


The book was found

Machine Learning With R



Synopsis

R gives you access to the cutting-edge software you need to prepare data for machine learning. No previous knowledge required - this book will take you methodically through every stage of applying machine learning. Overview Harness the power of R for statistical computing and data science Use R to apply common machine learning algorithms with real-world applications Prepare, examine, and visualize data for analysis Understand how to choose between machine learning models Packed with clear instructions to explore, forecast, and classify data In Detail Machine learning, at its core, is concerned with transforming data into actionable knowledge. This fact makes machine learning well-suited to the present-day era of "big data" and "data science". Given the growing prominence of R; a cross-platform, zero-cost statistical programming environment; there has never been a better time to start applying machine learning. Whether you are new to data science or a veteran, machine learning with R offers a powerful set of methods for quickly and easily gaining insight from your data. "Machine Learning with R" is a practical tutorial that uses hands-on examples to step through real-world application of machine learning. Without shying away from the technical details, we will explore Machine Learning with R using clear and practical examples.

Well-suited to machine learning beginners or those with experience. Explore R to find the answer to all of your questions. How can we use machine learning to transform data into action? Using practical examples, we will explore how to prepare data for analysis, choose a machine learning method, and measure the success of the process. We will learn how to apply machine learning methods to a variety of common tasks including classification, prediction, forecasting, market basket analysis, and clustering. By applying the most effective machine learning methods to real-world problems, you will gain hands-on experience that will transform the way you think about data.

"Machine Learning with R" will provide you with the analytical tools you need to quickly gain insight from complex data. What you will learn from this book Understand the basic terminology of machine learning and how to differentiate among various machine learning approaches Use R to prepare data for machine learning Explore and visualize data with R Classify data using nearest neighbor methods Learn about Bayesian methods for classifying data Predict values using decision trees, rules, and support vector machines Forecast numeric values using linear regression Model data using neural networks Find patterns in data using association rules for market basket analysis Group data into clusters for segmentation Evaluate and improve the performance of machine learning models Learn specialized machine learning techniques for text mining, social network data, and ;big; data Approach Written as a tutorial to explore and understand the power of R for machine learning. This practical guide that covers all of the need to know topics in a very

systematic way. For each machine learning approach, each step in the process is detailed, from preparing the data for analysis to evaluating the results. These steps will build the knowledge you need to apply them to your own data science tasks.

Book Information

Paperback: 396 pages

Publisher: Packt Publishing (October 25, 2013)

Language: English

ISBN-10: 1782162143

ISBN-13: 978-1782162148

Product Dimensions: 7.5 x 0.9 x 9.2 inches

Shipping Weight: 1.9 pounds (View shipping rates and policies)

Average Customer Review: 4.6 out of 5 starsÂ See all reviewsÂ (47 customer reviews)

Best Sellers Rank: #475,642 in Books (See Top 100 in Books) #79 inÂ Books > Computers & Technology > Computer Science > AI & Machine Learning > Machine Theory #289 inÂ Books > Computers & Technology > Programming > Algorithms #318 inÂ Books > Computers & Technology > Software > Mathematical & Statistical

Customer Reviews

Some of the other reviewers seem to be unclear on the concept of "introduction", demanding rigor and depth - and, unfortunately, not suggesting any alternatives. If you want deep understanding of the algorithms, you will need to start with a proper textbook, like "Pattern matching" by Bishop, or "Elements of statistical learning" by Hastie and Tibshirani. If you want a more accessible, high-level presentation with examples that can be reproduced, in R, you want to get "Introduction to statistical learning" by James, Witten, Hastie and Tibshirani. If you want a *really* accessible introduction to the techniques - again, with examples that you can try and build on - well, "Machine learning with R" just might be the best choice today. (On R side; if you have invested in Python, Peter Harrington's "Machine learning in action" and books by Wes McKinney are a good bet). OK, this is not an outstanding book, it is under-edited and plain-looking - unfortunately, Packt follow the no-frills approach of O'Reilly - but it is friendly, reasonably well written, and offers a good deal of content. (Extra brownie points for Chapter 11). Let me put it this way: you want to read "Introduction to statistical learning", but "Machine learning with R" is a good warm-up.UPD. With the benefit of a little more life experience, I would say: don't spend your time on *any* R book. Python is the way to go.

First off, I am newbie to both machine learning and R and wanted find a starting point somewhere. I browsed around many books before deciding on this one. The writing style of Mr. Lantz is provided in a very understandable/readable manner. It's akin to someone sitting next to you and explaining things in a down to earth, layman's fashion rather than try to "tech speak" you to death with complicated explanations (aka formal textbook). Just the right amount of hand holding for me. I highlight quite bit and it's actually difficult with this book as there isn't much fluff. He's very succinct. The books states that it's for someone who know some ML and no R or R and no ML. I don't know either and the material is digestible except for one thing: review your stats! I took statistics long ago in college and never really learned it well the first time so I had stop and reread core concepts before continuing. Do yourself a favor and review basic statistics and probability before you start this book. I read both "Naked Statistics" and "Statistics in Plain English" and it helped me a great deal (and probably will continue to do so since it appears a bulk of machine learning is stats and prob). Currently into about a third of the way in and I am finding it to be very enjoyable and practical. Other reviewers point out that this book is too basic and this may be the case, but for someone like me who is starting from absolute scratch and who needs to understand basic ML concepts (AND basic R) I find it a great book. Will post an addendum once I complete it.

This book uses R packages that are have been updated since its publication and no longer work with the code given in the book. I contacted the publisher, but because the code works fine with the package versions it was written for, they will not offer updates on their website. If you know machine learning and R well, you can probably figure out a workaround, but you're also not the intended audience for this book.

If you are new to both machine learning and R and want to learn both at the same time, I can't imagine there being a better book. I needed to figure out how to implement nearest neighbors, decision trees, SVM, neural networks, and boosting on two data sets, in a short amount of time. I had no experience with R and my only prior experience with machine learning was neural networks. Using this book I was able to implement four algorithms in R. For each topic the book describes an application, the algorithm, provides code to implement the algorithm. You can download the data set from the publisher's website so you can try it out.

This is a great book. I liked the way authors highlight syntax for models and discuss strengths and weaknesses. It has a nice balance of theory and hands-on training. However, I would need to use a

R book, such as R in action, in conjunction with it. I have looked at many books on the topic. I will put my review for all of these. Perhaps this can save you some time. 1) <http://www..com/dp/0470650931> : Good theoretical book, but badly written and does not have any hands on exercise. 2) <http://www..com/dp/1466503963> : This is another great book. Good balance of theory and hands-on exercise. This is an excellent book to start learning data mining and R. However, this book relies on a GUI RCommander. It does a good job and one can do a lot with it but it has its limitations. However, I will still use this book. 3) <http://www..com/dp/1439810184> : This is an advanced book and heavily entrenched in cases. This makes it difficult to replicate things unless your work is directly related to one of the case studies covered. 4) <http://www..com/dp/0133412938> : Good examples, but does not explain much about the interpretation. This leaves one wondering what is the purpose of certain graph, what are the axis and how to interpret it. If appropriate explanation is added, this would be an excellent book. 5) <http://www..com/dp/111844714X> : This book is very expensive and almost totally devoid of any theory or discussion. I would not use it. 6)

<http://www..com/dp/1441998896>: This is a decent book. It relies on another GUI, Rattle. It is a strong contender to the book 2 in this list.

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

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) Machine Learning: A Probabilistic Perspective (Adaptive Computation and Machine Learning series) Unsupervised Machine Learning in Python: Master Data Science and Machine Learning with Cluster Analysis, Gaussian Mixture Models, and Principal Components Analysis 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) Bioinformatics: The Machine Learning Approach, Second Edition (Adaptive Computation and

Machine Learning) First-Time Machine Applique: Learning to Machine Applique in Nine Easy Lessons 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) Introduction to Statistical Relational Learning (Adaptive Computation and Machine Learning series) Reinforcement Learning: An Introduction (Adaptive Computation and Machine Learning series) Learning Deep Architectures for AI (Foundations and Trends(r) in Machine Learning) The Bread Lover's Bread Machine Cookbook: A Master Baker's 300 Favorite Recipes for Perfect-Every-Time Bread-From Every Kind of Machine More Bread Machine Magic : More Than 140 New Recipes From the Authors of Bread Machine Magic for Use in All Types of Sizes of Bread Machines Bread Machine Cookbook: Delicious And Simple Bread Machine Recipes

[Dmca](#)