

# DeepMind's AI team explores navigation powers with 3-D maze

February 7 2016, by Nancy Owano

---



Two players, an empty board, white and black circles that look like stones as playing pieces, and [AI news around the world](#) is made.

Last month there were reports of a Google Go triumph as a milestone for [artificial intelligence](#) (AI) research.

Peter Cowling and Sam Devlin, University of York, wrote an article in which they reported that the DeepMind team developed a computer able to beat a human at the [game](#) Go.

Why should Google's DeepMind team bother to the extent that they have in their AI research activities to boardgame computing? Many researchers know the answer: some games like chess play the roles of lab rats in researchers seeing how far they can take artificial intelligence to come up with moves and decisions which we normally assume can be achieved only with the human brain.

Go is a 2,000-year-old [game](#) played by over 60 million people on Planet Earth, said the two, and another reason to turn to Go is the satisfaction in having met a very tough challenge.

"Creating a superhuman computer Go player able to beat these top pros has been one of the most challenging targets of AI research for decades."

What's the big mental challenge? Can it be as rigorous as chess? Cowling and Devlin had an answer for that. "Go has many more possible positions than even chess – in fact, there are more possibilities in a game of Go than we would get by considering a separate chess game played on every atom in the universe."

How DeepMind researchers pulled this off involved "analyzing millions of past games by professional human players and simulating thousands of possible future game states per second," said Cowling and Devlin. The team trained "convolutional neural networks," algorithms that mimic the brain structure and visual system. Cowling and Devlin said Monte Carlo tree search approaches were also part of the picture.

Cowling and Devlin offered an interesting observation, too, in their January article. They said, "Now that Go has seemingly been cracked, AI needs a new grand challenge – a new 'lab rat' – and it seems likely that many of these challenges will come from the \$100 billion digital games industry."

Well, on to February: This month, the story is more about Google DeepMind and their explorations than about their win in the game Go. *New Scientist* reported on Friday that "DeepMind's latest artificial intelligence can navigate a 3D maze reminiscent of the 1993 shooter game *Doom*." What is that game like? *New Scientist* reported this as a 3D maze game, Labyrinth, "a test bed for DeepMind's tech that resembles *Doom* without the shooting." The system is [rewarded](#) for finding apples and portals, the latter of which teleport it elsewhere in the maze, he said, and must score as high as possible in 60 seconds.

*New Scientist* said the team's system plays just as a human would, by looking at the screen and deciding how to proceed. "This ability to navigate a 3D space by 'sight,'" said Jacob Aron, "could be useful for AIs operating in the real world." *Engadget* also recognized that Google DeepMind AI can make its way through a 3D maze by 'sight'. "We're one step closer to AI machines that can [navigate](#) the real world as humans do," said *Engadget*'s Jessica Conditt.

The DeepMind team is working with a technique called asynchronous [reinforcement learning](#), which sees multiple versions of an AI tackling a

problem in parallel and comparing their experiences.

Their paper on this topic is on arXiv, titled "Asynchronous Methods for Deep Reinforcement Learning." The eight DeepMind researchers set out to confront limitations which they saw in Deep RL algorithms based on experience replay even though they had achieved unprecedented success in challenging domains such as Atari 2600.

The drawbacks they saw in experience replay: it uses more memory and more computation per real interaction; and it needs off-policy learning algorithms that can update from data generated by an older policy.

Instead of experience replay, said the authors, "we asynchronously execute multiple agents in parallel, multiple instances of the environment."

They presented asynchronous versions of four standard reinforcement learning algorithms. They said in the paper that they showed they can train neural network controllers on a variety of domains in a stable way.

**More information:** Asynchronous Methods for Deep Reinforcement Learning, arXiv:1602.01783 [cs.LG] [arxiv.org/abs/1602.01783](https://arxiv.org/abs/1602.01783)

## Abstract

We propose a conceptually simple and lightweight framework for deep reinforcement learning that uses asynchronous gradient descent for optimization of deep neural network controllers. We present asynchronous variants of four standard reinforcement learning algorithms and show that parallel actor-learners have a stabilizing effect on training allowing all four methods to successfully train neural network controllers. The best performing method, an asynchronous variant of actor-critic, surpasses the current state-of-the-art on the Atari domain while training for half the time on a single multi-core CPU instead of a

GPU. Furthermore, we show that asynchronous actor-critic succeeds on a wide variety of continuous motor control problems as well as on a new task involving finding rewards in random 3D mazes using a visual input.

© 2016 Tech Xplore

Citation: DeepMind's AI team explores navigation powers with 3-D maze (2016, February 7)  
retrieved 28 April 2024 from  
<https://techxplore.com/news/2016-02-deepmind-ai-team-explores-powers.html>

<p>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.</p>
--