Machine Learning With R - Second Edition

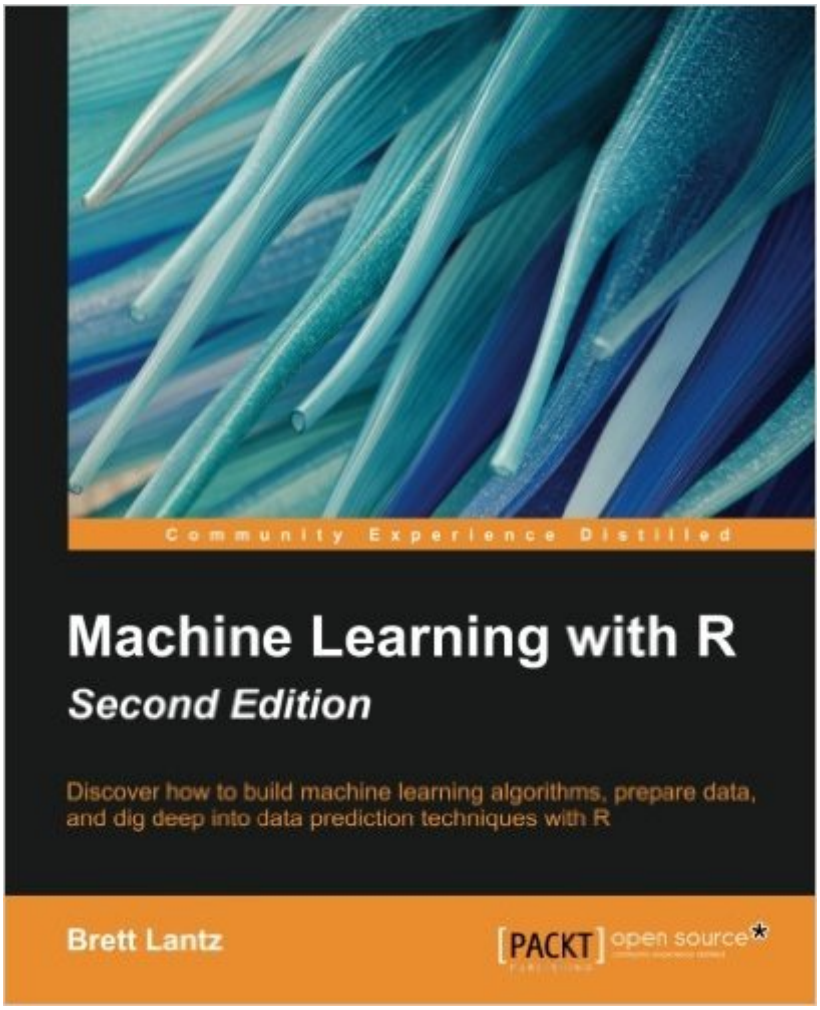

Synopsis

Key FeaturesHarness the power of R for statistical computing and data scienceExplore, forecast, and classify data with RUse R to apply common machine learning algorithms to real-world scenariosBook DescriptionMachine learning, at its core, is concerned with transforming data into actionable knowledge. This makes machine learning well suited to the present-day era of big data. Given the growing prominence of RÃ¢â ¬â •a cross-platform, zero-cost statistical programming environmentÃ¢â ¬â •there has never been a better time to start applying machine learning to your data. Whether you are new to data analytics or a veteran, machine learning with R offers a powerful set of methods to quickly and easily gain insights from your data.Want to turn your data into actionable knowledge, predict outcomes that make real impact, and have constantly developing insights? R gives you access to the cutting-edge power you need to master exceptional machine learning techniques.Updated and upgraded to the latest libraries and most modern thinking, the second edition of Machine Learning with R provides you with a rigorous introduction to this essential skill of professional data science. Without shying away from technical theory, it is written to provide focused and practical knowledge to get you building algorithms and crunching your data, with minimal previous experience.With this book youÃ¢â ¬â„¢ll discover all the analytical tools you need to gain insights from complex data and learn how to to choose the correct algorithm for your specific needs. Through full engagement with the sort of real-world problems data-wranglers face, youÃ¢â ¬â„¢ll learn to apply machine learning methods to deal with common tasks, including classification, prediction, forecasting, market analysis, and clustering. Transform the way you think about data; discover machine learning with R.What you will learnHarness the power of R to build common machine learning algorithms with real-world data science applicationsGet to grips with R techniques to clean and prepare your data for analysis, and visualize your resultsDiscover the different types of machine learning models and learn which is best to meet your data needs and solve your analysis problemsClassify your data with Bayesian and nearest neighbour methodsPredict values by using R to build decision trees, rules, and support vector machinesForecast numeric values with linear regression, and model your data with neural networksEvaluate and improve the performance of machine learning modelsLearn specialized machine learning techniques for text mining, social network data, big data, and moreAbout the AuthorBrett Lantz has used innovative data methods to understand human behavior for more than 10 years. A sociologist by training, he was first enchanted by machine learning while studying a large database of teenagers' social networking website profiles. Since then, he has worked on the interdisciplinary studies of cellular telephone calls, medical billing data, and philanthropic activity,

among others.Table of ContentsIntroducing Machine LearningManaging and Understanding DataLazy Learning Ã¢â ¬â œ Classification Using Nearest NeighborsProbabilistic Learning Ã¢â ¬â œ Classification Using Naive BayesDivide and Conquer Ã¢â ¬â œ Classification Using Decision Trees and RulesForecasting Numeric Data Ã¢â ¬â œ Regression MethodsBlack Box Methods Ã¢â ¬â œ Neural Networks and Support Vector MachinesFinding Patterns Ã¢â ¬â œ Market Basket Analysis Using Association RulesFinding Groups of Data Ã¢â ¬â œ Clustering with K-meansEvaluating Model PerformanceImproving Model PerformanceSpecialized Machine Learning Topics

Book Information

Paperback: 454 pages

Publisher: Packt Publishing - ebooks Account; 2 edition (August 3, 2015)

Language: English

ISBN-10: 1784393908

ISBN-13: 978-1784393908

Product Dimensions: 7.5 x 1 x 9.2 inches

Shipping Weight: 2.1 pounds (View shipping rates and policies)

Average Customer Review: 4.7 out of 5 starsÂ Â See all reviewsÂ (17 customer reviews)

Best Sellers Rank: #64,013 in Books (See Top 100 in Books) #9 inÂ Books > Computers & Technology > Computer Science > AI & Machine Learning > Machine Theory #35 inÂ Books > Computers & Technology > Programming > Algorithms #48 inÂ Books > Computers & Technology > Software > Mathematical & Statistical

Customer Reviews

I’m torn. There are some useful gems in this book, and for the most part, the presentation is simple, albeit a bit pedantic and cartoonish at times. If I was trying to get up to snuff on a new machine learning method, I might start here, since it *does* provide starter code for a variety of problems. That’s quite handy. It doesn’t, however, go into much depth at all on any one topic. You can’t read this book and expect to know how to do any one of these methods well. Certainly, it’s a tall order to ask any one book to cover all ML topics in depth, but any potential reader should be aware that this just skims the surface of a whole bunch of topics.On top of this, who in the world edited this book? Every other page has horrible typos, missing words, repeated sentences. These are not trivial errors either. This is a book about data analysis and yet the reported data are clearly wrong in places, e.g., a result is listed as .06 percent in one spot and then .0006 in another (p. 271). Basic subject-verb

agreement errors riddle the text, e.g. "These output is shown as follows". Sometimes these are trivial errors, but other times you have absolutely no idea what the intended meaning is. I have about 100 pages more to read but I'm starting to wonder if I'm just wasting my time.

Excellent overview for the different classical ways of machine learning. The writer clearly knows his way around. Explanations are down-to-earth, light on the math and theory but with references when they are needed. Glossing over the theory enables the author to condense a lot of information into those ~350 pages. Heavy on practical advice and good practices.Good accounts of the different algorithms are of course available online, but the big advantage of this book in my opinion is its eye for the end-game. What actionable insights can you extract from your dataset and how, using clear examples accompanied with R code. A book much appreciated.

I am pleased to have bought this book (directly from Packt, the publisher) based on positive reviews of the first edition. My background is as a SQL programmer and CRM data analyst, and although I had some experience of data mining algorithms in other software, I do not have a lot of prior experience in R.Before jumping into descriptions of the various data mining algorithms, there is some useful material on the basics of data handling in R, which was a useful refresher for me as an R novice.After this, the author describes clearly and concisely the use of the various algorithms, together with discussion on the strengths and weaknesses of each. There are examples given, using mostly real world data (which is available to download). These are easy to follow, giving enough detail to understand the concepts without getting bogged down in too much statistical detail.I found it useful to have some understanding of the concepts of some of the mining models, but this is not essential as the book gives a good grounding in both the concepts, and how to apply them in R.

I have read this book cover to cover. The readers can learn machine learning in practice without any prior knowledge in R and statistics.After an introduction to machine learning in chapter one the author explains essential concepts of R in second chapter. From chapter 3 to 9, each chapter covers a machine learning algorithm with its related extensions as well as introducing corresponding packages and commands in R. In addition for each algorithm a practical project in real world has been analyzed and discussed from preparing data to achieving final goal.The last three chapters are essential for using machine learning in data analysis. Chapter 10 introduces various tools to measure the performance of particular model applied on given data. It contains useful discussions

which hint to the researchers to judge correctly the efficiency of employed model. In chapter 11 the reader learn different concepts and techniques for improving the efficiency of machine learning models. The author introduced briefly various topics in chapter 12 which are important in real life analysis when one need obtain data from Internet, dealing with big data, databases, parallel programming and etc.In conclusion "Machine Learning with R - Second Edition" by Brett Lantz is a valuable book for learning machine learning practically with using R which expects the readers without background in field.

Excellent book. It covers all you need to get started with a solid foundation in machine learning. I would highly recommend it to any programmer (or reasonably logically minded individual) who wants to get into machine learning.EDIT: After finishing the book, I'd still recommend it, and everything I wrote previously still applies. I am now actively using the knowledge I gained from this book on a project. However, I am reducing my rating to 4 stars because the publisher/editor is absolutely atrocious. I found at least a dozen minor errors (mostly typographical such as repeating a section of a sentence) that never should have made it past a simple proofreading.

series) 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) Machine Learning with R - Second Edition Machine Learning with R - Second Edition - Deliver Data Insights with R and Predictive Analytics A First Course in Machine Learning, Second Edition 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)