

Machine Learning and the Prospect of a Master Algorithm

The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World. By Pedro Domingos, Basic Books, New York, 2015, 352 pages, \$29.99.

Machine learning (ML) has suddenly become very hot. It feels like every week, a machine learning algorithm achieves a major advance in some exciting application; and every other day the popular media publishes exaggerated accounts of these advances, or breathless predictions of what the future has in store. Over the last thirty-five years, ML has gone from being a largely-abandoned niche area of artificial intelligence to a central paradigm within computer science. Machine learning completely dominates applications such as computer vision, speech understanding, and recommender systems in addition to all forms of natural language processing.

Pedro Domingos' new book, *The Master Algorithm*, presents an introduction to the techniques of machine learning (ML), written for a popular audience. It seems to be the first popular book on ML, so it is certainly an important contribution toward general understanding of this enormously important technology.

Chapters 3-7 of Domingos' book survey the five major approaches to supervised ML: symbolic, connectionist, evolutionary, probabilistic, and exemplar-based. Chapter 8 covers unsupervised learning and meta-learning. These chapters are well-written, clear, and balanced; he carefully describes the intuitions behind each approach, the practical successes the approach has attained, its strengths, and its inherent weaknesses. The true believers in each of the approaches will undoubtedly complain that Domingos shortchanged their own approach and handled the other approaches much too politely. The author gives very clear accounts of overfitting and of "the curse of dimensionality," the two great hazards of ML. He almost entirely avoids the use of mathematics.

At his best, Domingos can be extremely good. He offers sharp, insightful statements of points easily overlooked. For instance, he writes that,

"The most important thing about an equation is all the quantities that don't appear in it; once we know what the essentials are, figuring out how they depend on each other is often the easier part."

Overall, these six chapters are a fine introduction to the state of the art of ML.

If only Domingos had been content with that! Unfortunately he has other things in mind, and these seriously degrade the book's overall quality. I found the prologue and chapters 1 and 2 so unengaging that I almost didn't make it to chapter 3, and my objections returned in full force upon reaching chapters 9 and 10.

In addition to explaining the actual state of the art, Domingos has two additional goals. The first, explicit in the book's title, is to develop the idea of a "Master Algorithm" for machine learning, which will subsume all other forms of learning and be able to learn anything that can be learned, and to argue that Domingos' own "Alchemy" system is a major step in that direction. The second objective, implicit but pervasive and very conspicuous, is to hype the present accomplishments and future impacts of ML as raucously as possible. In pursuit of this goal, the book is generally overwrought, sometimes seriously misleading or simply wrong, and occasionally really quite strange.

For example, in presenting connectionist models, Domingos gets carried away with phase transitions and the idea of an S-curve, which he characteristically calls "the most important curve in the world." He spends an entire page enumerating S-curves and phase transitions, including some very doubtful examples. Paradigm shifts in science and the fall of empires are supposedly S-curves. Falling in love, getting a job, and losing a job are allegedly phase transitions; clearly at this point "phase transition" just translates to "important change."

Domingos' view of computerization's past impact on science is wildly exaggerated. He writes that

"If computers hadn't been invented, science would have ground to a halt in the second half of the twentieth century."

It certainly would have been impeded, and some discoveries would have been impossible, but it seems safe to say that only a small fraction of scientific discoveries before about 1990 relied critically on computers.

Moreover, Domingos' view of the future impact of the Master Algorithm is fantasy:

"Science today is thoroughly balkanized, a Tower of Babel where each subcommunity can see only into a few adjacent subcommunities."

BOOK REVIEW

By Ernest Davis

He then proceeds to state that the Master Algorithm would offer a unified perspective of all science, one that could even lead to a new theory of everything. Even if a Master Algorithm could derive each scientific theory from data, there is no reason to expect that would provide a unifying view, rather than a collection of separate theories. Domingos continues by saying that

"The Master Algorithm is the germ of every theory: all we need to add to it to obtain theory X is the minimum amount of data required to induce it. (In the case of physics, that would be the results of perhaps a few hundred key experiments)."

That estimate is certainly too low by at least a factor of one hundred.

In chapter 10, Domingos presents his vision of the future. Some aspects of this seem reasonable; in particular, I agree with his belief that there is no danger of computers developing their own goals à la Skynet and exterminating or enslaving us. Other thoughts, however, appear both weird and dreary, like the following Baudrillardesque nightmare:

"In this rapidly approaching future, you're not going to be the only one with a 'digital half' doing your bidding twenty-four hours a day. Everyone will have a detailed model of him- or herself, and these models will talk to each other all the time."

He proceeds to give examples for how this will happen:

"If company X is looking to hire, its model will interview your model. It will be a lot like a real, flesh-and-blood interview . . . but it will take only a fraction of a second . . . Same with dating."¹

Domingos does not explain why a company would trust that your avatar is an accurate portrayal of yourself; nor why the company should not hire the avatar instead, since it presumably does the same things thousands of times faster.

¹ Felicia Day's (2009) "Do you wanna date my avatar?" is a particularly insightful discussion of this issue: <https://www.youtube.com/watch?v=urNyg1ftMIU>.

The issue most important to Domingos is his Master Algorithm conjecture:

"All knowledge — past, present, and future — can be derived from data by a single, universal learning algorithm."

In chapter 2, he presents many different arguments for the hypothesis, drawing evidence from fields such as neuroscience, evolution, physics, statistics, and computer science.

In my opinion, these assertions are incoherent. The different arguments rely on varied interpretations of the terms "data," "single," "universal," "learning," and "algorithm," and therefore point in different directions. For instance, Domingos bases his case for the neuroscience perspective on the uniform structure of the cortex. Yet in chapter 9, he is willing to consider a Master Algorithm which calls on a variety of very different algorithms and combines their votes. The uniform structure of the cortex is, if anything, evidence *against* a Master Algorithm of this kind. The kind of algorithm he proposes would be consistent with a brain that combines pieces of very diverse structures.

As Domingos himself observes, the Master Algorithm is very unlike existing trends in machine learning, which generally involve highly-handcrafted algorithms for fairly narrowly-defined tasks. Indeed, it is unlike anything that we have seen in computer science. Domingos' argument from the computer science viewpoint rests on the fact that the many NP-complete problems are, in a certain sense, actually the same problem. However, it seems to me that this points in exactly the opposite direction. NP-complete problems are all reducible to the problem of Boolean satisfiability, but there are, probably, tens of thousands of different algorithms for solving specific NP-complete problems, depending on the specifics of the problem and on the desired features of the solution and the algorithm.

The Master Algorithm does a fine job of explaining machine learning techniques to the general public. Unfortunately, it is also often extremely misleading. If you have a lay friend who wants to understand ML, by all means recommend it, but do so with a warning.

Ernest Davis is a professor of computer science at the Courant Institute of Mathematical Sciences, NYU.