


A Grand Unified Theory of Artificial Intelligence

March 30 2010



In the 1950s and '60s, artificial-intelligence researchers saw themselves as trying to uncover the rules of thought. But those rules turned out to be way more complicated than anyone had imagined. Since then, artificial-intelligence (AI) research has come to rely, instead, on probabilities -- statistical patterns that computers can learn from large sets of training data.

The probabilistic approach has been responsible for most of the recent progress in artificial intelligence, such as voice recognition systems, or

the system that recommends movies to Netflix subscribers. But Noah Goodman, an MIT research scientist whose department is Brain and Cognitive Sciences but whose lab is [Computer Science](#) and Artificial Intelligence, thinks that AI gave up too much when it gave up rules. By combining the old rule-based systems with insights from the new probabilistic systems, Goodman has found a way to model thought that could have broad implications for both AI and cognitive science.

Early AI researchers saw thinking as logical inference: if you know that birds can fly and are told that the waxwing is a bird, you can infer that waxwings can fly. One of AI's first projects was the development of a mathematical language — much like a computer language — in which researchers could encode assertions like “birds can fly” and “waxwings are birds.” If the language was rigorous enough, computer algorithms would be able to comb through assertions written in it and calculate all the logically valid inferences. Once they'd developed such languages, AI researchers started using them to encode lots of commonsense assertions, which they stored in huge databases.

The problem with this approach is, roughly speaking, that not all birds can fly. And among birds that can't fly, there's a distinction between a robin in a cage and a robin with a broken wing, and another distinction between any kind of robin and a penguin. The mathematical languages that the early AI researchers developed were flexible enough to represent such conceptual distinctions, but writing down all the distinctions necessary for even the most rudimentary cognitive tasks proved much harder than anticipated.

Embracing uncertainty

In probabilistic AI, by contrast, a computer is fed lots of examples of something — like pictures of birds — and is left to infer, on its own, what those examples have in common. This approach works fairly well

with concrete concepts like “bird,” but it has trouble with more abstract concepts — for example, flight, a capacity shared by birds, helicopters, kites and superheroes. You could show a probabilistic system lots of pictures of things in flight, but even if it figured out what they all had in common, it would be very likely to misidentify clouds, or the sun, or the antennas on top of buildings as instances of flight. And even flight is a concrete concept compared to, say, “grammar,” or “motherhood.”

As a research tool, Goodman has developed a computer programming language called Church — after the great American logician Alonzo Church — that, like the early AI languages, includes rules of inference. But those rules are probabilistic. Told that the cassowary is a bird, a program written in Church might conclude that cassowaries can probably fly. But if the program was then told that cassowaries can weigh almost 200 pounds, it might revise its initial probability estimate, concluding that, actually, cassowaries probably can’t fly.

“With probabilistic reasoning, you get all that structure for free,” Goodman says. A Church program that has never encountered a flightless bird might, initially, set the probability that any bird can fly at 99.99 percent. But as it learns more about cassowaries — and penguins, and caged and broken-winged robins — it revises its probabilities accordingly. Ultimately, the probabilities represent all the conceptual distinctions that early AI researchers would have had to code by hand. But the system learns those distinctions itself, over time — much the way humans learn new concepts and revise old ones.

“What’s brilliant about this is that it allows you to build a cognitive model in a fantastically much more straightforward and transparent way than you could do before,” says Nick Chater, a professor of cognitive and decision sciences at University College London. “You can imagine all the things that a human knows, and trying to list those would just be an endless task, and it might even be an infinite task. But the magic trick

is saying, ‘No, no, just tell me a few things,’ and then the brain — or in this case the Church system, hopefully somewhat analogous to the way the mind does it — can churn out, using its probabilistic calculation, all the consequences and inferences. And also, when you give the system new information, it can figure out the consequences of that.”

Modeling minds

Programs that use probabilistic inference seem to be able to model a wider range of human cognitive capacities than traditional cognitive models can. At the 2008 conference of the [Cognitive Science](#) Society, for instance, Goodman and Charles Kemp, who was a PhD student in BCS at the time, presented work in which they’d given human subjects a list of seven or eight employees at a fictitious company and told them which employees sent e-mail to which others. Then they gave the subjects a short list of employees at another fictitious company. Without any additional data, the subjects were asked to create a chart depicting who sent e-mail to whom at the second company.

If the e-mail patterns in the sample case formed a chain — Alice sent mail to Bob who sent mail to Carol, all the way to, say, Henry — the human subjects were very likely to predict that the e-mail patterns in the test case would also form a chain. If the e-mail patterns in the sample case formed a loop — Alice sent mail to Bob who sent mail to Carol, and so on, but Henry sent mail to Alice — the subjects predicted a loop in the test case, too.

A program that used probabilistic inference, asked to perform the same task, behaved almost exactly like a human subject, inferring chains from chains and loops from loops. But conventional cognitive models predicted totally random e-mail patterns in the test case: they were unable to extract the higher-level concepts of loops and chains. With a range of collaborators in the Department of Brain and Cognitive

Sciences, Goodman has conducted similar experiments in which subjects were asked to sort stylized drawings of bugs or trees into different categories, or to make inferences that required guessing what another person was thinking. In all these cases — several of which were also presented at the [Cognitive Science](#) Society's conference — Church programs did a significantly better job of modeling human thought than traditional [artificial-intelligence](#) algorithms did.

Chater cautions that, while Church programs perform well on such targeted tasks, they're currently too computationally intensive to serve as general-purpose mind simulators. "It's a serious issue if you're going to wheel it out to solve every problem under the sun," Chater says. "But it's just been built, and these things are always very poorly optimized when they've just been built." And Chater emphasizes that getting the system to work at all is an achievement in itself: "It's the kind of thing that somebody might produce as a theoretical suggestion, and you'd think, 'Wow, that's fantastically clever, but I'm sure you'll never make it run, really.' And the miracle is that it does run, and it works."

Provided by Massachusetts Institute of Technology

Citation: A Grand Unified Theory of Artificial Intelligence (2010, March 30) retrieved 24 April 2024 from <https://phys.org/news/2010-03-grand-theory-artificial-intelligence.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.