



Google Deep Mind's 'alien' chess computer reveals game's deeper truths

The self-taught AlphaZero's performance was the biggest paradigm shift in chess computing since Deep Blue defeated Garry Kasparov in 1997.



[Home](#) » [Innovation](#) » Google Deep Mind's 'alien' chess computer reveals game's deeper truths



March 20, 2018 by **David W. Smith**

The world's strongest human chess grandmasters have become resigned to being defeated by chess computers, but they have taken comfort in the knowledge that artificial intelligence (AI) plays what world champion Magnus Carlsen describes as "mechanical, dry and bland" moves. Chess computers, he says, have "no style" and when you play them, not only will you inevitably lose, you will also be bored.

Carlsen may have to revise his opinion, however, after the unprecedented performance of Google Deep Mind's AlphaZero chess-playing computer. **AlphaZero taught itself to play chess in four hours by playing millions of games**, then overwhelmed one of the world's strongest chess

computers, Stockfish 8, in a 100-game match in December. AlphaZero won 28 games and lost none against a chess engine that routinely dismantles human players.

But what was even more striking than the result was that AlphaZero's playing style was a world away from the boring, mechanical computer chess Carlsen abhors. Having learned the game without human input – apart from the rules – AlphaZero produced strategic masterpieces that stunned both the chess and AI worlds. "It surprised the hell out of me," says Professor **Jonathan Schaeffer**, an AI researcher at the University of Alberta, who built one of the world's strongest chess programmes in the 1990s and designed the Shinook programme that defeated the world checkers champion in 1994. "The games were beautiful and creative. AlphaZero made apparently crazy sacrifices that humans would not even consider in order to get more freedom of movement. But it also played differently to all other chess programmes which rely on human input."

Chess grandmasters were equally impressed. Russian champion **Peter Svidler** said that AlphaZero's play was "absolutely fantastic, phenomenal" and he felt in "awe" of its play. Carlsen has not reacted, but his coach Peter Nielsen said "the aliens came and showed us how to play chess", implying that AlphaZero represents a watershed in chess computing and, perhaps, in AI in general.

By 2015, the chess engine Komodo was offering grandmasters material advantages, such as an extra pawn, or a rook for a knight, and still beating them.



It is arguable that AlphaZero is the second paradigm shift in chess computing. The first came in 1997, when Garry Kasparov lost to IBM's Deep Blue. A bruised Kasparov, accustomed to demolishing humans, accused IBM's developers of cheating. He claimed he had seen signs of "human creativity" in the machine's moves.

The reality was more mundane. The 1997 version of Deep Blue derived its strength from brute force, calculating 200 million positions per second. Kasparov, who lost narrowly, may have underperformed and still competed with computers for years. In 2003, he took on Deep Junior for a US\$1 million prize billed as the man v machine World Championship, and the match ended in a draw.

China adopts competitive model

But post-Deep Blue it was clear that humans could not keep up for long with advances in computing power. By 2015, the chess engine Komodo was offering grandmasters material advantages, such as an extra pawn, or a rook for a knight, and still beating them. Today, grandmasters rely on computers to analyse positions, but don't dream of defeating them.

Despite increases in strength, methods of chess computing, from Deep Blue onwards, changed little until the advent of AlphaZero. Traditional engines evaluate positions using features handcrafted by grandmasters. The humans teach the machines to weigh up the relative importance of material, space and time in a position. They supply “opening books”, which suggest strong moves at the beginning of a game. Meanwhile, AI developers provide the engine with a vast tree search, allowing it to choose between many possible moves. For Stockfish, this was known as “alpha beta” search.

“Since Deep Blue there have been two sources of improvement in chess engines,” says **Kanwal Bhatia, an AI researcher into Machine Learning at Visulytix**, in London, as well as a strong chess player. “The first source is the increase in power. A standard laptop now computes more powerfully than a desktop did two decades ago. The second source is human updates to algorithms that allow computers to evaluate positions more accurately.”

The reliance on raw calculating power has limitations, however. Chess engines lack the intuitive understanding of certain positions that comes naturally to human grandmasters, she says. Garry Kasparov, and other strong players, have claimed that a human grandmaster-chess engine team could defeat an unassisted chess engine. “The human-computer team would have the benefit of the phenomenal, unemotional calculating powers of the machine, but would not make ‘computer mistakes,’” says Bhatia.

Traditional chess computers, she says, have a weakness in “closed positions”, when the pawns are locked together and it takes a lot of subtle manoeuvring to make something happen. “Computers are good at calculating eight or nine moves ahead, but what becomes important in a closed position is an intuitive understanding of long-term goals, which favours human players. The ex-world champion Vladimir Kramnik has had some success against computers exploiting this issue, which is known as ‘the horizon problem,’” she says. AlphaZero, however, doesn’t play chess like other chess computers and any weaknesses in such positions are not yet apparent. Learning through self-play means AlphaZero is free of preconceptions about how to win.

AlphaZero’s algorithm was a more generic version of the AlphaGo Zero algorithm, which last year defeated the leading human Go champions after playing itself 4.9 million times in a few days. For the chess-playing version, Deep Mind replaced the handcrafted knowledge of traditional engines with “**deep neural networks**”, a way of processing inspired by the human brain. The neural networks are layers of connected nodes that change as the system learns through self-play to play stronger moves. AlphaZero used a reinforcement learning algorithm and a general-purpose Monte-Carlo tree search, widely used in the development of games-playing AI. Before taking on Stockfish, AlphaZero played itself four million times in four hours on 5,000 of Google’s TPUS (tensor processing units), powerful chips that are not commercially available.

“AlphaZero learned like I did as a kid. I knew the basic rules and lost my first 20 or 30 games against simple strategies. After a while, I learned how not to fall for Scholar’s mate. For years, I continued to get better mainly through playing,” explains Bhatia. “AlphaZero uses reinforcement learning to retain moves and good positions that lead to wins. Moves that lead to losses are weighted down immediately. What makes such an approach viable is that it can play millions of games in a short space of time.”

Deep Mind CEO Demis Hassabis says AlphaZero compensates for the lower number of evaluations by focusing more selectively on promising options.



Meanwhile, how Bhatia develops her game as an experienced player resembles how Stockfish gets better. “To improve, Stockfish needs the superior understanding of positions which comes from external human input. Similarly, I improve now by studying books about the openings, or other aspects of chess, so I rely on human chess knowledge that has accumulated over decades. One of the amazing things about AlphaZero was that it worked out how to play the strongest chess openings all by itself.”

The difference in how **Stockfish** and AlphaZero select chess moves is striking. AlphaZero searches just **80,000 positions** per second, compared to 70 million for Stockfish. Deep Mind CEO Demis Hassabis says AlphaZero compensates for the lower number of evaluations by focusing more selectively on promising options. The understanding gained from deep neural networks allows it to take a more “human-like” approach by eliminating weak moves.

Despite the generally enthusiastic reaction of the chess world, there were criticisms of the match conditions. World number seven, the American Hikaru Nakamura, pointed out that AlphaZero had used the Google super-computer, whereas Stockfish had been running on the equivalent of a laptop. He reserved judgment until Stockfish had the same advantages. Other chess experts felt Stockfish had been deprived of its strongest opening book.

Meanwhile, **Stockfish author Tord Romstad** told chess.com that the results were “not particularly meaningful” because his programme had operated with some disadvantages. The games, he argued, were played at a fixed time of one minute per move, but Stockfish had been programmed to take longer to evaluate critical positions. And he added: “The version of Stockfish used is one year old, was playing with far more search threads than has ever received any significant amount of testing and had way too small hash tables for the number of threads.”

Professor Schaeffer, however, described the reservations as “nitpicking”. “There’s no question the result was enormously impressive,” he says. “It was less important to be completely sensitive to match conditions as there was no human involved, unlike with the **AlphaGo Zero match**. The aim was to prove a point, which they did, then retire the machine.”

As for the game of chess, AI advances have transformed our conception of it over the past few decades. Johann von Goethe once described chess “as the touchstone of the intellect” and many of the greatest AI researchers, including Charles Babbage, Alan Turing, Claude Shannon and John von Neumann, devised hardware, algorithms and theory to play chess. It later became the principal

challenge for a generation of AI researchers, leading to the Deep Blue revolution.

But chess was never going to be the ultimate test of AI. Professor Schaeffer says DeepBlue was an “idiot savant”, who played chess brilliantly, but could do nothing else. Stockfish, he argues, has the same limitations. “If you ask it to add up one and one, it wouldn’t have a clue,” he says.

AlphaZero’s intelligence is more generalisable. The same approach produced extraordinary performances in Go, chess, as well as the Japanese game of Shogi. But Professor Schaeffer says the achievement is still only a “small step” towards reaching the broader goal of solving AI, which means creating a general artificial intelligence capable of tackling any problem.

“AlphaZero is already a bit more generalisable because the neural networks allow it to mimic some characteristics of the human brain and learn for itself. But it’s far from being a general problem solver like a human. Humans can tackle any problem, from driving a car, to writing a sonnet, or playing chess,” he says.

“We have one thing that AI doesn’t have, which is human sentience. It means we can interact with the environment and feel pain, love and other emotions. We don’t know how to programme that, but one day we will figure it out and I would be surprised if that day is not in my lifetime.”



RELATED



GDPR is coming... and it's not all bad


True, there are some rather frightening bits. Like the fines for data breaches, which can be up to e20M or 4% of annual global turnover. Levels which one commentator has already called “existentially threatening” to some firms. Then there’s the cost in money and time of becoming compliant, with organisations being forced to streamline their personal data – knowing where every piece is, what it’s being used for and why, and deleting any that is no longer relevant.





Meet the man building the world's first AI recruitment agency

In 2013, tech entrepreneur and TED.com ambassador Jim Stolze got sick of discussing the potential of AI and decided he wanted to get his hands dirty. He founded Aigency, a recruitment company that connects businesses with the algorithms they need to streamline and automate their operations.



Going beyond the Internet of Things sales hype

Every Internet of Things (IoT) salesman is hyping their products. Businesses have to be clear about their goals before they decide what to install and how. But if they get it right, it can be transformative. Supply chain 'Ninja' Patrick Strauss always asks the same question when he comes into a business as an Internet of Things (IoT) supply chain consultant. The simple question is 'why do you need this IoT technology?'

DON'T MISS



ABOUT EUREKA

Eureka means “I found it!” and was the phrase that exclaimed Archimedes after discovering that the volume of water that ascends is equal to the volume of the submerged body. It is about problem solving, learning, and discovery. So that is precisely the purpose of this website: to understand, to learn. A tribute to our ancient history. From Europe to the world.

Powered by



[About Us](#) – [Legal Notices & Terms](#) – [EUREKA Cookie Usage](#) © KYOCERA Document Solutions Europe B.V.

