This book will get you up and running with R 3.5.x to help you implement DL techniques. The book starts with the various DL techniques that you can implement in your apps. A unique set of recipes will help you solve binomial and multinomial classification problems, and perform regression and hyperparameter optimization. To help you gain hands-on experience of concepts, the book features recipes for implementing convolutional neural networks (CNNs), recurrent neural networks (RNNs), and Long short-term memory (LSTMs) networks, as well as sequence-to-sequence models and reinforcement learning. You'll then learn about high-performance computation using GPUs, along with learning about parallel computation capabilities in R. Later, you'll explore libraries, such as MXNet, that are designed for GPU computing and state-of-the-art DL. Finally, you'll discover how to solve different problems in NLP, object detection, and action identification, before understanding how to use pre-trained models in DL apps. By the end of this book, you'll have comprehensive knowledge of DL and DL packages, and be able to develop effective solutions for different DL problems. What you will learn

Work with different datasets for image classification using CNNs Apply transfer learning to solve complex computer vision problems Use RNNs and their variants such as LSTMs and Gated Recurrent Units (GRUs) for sequence data generation and classification Implement autoencoders for DL tasks such as dimensionality reduction, denoising, and image colorization Build deep generative models to create photorealistic images using GANs and VAEs Use MXNet to accelerate the training of DL models through distributed computing

Who this book is for This deep learning book is for data scientists, machine learning practitioners, deep learning researchers and AI enthusiasts who want to learn key tasks in deep learning domains using a recipe-based approach. A strong understanding of machine learning and working knowledge of the R programming language is mandatory.

Python Machine Learning-Wei-Meng Lee 2019-04-04 Python makes machine learning easy for beginners and experienced developers With computing power increasing exponentially and costs decreasing at the same time, there is no better time to learn machine learning using Python. Machine learning tasks that once required enormous processing power are now possible on desktop machines. However, machine learning is not for the faint of heart—it requires a good foundation in statistics, as well as programming knowledge. Python Machine Learning will help coders of all levels master one of the most in-demand programming skillsets in use today. Readers will get started by following fundamental topics such as an introduction to Machine Learning and Data Science. For each learning algorithm, readers will use a real-life scenario to show how Python is used to solve the problem at hand.

- Python data science—manipulating data and data visualization
- Data cleansing
- Understanding Machine learning algorithms
- Supervised learning algorithms
- Unsupervised learning algorithms
- Deploying machine learning models

Python Machine Learning is essential reading for students, developers, or anyone with a keen interest in taking their coding skills to the next level.

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.

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.

Practical Machine Learning Cookbook-Atul Tripathi 2017-04-14 Resolving and offering solutions to your machine learning problems with R

About This Book Implement a wide range of algorithms and techniques for tackling complex data Improve predictions and recommendations to have better levels of accuracy Optimize performance of your machine-learning systems

Who This Book Is For This book is for analysts, statisticians, and data scientists with knowledge of fundamentals of machine learning and statistics, who need help in dealing with challenging scenarios faced every day of working in the field of machine learning and improving system performance and accuracy. It is assumed that as a reader you have a good understanding of mathematics. Working knowledge of R is expected. What You Will Learn Get equipped with a deeper understanding of how to apply machine-learning techniques Implement each of the advanced machine-learning techniques Solve real-life problems that are encountered in order to make your applications produce improved results Gain hands-on experience in problem solving for your machine-learning systems Understand the methods of collecting data, preparing data for usage, training the model, evaluating the model's performance, and improving the model's performance

In Detail Machine learning has become the new black. The challenge in today's world is the explosion of data from existing legacy data and incoming new structured and unstructured data. The complexity of discovering, understanding, performing analysis, and predicting outcomes on the data using machine learning algorithms is a challenge. This cookbook will help solve everyday challenges you face as a data scientist. The application of various data science techniques and on multiple data sets based on real-world challenges you face will help you appreciate a variety of techniques used in various situations. The first half of the book provides recipes on fairly complex machine-learning systems, where you'll learn to explore new areas of applications of machine learning and improve its efficiency. That includes recipes on classifications, neural networks, unsupervised and supervised learning, deep learning, reinforcement

learning, and more. The second half of the book focuses on three different machine learning case studies, all based on real-world data, and offers solutions and solves specific machine-learning issues in each one. Style and approach Following a cookbook approach, we'll teach you how to solve everyday difficulties and struggles you encounter.

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.

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

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

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.

Modern Python Cookbook-Steven F. Lott 2016-11-30 The latest in modern Python recipes for the busy modern programmer About This Book Develop succinct, expressive programs in Python Learn the best practices and common idioms through carefully explained and structured recipes Discover new ways to apply Python for the new age of development Who This Book Is For The book is for web developers, programmers, enterprise programmers, engineers, big data scientist, and so on. If you are a beginner, Python Cookbook will get you started. If you are experienced, it will expand your knowledge base. A basic knowledge of programming would help. What You Will Learn See the intricate details of the Python syntax and how to use it to your advantage Improve your code readability through functions in Python Manipulate data effectively using built-in data structures Get acquainted with advanced programming techniques in Python Equip yourself with functional and statistical programming features Write proper tests to be sure a program works as advertised Integrate application software using Python In Detail Python is the preferred choice of developers, engineers, data scientists, and hobbyists everywhere. It is a great scripting language that can power your applications and provide great speed, safety, and scalability. By exposing Python as a series of simple recipes, you can gain insight into specific language features in a particular context. Having a tangible context helps make the language or standard library feature easier to understand. This book comes with over 100 recipes on the latest version of Python. The recipes will benefit everyone ranging from beginner to an expert. The book is broken down into 13 chapters that build from simple language concepts to more complex applications of the language. The recipes will touch upon all the necessary Python concepts related to data structures, OOP, functional programming, as well as statistical programming. You will get acquainted with the nuances of Python syntax and how to effectively use the advantages that it offers. You will end the book equipped with the knowledge of testing, web services, and configuration and application integration tips and tricks. The recipes take a problem-solution approach to resolve issues commonly faced by Python programmers across the globe. You will be armed with the knowledge of creating applications with flexible logging, powerful configuration, and command-line options, automated unit tests, and good documentation. Style and approach This book takes a recipe-based approach, where each recipe addresses specific problems and issues. The recipes provide discussions and insights and an explanation of the problems.

Bioinformatics with Python Cookbook-Tiago Antao 2015-06-25 If you are either a computational biologist or a Python programmer, you will probably relate to the expression "explosive growth, exciting times". Python is arguably the main programming language for big data, and the deluge of data in biology, mostly from genomics and proteomics, makes bioinformatics one of the most exciting fields in data science. Using the hands-on recipes in this book, you'll be able to do practical research and analysis in computational biology with Python. We cover modern, next-generation sequencing libraries and explore real-world examples on how to handle real data. The main focus of the book is the practical application of bioinformatics, but we also cover modern programming techniques and frameworks to deal with the ever increasing deluge of bioinformatics data.

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.

Python Data Science Cookbook-Gopi Subramanian 2015-11-16 Over 60 practical recipes to help you explore Python and its robust data science capabilities About This Book The book is packed with simple and concise Python code examples to effectively demonstrate advanced concepts in action Explore concepts such as programming, data mining, data analysis, data visualization, and machine learning using Python Get up to speed on machine learning algorithms with the help of easy-to-follow, insightful recipes Who This Book Is For This book is intended for all levels of Data Science professionals, both students and practitioners, starting from novice to experts. Novices can spend their time in the first five chapters getting themselves acquainted with Data Science. Experts can refer to the chapters starting from 6 to understand how advanced techniques are implemented using Python. People from non-Python backgrounds can also effectively use this book, but it would be helpful if you have some prior basic programming experience. What You Will Learn Explore the complete range of Data Science algorithms Get to know the tricks used by industry engineers to create the most accurate data science models Manage and use Python libraries such as numpy, scipy, scikit learn, and matplotlib effectively Create meaningful features to solve real-world problems Take a look at Advanced Regression methods for model building and variable selection Get a thorough understanding of the underlying concepts and implementation of Ensemble methods Solve real-world problems using a variety of different datasets from numerical and text data modalities Get accustomed to modern state-of-the-art algorithms such as Gradient Boosting, Random Forest, Rotation Forest, and so on In Detail Python is increasingly becoming the language for data science. It is overtaking R in terms of adoption, it is widely known by many developers, and has a strong set of libraries such as Numpy, Pandas, scikit-learn, Matplotlib, Ipython and Scipy, to support its usage in this field. Data Science is the emerging new hot tech field, which is an amalgamation of different disciplines including statistics, machine learning, and computer science. It's a disruptive technology changing the face of today's business and altering the economy of various verticals including retail, manufacturing, online ventures, and hospitality, to name a few, in a big way. This book will walk you through the various steps, starting from simple to the most complex algorithms available in the Data Science arsenal, to effectively mine data and derive intelligence from it. At every step, we provide simple and efficient Python recipes that will not only show you how to implement these algorithms, but also clarify the underlying concept thoroughly. The book begins by introducing you to using Python for Data Science, followed by working with Python environments. You will then learn how to analyse your data with Python. The book then teaches you the concepts of data mining followed by an extensive coverage of machine learning methods. It introduces you to a number of Python libraries available to help implement machine learning and data mining routines effectively. It also covers the principles of shrinkage, ensemble methods, random forest, rotation forest, and extreme trees, which are a must-have for any successful Data Science Professional. Style and approach This is a step-by-step recipe-based approach to Data Science algorithms, introducing the math philosophy behind these algorithms.

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

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 FOR 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. TABLE OF CONTENTS 1. Boston Crime 2. World Happiness Report 3. Iris Species 4. Credit Card Fraud Detection 5. Heart Disease UCI

Building REST APIs with Flask-Kunal Relan 2019-09-12 Develop RESTful web services using the Flask micro-framework and integrate them using MySQL. Use Flask to develop, deploy, and manage REST APIs with easy-to-read and understand Python code. Solve your problem from a choice of libraries. Learn to use MySQL as the web services database for your Flask API using SQLAlchemy ORM. Building REST APIs with Flask provides a primer on Flask, RESTful services, and working with pip to set up your virtual environment. The key differences between NoSQL and SQL are covered, and you are taught how to connect MySQL and Flask using SQLAlchemy. Author Kunal Relan presents best practices for creating REST APIs and guides you in structuring your app and testing REST endpoints. He teaches you how to set up authentication and render HTML using views. You learn how to write unit tests for your REST APIs, and understand mocks, assertions, and integration testing. You will know how to document your REST APIs, deploy your Flask application on all of the major cloud platforms, and debug and monitor your Flask application. What You'll Learn Use MySQL to create Flask REST APIs Test REST endpoints Create CRUD endpoints with Flask and MySQL Deploy Flask on all of the major cloud platforms Monitor your Flask application Who This Book Is For Python developers interested in REST API development using Flask and web developers with basic programming knowledge who want to learn how Python and REST APIs work together. Readers should be familiar with Python (command line, or at least pip) and MySQL.

Cisco ACI Cookbook-Stuart Fordham 2017-05-25 Over 90 recipes to maximize automated solutions and policy-drive application profiles using Cisco ACI About This Book Confidently provision your virtual and physical infrastructure for application deployment Integrate Cisco ACI with hypervisors and other third party devices Packed with powerful recipes to automate your IT operations Who This Book Is For If you are a network administrator, system administrator, or engineer and are aware of the basics of Cisco ACI but want to start using it to automate your tasks, then this book is for you What You Will Learn Master the Cisco ACI architecture Discover the ACI fabric with easy-to-follow steps Set up quality of service within ACI Configure external networks with Cisco ACI Integrate with VMware and track VMware virtual machines Configure apply and verify access policies Extend or migrate a VMware virtual-machine LAN inside the ACI fabric Monitor ACI with third party tools and troubleshoot issues In Detail Cisco Application Centric Infrastructure (ACI) is a tough architecture that automates IT tasks and accelerates data-center application deployments. This book focuses on practical recipes to help you quickly

build, manage, and customize hybrid environment for your organization using Cisco ACI. You will begin by understanding the Cisco ACI architecture and its major components. You will then configure Cisco ACI policies and tenants. Next you will connect to hypervisors and other third-party devices. Moving on, you will configure routing to external networks and within ACI tenants and also learn to secure ACI through RBAC. Furthermore, you will understand how to set up quality of service and network programming with REST, XML, Python and so on. Finally you will learn to monitor and troubleshoot ACI in the event of any issues that arise. By the end of the book, you will gain have mastered automating your IT tasks and accelerating the deployment of your applications. Style and approach A set of exciting recipes to automate your IT operations related to datacenters, the Cloud, and networking tasks