

AlphaZero just wants to play

February 23 2018, by Prof. Roger Wattenhofer



Credit: CC0 Public Domain

Artificial intelligence is continually hyped up, but disappears from view again just as quickly. Roger Wattenhofer explains why that might soon change.

Over the last 60 years, artificial intelligence researchers have continually

made futuristic predictions. For example, they announced that a [computer](#) would become a chess grandmaster. It was expected by 1968; it happened around 30 years later. Even more optimistic were the predictions that computers would be able to perform any and every human task by 1985 – but even today, we're still a long way from that being the case.

A history of ups and downs

Disillusionment often followed, large-scale research programmes were scrapped, and a series of AI winters set in. Eventually, researchers didn't even want to talk about "artificial intelligence", preferring instead to study the more modest field of "machine learning".

At the turn of the millennium came a few successes: IBM's Deep Blue defeated Garry Kasparov at chess. Critics complained that chess was not a real test of intelligence, and repeated John Searle's [Chinese room](#) argument: Deep Blue didn't really understand the [game](#), it simply calculated very quickly. A few years later, a car from Stanford succeeded in navigating autonomously, but even that was only only in the desert. Finally, IBM's Watson won the quiz show Jeopardy! – but was it truly down to intelligence, or was it just a case of advanced googling?

Wolf cried too often?

I think that AI research and its advocates have cried wolf too often – meaning that now, people don't tend to see the value of real successes and genuine breakthroughs. But the wolf is really here, and it has a name: [AlphaZero](#).

AlphaZero, a computer programme from DeepMind, can teach itself skills such as chess. Since Deep Blue, chess programmes have been

continuously refined. The best programmes, such as Stockfish, reached such a high level years ago that humans stand no chance against them.

Stockfish knows a lot of chess theory, and is familiar with every game ever played. AlphaZero only knew the rules of chess; it learned the game by playing against itself for a few hours. It initially made a lot of beginner's mistakes, but after four hours of self-training, it defeated Stockfish. (There are a few [valid objections](#) that discuss the use of hardware and unfair competition rules.)

The professional game will change

I recommend that anyone who understands anything about chess watch the [tenth game against Stockfish](#). AlphaZero plays very unusually; not like a human, but also not like a typical computer. Instead, it plays with "real artificial" [intelligence](#). As AlphaZero began by knowing no chess theory, it had to develop its own theories. Grandmaster Daniel King thinks that AlphaZero will change chess theory and the professional game as soon as it is publicly accessible.

In a sense, AlphaZero is more human than a normal computer. The deep network structure is reminiscent of a human brain, and it calculates more slowly: per game, it assesses thousands of times fewer moves than Stockfish. It would be interesting to see how a game of human vs. AlphaZero would play out, if AlphaZero were slowed down to human speed. And AlphaZero can do more than just play [chess](#). It also knows other games that it has taught itself.

I, for one, am impressed. AlphaZero's performance shows that we have to take [artificial intelligence](#) seriously, and proves that we are witnessing an enormous change. I am convinced that AlphaZero & Co. will have a lasting effect on society.

Provided by ETH Zurich

Citation: AlphaZero just wants to play (2018, February 23) retrieved 1 May 2024 from <https://phys.org/news/2018-02-alphazero.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.