


The book was found

Introduction To Machine Learning With Python: A Guide For Data Scientists



Synopsis

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 will 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 will 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

Book Information

Paperback: 394 pages

Publisher: O'Reilly Media; 1 edition (October 21, 2016)

Language: English

ISBN-10: 1449369413

ISBN-13: 978-1449369415

Product Dimensions: 7 x 0.8 x 9.2 inches

Shipping Weight: 1.5 pounds (View shipping rates and policies)

Average Customer Review: 4.2 out of 5 stars 23 customer reviews

Best Sellers Rank: #12,120 in Books (See Top 100 in Books) #3 in Books > Computers & Technology > Computer Science > AI & Machine Learning > Natural Language Processing #4 in Books > Textbooks > Computer Science > Algorithms #6 in Books > Computers & Technology > Programming > Algorithms

Customer Reviews

Andreas Müller received his PhD in machine learning from the University of Bonn. After working as a machine learning researcher on computer vision applications at for a year, he recently joined the Center for Data Science at the New York University. In the last four years, he has been

maintainer and one of the core contributors of scikit-learn, a machine learning toolkit widely used in industry and academia, and author and contributor to several other widely used machine learning packages. His mission is to create open tools to lower the barrier of entry for machine learning applications, promote reproducible science and democratize the access to high-quality machine learning algorithms. Sarah is a data scientist who has spent a lot of time working in start-ups. She loves Python, machine learning, large quantities of data, and the tech world. She is an accomplished conference speaker, currently resides in New York City, and attended the University of Michigan for grad school.

Intriguing topic -- easy reading with lots of samples.

A practical and to-the-point introduction to many Machine Learning Algorithms and their applications. Andreas responds to the issues and pull requests on the book's github repo same day.

Very helpful reference.

Worth the money, if you just getting started.

Good for beginner

Fantastic introduction to machine learning in Python. The examples are well written, and do a very nice job of introducing both the implementation and the concept for each model. I'm halfway thru the book, and am really enjoying it. I have a background in math and wrote software professionally for a number of years, but haven't spent much time doing either for the past 5-10 years. This book is technical enough to keep me interested, and accessible enough to allow me to ramp up on the language and the scikit framework. An added bonus - the instructions actually allowed me to set up my development environment, and the code in the book actually runs! 100% recommend for someone looking to get started in ML with Python.

My current work revolves around using machine learning for the study of criminal behavior, so I read Introduction to Machine Learning with Python by Andreas Muller and Sarah Guido with great interest. The book comprises a complete documentation of the scikit-learn library, and provides a comprehensive overview of the machine learning models and the fundamental theory needed to get

started in applying ML tools in practice. Each chapter contains Python source code that cover a wide range of interesting and practical data science problems. In addition to the basic theory, scikit-learn tools and code samples, the book also includes many useful hints, tricks and words of wisdom that can save you a lot of time by avoiding issues that invariably arise in your learning process. This is an excellent book that I highly recommend both to machine learning experts who want to be proficient in scikit-learn and also to beginners who want to learn machine learning basics and how to apply them on data. The concepts are clearly described and their implementation is presented through useful and exciting data science problems, giving the reader a clear understanding of how to apply the ML tools on real problems. The code is very well organized and structured, and ready to be used as reference and as a starting point for future projects. The flow of the book is constructed such that it can serve two purposes: it can be read to familiarize one with the machine learning techniques and how they are being applied on data without actually having to get into coding with Python, or it can be read as a ML course for those who want to learn ML with scikit-learn by studying the theory and applying it on real data problems throughout the reading process. Therefore, I recommend this book not only for Python users, but for anyone who wants to learn the basics of ML, and to see their applications without delving too deeply into ML theory and math. The book is a comprehensive self-contained literature for those who want to learn the basics of ML and to try them out on data.

This is a great book, and I'd say it is even great for those that are not familiar with python (you just obviously won't be able to run the code). For anyone with some basic understanding of linear algebra/statistics, the authors are able to present to you all the important (and sometimes subtle but significant) details, without the usage of equations, and more importantly, how they all relate to one another. All the concepts mentioned here are heavily backed with well thought of and well presented figures, in such a way that again I'd suggest you don't even need python to understand. If you do know python, loading the data sets and reproducing the figures is just a few lines of easy to understand code away (with the exception of the mglearn library includes which does some "plotting magic" for you. However, I believe each of them were appropriate. You can ignore them and make the plots in your own way, or just print the variables, it just may not look as publication friendly). Normally, I hesitate purchasing books that claim they may explain algorithms without the need of equations, and I expect them rather to be cook books of lightly and disjointly explained techniques (like an encyclopedia). However, I do not think such is true of this book. The power of scikit-learn is demonstrated and the algorithms behind them explained intuitively, and are referred

as to how they fit together and complement each other. As with any introductory read, a supplement is needed from time to time and the authors' reference to Elements of Statistical Learning is a useful one (equation heavy). There are points in the book where the author defers to elements of statistical learning. I found these points suitable since further explanation would be out of scope. I read this book on my free time while on vacation, and much of the time I didn't have access to a computer. The concepts were so well presented that it was just a nice leisurely read. When I finally had time to access a computer, I was able to try the techniques on my data sets with some browsing back and forth through the book again, but otherwise with little effort. Finally, since I myself am a researcher, I would recommend this book to any other researcher willing to start delving into the world of machine learning. Further reading will always be necessary, but this book will give you such a good intuitive understanding and overview of the subject matter that you'll know what to do to proceed next, and how to do it without running in circles. Even better, you'll likely already have applied it to your research!

[Download to continue reading...](#)

Python: The Complete Python Quickstart Guide (For Beginner's) (Python, Python Programming, Python for Dummies, Python for Beginners) Python: Programming: Your Step By Step Guide To Easily Learn Python in 7 Days (Python for Beginners, Python Programming for Beginners, Learn Python, Python Language) Hacking with Python: Beginner's Guide to Ethical Hacking, Basic Security, Penetration Testing, and Python Hacking (Python Programming, Hacking, Python Coding, Python and Hacking Book 3) PYTHON: PYTHON'S COMPANION, A STEP BY STEP GUIDE FOR BEGINNERS TO START CODING TODAY! (INCLUDES A 6 PAGE PRINTABLE CHEAT SHEET)(PYTHON FOR BEGINNERS, PYTHON FOR DUMMIES, PYTHON PROGRAMMING) Data Analytics and Python Programming: 2 Bundle Manuscript: Beginners Guide to Learn Data Analytics, Predictive Analytics and Data Science with Python Programming Introduction to Machine Learning with Python: A Guide for Data Scientists Python Programming: Python Programming for Beginners, Python Programming for Intermediates, Python Programming for Advanced PYTHON: LEARN PYTHON in A Day and MASTER IT WELL. The Only Essential Book You Need To Start Programming in Python Now. Hands On Challenges INCLUDED! (Programming for Beginners, Python) Statistics, Data Mining, and Machine Learning in Astronomy: A Practical Python Guide for the Analysis of Survey Data (Princeton Series in Modern Observational Astronomy) Unsupervised Machine Learning in Python: How to Find Distinct Patterns in Your Data Without Being at the Mercy of Data Labeling by Third-Party Workers Big Data For Business: Your Comprehensive Guide to Understand Data Science, Data Analytics and Data Mining to Boost More Growth and Improve

Business - Data Analytics Book, Series 2 Data Analytics: What Every Business Must Know About Big Data And Data Science (Data Analytics for Business, Predictive Analysis, Big Data Book 1) Data Analytics: Applicable Data Analysis to Advance Any Business Using the Power of Data Driven Analytics (Big Data Analytics, Data Science, Business Intelligence Book 6) Python Programming: An In-Depth Guide Into The Essentials Of Python Programming (Included: 30+ Exercises To Master Python in No Time!) Python: The Fundamentals Of Python Programming: A Complete Beginners Guide To Python Mastery. Python: Learn Python in a Day and Master It Well: The Only Essential Book You Need to Start Programming in Python Now Maya Python for Games and Film: A Complete Reference for Maya Python and the Maya Python API Reinforcement Learning with Python: An Introduction (Adaptive Computation and Machine Learning series) Markov Models: Understanding Data Science, Markov Models, and Unsupervised Machine Learning in Python Analytics: Data Science, Data Analysis and Predictive Analytics for Business (Algorithms, Business Intelligence, Statistical Analysis, Decision Analysis, Business Analytics, Data Mining, Big Data)

[Contact Us](#)

[DMCA](#)

[Privacy](#)

[FAQ & Help](#)