

**BREAKING NEWS**

Porsche CEO Matthias Muller to become new Volkswagen CEO, person familiar with the matter says

# THE WALL STREET JOURNAL.

This copy is for your personal, non-commercial use only. To order presentation-ready copies for distribution to your colleagues, clients or customers visit <http://www.djreprints.com>.

<http://www.wsj.com/articles/the-sum-of-human-knowledge-1442610803>

ARTS | BOOKS | BOOKSHELF


## The Sum of Human Knowledge

Computers govern how long the microwave heats food or the dryer spins clothes. Can they learn to form ideas and theories about the world around them as well?

Russia's pavilion at the Architecture Exhibition of the 2012 Venice Biennale incorporated hundreds of QR codes.


*PHOTO: ULLSTEIN BILD VIA GETTY IMAGES*

Hide



Option

Dismiss



By **DANIEL LEVITIN**

Sept. 18, 2015 5:13 p.m. ET

**In a particularly** memorable episode of CBS’s “The Big Bang Theory,” physicist Sheldon Cooper and neurobiologist Amy Farrah Fowler get into an argument, a game of intellectual one-upmanship that threatens their relationship. Sheldon claims that “a grand unified theory, insofar as it explains everything, will ipso facto explain neurobiology.” Amy counters: “Yes, but if I’m successful, I will be able to map and reproduce your thought process in deriving a grand unified theory and therefore subsume your conclusions under my paradigm.”

The first contention is a familiar one—the second more surprising. But could it be true? Pedro Domingos, a computer scientist at the University of Washington, believes that a version of Amy’s notion is indeed true. All knowledge could be reproduced—and new knowledge produced—by “subsuming” human thought processes. And he thinks computer scientists are well on their way to doing it.

The central hypothesis of his new book is that “all knowledge—past, present, and future—can be derived from data by a single, universal learning algorithm.” Mr. Domingos calls this the Master Algorithm. He states that the discovery of the Master Algorithm, and its implementation by learning computers, will replace string theory, genetics and psychology, because it will be able to generate all hypotheses as well as prove theories that have yet to be formulated.

---

THE MASTER ALGORITHM

---

*By Pedro Domingos*

*Basic, 329 pages, \$29.99*

While a Master Algorithm is at best far off in the future, we’ve begun to see some applications of computer learning in our daily lives, and a measure of their success is that most of us aren’t aware of this. The Nest “learning thermostat” is an example.

It learns when you’re home and when you’re out and

what temperature you like to keep the house at various times of day and days of the week. Porsche’s latest Tiptronic transmission learns a driver’s shifting patterns and adapts itself to different driving styles using individual car keys as a means of assigning the correct profile to different drivers.

Indeed, one of the largest advances of recent years is that computers have become invisible in a great many cases, performing their functions in the background instead of being these big clunky things on our desks. This is mostly seen (or not seen) in home appliances and automotive applications. Computers govern how long the microwave heats food and the dryer spins clothes and how

much air to mix with fuel in an internal combustion engine. Most of us don't even know that there are computers inside these devices. They do their work by sensing environmental conditions, such as the level of humidity inside the tumbling dryer, but they don't learn anything. If they could, the dryer would know that when you put in jeans it's going to take longer than when you put in silk undies, and it would adjust the heat accordingly. It does not.

## Find your new home now ...

To teach a computer to learn takes inductive reasoning—that is, using data from a small number of instances to generate hypotheses and theories that apply to a very large number of cases. “The Master Algorithm” is, in part, an account of how the known and the unknown are duking it out in laboratories across the world as computer scientists, neuroscientists and engineers try to develop computers that can effectively program themselves, that can learn (more or less the way humans do) and effectively reprogram themselves. The author reviews the major streams within artificial intelligence and provides a bit of history to go along with it, framing the development of the field. But he—like the characters of “The Big Bang Theory”—goes too far in contending that one science will ever be able to replace all the others or that a single algorithm will be able to replace the sum total of human creativity and reasoning.

It's trivially obvious that to generate new knowledge, computers will have to learn. It's far-fetched to think that a single algorithm will accomplish all that, unless we stretch the notion of either “single” or “algorithm” beyond reasonable limits. The most sophisticated and accomplished computer we know of—the

brain—does not appear to do its work with a single algorithm. Rather, it uses a hodgepodge of heuristics, statistical inferences and special-purpose faculties to learn about how the world works. There is no evidence that learning a language follows the same algorithm as, say, learning to do long division, shoot a free throw or solve “Where’s Waldo?” in a picture book.

One recent big success of a computer algorithm has been collaborative filtering, the algorithm used by Amazon to predict consumer preferences based on past preferences. It functions like B.F. Skinner’s model of animal behavior, in which actions are broken down as a system of associational behaviors: See food, salivate. Liked Ray Kurzweil’s “The Singularity Is Near”? Try “The Master Algorithm.” Noam Chomsky famously argued that the capacity to learn language couldn’t follow this kind of algorithm, due to the complexity of language, the speed with which children learn it, and the fact that all of us can generate new sentences that have never before been uttered and may never be (such as “President Trump announced today that Jorge Ramos will be his press secretary”).

Our genes carry instructions for brain development, including possibly a language faculty, Mr. Chomsky has suggested. But though it is true that the genetic code is a set of instructions (as are computer programs), there is no evidence that the genetic code follows a single “master algorithm.”

Collaborative filtering can do many things, but it cannot teach a computer to learn language. Going through large amounts of data may well cure cancer. But that’s not the same as solving problems the way humans do or coming up with a master algorithm that solves all problems.

Mr. Domingos believes that computers should be able to perform induction, allowing them to go beyond what they’ve observed and discover generalized laws

that apply to all instances of something, even though they've only observed a few. For example, Newton's law that "an object at rest will tend to remain at rest, unless acted upon by outside forces" is a leap of induction, a generalization of specific instances. None of us can observe all possible objects and all possible outcomes, but based on observation and an understanding of natural processes, Newton formulated a general principle.

We do this in hundreds of ways, small and large, in our daily lives: I don't buy milk from that store on the corner because their prices are much higher than the supermarket a little further down the street; I keep milk refrigerated so that it won't turn sour. Now it's possible that the induction is wrong, or incomplete. The convenience store may have lowered its prices. Someone else might solve the same problem by purchasing European-style boxed milk that doesn't require refrigeration. Induction is not infallible in the way that a valid deduction is. And in fact this is all part of learning: Newton's induction about physical laws didn't hold at speeds close to light or in the subatomic realm, and it took Einstein, Heisenberg, Schrödinger and others to perform a new induction and generate new laws.

One crucial way of teaching computers induction is through the use of back-propagation neural networks, of which "The Master Algorithm" provides a compact history. In back-prop systems, artificial "neurons" take information from some kind of input (say, a music sound file) and then generate an output (such as an artist classification for that sound file). Back-prop networks receive supervised training on a data set and adjust themselves internally, without intervention from the computer scientist, to correct any errors in their output.

Back-prop's first significant real-world applications, beginning in the early 1990s, came in predicting the stock market. Because such systems could detect

small patterns in very noisy data, they beat the linear models that were then prevalent in finance, and so their use spread. A typical investment fund would train a separate network for each of a large number of stocks, let the networks pick the most promising ones and then have human analysts decide which of those to invest in. (Does this method work? It's not a secret that doctoral students who are studying machine learning tend not to stay in academia. I lost one myself to Wall Street, when, prior to even starting his thesis, he was offered twice as much money as he would ever make as an assistant professor.)

Back-prop again underscores the tension between the known and the unknown: By design, the way a back-prop system works is left hidden. Engineers and computer scientists are often content that something works. Scientists, including Mr. Chomsky, say that science won't advance if we only know that something works; we have to know how and why.

"The Master Algorithm" is in that category of book that promises to teach you about "the next big thing." But there's not all that much news in "The Master Algorithm," other than that the same family of algorithms we've been using for 20 years are doing a better job because they have access to big data. We are beginning to see this in cancer research. The most promising treatments involve computers sifting through large amounts of data (such as the repository found in David Haussler's UCSC Genome Browser) and matching cancer biopsy genomes to help craft treatments customized to the particular individual.

Donald Rumsfeld had it right: There are known unknowns and unknown unknowns. Although he used fractured language to express himself, he hit on a fundamental problem: That is, our biggest impediment to sound decision-making is our own ignorance of what we do and don't know. Some of what we don't know we can learn, and as yet the way by which we do this remains largely

hidden from us. And some of what we don't know we will probably never learn. This could be because some things are probably unknowable (such as "Where did all the matter that existed before the big bang come from?") or because our questions are malformed (such as "Will neurobiology replace physics or will physics replace neurobiology?").

A third option for the unknowable is that some phenomena are simply too complex to understand thoroughly. Physicists and neurobiologists agree that both the universe without and the universe within—the human brain—contain elements that are deterministic but not predictable. That is, if we could measure every single factor from the macroscopic to the microscopic and do so precisely, we could predict everything from when the next hurricane will be to whom Donald Trump will kick out of his next press conference. Such an approach would finally unite neurobiology and physics into a gigantic set of measurements. But it is the manifest impossibility of measuring all those factors that renders a deterministic system unpredictable.

Humans don't learn language, music, chess, basketball or mathematics by measuring everything or by applying a single algorithm. Computers don't either. Mr. Domingos believes that humans do and that we just haven't discovered how yet, and he believes that we humans have the capacity to invent computers that will discover how on their own. If he's right, neither Amy nor Sheldon win the argument about the primacy of their respective scientific endeavors; Pedro Domingos does. Who's right? So far, that is another known unknown.

*—Mr. Levitin is the author of, most recently, "The Organized Mind: Thinking Straight in the Age of Information Overload."*



Copyright 2014 Dow Jones & Company, Inc. All Rights Reserved

This copy is for your personal, non-commercial use only. Distribution and use of this material are governed by our Subscriber Agreement and by copyright law. For non-personal use or to order multiple copies, please contact Dow Jones Reprints at 1-800-843-0008 or visit [www.djreprints.com](http://www.djreprints.com).