

## Game-playing software holds lessons for neuroscience

26 February 2015 by Elizabeth Gibney - Nature

### ***DeepMind computer provides new way to investigate how the brain works.***

DeepMind, the Google-owned artificial-intelligence company, has revealed how it created a single computer algorithm that can learn how to play 49 different arcade games, including the 1970s classics *Pong* and *Space Invaders*. In more than half of those games, the computer became skilled enough to beat a professional human player.

The algorithm is the first artificial-intelligence (AI) system that can learn a variety of tasks from scratch given only the same, minimal starting information. “The fact that you have one system that can learn several games, without any tweaking from game to game, is surprising and pretty impressive,” says Nathan Sprague, a machine-learning scientist at James Madison University in Harrisonburg, Virginia.

DeepMind, which is based in London, says that the brain-inspired system could also provide insights into human intelligence. “Neuroscientists are studying intelligence and decision-making, and here’s a very clean test bed for those ideas,” says Demis Hassabis, co-founder of DeepMind.

Software that teaches itself to play *Space Invaders* could help scientists study the mind.

Games are to AI researchers what fruit flies are to biology — a stripped-back system in which to test theories, says Richard Sutton, a computer scientist who studies reinforcement learning at the University of Alberta in Edmonton, Canada. “Understanding the mind is an incredibly difficult problem, but games allow you to break it down into parts that you can study,” he says. But so far, most human-beating computers — such as IBM’s Deep Blue, which beat chess world champion Garry Kasparov in 1997— excel at only one game.

DeepMind’s versatility comes from joining two types of machine learning— an achievement that Sutton calls “a big deal”. The first, called deep learning, uses a brain-inspired architecture in which connections between layers of simulated neurons are strengthened on the basis of experience. Deep-learning systems can then draw complex information from reams of unstructured data. Google, of Mountain View, California, uses such algorithms to automatically classify

photographs and aims to use them for machine translation.

The second is reinforcement learning, a decision-making system inspired by the neurotransmitter dopamine reward system in the animal brain. Using only the screen’s pixels and game score as input, the algorithm learned by trial and error which actions — such as go left, go right or fire — to take at any given time to bring the greatest rewards. After spending several hours on each game, it mastered a range of arcade classics, including car racing, boxing and *Space Invaders*.

Sprague suggests that the technique could enable robots to solve problems by interacting with their environments.

But a major driver is science itself, says Hassabis, because building smarter systems means gaining a greater understanding of intelligence. Many in computational neuroscience agree. Sprague, who has created his own version of DeepMind’s algorithm, explains that whereas AI is largely irrelevant to neuroscience at the level of anatomical connections among neurons, it can bring insight at the higher level of computational principles.

“The tricks we use for training a system might lead to new ideas about the brain.”

DeepMind was bought by Google in 2014 for a reported £400 million (US\$617 million), and has been poaching leading computer scientists and neuroscientists from academia, growing from 80 to 140 researchers so far.

Its next steps are again likely to be influenced by neuroscience. One project could be building a memory into its algorithm, allowing the system to transfer its learning to new tasks. Unlike humans, when the current system masters one game, it is no better at tackling the next. Another challenge is to mimic the brain’s way of breaking problems down into smaller tasks. Currently, DeepMind’s system struggles to link actions with distant consequences — a limitation that, for example, prevented it from mastering maze games such as *Ms. Pac-Man*.