

Google's Go triumph is a milestone for artificial intelligence research

January 28 2016, by Peter Cowling And Sam Devlin, University Of York



Credit: AI-generated image ([disclaimer](#))

Researchers from Google DeepMind have developed the first computer able to defeat a human champion at the board game Go. But why has the online giant invested millions of dollars and some of the finest minds in Artificial Intelligence (AI) research to create a computer board game player?

Go is not just any [board game](#). It's more than 2,000 years old and is played by more than [60m people](#) across the world – including a thousand professionals. 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.

The rules are deceptively simple: two players take turns to place white and black "stones" on an empty 19x19 board, each aiming to encircle the most territory. Yet these basics yield a game of extraordinary beauty and complexity, full of patterns and flow. 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.

AI researchers have therefore long regarded Go as a "grand challenge". Whereas even the best human chess players had fallen to computers by the 1990s, Go remained unbeaten. This is a truly historic breakthrough.

Games are the 'lab rats' of AI research

Since the term "[artificial intelligence](#)" or "AI" was first coined in the 1950s, the range of problems which it can solve has been increasing at an accelerating rate. We take it for granted that Amazon has a pretty good idea of what we might want to buy, for instance, or that Google can complete our partially typed search term, though these are both due to [recent advances in AI](#).

Computer games have been a crucial test bed for developing and testing new AI techniques – the "lab rat" of our research. This has led to superhuman players in [checkers](#), [chess](#), [Scrabble](#), [backgammon](#) and more recently, simple forms of [poker](#).

Games provide a fascinating source of tough problems – they have well-

defined rules and a clear target: to win. To beat these games the AIs were programmed to search forward into possible futures and choose the move which leads to the best outcome – which is similar to how good human players make decisions.

Yet Go proved hardest to beat because of its enormous search space and the difficulty of working out who is winning from an unfinished game position. Back in 2001, Jonathan Schaeffer, a brilliant researcher who created a perfect AI checkers player, [said it would](#) "take many decades of research and development before world-championship-caliber Go programs exist". Until now, even with recent advances, it still seemed at least ten years out of reach.

The breakthrough

Google's announcement, in the journal [Nature](#), details [how its machine "learned" to play Go](#) by analysing millions of past games by professional human players and simulating thousands of possible future game states per second. Specifically, the researchers at DeepMind trained "convolutional neural networks", algorithms that mimic the high-level structure of the brain and visual system and which have recently seen [an explosion in their effectiveness](#), to predict expert moves.

This learning was combined with [Monte Carlo tree search](#) approaches which use randomness and machine learning to intelligently search the "tree" of possible future board states. These searches have massively increased the strength of computer Go players since their invention less than ten years ago, as well as finding applications in many other domains.

The resulting "player" significantly outperformed all existing state-of-the-art AI players and went on to beat the current European champion, Fan Hui, 5-0 under tournament conditions.

AI passes 'Go'

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. The ability to play alongside or against millions of engaged human players provides unique opportunities for AI research. At York's centre for [Intelligent Games and Game Intelligence](#), we're working on projects such as building an AI aimed at player fun (rather than playing strength), for instance, or using games to improve well-being of people with Alzheimer's. Collaborations between multidisciplinary labs like ours, the games industry and big business are likely to yield the next big AI breakthroughs.

However the real world is a step up, full of ill-defined questions that are far more complex than even the trickiest of board games. The techniques which conquered Go can certainly be applied in [medicine](#), [education](#), science or any other domain where data is available and outcomes can be evaluated and understood.

The big question is whether Google just helped us towards the next generation of [Artificial General Intelligence](#), where machines learn to truly think like – and beyond – humans. Whether we'll see AlphaGo as a step towards Hollywood's dreams (and nightmares) of AI agents with self-awareness, emotion and motivation remains to be seen. However the latest breakthrough points to a brave new future where AI will continue to improve our lives by helping us to make better-informed decisions in a world of ever-increasing complexity.

This article was originally published on [The Conversation](#). Read the [original article](#).

Source: The Conversation

Citation: Google's Go triumph is a milestone for artificial intelligence research (2016, January 28) retrieved 19 April 2024 from <https://phys.org/news/2016-01-google-triumph-milestone-artificial-intelligence.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.