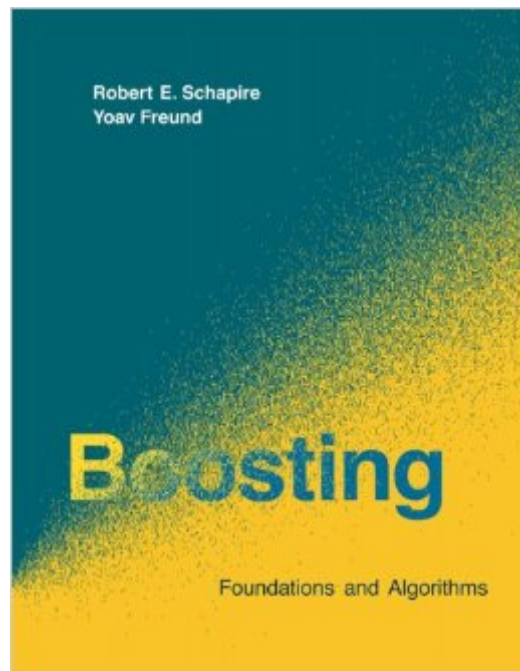# Boosting: Foundations And Algorithms (Adaptive Computation And Machine Learning Series)

Synopsis

Boosting is an approach to machine learning based on the idea of creating a highly accurate predictor by combining many weak and inaccurate "rules of thumb." A remarkably rich theory has evolved around boosting, with connections to a range of topics, including statistics, game theory, convex optimization, and information geometry. Boosting algorithms have also enjoyed practical success in such fields as biology, vision, and speech processing. At various times in its history, boosting has been perceived as mysterious, controversial, even paradoxical.This book, written by the inventors of the method, brings together, organizes, simplifies, and substantially extends two decades of research on boosting, presenting both theory and applications in a way that is accessible to readers from diverse backgrounds while also providing an authoritative reference for advanced researchers. With its introductory treatment of all material and its inclusion of exercises in every chapter, the book is appropriate for course use as well. The book begins with a general introduction to machine learning algorithms and their analysis; then explores the core theory of boosting, especially its ability to generalize; examines some of the myriad other theoretical viewpoints that help to explain and understand boosting; provides practical extensions of boosting for more complex learning problems; and finally presents a number of advanced theoretical topics. Numerous applications and practical illustrations are offered throughout.

Book Information

Series: Adaptive Computation and Machine Learning series

Paperback: 544 pages

Publisher: The MIT Press (January 10, 2014)

Language: English

ISBN-10: 0262526034

ISBN-13: 978-0262526036

Product Dimensions: 7 x 0.9 x 9 inches

Shipping Weight: 1.6 pounds (View shipping rates and policies)

Average Customer Review: 4.2 out of 5 stars See all reviews (15 customer reviews)

Best Sellers Rank: #326,311 in Books (See Top 100 in Books)   #55 in Books > Computers & Technology > Computer Science > AI & Machine Learning > Machine Theory   #202 in Books > Computers & Technology > Programming > Algorithms

Customer Reviews

A book on boosting coming from its inventors... do I need to say more? I was a CS student at

Princeton University, and I was once fortunate enough to take Prof. Schapire's course on theoretical machine learning. Prof. Schapire is an amazing teacher and researcher. When I was writing this review I tried to avoid any bias I may have because of my respect and personal admiration for him, and yet I have to say this book is simply a masterpiece.Reading this book was simply enjoyable. It is very well structured. Every chapter illustrates a differing perspective on boosting (either theoretical or practical), and as a whole they offer a complete view on this fantastic algorithm/mechanism. Chapter 1 is already good enough for normal practitioners, and if you are fascinated by the incredible performance of boosting and want to know more, please keep on reading and I believe when you finish the last page you will feel like everything is so clear. Theorems are rigorously proved. Algorithms are unequivocally laid out. Theories and practices work in perfect concert. The second and third parts of the book were the most interesting to me. Three distinct justifications of boosting's effectiveness are beautifully illustrated, and even more awesome, these theoretical underpinnings are directly related to ways to generalize the basic AdaBoost to other classification scenarios e.g. how to incorporate probability outputs from base classifiers, how to derive new variants by changing the underlying optimization problem, how to naturally extend to multi-class or multi-label classification, etc.

It's a kind of problem I come across way too often: in trying to determine whether a specific condition exists, a number of symptoms might help the diagnosis. But, some of the symptoms can also appear when something else is going on, instead. On the other hand, not all of the symptoms necessarily appear when the condition in fact is active. The question then becomes, given a number of indicators that have some diagnostic value, and given that all of them are inaccurate some of the time, how do I combine the indicators' answers to get the best diagnosis of that condition?This book proposes "Boosting" as the answer. Start with remarkably few technical requirements on the diagnostic indicators, plus some number of cases in which the condition's presence or absence is already known. Given that, Boosting iteratively determines the weight to assign each indicator. One requirement is that each indicator suggests presence or absence of the condition at least somewhat differently than random guessing - and being wrong most of the time is just as useful as being right more often than not, since the algorithm automatically assigns negative weights to such indicators.After that, the authors present rigorous development in a number of directions. I emphasize "rigor" - this text offers detailed development, analysis, and formal proof of the algorithm and its properties, far beyond the needs of someone who just wants to implement the technique. Implementable detail is there, but you'll spend a fair bit of time teasing it out of the dense notation

used here.Then, once basics have been established, the discussion branches out. The authors offer game-theoretic analysis of the algorithm, along with comparisons to related optimization techniques.

Boosting: Foundations and Algorithms (Adaptive Computation and Machine Learning series) Foundations of Machine Learning (Adaptive Computation and Machine Learning series) Machine Learning: A Probabilistic Perspective (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) Introduction to Statistical Relational Learning (Adaptive Computation and Machine Learning series) Reinforcement Learning: An Introduction (Adaptive Computation and Machine Learning series) Probabilistic Graphical Models: Principles and Techniques (Adaptive Computation and Machine Learning series) Machine Learning with Spark - Tackle Big Data with Powerful Spark Machine Learning Algorithms The Design of Innovation: Lessons from and for Competent Genetic Algorithms (Genetic Algorithms and Evolutionary Computation) 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) IntAR, Interventions Adaptive Reuse, Volume 03; Adaptive Reuse in Emerging Economies Learning Deep Architectures for AI (Foundations and Trends(r) in Machine Learning) Testosterone Boosting with Tongkat Ali and 7 Potent Herbs Unsupervised Machine Learning in Python: Master Data Science and Machine Learning with Cluster Analysis, Gaussian Mixture Models, and Principal Components Analysis