


The book was found

The Master Algorithm: How The Quest For The Ultimate Learning Machine Will Remake Our World



Synopsis

Under the aegis of machine learning in our data-driven machine age, computers are programming themselves and learning about - and solving - an extraordinary range of problems, from the mundane to the most daunting. Today it is machine learning programs that enable Amazon to predict what users will like, Netflix to predict what users will like, Apple to power Siri's ability to understand voices, and Google to pilot cars. These programs are already helping us fight the war on cancer and predict the movements of the stock market, and they are making great headway with instant language translation and discovering new laws of nature. But machine learning is incomplete, and its practitioners across the globe are seeking the most powerful algorithm of all. The Master Algorithm will not be limited to solving particular problems but will be able to learn anything and solve any problem, however difficult, and Pedro Domingos, a trailblazing computer scientist, is at the very forefront of the search for it. With the Master Algorithm in hand and data as its fuel, machine learning - essentially the automation of discovery, a kind of scientific method on steroids - will become the most powerful technology humanity has ever devised. And The Master Algorithm will be its bible.

Book Information

Audible Audio Edition

Listening Length: 13 hours and 3 minutes

Program Type: Audiobook

Version: Unabridged

Publisher: Brilliance Audio

Audible.com Release Date: September 22, 2015

Whispersync for Voice: Ready

Language: English

ASIN: B014X01SS0

Best Sellers Rank: #1 in Books > Computers & Technology > Computer Science > AI & Machine Learning > Expert Systems #8 in Books > Computers & Technology > Computer Science > Information Theory #9 in Books > Audible Audiobooks > Science > Mathematics

Customer Reviews

This book isn't perfect, but I'm still giving it 5 stars because it provides a better overview of the entire field of Machine Learning than any other book I've come across. My background is in computer science and software engineering and I've been interested ML as more of a hobbyist and outside observer for a few years (reading some books, taking Andrew NG's coursera course), just recently

dabbling in some applications professionally. What I was still missing before reading this book was as high a level understanding of where all of the models and technique in the field of ML fit. Other books describe the difference between supervised and unsupervised learning, but this book goes further in describing how, say, decisions trees, support vector machines and deep neural networks fit compared to each other and within which subfields statistics play a larger role than others. The book also puts many techniques in historical perspective that I found very helpful, such as the rise, fall and rise again of deep neural networks with support vector machines taking a lead as the hottest technique in between (while also making clear that SVMs are a useful technique with unique strengths today). Finally, it makes clear that these techniques are not all competing for being the best overall at everything, but that they can be used quite complementary and/or they have unique strengths within certain problem domains. The book accomplishes all of this through a survey of broad subfields of ML, how each has attempted to be *the* master algorithm, has fallen short in some ways, but remains the best at some things and could play a role in the state of the art master algorithm (while acknowledging we're not quite there yet). So while the term 'master algorithm' is somewhat of a gimmick (as he acknowledges), it's a good way to think about what ML is attempting to accomplish as a field: building working, adaptive software systems with less and less human assistance by learning from data, and to see how many specific techniques have played a role in progress. What I don't know is how accessible this book might be to someone who's less technical. I think the first couple chapters would be a great read for anyone with a general interest, making clear how ML differs from the traditional software / automation that has brought us so far, but it could be that the details within the rest of the chapters that go into more depth would be too in the weeds. I've also read some other reviews from technical readers that assert the book lacks enough depth to be helpful, but this wasn't the case for me, in fact the level of detail was perfect – just deep enough to match with details I'd skimmed before in previous surveys of the field yet not so deep that I couldn't get through and enjoy the chapters in a casual evening read. The author also explained some concepts better than I've read anywhere else before, such as the debate is between frequentist and bayesian statisticians.

This is a tremendous work for people that want a good overview of the major theories behind machine learning but aren't necessarily looking to dive deep into the underlying math. The book has a very humorous tone, when appropriate, and creates numerous examples of machine learning and AI in practical use in industry. The book coincides with the free online course Pedro Domingos teaches at [...] Well done!

The Master Algorithm by Pedro Domingos is a fascinating and accessible deep dive into artificial intelligence written by a giant in the field. It explains the history, functionality, limitations, promise, and implications of a technology that is reshaping nearly every industry. The first and last chapter are particularly important as they provide the conceptual framework for making sense of the machine learning universe. If you want to understand where computing is headed, this is required reading.

A word of warning: I'm not a machine learning practitioner and not even mathematically inclined. I am merely an average person with a (reasonable) degree of culture (I suppose I classify as a sociologist) looking out to understand the world around me a little better. Yet I loved this book, it opened my eyes on a world I suspected existed somewhere in the depths of my computer or when I did a Google search but that I never understood. If you are afraid to be engulfed in equations, don't be. The author prefers to use allegories and is very good at giving simple explanations, making everything (almost) crystal clear. What you get is a comprehensive overview of where machine learning is going. This is hugely important, considering the sometimes disturbing news you get in the media - not just from Snowden but also from "neurotechnologists" who guide political campaigns by embedding cameras in ads, cameras that detect the viewers' reactions and then adjust the campaign in function of these reactions. Perhaps my only quarrel with this book - but it does not detract from its excellence and the five stars I gave it - is the philosophical position taken by the author. I don't quite share his optimism about our robotic future. And this for 2 reasons: One, the connection between learning algorithms and data. Domingos does note at the outset: "...the more data they [the learners] have, the more intricate the algorithms can be." So without "big data", you don't get "good" learning algorithms - the "bigger" the data, the better the algorithm. Or as Domingos writes, "the more data we have, the more we can learn". True enough. But what about the quality and size of the big data? What if the data is error-filled and we're not aware of it? Can machine learning be aware of something their human masters are not? On what basis? The algorithms are taught to deal with an imperfect world and draw the most likely conclusions, where "most likely" is highly subjective. Or at least, it is "highly subjective" in my opinion and I realize that it is only one opinion and not necessarily one shared by the author. Two, the connection between you as you are and an expanded "digital you": a "model" of who you are, your work experience, your tastes, a model that the algorithm has learned - and thanks to algorithms, you will be able to do more things and faster than ever before. Domingos describes an extraordinary future where, for example, in your

LinkedIn account, "you'll immediately interview for every job in the universe that remotely fits your parameters (profession, location, pay, etc) LinkedIn will respond on the spot with a ranked list of the best prospects, and out of those, you'll pick the first company that you want to have a chat with. Same with dating: your model will go on millions of dates so you don't have to..."Big time-saver obviously. But it comes at a cost: you have to give your parameters (stuff about your work, your likes and dislikes) to the algorithms. Ye who go digital, leave behind any notion of privacy...So users of this algorithmic landscape will have few shaded areas in which to hide. But for me, privacy concerns aside (and I'm not that concerned, I have nothing to hide), there is yet another matter that I find more worrisome. Domingos (on p.283) tells us he is confident that the future can only get better: "In fact, it's the systems that have a slight edge in serving us better that will, generation after generation, multiply and take over the gene pool." Then he adds: "Of course, if we're so foolish as to deliberately program a computer to put itself above us, then maybe we'll get what we deserve".Indeed. That is precisely my worry: can we be sure that there won't be a Doctor Evil who will do so? Can we be sure that there won't be an Apprentice Sorcerer who might end up doing this even if unintentionally? A mistake can happen...So yes, I'm deeply worried about this algorithmic future and I don't share the author's unflagging optimism. But that doesn't mean it's a book you shouldn't read. On the contrary, it's a must read precisely because it raises all these fundamental questions about the future of humanity. And it is an easy read even for the non-cognoscenti like myself, I highly recommend the introductory chapters, and if you get bogged down in the middle with all the technicalities, no fear, skip to the last two chapters, they are well worth reading!

Great overview including the math involved, which is explained in a creative but intuitive way. Some parts require one to reread a couple of times in part because it takes some time to visualize what's going on, but overall very accessible. The last chapter really frames the AI debate over humanity's future very well. It also provides the next steps for people to learn and explore machine learning on their own-it looks fun!

Helped me put the subject into a broad perspective seeing how different aspects relate to each other. For non experts it's worth it to read sections more than once to improve understanding. Well written and a strong contribution.

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

The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World
Machine Learning: An Algorithmic Perspective, Second Edition (Chapman & Hall/Crc Machine

Learning & Pattern Recognition) Machine Learning: A Probabilistic Perspective (Adaptive Computation and Machine Learning series) Introduction to Machine Learning (Adaptive Computation and Machine Learning series) Machine Learning: For Beginners: Definitive Guide for Neural Networks, Algorithms, Random Forests and Decision Trees Made Simple (Machine Learning, Book 1) BREAD MACHINE COOKBOOK: 120 Most Delicious Bread Machine Recipes (bread, bread bible, bread makers, breakfast, bread machine cookbook, bread baking, bread making, healthy, healthy recipes) The Refashion Handbook: Refit, Redesign, Remake for Every Body Sneaky Green Uses for Everyday Things: How to Craft Eco-Garments and Sneaky Snack Kits, Create Green Cleaners, Remake Paper into Flying Toys, Assemble ... a Robot Recycle Bin with Everyday Things The Bread Lover's Bread Machine Cookbook: A Master Baker's 300 Favorite Recipes for Perfect-Every-Time Bread-From Every Kind of Machine Bread Lover's Bread Machine Cookbook: A Master Baker's 300 Favorite Recipes for Perfect-Every-Time Bread-From Every Kind of Machine Machine Learning: Fundamental Algorithms for Supervised and Unsupervised Learning With Real-World Applications Inheritance: How Our Genes Change Our Lives--and Our Lives Change Our Genes The Mind-Gut Connection: How the Hidden Conversation Within Our Bodies Impacts Our Mood, Our Choices, and Our Overall Health Cricut Tips the Ultimate Troubleshooting Guide: How to Master Your Cricut Machine Reinforcement Learning with Python: An Introduction (Adaptive Computation and Machine Learning series) Deep Learning (Adaptive Computation and Machine Learning series) Practical Machine Learning with H2O: Powerful, Scalable Techniques for Deep Learning and AI Reinforcement Learning: An Introduction (Adaptive Computation and Machine Learning) MATLAB Deep Learning: With Machine Learning, Neural Networks and Artificial Intelligence Semi-Supervised Learning (Adaptive Computation and Machine Learning series)

[Contact Us](#)

[DMCA](#)

[Privacy](#)

[FAQ & Help](#)