


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

Pedro Domingos

Download now

Read Online 

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

Pedro Domingos

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

Pedro Domingos

Algorithms increasingly run our lives. They find books, movies, jobs, and dates for us, manage our investments, and discover new drugs. More and more, these algorithms work by learning from the trails of data we leave in our newly digital world. Like curious children, they observe us, imitate, and experiment. And in the world's top research labs and universities, the race is on to invent the ultimate learning algorithm: one capable of discovering any knowledge from data, and doing anything we want, before we even ask.

Machine learning is the automation of discovery—the scientific method on steroids—that enables intelligent robots and computers to program themselves. No field of science today is more important yet more shrouded in mystery. Pedro Domingos, one of the field's leading lights, lifts the veil for the first time to give us a peek inside the learning machines that power Google, Amazon, and your smartphone. He charts a course through machine learning's five major schools of thought, showing how they turn ideas from neuroscience, evolution, psychology, physics, and statistics into algorithms ready to serve you. Step by step, he assembles a blueprint for the future universal learner—the Master Algorithm—and discusses what it means for you, and for the future of business, science, and society.

If data-ism is today's rising philosophy, this book will be its bible. The quest for universal learning is one of the most significant, fascinating, and revolutionary intellectual developments of all time. A groundbreaking book, *The Master Algorithm* is the essential guide for anyone and everyone wanting to understand not just how the revolution will happen, but how to be at its forefront.

The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World Details

Date : Published September 22nd 2015 by Basic Books (first published September 8th 2015)

ISBN :

Author : Pedro Domingos

Format : Kindle Edition 352 pages

Genre : Science, Nonfiction, Technology, Artificial Intelligence, Computer Science, Business, Programming, Abandoned, Computers, Software

 [Download The Master Algorithm: How the Quest for the Ultimate Le ...pdf](#)

 [Read Online The Master Algorithm: How the Quest for the Ultimate ...pdf](#)

Download and Read Free Online The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World Pedro Domingos

From Reader Review The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World for online ebook

Meghan says

Good overview of machine learning. The master algorithm seems like an overwhelming concept at first, but the book is very accessible for anyone who has a basic education in comp sci. However it's pretty clear that the intended audience is nerdy high school boys considering a career in machine learning, despite the author saying this book is written for everyone/anyone.

Also, I can't help but find his depiction of the post-master algorithm world creepy. Call me old fashioned.

Oliver Sampson says

While coached as a guidebook to help find "The Master Algorithm," the one AI algorithm "that will rule them all" (his words, not mine), this book is much, much more. At times written whimsically, and at times treating very advance material in a way that non-sophisticated readers can understand, the book is part history lesson, part cultural commentary, and part description of the scientific process. I work exactly in the field of Artificial Intelligence and Machine Learning, and I am definitely a member of the target audience. Maybe I'm just an ML fanboy, but I found the treatment of everything, including that of the author's own work, to be at just the right level to keep the non-specialist interested, while informing the specialist about those other areas, where he is not a specialist.

The book finishes with musings about what it means to be a source of data for the corporations (and governments) that would use that data for good and/or evil. This alone makes the book worth reading. Such candor from those in the field of machine learning is really refreshing.

Ismail says

I like Pedro Domingos. He has some very nice accessible papers, and he seems like a nice guy (having done an online course, being open source fan, etc etc).

But, this book is a pile of crap. Despite his best efforts, Domingos isn't a novelist, which makes the writing a bit cheesy. Putting that aside, I think that the book has several problems:

- The entire premise of the book is that a master algorithm exist. I don't think that we have any idea about that yet.
- The separation of machine learning people into groups looked to me extremely superficial. The author gives the expression that the communities are almost completely separated from each other which is far from true. Michael Jordan (arguably the biggest machine learning scientist of our time) has given contribute in neural networks in the past, as he has done in clustering, but now works more on graphical models. Andrew Ng did his college as a student of Mitchell, did his PhD with Jordan, but his current work is on the same

'group' as Hinton.

- The book is confusing at times, and jumps from topic to topic. I don't know if not machine learning scientists were able to understand much from it, while on the same time, there wasn't anything interesting for people who are familiar with machine learning.
- 'No free lunch' relation with 'no classifier can do better than the coin flip classifier' is arguably the worst explanation of 'no free lunch in statistics' 'theorem'. Same about Bayesian inference.
- For a book that mentions a lot of time how better the data is to intuition, concluding it with complete intuition about the future of machine learning looked a bit controversial to me.

On the bright side, I liked the Lord of the Ring-ish poetry about the Master Algorithm in the last chapter.

1.5 stars out of 5. Hoping that Domingos sticks to scientific papers.

Gwern says

Domingos wants to cover all of machine learning for the layman, but it winds up being a big mess.

The level of explanation veers wildly from ridiculously oversimplified to technical minutiae. It is more confusing than enlightening as it goes through topics in an almost random order, scattering them all throughout the book. (You would think that Hume's problem of induction, the underdetermination of data, Occam's razor, the curse of dimensionality, and overfitting, would all be discussed in one and the same place in order to set the stage for how the various 'tribes' work, but you would be wrong.) The manic stream-of-consciousness writing style also drives me nuts, and the little medieval-fantasy passages come off as puerile. (I smiled. Once.)

The explanations are almost uniformly terrible (another reviewer asks if this is the worst explanation of Bayesian inference one has ever read; I would have to say that at least for me, this is competitive for that distinction), and most are explained as briefly as decision trees are endlessly waxed upon. Major premises like there being any really universal algorithm are poorly presented; compare Domingo's argument for there being a neural algorithm with, say Jacob Cannell's "The Brain as a Universal Learning Machine", where Domingos is provoking rather than thought-provoking.

Content-wise, I have to seriously question the inclusion of evolutionary programming as a top-tier paradigm, and analogies hardly seem much more relevant a grouping either, and all that space comes at a huge cost of extreme superficiality about what *deep learning is doing right now*. Let me remark on how astounding it is to read a book whose self-proclaimed goal is to de-mystify machine learning for the layman, explain recent advances in deep learning that have created such media hype and sparked so much commercial & public & research interest, and which seems to only go from strength to strength to the point where sometimes it feels like one can hardly even skim a fascinating paper before yet another one has been uploaded to Arxiv, and which winds up doing little but explaining what backpropagation is and then passing grandiosely onto other topics and not, y'know, covering anything like solving ImageNet, caption generation, logical inference using reading of passages, etc. Or to read a decent capsule description of the general paradigm of reinforcement learning... and then see deep reinforcement learning described in a few sentences mostly to the effect that learning can be unstable - really? *That* is what laymen need to know about deep reinforcement learning, that - whatever it is - it can be unstable?

Oh, and he offers us his thoughts on AI risk, the fruit of his decades of experience with machine learning:

Relax. The chances that an AI equipped with the Master Algorithm will take over the world are zero. The reason is simple: unlike humans, computers don't have a will of their own. They're products of engineering, not evolution. Even an infinitely powerful computer would still be only an extension of our will and nothing to fear...The optimizer then does everything in its power to maximize the evaluation function—no more and no less—and the evaluation function is *determined by us*. A more powerful computer will just optimize it better. There's no risk of it getting out of control, even if it's a genetic algorithm. A learned system that didn't do what we want would be severely unfit and soon die out. In fact, it's the systems that have even a slight edge in serving us better that will, generation after generation, multiply and take over the gene pool. 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. The same reasoning applies to all AI systems because they all—explicitly or implicitly—have the same three components. They can vary what they do, even come up with surprising plans, but only in service of the goals we set them. A robot whose programmed goal is “make a good dinner” may decide to cook a steak, a bouillabaisse, or even a delicious new dish of its own creation, but it can't decide to murder its owner any more than a car can decide to fly away. The purpose of AI systems is to solve NP-complete problems, which, as you may recall from Chapter 2, may take exponential time, but the solutions can always be checked efficiently. We should therefore welcome with open arms computers that are vastly more powerful than our brains, safe in the knowledge that our job is exponentially easier than theirs.

How I wish I was making up these arguments. (Aside from the invocation of complexity theory which is not even wrong as many problems we want AI to solve are not expressible as decision problems, the ones which are can fall into anything of 'much easier than NP-complete' or 'much harder', and a problem falling into a particular complexity class is no guarantee of safety in the first place, this sort of naiveté is sad coming from someone so enthusiastic about genetic algorithms - where researchers routinely discover that defining a good reward/fitness/evaluation function is quite difficult and they have to fight their algorithms to get a useful rather than a hilariously perversely correct answer!)

Overall, I would absolutely recommend against this book for any laymen interested in statistics or machine learning. The explanations are so poor and garbled that you will either not learn anything or what you take away will be as likely to be misleading as not. You will be better off with Silver's *The Signal and the Noise*, reading random presentations on Schmidhuber's website, Bostrom's *Superintelligence* or even Hutter's *Machine Super Intelligence* or Domingos's own "A Few Useful Things to Know about Machine Learning" (which was really good) or anything really. (Suggestions are welcomed on things I can recommend for laymen *instead* of this...)

Brian Nwokedi says

From a content standpoint, *The Master Algorithm* by Pedro Domingos is a great crash course for anyone who is interested in learning more about machine learning. But from an “ease of comprehension” standpoint, this book is far from the layman’s journey that Domingos claims it is.

I found myself able to follow roughly 70% or so of the technical content of this book, and there were definitely some times that it was a bit too technical for me to completely grasp what he was trying to say. The writing at times can become circular in nature and to use a nerd joke... feels like a bit of a “circular

reference” J. Based on these two components (content vs. comprehension), I believe that the Master Algorithm is a 2.5 to 3.0 star book that is worth reading if you are remotely interested in A.I., machine learning, and computer science. You just have to be comfortable with some of the complex concepts that he covers (maybe that is just me)

More about the book itself... The main point that I grabbed onto and took away is that there are five distinct methods/tribes for machine learning that will have to come together to create the Master Algorithm. The five tribes/methods are as follows:

1. Symbolologists (inverse deduction);
2. Connectionists (backpropagation);
3. Evolutionaries (genetic algorithms);
4. Bayesians (probabilistic inference);
5. Analogizers (support vector machines))

Each of these five tribes has a piece of the puzzle that will ultimately help to create the Master Algorithm. And with this algorithm, true and grand machine learning will be able to flourish in a way that we have yet to see. The Master Algorithm is the unifier of machine learning: it lets any application use any learner, by abstracting the learners into a common form that is all the applications need to know.

But as Domingos states, the future of machine learning, although bright, is very much dependent on cracking the code of this Master Algorithm. Domingos implores each of us, novice and expert alike, to play our part in the development of this future and I suggest anyone interested in this material go to the websites that are in the book and take a look. There is some pretty cool stuff out there!

Maria Espadinha says

Remembering Sophia...

While reading this book, the image of Sophia was constantly assaulting my mental screen!

I'm sure most of you remember that gorgeous social robot, that could blink, smile, raise an eyebrow, and was capable of 59 more facial expressions.

She has been a media darling, showing herself in magazines, newspapers, tv news, talk shows,... spreading her charm all over the world.

She could handle a clever conversation, make eye contact and even show some sense of humour!

Sophia dreams about helping humanity building a better world, and trembles with the thought of taking a shower ;)

She belongs to a project of genius-androids -- highly intelligent robots, whose main skill is solving all those mind cracking problems, human brains find so hard to unravel!

According to their creators, in something like twenty years time, androids will be able to perform most of human jobs.

Scary?!...

My one and only fear is, by the time it happens, the ones who'll be ending like robots, will be us, common

mortals!

HeHe!... Just kidding!

Robots will always need us around! I'm sure they will demonstrate great abilities as collaborators, but couldn't possibly 100% substitute us.

The creation never surpasses its Master!

Since I'm crazy about happy endings, I have a strong feeling this Humanoids Creation story will grow into a happily ever after -- Sophia's dream will come true and we'll all share a better world. ;)

Can't get worse, can it?!...

Androids intelligence is possible, thanks to machine learning -- this book's core.

Hence, if you are curious about knowing all possible scopes of machine learning, when applied to human society, this book is for you :)

I'm ending this review with a video, where you'll be watching a conversation between Sophia and Jimmy Fallon, on Tonight's Show. I hope you'll enjoy it:

https://m.youtube.com/watch?v=Bg_tJvC...

Gary says

The author states that "intuition is what you use when you don't have enough data". The author will show heuristically how intuition is slowly being taken out of analyzing big data and being replaced with algorithms which teach themselves how to make the data speak for themselves. "All learning starts with some knowledge" (a quote from Hume, that the author invokes), and from Hume we know that there is a problem with induction, no matter what the particular can not prove the universal. The trick is to get from the data (the particular) to the universal and the author explains in detail the five general ways we learn and shows how they work in practice. The five ways are Symbolic (think: rational thought), Connective (modeling like the Network in the brain), Bayesian (nothing is certain and all is contingent), Evolutionary (see "The Selfish Gene" by Dawkins), and by Analogy.

The key is to use some variations of the ways ('tribes') and have the method (algorithm) use the data to exploit the information that is within the data set and do it recursively (and as Douglas Hofstadter says "I am a Strange Loop"). The computers are becoming faster, cheaper and can manipulate ever larger and more easily accessible data sets, and the methods have become more refined and usable. For example, brute force Bayesian methods are not used since the whole decision tree necessary for learning complex solutions are never practical and are now replaced by naive Bayesian techniques (only some of the dependent states need to be computed) giving only a small loss in overall accuracy.

The overall point of the book is to show that there is evolutionary thinking going on in writing smart algorithms which are able to let the data speak for themselves and the computer scientists have a tool box of techniques which enable real objective knowledge to be extracted from the data.

I like the TV show Person of Interest. Everything that "The Machine" does on that show can be explained by the techniques discussed in this book. This author doesn't think the computer will ever be able to think or have its own "will". I think this book would be an excellent lead in to the Nick Bostrom book "Superintelligence: Paths, Dangers and Strategies". That book does think super AI will happen and a computer will develop a 'will'. This book, "Master Algorithm" is an excellent primer for someone who

believes the "singularity is near" even though the author disagrees (It's odd this author thinks the super AI is not possible because the way he starts off the book by explaining the P=NP problem and how solving that could create a master algorithm which in my way of thinking would lead to a super AI).

Charlene says

Most of this book was great because it read like a short summary of what is taught in an introduction to cognitive science class. While spouting about how Bayesian stats decidedly kick frequentist stat's ass, to which I agree, the author showed how to look at the world itself, and everything in it, through a more Bayesian lens. He hammered home the central point that nothing can be understood in isolation and must rather be understood through its connection to the things around it. One more book on my list that belongs on the complexity/emergence/network shelf. Any book in that category, if it's any good at all, will be among my favorites. Adding to that, this author did a great job of covering one of my favorite subjects, Markov chains.

Every time I dictate into my phone and realize that Apple has corrected "thus" to DOS and can recognize any computer programming term but not the science terms I dictate, I laugh to myself at the coders who create the Markov chains. They are seriously biased in favor of programmers. But in the end, that is because they are programmers:)

This book does an exceptional job of explaining Bayes, providing a brief history of how it came to be used, and showing how it is at work naturally, all around us-- neurons, self driving cars, electrons, etc.

In my opinion, he should have stopped before writing the last chapter. He basically ruined his beautiful book by merging it with the Selfish Gene theory and going on a rant about how humans are special. Does he know that any and every time we humans thought we were special, it turned out we were wrong? For some reason, he thinks humans will always control AI. He doesn't seem to be of the opinion that AI itself will be a mashup of human and machine, with no differentiation. He also went on to trash any advances that arise from Moore's Law, which he said was on it's last leg. Tell that to Stephen Hawking who just used tech derived from Moore's Law to create AI nanoships that will literally sail to Alpha Centauri to search for life on the earth-like planet in that system.

Brian Clegg says

I am really struggling to remember a book that has irritated me as much as this one, which is a shame because it's on a very interesting and significant subject. Pedro Domingos takes us into the world of computer programs that solve problems through learning, exploring everything from back propagating neural networks to Bayesian algorithms, looking for the direction in which we might spot the computing equivalent of the theory of everything, the master algorithm that can do pretty much anything that can be done with a computer (Turing proved a long time ago that there will always be some things that can't). As the subtitle puts it, this is the quest for the ultimate learning machine that will remake our world.

So far, so good. Not only an interesting subject but one I have a personal interest in as I had some involvement in artificial intelligence many moons ago. But just reading the prologue put my hackles up. It was one of those descriptions of how a technology influences every moment of your life, as the author takes us through a typical day. Except 90% of his examples have only ever been experienced by a Silicon Valley

geek, and those that the rest of us have come across, like algorithms to make recommendations to you on shopping websites and video streaming sites, in my experience, are always so terrible that they are almost funny.

The pain carries on in part because of a kind of messianic fervour for the topic that means that the author seems convinced it is about to totally takeover the world - and like most fanatics, he presents this view while viciously attacking everyone who disagrees, from the likes of Marvin Minsky and Noam Chomsky to Black Swan author Nassim Nicholas Taleb. It's interesting that Domingos is totally dismissive of the early knowledge engineers who thought their methodology would take over the world, but can't see that his own pursuit of the 'master algorithm' (think of Lord of the Rings, but substitute 'algorithm' for 'ring') is equally likely to be a pursuit that is much easier to theorise about than to bring to success.

To make matters worse, Domingos repeatedly claims, for instance, that thanks to learning algorithms it's possible to predict the movement of the stock market, or to predict the kind of 'black swan' events that Taleb shows so convincingly are unpredictable. Yet I have never seen any evidence that this is true, it seems to go totally against what we know from chaos theory, and Domingos doesn't present any evidence, he just states it as fact. (Could you really have predicted the existence of black swans before they were discovered? How about blue ones?)

One other problem I have with the book is that the author isn't very good at explaining the complexities he is dealing with. I've seen many explanations of Bayesian statistics over the years, for instance, and this was one of the most impenetrable I've ever seen.

I can't tell you to avoid this book, because I've not come across another that introduces the whole range of machine learning options in the way that Domingos does. But any recommendation has to be made through gritted teeth because I did not like the way that information was put across.

Maria Ferreira says

Pedro Domingos é professor de Ciências da computação na Universidade de Washington. Formado no Instituto Superior Técnico em Lisboa, ganhou o prémio de inovação da SIGKDD, o mais prestigiado na área das ciências de dados.

Domingos termina a Revolução do Algoritmo Mestre com uma mensagem que encerra em si o grande propósito do mesmo, dar a conhecer aos leitores o significado de Inteligência Artificial.

”A aprendizagem automática toca a vida a cada um de nós e cabe-nos a todos decidir o que queremos fazer com ela. Com o entendimento sobre a aprendizagem automática, ficamos numa posição muito melhor para decidir sobre a partilha de dados, o futuro do trabalho, a guerra robotizada, as promessas e os perigos da IA; e quanto mais forem aqueles que tiverem este entendimento, mais provável será que evitemos armadilhas e encontremos os caminhos certos.”

As máquinas são criação do homem. Enquanto algumas pessoas ficam deslumbradas com a Inteligência Artificial das máquinas (por exemplo: eu), outras pessoas se sentem atemorizadas achando que as máquinas irão ter capacidades subversivas e atacar o homem, e até mesmo extinguir os humanos do planeta, tal como NÓS humano extinguimos algumas espécies de animais que coabitavam connosco.

Nada disso: **“As hipóteses de uma IA equipada com o Algoritmo Mestre assumir o controlo do mundo são nulas. O motivo é simples: ao contrário dos seres humanos os computadores não têm vontade própria. São produtos da engenharia e não da evolução. Mesmo um computador infinitamente potente não seria mais do que uma extensão da nossa vontade, não sendo, nada a rezear. Os Algoritmos de aprendizagem são compostos por três componentes: representação, avaliação e otimização. A representação circunscreve o que ele pode aprender, a otimização faz tudo ao seu alcance para otimizar a aprendizagem – nem mais, nem menos -maximiza a função de avaliação. A função da avaliação é determinada por nós. Um computador mais potente apenas irá otimiza-la melhor. Não existe o risco de perder o controlo, mesmo que seja um algoritmo genético.”**

Frequentemente ouvimos falar em “erros informáticos”, pois bem... isso não existe... o que existe são erros humanos, e estes, claro, culpam as máquinas por fazerem o trabalho errado, mas elas apenas fazem o que lhes mandam, e quem manda somos nós, quem cria e implementa os algoritmos são os humanos.

Impressionante a quantidade de citações que retirei do livro, deixo apenas algumas, as que considero mais relevantes:

84.17% "A empresa Narrative Science possui um sistema de IA que é capaz de escrever resumos bastante bons de um jogo de basebol, mas não romances, porque - lamento, George F. Will - a vida é muito mais que um jogo de basebol.

O senso comum é importante não só porque a nossa mãe assim nos ensinou, mas porque os computadores não o têm."

Com a rápida e crescente evolução dos tradutores automáticos, cada vez mais eficientes, facilmente adquirimos na Internet qualquer livro, em qualquer língua, que os sistemas conseguirão traduzir para a nossa língua em poucos minutos.

72.5% "Á semelhança da memória humana, a aprendizagem relacional tece uma abundante rede de associações. Ela conecta objetos percebidos...

A aprendizagem relacional é a última peça do puzzle, o ingrediente final que nos falta para a nossa alquimia, o Algoritmo-Mestre".

70.83% "Se equiparmos o robô Robby com todas as capacidades de aprendizagem, que já vimos neste livro, ele será bastante inteligente, mas ainda um pouco autista. Há de ver o mundo como uma série de objetos separados, que ele pode identificar, manipular, e sobre os quais pode até fazer previsões, mas não perceberá que o mundo é uma rede de interconexões."

45.0% "O cérebro consegue aprender qualquer coisa, mas não consegue desenvolver um cérebro. Se compreendêssemos perfeitamente a sua arquitetura, poderíamos simplesmente implementá-la em hardware, mas estamos muito longe disso."

24.72% "Aristóteles afirmou que não existe nada no intelecto que não existisse primeiro nos sentidos. Leibniz acrescentou: «Exceto o próprio intelecto.» O cérebro humano não é uma folha em branco porque não é uma folha. Uma folha é algo passiva, algo no qual escrevemos, mas o cérebro processa ativamente a informação que recebe. A memória é a folha no qual ela escreve, e sim, começa por estar em branco."

19.44% "...no nosso trabalho, o que pode ser feito por um algoritmo de aprendizagem, e o que é que não pode, e - o mais importante - **como é que posso tirar partido da aprendizagem automática para o fazer**

melhor? O computador é a nossa ferramenta, não o nosso adversário,"

9.44% "A Revolução Industrial automatizou o trabalho manual e a Revolução da Informação fez o mesmo pelo trabalho mental, mas a aprendizagem automática automatiza a própria automatização. Sem ela, os programadores tornam-se o estrangulamento que atrasa o progresso. Com ela, o ritmo do progresso aumenta.

... a revolução da aprendizagem automática causará extensas alterações económicas e sociais." Estima-se que daqui por 10 anos a IA seja vulgar nas nossas vidas: na nossa casa, no nosso trabalho e no lazer, tal como a eletricidade, a IA será imprescindível para nós.

October 25, 2017 – 0.0% "Qual é o melhor modelo para o Algoritmo-Mestre: a evolução ou o cérebro? Esta é a versão da aprendizagem automática do debate **inato versus adquirido**. E, do mesmo modo que a natureza e a educação se combinam para nos produzir a nós, seres humanos, talvez o verdadeiro Algoritmo-Mestre contenha elementos de ambas."

Afinal o que é o Algoritmo – Mestre?

É a consciência da máquina. (Estamos no bom caminho, mas ainda não a alcançamos)

Nota: este livro, de todo, não é de fácil leitura, é necessário que o leitor domine as ciências exatas. O livro aborda vários modelos matemáticos que estão na base da criação dos algoritmos e a linguagem matemática para quem não é da área torna o livro enfadonho e incompreensível. Talvez por isso o autor afirma ser um livro para profissionais e pessoas que adorem esta área do conhecimento.

Por fim uma nota para a escrita, na minha opinião está pouco cuidada, algumas lacunas, sei que o autor não domina a ciência da escrita, que provém de outra área de conhecimento, contudo falta nesta obra, duas a três leituras de aperfeiçoamento da escrita.

Yazeed says

I've finally come around to finishing this book after I started reading it more than a year ago. The Master Algorithm attempts to present a high-level overview of machine learning for the non-technical reader. The author describes the five different 'tribes' of machine learning (analogizers, evolutionaries, Bayesians, connectionists, and symbolists). The author also talks about unsupervised learning and attempts (although in a very superficial way) to combine the five different tribes into one that uses a universal machine learning algorithm that he calls 'the Master Algorithm'.

The book has several interesting anecdotes. I really enjoyed the last chapter that was mainly a discussion about (among other things) the future of AI, privacy, digital ethics, and the author's view of Kurzweil's looming 'singularity'.

Roxanne Russell says

Bill Gates put this book on his list of recommended reads this year. It interested me because of my work on

an ed tech tool to help young adults read better and enjoy their reading experience more. To get the tool right, we have had to integrate artificial intelligence and to make it better we will need machine learning. We have experts on the team for that, but I like to know what's going on around me. I can't pretend I now fully understand machine learning, but Domingos did an excellent job surveying the field so I could get a gist of it.

He lays the foundations for machine learning and then identifies 5 approaches that are being pursued now and details the formal languages they rely on:

Symbolists- Logic

Connectionists- Neural Networks

Evolutionaries- Genetic Programs

Bayesians- Graphical Modeling

Analogizers- Specific Instances

Throughout these explanations he explores the challenges:

Complexity monster

Overfitting problem

Curse of dimensionality

Exploration/Exploitation dilemma

My background in educational philosophy and epistemology helped me follow along with ease and enthusiasm as he discussed the concepts that were related to learning theories, but I could only just barely hang on for the rides through mathematical foundations.

Points of interest:

Machine learning helped Obama's campaign make strategic ad buying decisions in 2012.

Machine learning's version of the nature v nurture debate is whether or not the brain or evolution are better models for the master algorithm.

Mastering Tetris is a great step towards solving 1000s of problems because it is a basic NP completeness problem.

Maria Espadinha says

A Era dos Humanóides

Que sensação experimentariam, se algum dia entrassem de urgência num hospital, para uma intervenção cirúrgica, e dessem de caras com um robot cirurgião?!

E se numa ida ao tribunal, a ocupar a cadeira do respeitável juiz, se encontrasse um humanóide?!

Para já, tais eventos ainda figuram na lista dos incríveis impossíveis, mas...não estaremos a enveredar por aí ?!...

A verdade é que já se conhecem softwares capazes de diagnosticar doenças, e robots a colaborar em cirurgias. Estes últimos são braços e mãos capazes dum alcance e precisão inatingíveis por um ser humano. Porém, carecem de autonomia, pois ainda são manipulados por humanos!

E também já há por aí um robot advogado -- o famoso DoNotPay, cuja especialidade consiste em anular multas de estacionamento, e que já venceu uns bons milhares de casos.

E não podemos obliterar o projecto de robots geniais da Hanson Robotics, que foram especialmente criados para ensinar, entreter e servir, e assim cooperar com os humanos na construção dum mundo melhor.

As expectativas apontam para que dentro de dezenas de anos, muitas (todas?!) das funções desempenhadas por humanos, passem a sê-lo por robots.

Pois!... Ao que parece, a grande Revolução da Inteligência Artificial, está prestes a rebentar!

Para já, temos a semente, mas ainda há por aí algumas cartas para dar!...

E o seu sucesso depende directamente do Algoritmo Mestre -- o super herói deste livro! ;)

Para conhecê-lo de perto, só lendo "a revolução do algoritmo mestre" -- este mesmo livro, pois!...

“Todo o conhecimento — passado, presente e futuro — pode ser deduzido de dados por um único algoritmo de aprendizagem universal.

Eu chamo a este algoritmo o Algoritmo-Mestre.”

Este fabuloso algoritmo é capaz de aprender o que quer que seja, a partir dos dados fornecidos:

“Tudo o que temos de fazer é fornecer-lhe a quantidade suficiente do tipo certo de dados, e ele há de descobrir o conhecimento correspondente. Se lhe dermos um vídeo, ele aprende a ver. Se lhe dermos uma biblioteca, ele aprende a ler. Se lhe dermos os resultados de experiências de física, ele descobre as leis da física. Se lhe dermos dados de cristalografia do ADN, ele descobre a estrutura do ADN.”

Preparem-se para uma vida de papo para o ar!!!... ;)

Atila Iamarino says

Um livro bem denso que usa a inteligência artificial como contexto para explicar algoritmos. Passando pelos vários tipos de algoritmos de machine learning: conexionistas, bayesianos, evolutivos, simbolistas e analogistas.

Pesado, mas didático e bem humorado. Não recomendo para o público geral. Recomendo muito para quem já tem um background em matemática ou ciências em geral, ou quem realmente quer entender como funcionam os métodos de análise de big data atuais, porque depende de vários conceitos prévios. E recomendo especialmente para quem trabalha com bioinformática, pois achei a didática do autor sensacional para explicar métodos de análise como bayesiana, MCMC, cadeias de Markov, scale vector e afins. Talvez sirva para várias outras áreas com as quais não sou familiar. E mesmo para quem já tem um bom conhecimento na área, pode ser relevante, dado o tanto que o autor conhece, explica e relaciona.

Bastante atual, com vários exemplos interessantes e relevantes de companhias e problemas onde os algoritmos são aplicados. Além de um escopo muito maior do que só algoritmos, a discussão é bem situada no presente. E a discussão final sobre o que fazemos com nossos dados pessoais e como (possivelmente) lidar com a relação perda de privacidade/ganho de personalização é bastante sóbria e relevante. Me lembra um pouco a densidade e compreensão do autor do Superintelligence: Paths, Dangers, Strategies.

Só cansei de ler "cura do câncer" como o exemplo de objetivo.

Manuel Antão says

If you're into stuff like this, you can read the full review.

Machine Learning Made Easier (or NOT!): "The Master Algorithm" by Pedro Domingos Published September 22nd 2015.

How can one become an expert in ML? All one needs is a basic background in (multivariate) Calculus, Linear Algebra, and Probability. ML is math. If one wants to understand the techniques, one has to understand the math. No shortcut. If one wants to start looking into the field of ML, this book is for you. If not, stay well clear.

My background is in computer science and software engineering and I've been interested in ML since I can remember. In 2013 I took Andrew NG's ML class at Stanford University (for those of you who want to dive into stuff like this here are mynotes of the class; while learning the needed math can look daunting at first it is actually quite fun once you get into it), and I was never literally the same...After that I made some Python coding to get a feel for the real thing, which I'm still doing to this day.

Humans ARE machines, albeit biologically-based. Billions of highly interconnected neurons receiving sensory input, lots of internal feedback, and signals that go out to motors, etc. Emotions, feelings, consciousness, are all just "concepts" we've constructed through a mixture of self-introspection and communicating with other self-introspecting machines (humans).

Read on, if learning comes as second nature to you.
