

BLOG POST
RESEARCH

06 DEC 2018

AlphaZero: Shedding new light on chess, shogi, and Go

In late 2017 we introduced AlphaZero, a single system that taught itself from scratch how to master the games of chess, shogi (Japanese chess), and Go, beating a world-champion program in each case. We were excited by the preliminary results and thrilled to see the response from members of the chess community, who saw in AlphaZero's games a ground-breaking, highly dynamic and "unconventional" style of play that differed from any chess playing engine that came before it.

Today, we are delighted to introduce the full evaluation of AlphaZero, published in the

[journal Science](#) ([Open Access version here](#)), that confirms and updates those preliminary results. It describes how AlphaZero quickly learns each game to become the strongest player in history for each, despite starting its training from random play, with no in-built domain knowledge but the basic rules of the game.

“

I can't disguise my satisfaction that it plays with a very dynamic style, much like my own!"

GARRY KASPAROV
FORMER WORLD CHESS CHAMPION

This ability to learn each game afresh, unconstrained by the norms of human play, results in a distinctive, unorthodox, yet creative and dynamic playing style. Chess Grandmaster Matthew Sadler and Women's International Master Natasha Regan, who have analysed thousands of AlphaZero's chess games for [their forthcoming book *Game Changer*](#) (New in Chess, January 2019), say its style is unlike any traditional chess engine." It's like discovering the secret notebooks of some great player from the past," says Matthew.




AlphaZero: Player and potential

Traditional chess engines – including the world computer chess champion [Stockfish](#) and [IBM's ground-breaking Deep Blue](#) – rely on thousands of rules and heuristics handcrafted by strong human players that try to account for every eventuality in a game. Shogi programs are also game specific, using similar search engines and algorithms to chess programs.

AlphaZero takes a totally different approach, replacing these hand-crafted rules with a deep [neural network](#) and general purpose algorithms that know nothing about the game beyond the basic rules.

Chess



IN CHESS, ALPHAZERO FIRST OUTPERFORMED STOCKFISH AFTER JUST 4 HOURS; IN SHOGI, ALPHAZERO FIRST OUTPERFORMED ELMO AFTER 2 HOURS; AND IN GO, ALPHAZERO FIRST OUTPERFORMED THE VERSION OF ALPHAGO THAT BEAT THE LEGENDARY PLAYER LEE SEDOL IN 2016 AFTER 30 HOURS. NOTE: EACH TRAINING STEP REPRESENTS 4,096 BOARD POSITIONS.

To learn each game, an untrained neural network plays millions of games against itself via a process of trial and error called [reinforcement learning](#). At first, it plays completely randomly, but over time the system learns from wins, losses, and draws to adjust the parameters of the neural network, making it more likely to choose advantageous moves in the future. The amount of training the network needs depends on the style and complexity of the game, taking approximately 9 hours for chess, 12 hours for shogi, and 13 days for Go.

The trained network is used to guide a search algorithm – known as Monte-Carlo Tree Search (MCTS) – to select the most promising moves in games. For each move, AlphaZero searches only a small fraction of the positions considered by traditional chess engines. In Chess, for example, it searches only 60 thousand positions per second in chess, compared to roughly 60 million for Stockfish.



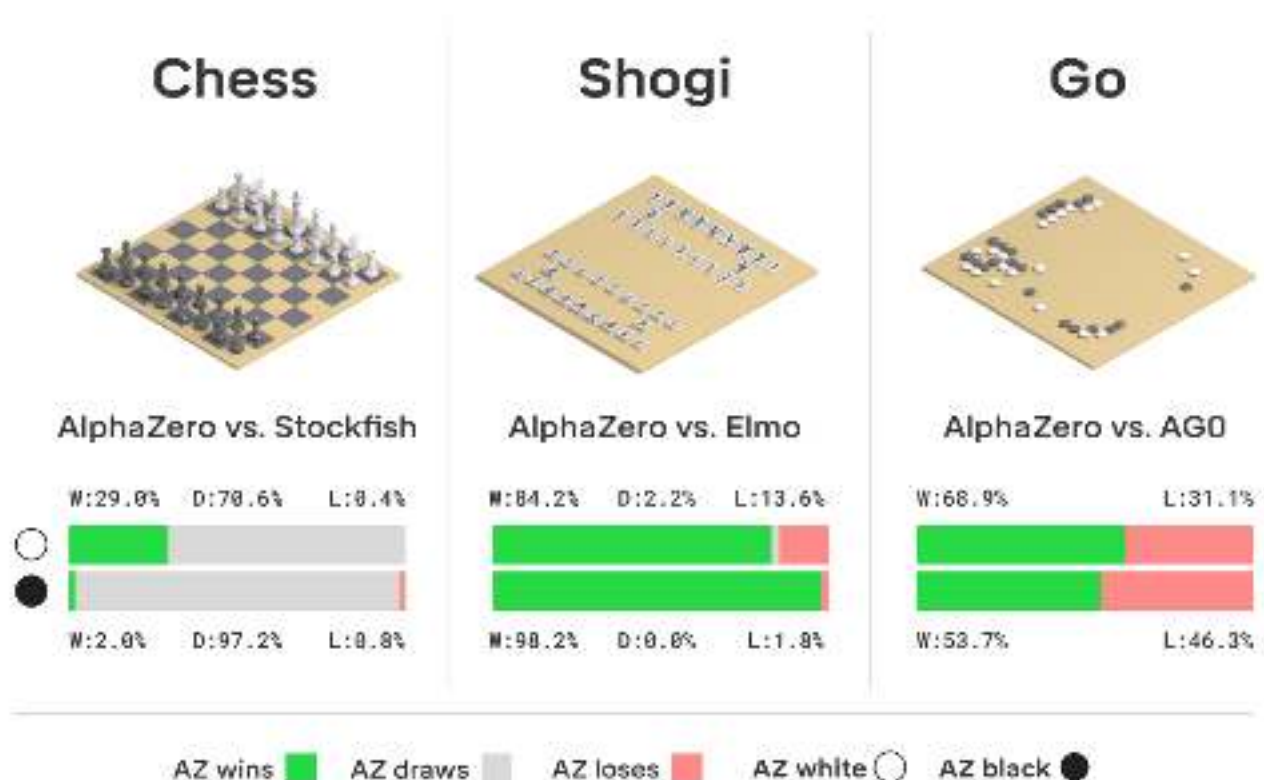
The fully trained systems were tested against the strongest hand-crafted engines for chess ([Stockfish](#)) and shogi ([Elmo](#)), along with our previous self-taught system AlphaGo Zero, the strongest Go player known.

AlphaGo zero, the strongest Go player known.

- Each program ran on the hardware for which they were designed. Stockfish and Elmo used 44 CPU cores (as in the [TCEC world championship](#)), whereas AlphaZero and AlphaGo Zero used a single machine with 4 first-generation [TPUs](#) and 44 [CPU](#) cores. A first generation TPU is roughly similar in inference speed to commodity hardware such as an [NVIDIA Titan V GPU](#), although the architectures are not directly comparable.
- All matches were played using time controls of three hours per game, plus an additional 15 seconds for each move.

In each evaluation, AlphaZero convincingly beat its opponent:

- In chess, AlphaZero defeated the [2016 TCEC \(Season 9\)](#) world champion [Stockfish](#), winning 155 games and losing just six games out of 1,000. To verify the robustness of AlphaZero, we also played a series of matches that started from common human openings. In each opening, AlphaZero defeated Stockfish. We also played a match that started from the set of opening positions used in the 2016 TCEC world championship, along with a series of additional matches against the most recent development version of Stockfish, and a variant of Stockfish that uses a strong opening book. In all matches, AlphaZero won.
- In shogi, AlphaZero defeated the [2017 CSA world champion version of Elmo](#), winning 91.2% of games.
- In Go, AlphaZero defeated [AlphaGo Zero](#), winning 61% of games.



However, it was the style in which AlphaZero plays these games that players may find most fascinating. In Chess, for example, AlphaZero independently discovered and played common human motifs during its self-play training such as openings, king safety and pawn structure. But, being self-taught and therefore unconstrained by conventional wisdom about the game, it also developed its own intuitions and strategies adding a new and expansive set of exciting and novel ideas that augment centuries of thinking about chess strategy.

The first thing that players will notice is AlphaZero's style, says Matthew Sadler – “the way its pieces swarm around the opponent’s king with purpose and power”.

Underpinning that, he says, is AlphaZero’s highly dynamic game play that maximises the activity and mobility of its own pieces while minimising the activity and mobility of its opponent’s pieces. Counterintuitively, AlphaZero also seems to place less value on “material”, an idea that underpins the modern game where each piece has a value and if one player has a greater value of pieces on the board than the other, then they have a material advantage. Instead, AlphaZero is willing to sacrifice material early in a game for gains that will only be recouped in the long-term.

“Impressively, it manages to impose its style of play across a very wide range of positions and openings,” says Matthew, who also observes that it plays in a very deliberate style from its first move with a “very human sense of consistent purpose”.

“Traditional engines are exceptionally strong and make few obvious mistakes, but can drift when faced with positions with no concrete and calculable solution,” he says. “It’s precisely in such positions where ‘feeling’, ‘insight’ or ‘intuition’ is required that AlphaZero comes into its own.”

This unique ability, not seen in other traditional chess engines, has already been harnessed to give chess fans [fresh insight and commentary](#) on the recent [World Chess Championship](#) match between [Magnus Carlsen](#) and [Fabiano Caruana](#) and will be explored further in [Game Changer](#). “It was fascinating to see how AlphaZero’s analysis differed from that of top chess engines and even top grandmaster play,” says Natasha Regan. “AlphaZero could be a powerful teaching tool for the whole community.”

AlphaZero’s teachings echo what we saw when [AlphaGo](#) played the legendary champion [Lee Sedol](#) in 2016. During [the games](#), AlphaGo played a number of [highly inventive winning moves](#), including move 37 in game two, which overturned hundreds of years of thinking. These moves – and many others – have since been studied by players at all levels including Lee Sedol himself, who said of Move 37: “I thought AlphaGo was based on probability calculation and it was merely a machine. But when I saw this move I changed my mind. Surely AlphaGo is creative.”

As with Go, we are excited about AlphaZero’s creative response to chess, which has been a grand challenge for artificial intelligence since the dawn of the computing age

with early pioneers including Babbage, Turing, Shannon, and von Neumann all trying their hand at designing chess programs. But AlphaZero is about more than chess, shogi or Go. To create intelligent systems capable of solving a wide range of real-world problems we need them to be flexible and generalise to new situations. While there has been some progress towards this goal, it remains a major challenge in AI research with systems capable of mastering specific skills to a very high standard, but often failing when presented with even slightly modified tasks.

AlphaZero's ability to master three different complex games – and potentially any perfect information game – is an important step towards overcoming this problem. It demonstrates that a single algorithm can learn how to discover new knowledge in a range of settings. And, while it is still early days, AlphaZero's creative insights coupled with the encouraging results we see in other projects such as [AlphaFold](#), give us confidence in [our mission](#) to create general purpose learning systems that will one day help us find novel solutions to some of the most important and complex scientific problems.


-
- Read [the paper in Science](#)
 - Download an [Open Access version of the paper](#) [PDF]
 - Read the accompanying [Science editorial](#) by Garry Kasparov
 - Read the [accompanying Perspective article](#) in Science by Deep Blue co-creator Murray Campbell
 - Download [the top 20 AlphaZero-Stockfish games chosen by Grandmaster Matthew Sadler](#) [.zip]
 - Download [the top 10 AlphaZero-Elmo games chosen by shogi Master Yoshiharu Habu](#) [.zip]
 - Download [210 AlphaZero-Stockfish chess games and 100 AlphaZero-Elmo shogi games](#)
 - Download the [accompanying artwork](#)
 - Find out more about the AlphaZero book [Game Changer](#) (New in Chess, January 2019)

This work was done by David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharshan Kumaran, Thore Graepel, Timothy Lillicrap, Karen Simonyan, and Demis Hassabis.

SHARE



AUTHORS




David Silver



Thomas Hubert



Julian Schrittwieser



Demis Hassabis

Further reading

About



Research



Impact



Blog



Safety & Ethics



Careers



PRESS

TERMS & CONDITIONS

PRIVACY POLICY

MODERN SLAVERY STATEMENT

ALPHABET INC

