

Machine Learning With Python Cookbook: Practical Solutions From Preprocessing To Deep Learning

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 large 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 learn how to solve the credit card fraud and default problems using advanced classification techniques 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 learn on learning how to use deep learning (PyTorch) for approaching financial tasks. At 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 the 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 this book effectively.

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 such as classifying and generating text, images, and music. Each chapter consists of several recipes needed to complete a single project, such as training a music recommendation system. Author Douwe Osinga also provides a chapter with half a dozen techniques

Access PDF Machine Learning With Python Cookbook: Practical Solutions From Preprocessing To Deep Learning

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 recommendation 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 can generate icons Detect music styles and index song collections

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 graphically 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, and work with K-Nearest Neighbors, neural networks, clustering, and other algorithms

Description This book provides the concept of machine learning with mathematical explanation and programming examples. Every chapter starts with fundamentals of a technique and working example on the real-world dataset. Along with the advice on applying algorithms, each technique is provided with advantages and disadvantages of the data. In this book we provide code examples in Python. Python is the most widely accepted language for this. First, it is free and open source. It contains good support from open community. It contains a lot of library, so you don't need everything. Also, it is scalable for large amount of data and suitable for big data technologies. This book: Covers all major areas in Machine Learning. Topics are discussed with graphical explanations. Comparison of different Machine Learning methods to solve any problem. Methods to handle real-world noisy data before applying any Machine Learning algorithm. Python code example for each concept discussed. Jupyter notebook scripts are provided with dataset used to test and compare algorithms Contents Introduction to Machine Learning Understanding Python Feature Engineering Data Visualisation Basic and Advanced Regression techniques Classification Un Supervised Learning Text Analysis Neural Network and Deep Learning Recommendation System Time Series Analysis

Discover powerful ways to effectively solve real-world machine learning problems using key libraries including scikit-learn, TensorFlow, and PyTorch Key Features Learn how to implement machine learning algorithms in a variety of real-life scenarios Cover a wide range of tasks catering to supervised, unsupervised and reinforcement learning techniques easy-to-follow code solutions for tackling common and not-so-common challenges 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

Access PDF Machine Learning With Python Cookbook: Practical Solutions From Preprocessing To Deep Learning

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 this 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 preprocessing techniques to clean and prepare data for machine learning
- Use predictive modeling and apply it to real-world problems
- Explore data visualization techniques to understand and 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 build models to 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 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, this book is what you need. Familiarity with Python programming and machine learning concepts will be useful.

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 learn with the best tools to streamline your feature engineering pipelines and techniques to 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, polynomial 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. You will learn

- Simplify your feature engineering pipelines with powerful Python pandas
- Get to grips with imputing missing values
- Encode categorical variables with a wide range of techniques

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 in understanding the concepts covered in this book.

If you're a data scientist already familiar with Python but not Scikit-Learn, or are familiar with other programming languages like R and want to take the plunge with the gold standard of Python machine learning libraries, then this is the book for you. Every other day we hear about new ways to put deep learning to good use: improved medical imaging, accurate credit card fraud detection, long range weather forecasting, and more. PyTorch puts these superpowers in your hands, providing a comfortable Python experience that gets you started quickly and then grows with you as your deep learning skills—become more sophisticated. Deep Learning with PyTorch will make that journey engaging and fun. Summary Every other day we hear about new ways to put deep learning to good use: improved medical imaging, accurate credit card fraud detection, long range weather forecasting, and more. PyTorch puts these superpowers in your hands, providing a comfortable Python experience that gets you started quickly and then grows with you as you—and your deep learning skills—become more sophisticated. Deep Learning with PyTorch will make that journey engaging and fun. Foreword by Soumith Chintala, Cocreator of PyTorch. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Although many deep learning tools use Python, the PyTorch library is Pythonic. Instantly familiar to anyone who knows PyData tools like NumPy and SciPy, PyTorch simplifies deep learning without sacrificing advanced features. It's excellent for building quick models, and it scales smoothly from laptop to enterprise. Because companies like Apple, Facebook, and JPMorgan Chase rely on PyTorch, it's a great skill to have as you expand your career options. It's easy to get started with PyTorch. It minimizes cognitive overhead without sacrificing the access to advanced features, meaning you can focus on what matters the most - building and training the latest and greatest deep learning models and contribute to making a dent in the world. PyTorch is also a snap to scale and extend, and it partners well with other Python tooling. PyTorch has been adopted by hundreds of deep learning practitioners and several first-class players like FAIR, OpenAI, FastAI and Purdue. About the book Deep Learning with PyTorch teaches you to create neural networks and deep learning models with PyTorch. This practical book quickly gets you to work building a real-world example from scratch: a tumor image classifier. Along the way, it covers best practices for the entire DL pipeline, including the PyTorch Tensor API, loading data in Python, monitoring training, and visualizing results. After covering the basics, the book will take you on a journey through larger projects. The centerpiece of the book is a neural network designed for cancer detection. You'll discover ways for training networks with li

Access PDF Machine Learning With Python Cookbook: Practical Solutions From Preprocessing To Deep Learning

inputs and start processing data to get some results. You'll sift through the uninitial results and focus on how to diagnose and fix the problems in your neural. Finally, you'll look at ways to improve your results by training with augmented data, make improvements to the model architecture, and perform other fine tuning. What's inside Training deep neural networks Implementing modules and loss functions Using pretrained models from PyTorch Hub Exploring code samples in Jupyter Notebooks About the reader For Python programmers with an interest in machine learning. the author Eli Stevens had roles from software engineer to CTO, and is currently on machine learning in the self-driving-car industry. Luca Antiga is cofounder of engineering company and an AI tech startup, as well as a former PyTorch contributor. Thomas Viehmann is a PyTorch core developer and machine learning trainer and consultant. consultant based in Munich, Germany and a PyTorch core developer.

of Contents PART 1 - CORE PYTORCH 1 Introducing deep learning and the PyTorch Library 2 Pretrained networks 3 It starts with a tensor 4 Real-world data representation using tensors 5 The mechanics of learning 6 Using a neural network to fit the data Telling birds from airplanes: Learning from images 8 Using convolutions to generate PART 2 - LEARNING FROM IMAGES IN THE REAL WORLD: EARLY DETECTION OF LUNG CANCER 9 Using PyTorch to fight cancer 10 Combining data sources into a unified dataset 11 Training a classification model to detect suspected tumors 12 Improving training with metrics and augmentation 13 Using segmentation to find suspected nodules 14 End-to-end nodule analysis, and where to go next PART 3 - DEPLOYMENT 15 Deploying to production

[TensorFlow 1.x Deep Learning Cookbook](#)

[Machine Learning with Python for Everyone](#)

[Over 70 recipes for creating, engineering, and transforming features to build machine learning models](#)

[Recipes for Mastering Python 3](#)

[Introducing Data Science](#)

[Over 80 recipes on how to implement machine learning algorithms for building systems using Python](#)

[Python Deep Learning Cookbook](#)

[Python Machine Learning By Example](#)

[Machine Learning for Time Series Forecasting with Python](#)

[Machine Learning with Python Cookbook](#)

[Practical Solutions from Preprocessing to Deep Learning](#)

[Machine Learning Cookbook with Python](#)

[A Guide for Data Scientists](#)

This book focuses on expert-level explanations and implementations of scalable reinforcement learning algorithms and approaches. Starting with the fundamentals, the book covers state-of-the-art methods from bandit problems to meta-reinforcement learning. You'll also explore practical examples inspired by real-life problems from the industry.

If you need help writing programs in Python 3, or want to update older Python 2

Access PDF Machine Learning With Python Cookbook: Practical Solutions From Preprocessing To Deep Learning

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

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), naive Bayes, clustering, and neural networks Saving and loading trained models

The book shows you how to view data from multiple perspectives, including data frame and column attributes. You will cover common and not-so-common challenges that are faced while cleaning messy data for complex situations. You learn to manipulate data and get them down to a form that can be useful for making the right decisions.

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.

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.

Unlock the groundbreaking advances of deep learning with this extensively revised new edition of the bestselling original. Learn directly from the creator of Keras a master practical Python deep learning techniques that are easy to apply in the real world. In Deep Learning with Python, Second Edition you will learn: Deep learning from first principles Image classification and image segmentation Timeseries forecasting Text classification and machine translation Text generation, neural style transfer, and image generation Deep Learning with Python has taught thousands of readers how to put the full capabilities of deep learning into action. This extensively revised second edition introduces deep learning using Python and Keras, and is loaded with insights for both novice and experienced ML practitioners. You'll learn practical techniques that are easy to apply in the real world, and important theory for perfecting neural networks. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology Recent innovations in deep learning unlock exciting new software capabilities like

automated language translation, image recognition, and more. Deep learning is quickly becoming essential knowledge for every software developer, and modern tools like Keras and TensorFlow put it within your reach—even if you have no background in mathematics or data science. This book shows you how to get started. About the book *Deep Learning with Python, Second Edition* introduces the field of deep learning using Python and the powerful Keras library. In this revised and expanded new edition, Keras creator François Chollet offers insights for both novice and experienced machine learning practitioners. As you move through this book, you'll build your understanding through intuitive explanations, crisp illustrations, and clear examples. You'll quickly pick up the skills you need to start developing deep-learning applications. What's inside

- Deep learning from first principles
- Image classification and image segmentation
- Time series forecasting
- Text classification
- machine translation
- Text generation, neural style transfer, and image generation

About the reader For readers with intermediate Python skills. No previous experience with Keras, TensorFlow, or machine learning is required. About the author François Chollet is a software engineer at Google and creator of the Keras deep-learning library. Table of Contents

- 1 What is deep learning?
- 2 The mathematical building blocks of neural networks
- 3 Introduction to Keras and TensorFlow
- 4 Getting started with neural networks: Classification and regression
- 5 Fundamentals of machine learning
- 6 The universal workflow of machine learning
- 7 Working with Keras: A deep dive
- 8 Introduction to deep learning for computer vision
- 9 Advanced deep learning for computer vision
- 10 Deep learning for timeseries
- 11 Deep learning for text
- 12 Generative deep learning
- 13 Best practices for the real world
- 14 Conclusions

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

Access PDF Machine Learning With Python Cookbook: Practical Solutions From Preprocessing To Deep Learning

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 solution. 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, 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 Cookbook](#)

[Concepts, Tools, and Techniques to Build Intelligent Systems](#)

[Modern techniques and Python tools to detect and remove dirty data and extract key insights](#)

[Deep Learning Cookbook](#)

[Practical Recipes to Get Started Quickly](#)

[TensorFlow Machine Learning Cookbook](#)

[Keras Deep Learning Cookbook](#)

[Scikit-Learn Cookbook](#)

[Machine Learning Using TensorFlow Cookbook](#)

[Python Machine Learning Cookbook](#)

[Over 90 unique recipes to solve artificial-intelligence driven problems with Python](#)

[Deep Learning with Python](#)

[Essential Tools for Working with Data](#)

This book presents practical solutions to the most common reinforcement learning problems. The recipes in this book will help you understand the fundamental concepts to develop popular RL algorithms. You will gain practical experience in the RL domain using the modern offerings of the PyTorch 1.x library.

Summary Introducing Data Science teaches you how to accomplish the fundamental tasks that occupy data scientists. Using the Python language and common Python libraries, you'll experience firsthand the challenges of dealing with data at scale and gain a solid foundation in data science.

Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Many companies need developers with data science skills to work on projects ranging from social media marketing to machine learning. Discovering what you need to learn to begin a career as a data scientist can seem bewildering. This book is designed to help you get started. About the Book Introducing Data Science Introducing Data Science explains vital data

Access PDF Machine Learning With Python Cookbook: Practical Solutions From Preprocessing To Deep Learning

science concepts and teaches you how to accomplish the fundamental tasks that occupy data scientists. You'll explore data visualization, graph databases, the use of NoSQL, and the data science process. You'll use the Python language and common Python libraries as you experience firsthand the challenges of dealing with data at scale. Discover how Python allows you to gain insights from data sets so big that they need to be stored on multiple machines, or from data moving so quickly that no single machine can handle it. This book gives you hands-on experience with the most popular Python data science libraries, Scikit-learn and StatsModels. After reading this book, you'll have the solid foundation you need to start a career in data science. What's Inside Handling large data Introduction to machine learning Using Python to work with data Writing data science algorithms About the Reader This book assumes you're comfortable reading code in Python or a similar language, such as C, Ruby, or JavaScript. No prior experience with data science is required. About the Authors Davy Cielen, Arno D. B. Meysman, and Mohamed Ali are the founders and managing partners of Optimately and Maiton, where they focus on developing data science projects and solutions in various sectors. Table of Contents Data science in a big data world The data science process Machine learning Handling large data on a single computer First steps in big data Join the NoSQL movement The rise of graph databases Text mining and text analytics Data visualization to the end user

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. 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.

This book uses a recipe-based approach to showcase the power of machine learning algorithms to build ensemble models using Python libraries. Through this book, you will be able to pick up the code, understand in depth how it works, execute and implement it efficiently. This will be a desk reference to implement a wide range of tasks and solve ... Learn how to apply the principles of machine learning to time series modeling with this indispensable resource Machine Learning for Time Series Forecasting with Python is an incisive and straightforward examination of one of the most crucial elements of decision-making in finance, marketing, education, and healthcare: time series modeling. Despite the centrality of time series forecasting, few business analysts are

familiar with the power or utility of applying machine learning to time series modeling. Author Francesca Lazeri, a distinguished machine learning scientist and economist, corrects that deficiency by providing readers with comprehensive and approachable explanation and treatment of the application of machine learning to time series forecasting. Written for readers who have little to no experience in time series forecasting or machine learning, the book comprehensively covers all the topics necessary to: Understand time series forecasting concepts, such as stationarity, horizon, trend, and seasonality Prepare time series data for modeling Evaluate time series forecasting models' performance and accuracy Understand when to use neural networks instead of traditional time series models in time series forecasting Machine Learning for Time Series Forecasting with Python is full real-world examples, resources and concrete strategies to help readers explore and transform data and develop usable, practical time series forecasts. Perfect for entry-level data scientists, business analysts, developers, and researchers, this book is an invaluable and indispensable guide to the fundamental and advanced concepts of machine learning applied to time series modeling.

This book shows readers how they can successfully analyze data using only two core machine learning algorithms---and how to do so using the popular Python programming language. These algorithms deal with common scenarios faced by all data analysts and data scientists. This book focuses on two algorithm families (linear methods and ensemble methods) that effectively predict outcomes. This type of problem covers a multitude of use cases (what ad to place on a web page, predicting prices in securities markets, detecting credit card fraud, etc.). The focus on two families gives enough room for full descriptions of the mechanisms at work in the algorithms. Then the code examples serve to illustrate the workings of the machinery with specific hackable code. The author will explain in simple terms, using no complex math, how these algorithms work, and will then show how to apply them in Python. He will also provide advice on how to select from among these algorithms, and will show how to prepare the data, and how to use the trained models in practice. The author begins with an overview of the two core algorithms, explaining the types of problems solved by each one. He then introduces a core set of Python programming techniques that can be used to apply these algorithms. The author shows various techniques for building predictive models that solve a range of problems, from simple to complex; he also shows how to measure the performance of each model to ensure you use the right one. The following chapters provide a deep dive into each of the two algorithms: penalized linear regression and ensemble methods. Chapters will show how to apply each algorithm in Python. Readers can directly use the sample code to build their own solutions.

Access PDF Machine Learning With Python Cookbook: Practical Solutions From Preprocessing To Deep Learning

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

[Essential Techniques for Predictive Analysis](#)

[Over 50 recipes for applying modern Python libraries to financial data analysis](#)

[Build next-generation, self-learning models using reinforcement learning techniques and best practices](#)

[Introduction to Machine Learning with Python](#)

[Create ML and Data Analytics Projects Using Some Amazing Open Datasets](#)

[Python Feature Engineering Cookbook](#)

[Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow A Test-Driven Approach](#)

[Over 35 practical recipes to explore ensemble machine learning techniques using Python](#)

[Over 75 practical recipes on neural network modeling, reinforcement learning, and transfer learning using Python](#)

[IPython Interactive Computing and Visualization Cookbook](#)

[Create powerful machine learning algorithms with TensorFlow](#)

[Over 60 recipes to design, develop, and deploy self-learning AI models using Python](#)

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.

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

Access PDF Machine Learning With Python Cookbook: Practical Solutions From Preprocessing To Deep Learning

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

This book is designed to guide you through TensorFlow and how to use it effectively. Throughout the book, you will work through recipes and get hands-on experience to perform complex data computations, gain insights into your data, and more.

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.

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

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 algorithm Who This Book Is For Data 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-learn In Detail Python 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 Approach This 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.

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.

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

Access PDF Machine Learning With Python Cookbook: Practical Solutions From Preprocessing To Deep Learning

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.

[Machine Learning for Cybersecurity Cookbook](#)

[Deep Learning with PyTorch](#)

[PyTorch 1.x Reinforcement Learning Cookbook](#)

[Python Data Cleaning Cookbook](#)

[Python for Finance Cookbook](#)

[Ensemble Machine Learning Cookbook](#)

[Proven recipes for applying AI algorithms and deep learning techniques using TensorFlow 2.x and PyTorch 1.6](#)

[Practical Machine Learning with Python](#)

[Over 30 recipes for implementing deep neural networks in Python](#)

[Python Data Science Handbook](#)

[Artificial Intelligence with Python Cookbook](#)

[A Problem-Solver's Guide to Building Real-World Intelligent Systems](#)

[Scikit-Learn Cookbook - Second Edition](#)

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

If you are looking to build next-generation AI solutions for work or even for your pet projects, you'll find this cookbook useful. With the help of easy-to-follow recipes, this book will take you through the advanced AI and machine learning approaches and algorithms that are required to build smart models for problem-solving.

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

The 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.

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

Portable, powerful, and a breeze to use, Python is the popular open source object-oriented programming language used for both standalone programs and scripting applications. It is now being used by an increasing number of major organizations, including NASA and Google. Updated for Python 2.4, The Python Cookbook, 2nd Edition offers a wealth of useful code for all Python programmers, not just advanced practitioners. Like its predecessor, the new edition provides solutions to problems that Python programmers face everyday. It now includes over 200 recipes that range from simple tasks, such as working with dictionaries and list comprehensions, to complex tasks, such as monitoring a network and building a templating system. This revised version also includes new chapters on topics such as time, money, and metaprogramming. Here's a list of additional topics covered: Manipulating text Searching and sorting Working with files and the filesystem Object-oriented programming Dealing with threads and processes System administration Interacting with databases Creating user interfaces Network and web programming Processing XML Distributed programming Debugging and testing Another advantage of The Python Cookbook, 2nd Edition is its trio of authors--three well-known Python programming experts, who are highly visible on email lists and in newsgroups, and speak often at Python conferences. With scores of practical examples and pertinent background information, The Python Cookbook, 2nd Edition is the one source you need if you're looking to build efficient, flexible, scalable, and well-integrated systems.

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 FORÊ 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

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.

[Big data, machine learning, and more, using Python tools](#)

[Over 100 recipes to progress from smart data analytics to deep learning using real-world datasets, 2nd Edition](#)

[TensorFlow 2 Reinforcement Learning Cookbook](#)

[Over 50 recipes to help you build, train, and deploy learning agents for real-world applications](#)

[MACHINE LEARNING WITH PYTHON](#)

[Thoughtful Machine Learning with Python](#)

[Mastering Reinforcement Learning with Python](#)

[Machine Learning in Python](#)

[Deep Learning with Python, Second Edition](#)