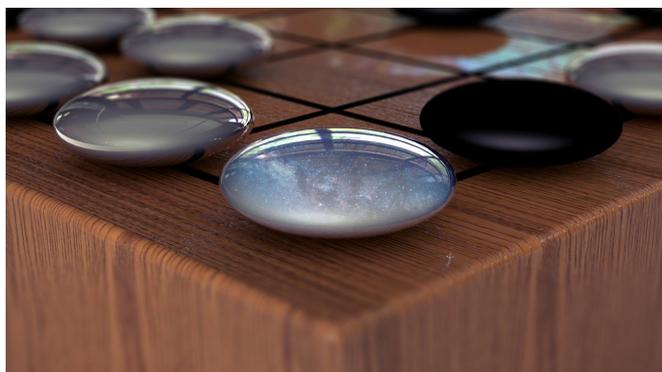


Self-taught, 'superhuman' AI now even smarter: makers

18 October 2017, by Mariëtte Le Roux



There are an astonishing 10 to the power of 170 possible board configurations in Go - more than the number of atoms in the known universe. Credit: DeepMind

The computer that stunned humanity by beating the best mortal players at a strategy board game requiring "intuition" has become even smarter, its makers said Wednesday.

Even more startling, the updated version of AlphaGo is entirely self-taught—a major step towards the rise of machines that achieve superhuman abilities "with no human input", they reported in the science journal *Nature*.

Dubbed AlphaGo Zero, the Artificial Intelligence (AI) system learnt by itself, within days, to master the ancient Chinese board [game](#) known as "Go"—said to be the most complex two-person challenge ever invented.

It came up with its own, novel moves to eclipse all the Go acumen humans have acquired over thousands of years.

After just three days of self-training it was put to the ultimate test against AlphaGo, its forerunner which previously dethroned the top human

champs.

AlphaGo Zero won by 100 games to zero.

"AlphaGo Zero not only rediscovered the common patterns and openings that humans tend to play... it ultimately discarded them in preference for its own variants which humans don't even know about or play at the moment," said AlphaGo lead researcher David Silver.

The 3,000-year-old Chinese game played with black and white stones on a board has more move configurations possible than there are atoms in the Universe.

AlphaGo made world headlines with its shock 4-1 victory in March 2016 over 18-time Go champion Lee Se-Dol, one of the game's all-time masters.

Lee's defeat showed that AI was progressing faster than widely thought, said experts at the time who called for rules to make sure powerful AI always remains completely under human control.

In May this year, an updated AlphaGo Master programme beat world Number One Ke Jie in three matches out of three.

Not constrained by humans

Unlike its predecessors which trained on data from thousands of human games before practising by playing against itself, AlphaGo Zero did not learn from humans, or by playing against them, according to researchers at DeepMind, the British [artificial intelligence](#) (AI) company developing the system.

"All previous versions of AlphaGo... were told: 'Well, in this position the human expert played this particular move, and in this other position the human expert played here'," Silver said in a video explaining the advance.

AlphaGo Zero skipped this step.

Instead, it was programmed to respond to reward—a positive point for a win versus a negative point for a loss.

Starting with just the rules of Go and no instructions, the system learnt the game, devised strategy and improved as it competed against itself—starting with "completely random play" to figure out how the reward is earned.

This is a trial-and-error process known as "reinforcement learning".

Unlike its predecessors, AlphaGo Zero "is no longer constrained by the limits of human knowledge," Silver and DeepMind CEO Demis Hassabis wrote in a blog.

Amazingly, AlphaGo Zero used a single machine—a human brain-mimicking "neural network"—compared to the multiple-machine "brain" that beat Lee.

It had four data processing units compared to AlphaGo's 48, and played 4.9 million training games over three days compared to 30 million over several months. © 2017 AFP

Beginning of the end?

"People tend to assume that machine learning is all about big data and massive amounts of computation but actually what we saw with AlphaGo Zero is that algorithms matter much more," said Silver.

The findings suggested that AI based on reinforcement learning performed better than those that rely on human expertise, Satinder Singh of the University of Michigan wrote in a commentary also carried by Nature.

"However, this is not the beginning of any end because AlphaGo Zero, like all other successful AI so far, is extremely limited in what it knows and in what it can do compared with humans and even other animals," he said.

AlphaGo Zero's ability to learn on its own "might

appear creepily autonomous", added Anders Sandberg of the Future of Humanity Institute at Oxford University.

But there was an important difference, he told AFP, "between the general-purpose smarts humans have and the specialised smarts" of computer software.

"What DeepMind has demonstrated over the past years is that one can make software that can be turned into experts in different domains... but it does not become generally intelligent."

It was also worth noting that AlphaGo was not programming itself, said Sandberg.

"The clever insights making Zero better was due to humans, not any piece of software suggesting that this approach would be good. I would start to get worried when that happens."

More information: Mastering the game of Go without human knowledge, *Nature* (2017). [nature.com/articles/doi:10.1038/nature24270](https://www.nature.com/articles/doi:10.1038/nature24270)

APA citation: Self-taught, 'superhuman' AI now even smarter: makers (2017, October 18) retrieved 30 November 2020 from <https://phys.org/news/2017-10-self-taught-superhuman-ai-smarter-makers.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.