


The book was found

Large Scale Machine Learning With Python



Synopsis

Learn to build powerful machine learning models quickly and deploy large-scale predictive applications

About This Book

Design, engineer and deploy scalable machine learning solutions with the power of Python

Take command of Hadoop and Spark with Python for effective machine learning on a map reduce framework

Build state-of-the-art models and develop personalized recommendations to perform machine learning at scale

Who This Book Is For

This book is for anyone who intends to work with large and complex data sets. Familiarity with basic Python and machine learning concepts is recommended. Working knowledge in statistics and computational mathematics would also be helpful.

What You Will Learn

- Apply the most scalable machine learning algorithms
- Work with modern state-of-the-art large-scale machine learning techniques
- Increase predictive accuracy with deep learning and scalable data-handling techniques
- Improve your work by combining the MapReduce framework with Spark
- Build powerful ensembles at scale
- Use data streams to train linear and non-linear predictive models from extremely large datasets using a single machine

In Detail

Large Python machine learning projects involve new problems associated with specialized machine learning architectures and designs that many data scientists have yet to tackle. But finding algorithms and designing and building platforms that deal with large sets of data is a growing need. Data scientists have to manage and maintain increasingly complex data projects, and with the rise of big data comes an increasing demand for computational and algorithmic efficiency.

Large Scale Machine Learning with Python uncovers a new wave of machine learning algorithms that meet scalability demands together with a high predictive accuracy.

Dive into scalable machine learning and the three forms of scalability. Speed up algorithms that can be used on a desktop computer with tips on parallelization and memory allocation. Get to grips with new algorithms that are specifically designed for large projects and can handle bigger files, and learn about machine learning in big data environments. We will also cover the most effective machine learning techniques on a map reduce framework in Hadoop and Spark in Python.

Style and Approach

This efficient and practical title is stuffed full of the techniques, tips and tools you need to ensure your large scale Python machine learning runs swiftly and seamlessly.

Large-scale machine learning tackles a different issue to what is currently on the market. Those working with Hadoop clusters and in data intensive environments can now learn effective ways of building powerful machine learning models from prototype to production.

This book is written in a style that programmers from other languages (R, Julia, Java, Matlab) can follow.

Book Information

Paperback: 420 pages

Publisher: Packt Publishing - ebooks Account (August 3, 2016)

Language: English

ISBN-10: 1785887211

ISBN-13: 978-1785887215

Product Dimensions: 7.5 x 1 x 9.2 inches

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

Average Customer Review: 5.0 out of 5 stars (See all reviews) (2 customer reviews)

Best Sellers Rank: #389,874 in Books (See Top 100 in Books) #60 in Books > Computers & Technology > Computer Science > AI & Machine Learning > Neural Networks #64 in Books > Computers & Technology > Computer Science > AI & Machine Learning > Machine Theory #232 in Books > Computers & Technology > Databases & Big Data > Data Mining

Customer Reviews

Disclosure: I was a technical reviewer of this book. Many books when their subject is Machine Learning with Python concentrate on a few most known and used libraries to explain Machine Learning tasks and solutions. Although I don't want to say that such books are useless for readers, they may still leave gaps in understanding of how a certain method or library would work in real-world scenarios. Authors of the book "Large Scale Machine Learning with Python" set up an ambitious goal to teach readers how to solve real-world Machine Learning problems by employing a variety of libraries, frameworks, and tools relying on Python. This advantageously differentiates a given book from many other books on the same subject. The following practical situations are considered and their solutions are presented:- Tall datasets when the number of cases is large, compared to the number of features.- Wide datasets when the number of features is large, compared to the number of cases.- Both tall and wide datasets when both the number of features and the number of cases are large.- Sparse datasets when there are many zero-valued elements. The book treats the problem of scalability from different angles, such as fast batch (offline) processing, incremental online processing (one instance at a time arrives), streaming processing (a chunk of instances at a time arrives) and distributed processing. Popular libraries and frameworks, such as Gensim, H2O, XGBoost, TensorFlow, Theano, Theanets, Keras, Vowpal Wabbit, and Spark and their applications are explained through numerous Python snippets. In my opinion, this is one of the first books presenting all these tools under one cover.

[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) Large Scale Machine Learning with Python Rand McNally 2017 Large Scale Road Atlas (Rand McNally Large Scale Road Atlas USA) 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 Pocket Neighborhoods: Creating Small-Scale Community in a Large-Scale World 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