


The book was found

Python Machine Learning



Synopsis

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 Pylearn 2 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 Pylearn2, 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.

Book Information

Paperback: 454 pages

Publisher: Packt Publishing - ebooks Account (September 2015)

Language: English

ISBN-10: 1783555130

ISBN-13: 978-1783555130

Product Dimensions: 7.5 x 1 x 9.2 inches

Shipping Weight: 1.8 pounds (View shipping rates and policies)

Average Customer Review: 4.5 out of 5 stars [See all reviews](#) (73 customer reviews)

Best Sellers Rank: #1,884 in Books (See Top 100 in Books) [#1 in Books > Computers & Technology > Computer Science > AI & Machine Learning > Neural Networks](#) [#1 in Books > Computers & Technology > Databases & Big Data > Data Modeling & Design](#) [#2 in Books > Computers & Technology > Databases & Big Data > Data Processing](#)

Customer Reviews

First some general, higher-level thoughts and comments before I dive into specifics:MY

BACKGROUND:Data Scientist; B.S. in Economics and M.S. in Business Analytics; experienced

(though by no means expert) user of Scikit-learnOVERALL THOUGHTS:I've purchased and read

(virtually) every Machine Learning book that aims to teach the reader the basics of ML using the

Scikit-learn library as the main focus. I've found them to be...less than satisfactory. The examples in

other books often use ML techniques in contexts for which they are not intended to be used and/or

contexts they are not used in out in the real world (among other issues I have found within them).In

stark contrast, Python Machine Learning by Sebastian Raschka is stunningly-impressive, not only

for the breadth and depth of coverage, but also in the manner the information is presented to the

reader.To date, I have not encountered a book on ML that incorporates multiple levels of learning in

a manner such as this. It is the textual equivalent of a Neural Network with hundreds of hidden

layers running on the latest NVIDIA GPU (if that comparison is lost on you, don't worry; it'll all

make sense by the time you finish the book).One of the underlying (though understated) themes in

the book is the importance of using visual aids where appropriate to gauge the performance of the

algorithms you're using as well as to understand exactly what is going on behind the scenes,

so-to-speak. If you're a novice user of the Matplotlib graphics library for Python, this book will

greatly improve your visualization skills by the time you're done which I found to be an added

bonus.Another underlying theme is basic optimization using the NumPy library.

This is a fantastic book, even for a relative beginner to machine learning such as myself. The first thing that comes to mind after reading this book is that it was the perfect blend (for me at least) of theory and practice, as well as breadth and depth. Let's face it, we know that machine learning isn't an easy subject. You need theory but you also need practice in the form of some serious coding before you really start understanding it. And this is one area where Sebastian's book shines: it contains a plethora of really good code examples that are illuminating and well explained, and which cover a very wide range of different machine learning algorithms. And, speaking of code, as another reviewer has pointed out, another huge plus is that, in many places, Sebastian shows you how to gauge the performance of your code and make it more efficient. For me, the best measure of any book such as this is how many 'ah ha!' moments I had while reading it. And I had more than a few while reading Sebastian's book. One such 'ah ha!' moment came while reading chapter 12 (and this also illustrates that nice blend of theory and practice I already mentioned above). In this particular chapter, he discusses training artificial neural networks for image recognition. At the heart of this approach is back propagation, which is pretty much THE bread and butter behind multilayered neural networks. He presents a detailed discussion of back propagation in two separate pieces: one that is intuitive and 'top down'; the other a more mathematical, 'bottoms up' approach that goes through the algorithm step by step, showing how the gradients are computed and the weights updated. His treatment of back propagation was one of the better explanations I've seen and really cleared things up for me.

Technical, but not too much. Let's face it, machine learning algorithms are technical in nature. However, this book allows you to gloss over the actual technical details if you don't really need to understand them right away and view the implementation of the logic in the code snippets. Though, I must say, the presentation of the technical subjects are explained clearly and with supporting graphs and images to help visualize the concepts. It was a wonderful experience to understand the code, even though the theory was also given. This allows most people to jump right in and start writing in python. For the mathematicians out there, you can take the equations and verify them if need be. The ideas build upon each other and just like teaching a child to talk, the quality of machine learning seems to be about getting good training sets for your algorithms. As such, Sebastian is good about giving in-depth, best practice steps on how to make sure your training data is clean and normalized, as well as your feature selection is relevant - which was great. You'll learn how to merge results from multiple data sets into a more thorough model in order to filter out weaknesses of various algorithms. You'll be able to predict future outcomes using regression analysis using

techniques from statistics to look for patterns and anomalies, again, all explained in very understandable words. Though the content and speed of the book is all very good and relevant, the icing on the cake is in the last two chapters (which you need to have worked up to in the previous chapters); understanding and then creating a layered neural network to solve complex problems like hand written digit recognition. And to top it all off he teaches us how to make it more powerful using the Theano tool.

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

Python: Python Programming Course: Learn the Crash Course to Learning the Basics of Python (Python Programming, Python Programming Course, Python Beginners Course) Deep Learning: Recurrent Neural Networks in Python: LSTM, GRU, and more RNN machine learning architectures in Python and Theano (Machine Learning in Python) Unsupervised Deep Learning in Python: Master Data Science and Machine Learning with Modern Neural Networks written in Python and Theano (Machine Learning in Python) Deep Learning in Python Prerequisites: Master Data Science and Machine Learning with Linear Regression and Logistic Regression in Python (Machine Learning in Python) Convolutional Neural Networks in Python: Master Data Science and Machine Learning with Modern Deep Learning in Python, Theano, and TensorFlow (Machine Learning in Python) Deep Learning in Python: Master Data Science and Machine Learning with Modern Neural Networks written in Python, Theano, and TensorFlow (Machine Learning in Python) Python: Python Programming For Beginners - The Comprehensive Guide To Python Programming: Computer Programming, Computer Language, Computer Science (Machine Language) Unsupervised Machine Learning in Python: Master Data Science and Machine Learning with Cluster Analysis, Gaussian Mixture Models, and Principal Components Analysis A collection of Advanced Data Science and Machine Learning Interview Questions Solved in Python and Spark (II): Hands-on Big Data and Machine ... Programming Interview Questions) (Volume 7) Beginning Python Programming: Learn Python Programming in 7 Days: Treading on Python, Book 1 Python: Python Programming For Beginners - The Comprehensive Guide To Python Programming: Computer Programming, Computer Language, Computer Science Learn Python in One Day and Learn It Well: Python for Beginners with Hands-on Project. The only book you need to start coding in Python immediately Maya Python for Games and Film: A Complete Reference for Maya Python and the Maya Python API Deep Learning: Natural Language Processing in Python with GLoVe: From Word2Vec to GLoVe in Python and Theano (Deep Learning and Natural Language Processing) Deep Learning: Natural Language Processing in Python with Word2Vec: Word2Vec and Word Embeddings in Python and Theano (Deep Learning and Natural Language Processing Book 1)

Machine Learning: A Probabilistic Perspective (Adaptive Computation and Machine Learning series)
Machine Learning with Spark - Tackle Big Data with Powerful Spark Machine Learning Algorithms
Foundations of Machine Learning (Adaptive Computation and Machine Learning series) Introduction
to Machine Learning (Adaptive Computation and Machine Learning series) Gaussian Processes for
Machine Learning (Adaptive Computation and Machine Learning series)

[Dmca](#)