


Scientists make a maze-running artificial intelligence program that learns to take shortcuts

May 11 2018, by Amina Khan, Los Angeles Times



An AI system learns to take shortcuts. Credit: *Nature* (2018). DOI: 10.1038/s41586-018-0102-6

Call it an a-MAZE-ing development: A U.K.-based team of researchers

has developed an artificial intelligence program that can learn to take shortcuts through a labyrinth to reach its goal. In the process, the program developed structures akin to those in the human brain.

The emergence of these computational "[grid cells](#)," described in the journal *Nature*, could help scientists design better navigational software for future robots and even offer a new window through which to probe the mysteries of the mammalian brain.

In recent years, AI researchers have developed and fine-tuned deep-learning networks—layered programs that can come up with novel solutions to achieve their assigned goal. For example, a deep-learning network can be told which face to identify in a series of different photos, and through several rounds of training, can tune its algorithms until it spots the right face virtually every time.

These networks are inspired by the brain, but they don't work quite like them, said Francesco Savelli, a neuroscientist at Johns Hopkins University who was not involved in the paper. So far, AI systems don't come close to emulating the brain's architecture, the diversity of real neurons, the complexity of individual neurons or even the rules by which they learn.

"Most of the learning is thought to occur with the strengthening and weakening of these synapses," Savelli said in an interview, referring to the connections between neurons. "And that's true of these AI systems too—but exactly how you do it, and the rules that govern that kind of learning, might be very different in the brain and in these systems."

Regardless, AI has been really useful for a number of functions, from facial recognition to deciphering handwriting and translating languages, Savelli said. But higher-level activities—such as navigating a complex environment—have proved far more challenging.

One aspect of navigation that our brains seem to perform without conscious effort is path integration. Mammals use this process to recalculate their position after every step they take by accounting for the distance they've traveled and the direction they're facing. It's thought to be key to the brain's ability to produce a map of its surroundings.

Among the neurons associated with these "cognitive maps": [place cells](#), which light up when their owner is in some particular spot in the environment; head-direction [cells](#), which tell their owner what direction they're facing; and grid cells, which appear to respond to an imaginary hexagonal grid mapped over the surrounding terrain. Every time a person steps on a "node" in this grid, the neuron fires.

"Grid cells are thought to endow the cognitive map with geometric properties that help in planning and following trajectories," Savelli and fellow Johns Hopkins neuroscientist James Knierim wrote in a commentary on the paper. The discovery of grid cells earned three scientists the 2014 Nobel Prize in physiology or medicine.

Humans and other animals seem to have very little trouble moving through space because all of these highly specialized neurons work together to tell us where we are and where we're going.

Scientists at DeepMind, which is owned by Google and University College London, wondered whether they could develop a program that could also perform path integration. So they trained the network with simulations of paths used by rodents looking for food. They also gave it data for a rodent's movement and speed as well as feedback from simulated place cells and head-direction cells.

During this training, the researchers noticed something strange: The simulated rodent appeared to develop patterns of activity that looked remarkably like grid cells—even though grid cells had not been part of

their training system.

"The emergence of grid-like units is an impressive example of deep learning doing what it does best: inventing an original, often unpredicted internal representation to help solve a task," Savelli and Knierim wrote.

Grid cells appear to be so useful for path integration that this faux-rodent came up with a solution eerily similar to a real rodent brain. The researchers then wondered: Could grid cells also be useful in another crucial aspect of mammal navigation?

That aspect, called vector-based navigation, is basically the ability to calculate the straight-shot, "as the crow flies" distance to a goal even if you originally took a longer, less-direct route. That's a useful skill for finding shortcuts to your destination, Savelli pointed out.

To test this, researchers challenged the grid-cell-enabled faux-rodent to solve a maze, but blocked off most of the doorways so the program would have to take the long route to its goal. They also modified the program so it was rewarded for actions that brought it closer to the goal. They trained the network on a given maze and then opened shortcuts to see what happened.

Sure enough, the simulated rodent with grid cells quickly found and used the shortcuts, even though those pathways were new and unknown. And it performed far better than a faux-rodent whose start point and goal point were tracked only by place cells and head-direction cells. It even beat out a "human expert," the study authors said.

The findings eventually could prove useful for robots making their way through unknown territory, Savelli said. And from a neuroscientific perspective, they could help researchers better understand how these neurons do their job in the mammalian brain.

Of course, this program was highly simplified compared to its biological counterpart, Savelli pointed out. In the simulated rodent, the "place cells" didn't change—even though place cells and grid cells influence each other in complex ways in real brains.

"By developing the network such that the place-cell layer can be modulated by [grid](#)-like inputs, we could begin to unpack this relationship," Savelli and Knierim wrote.

Developing this AI program further could help scientists start to understand all the complex relationships that come into play in living neural systems, they added.

But whether they want to hone the technology or use it to understand biology, scientists will have to get a better handle on their own deep-learning programs, whose solutions to problems are often hard to decipher even if they consistently get results, scientists said.

"Making deep-learning systems more intelligible to human reasoning is an exciting challenge for the future," Savelli and Knierim wrote.

More information: Andrea Banino et al. Vector-based navigation using grid-like representations in artificial agents, *Nature* (2018). [DOI: 10.1038/s41586-018-0102-6](https://doi.org/10.1038/s41586-018-0102-6)

©2018 Los Angeles Times

Distributed by Tribune Content Agency, LLC.

Citation: Scientists make a maze-running artificial intelligence program that learns to take shortcuts (2018, May 11) retrieved 24 April 2024 from <https://phys.org/news/2018-05-scientists-maze-running-artificial-intelligence-shortcuts.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.