Brain-inspired algorithm helps AI systems multitask and remember

October 17 2018

Credit: AI-generated image (disclaimer)

Behind most of today's artificial intelligence technologies, from self-driving cars to facial recognition and virtual assistants, lie artificial neural networks. Though based loosely on the way neurons communicate in the brain, these "deep learning" systems remain incapable of many basic functions that would be essential for primates and other organisms.

However, a new study from University of Chicago neuroscientists found that adapting a well-known [brain](link) mechanism can dramatically improve the ability of [artificial neural networks](link) to learn [multiple tasks](link) and avoid the persistent AI challenge of "catastrophic forgetting." The study, published in *Proceedings of the National Academy of Sciences*, provides a unique example of how neuroscience research can inform new computer science strategies, and, conversely, how AI technology can help scientists better understand the human brain.

When combined with previously reported methods for stabilizing synaptic connections in artificial neural networks, the new algorithm allowed single artificial neural networks to learn and perform hundreds of tasks with only minimal loss of accuracy, potentially enabling more powerful and efficient AI technologies.

"Intuitively, you might think the more tasks you want a [network](link) to know, the bigger the network might have to be," said David Freedman, professor of neurobiology at UChicago. "But the brain suggests there's probably some efficient way of packing in lots of knowledge into a fairly small network. When you look at parts of the brain involved in higher cognitive functions, you tend to find that the same areas, even the same cells, participate in many different functions. The idea was to draw inspiration from what the brain does in order to solve challenges with neural networks."

In artificial neural networks, "catastrophic forgetting" refers to the difficulty in teaching the system to perform new skills without losing previously learned functions. For example, if a network initially trained to distinguish between photos of dogs and cats is then re-trained to distinguish between dogs and horses, it will lose its earlier ability.

"If you show a trained neural network a new [task](link), it will forget about its previous task completely," said Gregory Grant, AB'18, who is now a

researcher in the Freedman lab. "It says, 'I don't need that information,' and overwrites it. That's catastrophic forgetting. It happens very quickly; within just a couple of iterations, your previous task could be utterly obliterated."

By contrast, the brain is capable of "continual learning," acquiring new knowledge without eliminating old memories, even when the same neurons are used for multiple tasks. One strategy the brain uses for this learning challenge is the selective activation of cells or cellular components for different tasks—essentially turning on smaller, overlapping sub-networks for each individual skill, or under different contexts.

The UChicago researchers adapted this neuroscientific mechanism to artificial neural networks through an algorithm they called "context-dependent gating." For each new task learned, only a random 20 percent of a neural network is activated. After the network is trained on hundreds of different tasks, a single node might be involved in dozens of operations, but with a unique set of peers for each individual skill.

When combined with methods previously developed by Google and Stanford researchers, context-dependent gating allowed networks to learn as many as 500 tasks with only a small decrease in accuracy.

"It was a little bit surprising that something this simple worked so well," said Nicolas Masse, a postdoctoral researcher in the Freedman lab. "But with this method, a fairly medium-sized network can be carved up a whole bunch of ways to be able to learn many different tasks if done properly."

As such, the approach likely has great potential in the growing AI industry, where companies developing autonomous vehicles, robotics and other smart technologies need to pack complex learning capabilities

into consumer-level computers. The UChicago team is currently working with the Polsky Center for Entrepreneurship and Innovation to explore commercialization options for the algorithm.

The computational research also benefits the laboratory's original focus on better understanding the primate brain by recording its activity as animals learn and behave. Modeling and testing strategies that enable learning, attention, sensory processing and other functions in a computer can motivate and suggest new biological experiments that probe the mechanisms of intelligence both natural and artificial, the researchers said.

"Adding in this component of research to the lab has really opened a lot of doors in terms of allowing us to think about new kinds of problems, new kinds of neuroscience topics and problems that we normally can't really address using the experimental techniques currently available to us in the lab," Freedman said. "We hope this is the starting point for more work in the lab to both identify those principles and to help create artificial networks that continue learning and building on prior knowledge."

More information: Nicolas Y. Masse et al. Alleviating catastrophic forgetting using context-dependent gating and synaptic stabilization, *Proceedings of the National Academy of Sciences* (2018). DOI: 10.1073/pnas.1803839115

Provided by University of Chicago