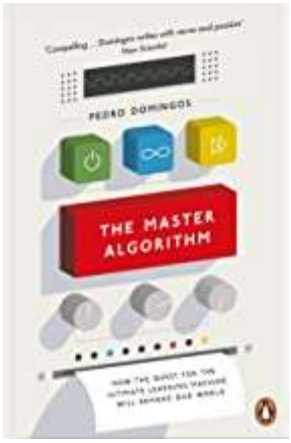

The Master Algorithm

How the Quest for the Ultimate Learning Machine will Remake Our World.
by Pedro Domingos

Penguin 26. Januar 2017
Englisch

ISBN-10: 9780141979243
ISBN-13: 978-0141979243

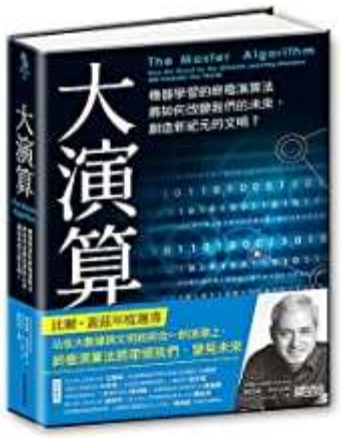

You can't learn without knowledge
("No free lunch"- Rule)

"You trust a middleman who is completely insane. Who doesn't know what he's saying."
Sandra Wachter, AI –Ethicist on trusting a chatbot. [1] Spiegel Nr. 10 March 4, 2023]

The Spiegel article warns that artificial intelligence as described in Pedro Domingos' book will be the greatest possible destabilizer of the political order. [1]

With the cover story in the Spiegel-magazine (Nr. 10, March 4, 2023, titled "The new world power"), and the hype about chatbots, I was motivated to read Paolo Domingos' standard book "The Master Algorithm" to see where we are heading.

As quoted in the above-mentioned Spiegel article, Paolo Domingos' book supposedly found its way onto Xi Jinping's bookshelf during his speech to the Chinese people on new year's eve 2018, thus setting the signal for intensive artificial intelligence (AI) research in China, heavily sponsored by the government.

Reading the book, the first thing you learn is the difference between *AI* and *Learning Machines*. The example given by Paolo is very instructive: "Learning algorithms are the seeds, data is the soil, and the learned programs are the grown plants. The machine-learning expert is like a farmer, sowing the seeds, irrigating and fertilizing the soil, and keeping an eye on the health of the crop but otherwise staying out of the way."
In other words, machine learning (ML) is a subset of AI that focuses on teaching machines how to learn from data, while AI is a broader field that encompasses various techniques, including machine learning, to create intelligent machines.[2]

Believe it or not, every algorithm, no matter how complex, can be reduced to just these three operations: AND, OR, and NOT explains Domingos.
Simple algorithms can be represented by diagrams, using different symbols for the AND, OR, and NOT operations. For example, if a fever can be caused by influenza or malaria, and you should take Tylenol for a fever and a headache, this can be expressed as follows:

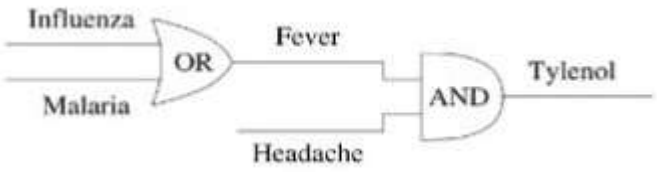

The main machine learning paradigms are *five* in number, and a chapter to each is devoted.

1. *Symbolists* view learning as the inverse of deduction and take ideas from philosophy, psychology, and logic.
2. *Connectionists* try to reverse engineer the brain and are inspired by physics.
3. *Evolutionaries* simulate evolution on the computer and draw on genetics and evolutionary biology.
4. *Bayesians* believe learning is a form of probabilistic inference and have their roots in statistics.
5. *Analogizers* learn by extrapolating from similarity judgments and are influenced by psychology and mathematical optimization.

Driven by the goal of building learning machines, Pedro Domingos tours a good portion of the intellectual history of the last hundred years and shows it in a new light.

He describes the above five principles in much detail explaining them with elaborated, understandable and plausible examples pointing out their merits, shortcomings, and flaws, working his way to the "citadel of the Master algorithm" calling it "part 1 of the master algorithm, leaving "part 2, the ultimate master algorithm" to the reader to be discovered.
Don't get discouraged when the explanations and reasoning get too specialized – you always get the message.

Here, then, is the central hypothesis of this book: All knowledge—past, present, and future—can be derived from data by a single, universal learning algorithm. The author calls this "learner" the Master Algorithm.

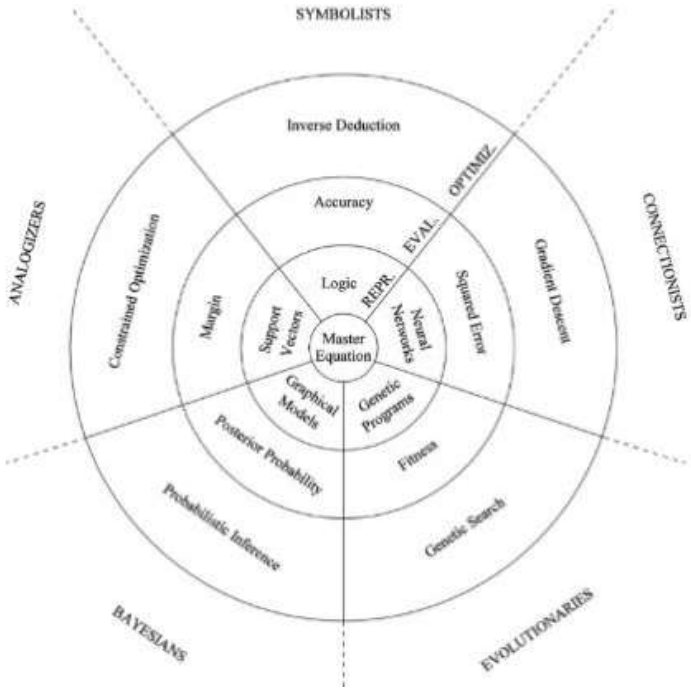

Domnigos' approach to the Master Algorithm

Pedro Domingos, professor emeritus of computer science and engineering at the University of Washington, and his co-workers have developed a unified learner which uses Markov Logic Networks (MLNs) as the representation, posterior probability as the evaluation function, and genetic search coupled with gradient descent as the optimizer.

Pedro Domingos says about his development: "I wouldn't be so rash as to call this learner the Master Algorithm. However, for one, the proof of the pudding is in the eating, and although over the last decade this algorithm (or variations of it) has been successfully applied in many areas, there are many more to which it hasn't, and so it's not yet clear just how general purpose it is. Second, there are some important problems that it doesn't solve."
You can download this learner from alchemy.cs.washington.edu. It was christened *Alchemy* to remind ourselves that, despite all its successes, machine learning is still in the alchemy stage of science.

In Chapter 10 the author presents his vision of the future world of machine learning: "Machine learning will not single-handedly determine the future, any more than any other technology; it's what we decide to do with it that counts, and now [we] have the tools to decide."

I personally wouldn't necessarily agree with all the author's conclusions e.g., that one should actively influence the "digital model" of yourself, which is created in slivers by learning machines (learners) today every time you interact with the internet anyways. Domingos' idea is, the learning machine should only know as much about you as you want, and not just learn it indirectly from what you do on the internet.
More than that, you should be able to inspect the learner's digital model of yourself and correct it as desired, and entrust it to a new type of company that is to your data what your bank is to your money, assuming no bank will misuse your money, or…?
The concept is alright, however, life experience has demonstrated often enough that this control is rather difficult to be exercised and you end up delivering even more accurate digital information about yourself prone to be misused.
The author is also convinced that it would be possible to have combat robots or other harmful robot tools learn basic ethical rules to eradicate cruelty and misuse by humans forever. Of course, a robot can only follow the ethics his human programmer tells him but they could contain some "unorthodox" human ethics which might grant him victory over his adversary.

On the other hand Domingos' plea for distributing an emerging final master algorithm "open source" is appreciated and his conviction that even the most sophisticated AI robots never will take over human intelligence because computers "have no will of their own" is reassuring!

Summary

Paolo Domingos tells you all-embracing by using illustrative examples of the basics of artificial intelligence (AI), machine learning (ML), and optimization algorithms.
After you have digested that, he embarks on a "tour de force" covering all aspects of achieving the ultimate tool, what he calls the "Master Algorithm", but admits that there is some way to go yet.
However, he invites and encourages every reader of his book to join the global community to achieve this goal based on the knowledge provided in his book and the *Alchemy* algorithm, developed by the author and his researchers as a unified algorithm combining suitable features of the five optimization and machine learning domains.

So, as he says in the prolog if you're curious what all the hubbub surrounding big data and machine learning is about, or if your main interest is in the business uses of machine learning, this is your book to read.

If you’re a citizen or policy maker concerned with the social and political issues raised by big data and machine learning, this book will give you a primer on the technology.
If you’re a scientist or engineer, machine learning is a powerful armory that you don’t want to be without. If you’re a machine-learning expert, you’re already familiar with much of what the book covers, but you’ll also find in it many fresh ideas.
If you’re a student of any age—a high schooler wondering what to major in, a college undergraduate deciding whether to go into research, or a seasoned professional considering a career change—this book might spark in you an interest in this fascinating field.
Last but not least, if you have an appetite for wonder, machine learning is an intellectual feast and you’re invited to participate—RSVP! …
…to which I wholeheartedly agree, and additionally you become highly aware of how precious your personal data are and that it might have global influence to consciously answer the question “to share or not to share” every day.
Go get the book and read it!
It makes you understand the internet world much better and cope with it in the future.

References:

[1] Spiegel Nr. 10 March 4, 2023
[2] Definition provided by chatGPT

March 2023, Joachim J. Kehr, Editor Journal of Space Operations & Communicator
https://opsjournal.org