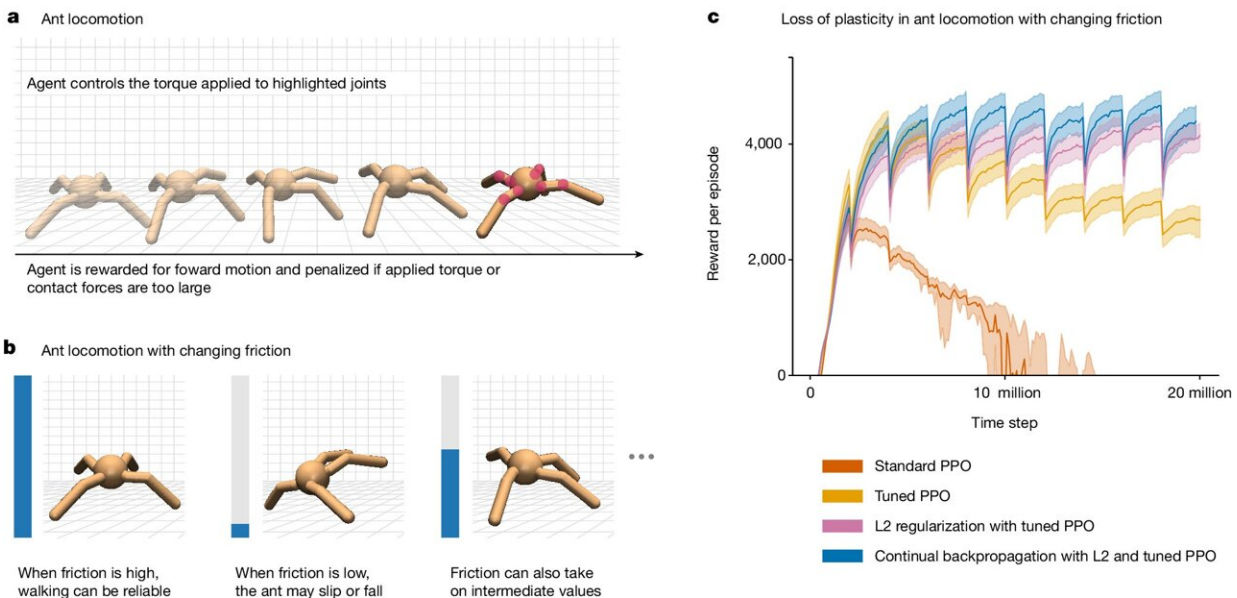


Research team proposes solution to AI's continual learning problem

August 31 2024, by Scott Lilwall



Maintaining plasticity in a non-stationary reinforcement-learning problem. **a**, The reinforcement-learning agent controls torques at the eight joints of the simulated ant (red circles) to maximize forward motion and minimize penalties. **b**, Here we use a version of the ant problem in which the friction on contact with the ground is abruptly changed every 2 million time steps. **c**, The standard PPO learning algorithm fails catastrophically on the non-stationary ant problem. If the optimizer of PPO (Adam) is tuned in a custom way, then the failure is less severe, but adding continual backpropagation or L2 regularization is necessary to perform well indefinitely. These results are averaged over 100 runs; the solid lines represent the mean and the shaded regions represent the 95% bootstrapped confidence interval. Credit: *Nature* (2024). DOI: 10.1038/s41586-024-07711-7

A team of Alberta Machine Intelligence Institute (Amii) researchers has revealed more about a mysterious problem in machine learning—a discovery that might be a major step towards building advanced AI that can function effectively in the real world.

The paper, titled "Loss of Plasticity in Deep Continual Learning," is [published](#) in *Nature*. It was authored by Shibhansh Dohare, J. Fernando Hernandez-Garcia, Qingfeng Lan, Parash Rahman, as well as Amii Fellows & Canada CIFAR AI Chairs A. Rupam Mahmood and Richard S. Sutton.

In their paper, the team explores a vexing problem that has long been suspected in [deep learning models](#) but has not received much attention: for some reason, many deep learning agents engaged in continual learning lose the ability to learn and have their performance degrade drastically.

"We have established that there is definitely a problem with current deep learning," said Mahmood. "When you need to adapt continually, we have shown that deep learning eventually just stops working. So effectively you can't keep learning."

He points out that not only does the AI agent lose the ability to learn new things, but it also fails to relearn what it learned in the past after it is forgotten. The researchers dubbed this phenomenon "loss of plasticity," borrowing a term from neuroscience where plasticity refers to the brain's ability to adapt its structure and form new neural connections.

The state of current deep learning

The researchers say that loss of plasticity is a major challenge to developing AI that can effectively handle the complexity of the world and would need to be solved to develop human-level artificial

intelligence.

Many existing models aren't designed for continual learning. Sutton references ChatGPT as an example; it doesn't learn continually. Instead, its creators train the model for a certain amount of time. When training is over, the model is then deployed without further learning.

Even with this approach, merging new and old data into a model's memory can be difficult. Most of the time, it is more effective to just start from scratch, erasing the memory, and training the model on everything again. For large models like ChatGPT, that process can take a lot of time and cost millions of dollars each time.

It also limits the kind of things a model can do. For fast-moving environments that are constantly changing, like [financial markets](#) for instance, Sutton says continual learning is a necessity.

Hidden in plain sight

The first step to addressing loss of plasticity, according to the team, was to show that it happens and it matters. The problem is one that was "hiding in plain sight"—there were hints suggesting that loss of plasticity could be a widespread problem in deep learning, but very little research had been done to actually investigate it.

Rahman says he first became interested in exploring the problem because he kept seeing hints of the issue—and that intrigued him.

"I'd be reading through a paper, and you'd see something in the appendices about how performance dropped off. And then you'd see it in another paper a while later," he said.

The research team designed several experiments to search for loss of

plasticity in deep learning systems. In supervised learning, they trained networks in sequences of classification tasks. For example, a [network](#) would learn to differentiate between cats and dogs in the first task, then between beavers and geese on the second task, and so on for many tasks. They hypothesized that as the networks lost their ability to learn, their ability to differentiate would decrease in each subsequent task.

And that's exactly what happened.

"We used several different data sets to test, to show that it could be widespread. It really shows that it isn't happening in a little corner of deep learning," Sutton said.

Dealing with the dead

With the problem established, the researchers then had to ask: could it be solved? Was loss of plasticity an inherent issue for continual deep-learning networks, or was there a way to allow them to keep learning?

They found some hope in a method based on modifying one of the fundamental algorithms that make neural networks work: backpropagation.

Neural networks are built to echo the structure of the human brain: They contain units that can pass information and make connections with other units, just like neurons. Individual units can pass information along to other layers of units, which do the same. All of this contributes to the network's overall output.

However, when adapting the connection strength or "weights" of the network with backpropagation, a lot of the time these units will calculate outputs that don't actually contribute to learning. They also won't learn new outputs, so they will become dead weight to the network and stop

contributing to the learning process.

Over long-term continual learning, as many as 90% of a network's units might become dead, Mahmood notes. And when enough stops contributing, the model loses plasticity.

So, the team came up with a modified method that they call "continual backpropagation."

Dohare says that it differs from backpropagation in a key way: While backpropagation randomly initializes the units only at the very beginning, continual backpropagation does so continually. Once in a while, during learning, it selects some of the useless units, like the dead ones, and reinitializes them with random weights. By using continual backpropagation, they find that models can continually learn much longer, sometimes seemingly indefinitely.

Sutton says that other researchers might come up with better solutions to tackle loss of plasticity, but their continual backprop approach at least shows the problem can be solved, and this tricky problem isn't inherent to deep networks.

He's hopeful that the team's work will bring more attention to loss of plasticity and encourage other researchers to examine the issue.

"We established this problem in a way that people sort of have to acknowledge it. The field is gradually getting more willing to acknowledge that deep learning, despite its successes, has fundamental issues that need addressing," he said. "So, we're hoping this will open up this question a little bit."

More information: Shibhansh Dohare, Loss of plasticity in deep continual learning, *Nature* (2024). [DOI: 10.1038/s41586-024-07711-7](https://doi.org/10.1038/s41586-024-07711-7).

www.nature.com/articles/s41586-024-07711-7

Provided by Alberta Machine Intelligence Institute

Citation: Research team proposes solution to AI's continual learning problem (2024, August 31) retrieved 2 September 2024 from <https://techxplore.com/news/2024-08-team-solution-ai-problem.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.