

[Home](#) \* [Engines](#) \* AlphaZero



The Krampus has come <sup>[3]</sup> <sup>[4]</sup>

### AlphaZero,

a chess and [Go](#) playing entity by [Google DeepMind](#) based on a general [reinforcement learning](#) algorithm with the same name. On [December 5, 2017](#), the DeepMind team around [David Silver](#), [Thomas Hubert](#), and [Julian Schrittwieser](#) along with former [Giraffe](#) author [Matthew Lai](#), reported on their generalized algorithm, combining [Deep learning](#) with [Monte-Carlo Tree Search](#) (MCTS) <sup>[1]</sup>.

A 100 game match versus [Stockfish 8](#) using 64 [threads](#) and a [transposition table](#) size of 1GiB, was won by AlphaZero using a single machine with 4 [Tensor processing units](#) (TPUs) with +28=72-0. Despite a possible hardware advantage of AlphaZero and criticized playing conditions <sup>[2]</sup>, this seems a tremendous achievement.

## Table of Contents

[Description](#)

[Network Architecture](#)

[Training](#)

[See also](#)

[Publications](#)

[Forum Posts](#)

[2017](#)

[2018](#)

[External Links](#)

[Reports](#)

[Stockfish Match](#)

[Misc](#)

[References](#)

[What links here?](#)

## Description

Starting from random play, and given no domain knowledge except the game rules, AlphaZero achieved a superhuman level of play in the games of chess and [Shogi](#) as well as in [Go](#). The algorithm is a more generic version of the [AlphaGo Zero](#) algorithm that was first introduced in the domain of Go <sup>[5]</sup>.

AlphaZero [evaluates positions](#) using non-linear function approximation based on a [deep neural network](#), rather than the [linear function approximation](#) as used in classical chess programs. This neural network takes the board position as input and outputs a vector of move probabilities. The MCTS consists of a series of simulated games of self-play whose move selection is controlled by the neural network. The search returns a vector representing a probability distribution over moves, either proportionally or greedily with respect to the visit counts at the root state.

## Network Architecture

The network is a [deep residual convolutional neural network](#) <sup>[6] [7]</sup> with many layers of spatial NxN planes - [8x8 board](#) arrays for chess. The input describes the [chess position](#) from [side's to move](#) point of view - that is [color flipped](#) for black to move. Each square cell consists of 12 [piece-type](#) and [color](#) bits, e.g. from the current [bitboard board definition](#), and to address [graph history](#) and [path-dependency](#) - times eight, that is up to seven predecessor positions as well - so that [en passant](#), immediate [repetitions](#), or some sense of progress is implicit. Additional inputs, redundant inside each square cell to be conform to the convolution net, consider [castling rights](#), [halfmove clock](#), total move count and side to move.

The deep hidden layers connect the pieces on different squares to each other due to consecutive 3x3 convolutions, where a cell of a layer is connected to the correspondent 3x3 [receptive field](#) of the previous layer, so that after 4 layers, each square is connected to every other cell in the original input layer <sup>[8]</sup>. The output of the neural network is finally represented as an 8x8 board array as well, for every [origin square](#) up to 73 [target square](#) possibilities ( $\text{NRayDirs} \times \text{MaxRayLength} + \text{NKnightDirs} + \text{NPawnDirs} *$

[NMinorPromotions](#)), encoding a probability distribution over  $64 \times 73 = 4,672$  possible moves, where illegal moves were masked out by setting their probabilities to zero, re-normalising the probabilities for remaining moves.

## Training

AlphaZero was trained in 700,000 steps or mini-batches of size 4096 each, starting from randomly initialized parameters, using 5,000 [first-generation TPUs](#) <sup>[9]</sup> to generate self-play games and 64 [second-generation TPUs](#) <sup>[10] [11] [12]</sup> to train the neural networks <sup>[13]</sup>.

## See also

- [Alpha-Beta](#)
- [Alpha I](#)
- [AlphaGo](#)
- [Chess Engines with Neural Networks](#)
- [Learning Chess Programs](#)
- [LCZero](#)

## Publications

- [David Silver](#), [Julian Schrittwieser](#), [Karen Simonyan](#), [Ioannis Antonoglou](#), [Aja Huang](#), [Arthur Guez](#), [Thomas Hubert](#), [Lucas Baker](#), [Matthew Lai](#), [Adrian Bolton](#), [Yutian Chen](#), [Timothy Lillicrap](#), [Fan Hui](#), [Laurent Sifre](#), [George van den Driessche](#), [Thore Graepel](#), [Demis Hassabis](#) (2017). [Mastering the game of Go without human knowledge](#). *Nature*, Vol. 550, [pdf](#)
- [David Silver](#), [Thomas Hubert](#), [Julian Schrittwieser](#), [Ioannis Antonoglou](#), [Matthew Lai](#), [Arthur Guez](#), [Marc Lanctot](#), [Laurent Sifre](#), [Dharshan Kumaran](#), [Thore Graepel](#), [Timothy Lillicrap](#), [Karen Simonyan](#), [Demis Hassabis](#) (2017). *Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm*. [arXiv:1712.01815](#)

## Forum Posts

### 2017

- [Google's AlphaGo team has been working on chess](#) by [Peter Kappler](#), [CCC](#), December 06, 2017
- [Historic Milestone: AlphaZero](#) by Miguel Castanuela, [CCC](#), December 06, 2017
- [AlphaZero beats AlphaGo Zero, Stockfish, and Elmo](#) by Carl Lumma, [CCC](#), December 06, 2017
- [AlphaZero vs Stockfish](#) by Bigler, [CCC](#), December 06, 2017
- [Deepmind drops the bomb](#) by Leebot, [FishCooking](#), December 06, 2017
- [AlphaZero beats Stockfish 8 by 64-36](#) by [Venator](#), [Rybka Forum](#), December 06, 2017

- [Alpha Zero](#) by [BB+](#), [OpenChess Forum](#), December 06, 2017
- [AlphaGo Zero And AlphaZero, RomiChess done better](#) by [Michael Sherwin](#), [CCC](#), December 07, 2017 » [RomiChess](#)
- [BBC News: 'Google's ... DeepMind AI claims chess crown'](#) by pennine22, [Hiarcs Forum](#), December 07, 2017
- [Press Release Stockfish vs. AlphaZero](#) by Michael Whiteley, [FishCooking](#), December 08, 2017
- [AlphaZero reinvents mobility and romanticism](#) by [Chris Whittington](#), [Rybka Forum](#), December 08, 2017 » [Alpha Zero's "Immortal Zugzwang Game"](#)
- [Reactions about AlphaZero from top GMs...](#) by [Norman Schmidt](#), [CCC](#), December 08, 2017 » [Reactions From Top GMs, Stockfish Author](#)
- [AlphaZero is not like other chess programs](#) by [Dann Corbit](#), [CCC](#), December 08, 2017
- [Re: AlphaZero is not like other chess programs](#) by [Rein Halbersma](#), [CCC](#), December 09, 2017
- [Photo of Google Cloud TPU cluster](#) by [Norman Schmidt](#), [CCC](#), December 09, 2017
- [Cerebellum analysis of the AlphaZero - Stockfish Games](#) by [Thomas Zippbroth](#), [CCC](#), December 11, 2017 » [Cerebellum](#)
- [Open letter to Google DeepMind](#) by Michael Stempera, [FishCooking](#), December 12, 2017
- [recent article on alphazero ... 12/11/2017 ...](#) by Dan Ellwein, [CCC](#), December 14, 2017
- [An AlphaZero inspired project](#) by [Truls Edvard Stokke](#), [CCC](#), December 14, 2017
- [AlphaZero - Tactical Abilities](#) by [David Rasmussen](#), [CCC](#), December 16, 2017
- [In chess, AlphaZero outperformed Stockfish after just 4 hours](#) by [Ed Schroder](#), [CCC](#), December 18, 2017
- [AlphaZero - Youtube Videos](#) by Christoph Fieberg, [CSS Forum](#), December 18, 2017
- [AlphaZero Chess is not that strong ...](#) by [Vincent Lejeune](#), [CCC](#), December 19, 2017
- [David Silver \(Deepmind\) inaccuracies](#) by [Ed Schroder](#), [CCC](#), December 21, 2017
- [AZ vs SF - game 99](#) by [Rebel](#), [Rybka Forum](#), December 23, 2017
- [AlphaZero performance](#) by [Martin Sedlak](#), [CCC](#), December 25, 2017
- [A Simple Alpha\(Go\) Zero Tutorial](#) by Oliver Roese, [CCC](#), December 30, 2017
- [AlphaZero: The 10 Top Shots](#) by [Walter Eigenmann](#), [CCC](#), December 30, 2017

## 2018

- [SF was more seriously handicapped than I thought](#) by [Kai Laskos](#), [CCC](#), January 02, 2018
- [Chess World to Google Deep Mind..Prove You beat Stockfish 8!](#) by AA Ross, [CCC](#), January 11, 2018
- [Article:"How Alpha Zero Sees/Wins"](#) by AA Ross, [CCC](#), January 17, 2018 » [How AlphaZero Wins](#)
- [Connect 4 AlphaZero implemented using Python...](#) by [Steve Maughan](#), [CCC](#), January 29, 2018 » [Connect Four](#), [Python](#)
- [Seeing Alphazero in perspective ...](#) by Dan Ellwein, [CCC](#), February 10, 2018

## External Links

- [AlphaZero from Wikipedia](#)
- [AlphaGo Zero - AlphaZero from Wikipedia](#)
- Keynote [David Silver NIPS 2017 Deep Reinforcement Learning Symposium AlphaZero](#),

December 06, 2017, [YouTube](#) Video <sup>[14]</sup>

- [A Simple Alpha\(Go\) Zero Tutorial](#) by [Surag Nair](#), [Stanford University](#), December 29, 2017 <sup>[15]</sup>  
[GitHub - suragnair/alpha-zero-general: A clean and simple implementation of a self-play learning algorithm based on AlphaGo Zero \(any game, any framework!\)](#)

## Reports

- [DeepMind's AI became a superhuman chess player in a few hours, just for fun](#) by [James Vincent](#), [The Verge](#), December 06, 2017
- [Entire human chess knowledge learned and surpassed by DeepMind's AlphaZero in four hours](#) by [Sarah Knapton](#), and [Leon Watson](#), [The Telegraph](#), December 06, 2017
- [Google's 'superhuman' DeepMind AI claims chess crown](#), [BBC News](#), December 06, 2017 <sup>[16]</sup>
- [DeepMind's AlphaZero crushes chess](#) by [Colin McGourty](#), [Chess24.com](#), December 06, 2017
- [One Small Step for Computers, One Giant Leap for Mankind](#) by [Dana Mackenzie](#), [Dana Blogs Chess](#), December 06, 2017
- [Google's AlphaZero Destroys Stockfish In 100-Game Match](#) by [Mike Klein](#), [Chess.com](#), December 06, 2017
- [The future is here – AlphaZero learns chess](#) by [Albert Silver](#), [ChessBase News](#), December 06, 2017
- [AlphaZero: Reactions From Top GMs, Stockfish Author](#) by [Peter Doggers](#), [Chess.com](#), December 08, 2017 » [Stockfish](#), [Tord Romstad](#) <sup>[17]</sup>
- [Is AlphaZero really a scientific breakthrough in AI?](#) by [Jose Camacho Collados](#), [Medium](#), December 11, 2017 <sup>[18]</sup>
- [Alpha Zero: Comparing "Orangutans and Apples"](#) by [André Schulz](#), [ChessBase News](#), December 13, 2017
- [Kasparov on Deep Learning in chess](#) by [Frederic Friedel](#), [ChessBase News](#), December 13, 2017

## Stockfish Match

- [AlphaZero vs Stockfish Games • lichess.org](#)
- [The chess games of AlphaZero \(Computer\)](#) from [chessgames.com](#)
- [Cerebellum AlphaZero Analysis](#) » [Cerebellum](#) <sup>[19]</sup>
- [Deep Mind Alpha Zero's "Immortal Zugzwang Game" against Stockfish](#) by [Antonio Radic](#), December 07, 2017, [YouTube](#) Video <sup>[20]</sup> <sup>[21]</sup> » [Zugzwang](#)
- [Deep Mind AI Alpha Zero Dismantles Stockfish's French Defense](#) by [Antonio Radic](#), December 08, 2017, [YouTube](#) Video
- [How AlphaZero Wins](#) by [Dana Mackenzie](#), [Dana Blogs Chess](#), December 15, 2017 <sup>[22]</sup>

## Misc

- [How to build your own AlphaZero AI using Python and Keras](#) by [David Foster](#), January 26, 2018 » [Connect Four](#), [Python](#) <sup>[23]</sup>
- [Can - Halleluwah](#), from [Tango Mago](#) 1971, [YouTube](#) Video

lineup: [Irmin Schmidt](#), [Michael Karoli](#), [Holger Czukay](#), [Damo Suzuki](#), [Jaki Liebezeit](#)

## References

1. [^](#) [David Silver](#), [Thomas Hubert](#), [Julian Schrittwieser](#), [Ioannis Antonoglou](#), [Matthew Lai](#), [Arthur Guez](#), [Marc Lanctot](#), [Laurent Sifre](#), [Dharshan Kumaran](#), [Thore Graepel](#), [Timothy Lillicrap](#), [Karen Simonyan](#), [Demis Hassabis](#) (2017). *Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm*. [arXiv:1712.01815](#)
2. [^](#) [Alpha Zero](#) by [BB+](#), [OpenChess Forum](#), December 06, 2017
3. [^](#) [Krampus](#) at [Perchtenlauf](#) in [Klagenfurt](#), Image by [Anita Martinz](#), December 02, 2006, [CC BY 2.0](#), [Wikimedia Commons](#)
4. [^](#) "5th of December - The [Krampus](#) has come" was suggested by [Michael Scheidl](#) in [AlphaZero](#) by Peter Martan, [CSS Forum](#), December 06, 2017, with further comments by [Ingo Althöfer](#)
5. [^](#) [AlphaGo Zero: Learning from scratch](#) by [Demis Hassabis](#) and [David Silver](#), [DeepMind](#), October 18, 2017
6. [^](#) The principle of residual nets is to add the input of the layer to the output of each layer. With this simple modification training is faster and enables deeper networks, see [Tristan Cazenave](#) (2017). [Residual Networks for Computer Go](#). [IEEE Transactions on Computational Intelligence and AI in Games](#), Vol. PP, No. 99, [pdf](#)
7. [^](#) [Residual Networks for Computer Go](#) by Brahim Hamadicharef, [CCC](#), December 07, 2017
8. [^](#) [Re: AlphaZero is not like other chess programs](#) by [Rein Halbersma](#), [CCC](#), December 09, 2017
9. [^](#) [First In-Depth Look at Google's TPU Architecture](#) by [Nicole Hemsoth](#), [The Next Platform](#), April 05, 2017
10. [^](#) [Photo of Google Cloud TPU cluster](#) by [Norman Schmidt](#), [CCC](#), December 09, 2017
11. [^](#) [First In-Depth Look at Google's New Second-Generation TPU](#) by [Nicole Hemsoth](#), [The Next Platform](#), May 17, 2017
12. [^](#) [Under The Hood Of Google's TPU2 Machine Learning Clusters](#) by Paul Teich, [The Next Platform](#), May 22, 2017
13. [^](#) [David Silver](#), [Thomas Hubert](#), [Julian Schrittwieser](#), [Ioannis Antonoglou](#), [Matthew Lai](#), [Arthur Guez](#), [Marc Lanctot](#), [Laurent Sifre](#), [Dharshan Kumaran](#), [Thore Graepel](#), [Timothy Lillicrap](#), [Karen Simonyan](#), [Demis Hassabis](#) (2017). *Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm*. [arXiv:1712.01815](#)
14. [^](#) [AlphaZero explained by one creator](#) by [Mario Carbonell Martinez](#), [CCC](#), December 19, 2017
15. [^](#) [A Simple Alpha\(Go\) Zero Tutorial](#) by Oliver Roesse, [CCC](#), December 30, 2017
16. [^](#) [BBC News: 'Google's ... DeepMind AI claims chess crown'](#) by pennine22, [Hiarcs Forum](#), December 07, 2017
17. [^](#) [Reactions about AlphaZero from top GMs...](#) by [Norman Schmidt](#), [CCC](#), December 08, 2017
18. [^](#) [recent article on alphazero ... 12/11/2017 ...](#) by Dan Ellwein, [CCC](#), December 14, 2017
19. [^](#) [Cerebellum analysis of the AlphaZero - Stockfish Games](#) by [Thomas Zipproth](#), [CCC](#), December 11, 2017
20. [^](#) [AlphaZero reinvents mobility and romanticism](#) by [Chris Whittington](#), [Rybka Forum](#), December 08, 2017
21. [^](#) [Immortal Zugzwang Game from Wikipedia](#)
22. [^](#) [Article:"How Alpha Zero Sees/Wins"](#) by AA Ross, [CCC](#), January 17, 2018

23. <sup>^</sup> [Connect 4 AlphaZero implemented using Python...](#) by [Steve Maughan](#), [CCC](#), January 29, 2018

## What links here?

Page	Date Edited
<a href="#">Albert Silver</a>	Feb 25, 2018
<a href="#">Alpha I</a>	Dec 6, 2017
<a href="#">AlphaZero</a>	Feb 10, 2018
<a href="#">Arthur Guez</a>	Dec 6, 2017
<a href="#">Chess</a>	Jan 21, 2018
<a href="#">Chris Whittington</a>	Dec 9, 2017
<a href="#">Connect Four</a>	Jan 31, 2018
<a href="#">Dana Mackenzie</a>	Jan 18, 2018
<a href="#">Dann Corbit</a>	Dec 9, 2017
<a href="#">David Rasmussen</a>	Dec 16, 2017
<a href="#">David Silver</a>	Feb 11, 2018
<a href="#">Deep Learning</a>	Feb 12, 2018
<a href="#">DeepMind</a>	Dec 9, 2017
<a href="#">Demis Hassabis</a>	Dec 8, 2017
<a href="#">Dharshan Kumaran</a>	Dec 9, 2017
<a href="#">Engines</a>	Mar 10, 2018
<a href="#">Frederic Friedel</a>	Jan 6, 2018
<a href="#">Go</a>	Jan 24, 2018
<a href="#">Ioannis Antonoglou</a>	Dec 6, 2017
<a href="#">Julian Schrittwieser</a>	Dec 7, 2017
<a href="#">Karen Simonyan</a>	Dec 10, 2017
<a href="#">Laurent Sifre</a>	Dec 7, 2017
<a href="#">LCZero</a>	Apr 18, 2018
<a href="#">Learning</a>	Feb 20, 2018
<a href="#">Marc Lanctot</a>	Jan 10, 2018
<a href="#">Mark Watkins</a>	Dec 6, 2017
<a href="#">Martin Sedlak</a>	Dec 25, 2017
<a href="#">Mathematician</a>	Apr 9, 2018
<a href="#">Matthew Lai</a>	Dec 6, 2017
<a href="#">Michael Sherwin</a>	Dec 18, 2017
<a href="#">Monte-Carlo Tree Search</a>	Apr 26, 2018
<a href="#">Neural Networks</a>	Mar 12, 2018
<a href="#">Othello</a>	Jan 4, 2018
<a href="#">Peter Kappler</a>	Feb 4, 2018
<a href="#">Python</a>	Jan 31, 2018
<a href="#">Reinforcement Learning</a>	Feb 12, 2018
<a href="#">RomiChess</a>	Jan 2, 2018
<a href="#">Shogi</a>	Feb 19, 2018
<a href="#">Steve Maughan</a>	Jan 31, 2018
<a href="#">Stockfish</a>	Apr 7, 2018

Page	Date Edited
<a href="#">Test-Positions</a>	Feb 25, 2018
<a href="#">Thomas Hubert</a>	Dec 7, 2017
<a href="#">Thore Graepel</a>	Jan 10, 2018
<a href="#">Timothy Lillicrap</a>	Dec 9, 2017
<a href="#">Tord Romstad</a>	Dec 9, 2017
<a href="#">University of Toronto</a>	Dec 8, 2017
<a href="#">Zugzwang</a>	Dec 8, 2017

[Up one Level](#)