

DANIEL ZENDER

Intelligent Machines

We analyzed 16,625 papers to figure out where AI is headed next

Our study of 25 years of artificial-intelligence research suggests the era of deep learning may come to an end.

by Karen Hao January 25, 2019

Almost everything you hear about artificial intelligence today is thanks to deep learning. This category of algorithms works by using statistics to find patterns in data, and it has proved immensely powerful in mimicking human skills such as our ability to see and hear. To a very narrow extent, it can even emulate our ability to reason. These capabilities power Google's search, Facebook's news feed, and Netflix's recommendation engine—and are transforming industries like health care and education.

Advertisement

But though deep learning has singlehandedly thrust AI into the public eye, it represents just a small blip in the history of humanity's quest to replicate our own intelligence. It's been at the forefront of that effort for less than 10 years. When you zoom out on the whole history of the field, it's easy to realize that it could soon be on its way out.

Sign up for **The Algorithm**
Artificial intelligence, demystified

Your email

Sign Up

Stay updated on MIT Technology Review initiatives and events? Yes No

“If somebody had written in 2011 that this was going to be on the front page of newspapers and magazines in a few years, we would’ve been like, ‘Wow, you’re smoking something really strong,’” says Pedro Domingos, a professor of computer science at the University of Washington and author of *The Master Algorithm*.

The sudden rise and fall of different techniques has characterized AI research for a long time, he says. Every decade has seen a heated competition between different ideas. Then, once in a while, a switch flips, and everyone in the community converges on a specific one.

At MIT Technology Review, we wanted to visualize these fits and starts. So we turned to one of the largest open-source databases of scientific papers, known as the arXiv (pronounced “archive”). We downloaded the abstracts of all 16,625 papers available in the “artificial intelligence” section through November 18, 2018, and tracked the words mentioned through the years to see how the field has evolved.

The number of papers we downloaded from the arXiv

All of the papers available in the “artificial intelligence” section through November 18, 2018



Chart: MIT Technology Review • Source: arXiv.org • Created with Datawrapper

Through our analysis, we found three major trends: a shift toward machine learning during the late 1990s and early 2000s, a rise in the popularity of neural networks beginning in the early 2010s, and growth in reinforcement learning in the past few years.

There are a couple of caveats. First, the arXiv's AI section goes back only to 1993, while the term "artificial intelligence" dates to the 1950s, so the database represents just the latest chapters of the field's history. Second, the papers added to the database each year represent a fraction of the work being done in the field at that moment. Nonetheless, the arXiv offers a great resource for gleaning some of the larger research trends and for seeing the push and pull of different ideas.

A MACHINE-LEARNING PARADIGM

The biggest shift we found was a transition away from knowledge-based systems by the early 2000s. These computer programs are based on the idea that you can use rules to encode all human knowledge. In their place, researchers turned to machine learning—the parent category of algorithms that includes deep learning.

Among the top 100 words mentioned, those related to knowledge-based systems—like "logic," "constraint," and "rule"—saw the greatest decline. Those related to machine learning—like "data," "network," and "performance"—saw the highest growth.

Machine learning eclipses knowledge-based reasoning

Change in mentions per 1,000 words for the top 100 words





Chart: MIT Technology Review • Source: arXiv.org • Created with Datawrapper



Chart: MIT Technology Review • Source: arXiv.org

The reason for this sea change is rather simple. In the ’80s, knowledge-based systems amassed a popular following thanks to the excitement surrounding ambitious projects that were attempting to re-create common sense within machines. But as those projects unfolded, researchers hit a major problem: there were simply too many rules that needed to be encoded for a system to do anything useful. This jacked up costs and significantly slowed ongoing efforts.

Machine learning became an answer to that problem. Instead of requiring people to manually encode hundreds of thousands of rules, this approach programs machines to extract those rules automatically from a pile of data. Just like that, the field abandoned knowledge-based systems and turned to refining machine learning.

THE NEURAL-NETWORK BOOM

Under the new machine-learning paradigm, the shift to deep learning didn't happen immediately. Instead, as our analysis of key terms shows, researchers tested a variety of methods in addition to neural networks, the core machinery of deep learning. Some of the other popular techniques included Bayesian networks, support vector machines, and evolutionary algorithms, all of which take different approaches to finding patterns in data.

Neural networks take over other machine-learning methods

Percentage of papers that mention each method



Hover over the chart areas to see their labels.

Chart: MIT Technology Review • Source: arXiv.org • Created with Datawrapper

Through the 1990s and 2000s, there was steady competition between all of these methods. Then, in 2012, a pivotal breakthrough led to another sea change. During the annual ImageNet competition, intended to spur progress in computer vision, a researcher named Geoffrey Hinton, along with his colleagues at the University of Toronto, achieved the best accuracy in image recognition by an astonishing margin of more than 10 percentage points.

The technique he used, deep learning, sparked a wave of new research—first within the vision community and then beyond. As more and more researchers began using it to achieve impressive results, its popularity—along with that of neural networks—exploded.

THE RISE OF REINFORCEMENT LEARNING

In the few years since the rise of deep learning, our analysis reveals, a third and final shift has taken place in AI research.

As well as the different techniques in machine learning, there are three different types: supervised, unsupervised, and reinforcement learning. Supervised learning, which involves feeding a machine labeled data, is the most commonly used and also has the most practical applications by far. In the last few years, however, reinforcement learning, which mimics the process of training animals through punishments and rewards, has seen a rapid uptick of mentions in paper abstracts.

Reinforcement learning is gaining momentum

Share of papers that mention it compared to any type of machine learning



Chart: MIT Technology Review • Source: arXiv.org • Created with Datawrapper

The idea isn't new, but for many decades it didn't really work. "The supervised-learning people would make fun of the reinforcement-learning people," Domingos says. But, just as with deep learning, one pivotal moment suddenly placed it on the map.

That moment came in October 2015, when DeepMind's AlphaGo, trained with reinforcement learning, defeated the world champion in the ancient game of Go. The effect on the research community was immediate.

THE NEXT DECADE

Our analysis provides only the most recent snapshot of the competition among ideas that characterizes AI research. But it illustrates the fickleness of the quest to duplicate intelligence. “The key thing to realize is that nobody knows how to solve this problem,” Domingos says.

Many of the techniques used in the last 25 years originated at around the same time, in the 1950s, and have fallen in and out of favor with the challenges and successes of each decade. Neural networks, for example, peaked in the ’60s and briefly in the ’80s but nearly died before regaining their current popularity through deep learning.

Every decade, in other words, has essentially seen the reign of a different technique: neural networks in the late ’50s and ’60s, various symbolic approaches in the ’70s, knowledge-based systems in the ’80s, Bayesian networks in the ’90s, support vector machines in the ’00s, and neural networks again in the ’10s.

The 2020s should be no different, says Domingos, meaning the era of deep learning may soon come to an end. But characteristically, the research community has competing ideas about what will come next—whether an older technique will regain favor or whether the field will create an entirely new paradigm.

“If you answer that question,” Domingos says, “I want to patent the answer.”

Learn from the humans leading the way in deep learning at EmTech Next. Register Today!
June 11-12, 2019
Cambridge, MA
Register now ⟶

Share

Tagged AI, deep learning, machine learning

Karen Hao

Karen Hao is the artificial intelligence reporter for *MIT Technology Review*. In particular she covers the ethics and social impact of the technology as well as its applications for social good. She also writes the AI newsletter, the Algorithm,...
More

Related Video

More videos →

Intelligent Machines

Accelerating Discovery with Self-Driving Laboratories 27:00

Intelligent Machines

How to Develop Your Own AI Playbook 27:32

Intelligent Machines

AI Transforming the Developing World 20:56

More from Intelligent Machines

Artificial intelligence and robots are transforming how we work and live.

01

A robot has figured out how to use tools

In a startling demonstration, the machine drew on experimentation, data, and observation of humans to learn how simple implements could help it achieve a task.

by Will Knight



02 This robot can sort recycling by giving it a squeeze

The robot, called RoCycle, uses pincers to pick through garbage and identify what materials each bit contains. It could help reduce how much waste gets sent to landfill.

by Douglas Heaven

03 This may be the Apple II of AI-driven robot arms

A new low-cost robot arm that can be controlled using a virtual-reality headset will make it easier to experiment with AI and robotics.

by Will Knight



More from Intelligent Machines ⟶

FROM OUR ADVERTISERS

In association with **VMware**

Digital transformation sparks innovation in networking

In association with **Google**

For data-savvy marketers, there's a new keyword: Intent

In association with **Google**

Machine learning teaches marketers to cultivate a growth mindset

In association with **Oracle and Intel**

Machine learning - driven analytics: Key to digital transformation

Want more award-winning journalism?
Subscribe to Print + All Access Digital.

Print + All Access Digital $79.95/ye

Subscribe

The best of MIT Technology Review in print and online, plus unlimited access to our online archive, an ad-free web experience, discounts to MIT Technology Review events, and The Download delivered to your email in-box each weekday.

See details+

*Prices are for U.S. residents only
See international prices

Follow us

The Algorithm Artificial intelligence, demystified

Enter your email. Get the newsletter

Sign Up

Stay updated on MIT Technology Review initiatives and events? Yes No

MIT Technology Review

The mission of MIT Technology Review is to bring about better-informed and more conscious decisions about technology through authoritative, influential, and trustworthy journalism.

Browse International Editions +

Company +

Your Account +

Customer Support +

More +

Policies +

MIT Technology Review © 2019 v.|e^IT|