A Machine Learning playing GO

TUM Data Innovation Lab

Sponsored by: Lehrstuhl für Geometrie und Visualisierung

Mentor: M.Sc. Bernhard Werner

TUM-DI-LAB Coordinator & Project Manager: Dr. Ricardo Acevedo Cabra

Head of TUM-DI-LAB: Prof. Dr. Massimo Fornasier

Nathanael Bosch, Benjamin Degenhart, Karthikeya Kaushik, Yushan Liu, Benedikt Mairhörmann, Stefan Peidli, Faruk Toy, Patrick Wilson

2016: **AlphaGo vs. Lee Sedol**

www.newscientist.com/article/2079871-im-in-shock-how-an-ai-beat-the-worlds-best-human-at-go/

The Korea Herald

Sales of Go-playing products soar

By Korea Herald

NBC NEWS

'Go' Matches Between Lee Sedol and AlphaGo Push AI Boundaries

by Stella Kim / Mar.09.2016 / 1:58 PM ET

NATURE | NEWS

South Korea trumpets $860-million AI fund after AlphaGo 'shock'

CADE METZ BUSINESS 03.16.16 07:00 AM

WIRED

IN TWO MOVES, ALPHAGO AND LEE SEDOL REDEFINED THE FUTURE

The Guardian International edition

DeepMind's artificial intelligence astonishes fans to defeat human opponent and offers evidence computer software has mastered a major challenge

2014: **“May take another ten years or so”**

ALAN LEVINOVITZ BUSINESS 05.12.14 06:30 AM

WIRED

THE MYSTERY OF GO, THE ANCIENT GAME THAT COMPUTERS STILL CAN'T WIN

www.wired.com/2014/05/the-world-of-computer-go/

Complexity of Go compared to Chess

https://research.googleblog.com/2016/01/alphago-mastering-ancient-game-of-go.htm
http://scienceblogs.com/startswithabang/2011/11/14/found-the-first-atoms-in-the-u

Outline

- Go & Computer Go
- Our Project
- Neural Networks
- Different Approaches & Results
- Learning from AI

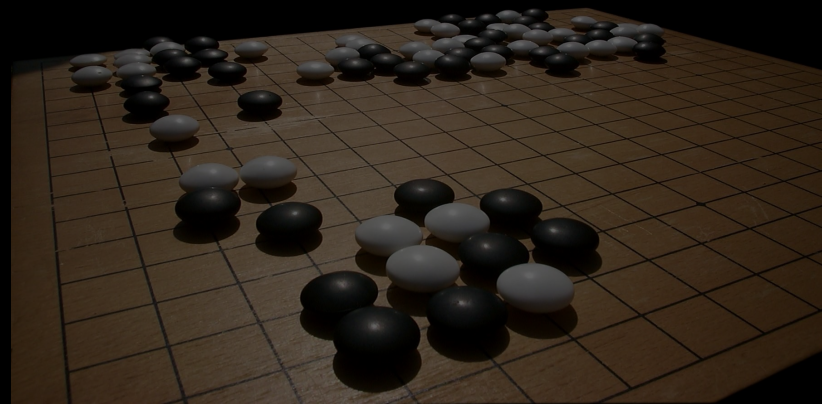

Go & Computer Go

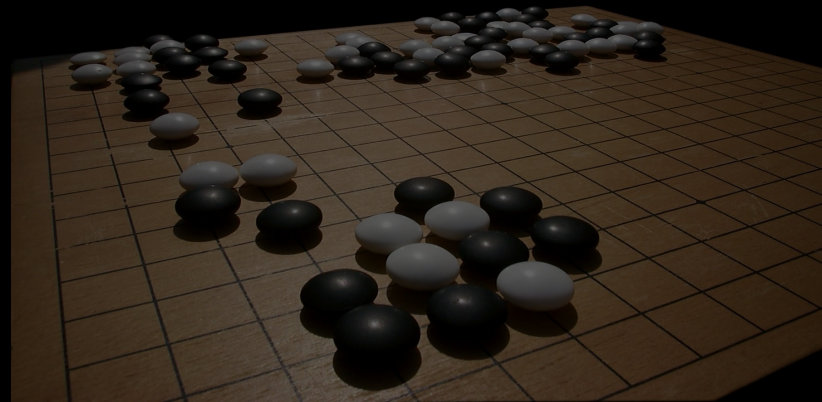

What is Go

- Go is an abstract strategy board game for two players, in which the aim is to surround more territory than the opponent.
- Origin: China ~4000 years ago
- Among the most played games in the world
- Top-level international players can become millionaires

Images from en.wikipedia.org/wiki/Go_(game)

Quotes

“The rules of go are so elegant, organic, and rigorously logical that if intelligent life forms exist elsewhere in the universe, they almost certainly play go”. - Edward Lasker

“The board is a mirror of the mind of the players as the moments pass. When a master studies the record of a game he can tell at what point greed overtook the pupil, when he became tired, when he fell into stupidity, and when the maid came by with tea.” - Anonymous

Proverbs

Go proverbs summarize wisdom in easy-to-remember phrases.

- Move where your opponent wants to move.
- If you have lost all four corners then you have lost.
- If you have secured all four corners then you have lost.
- Don't follow proverbs blindly.
- ...

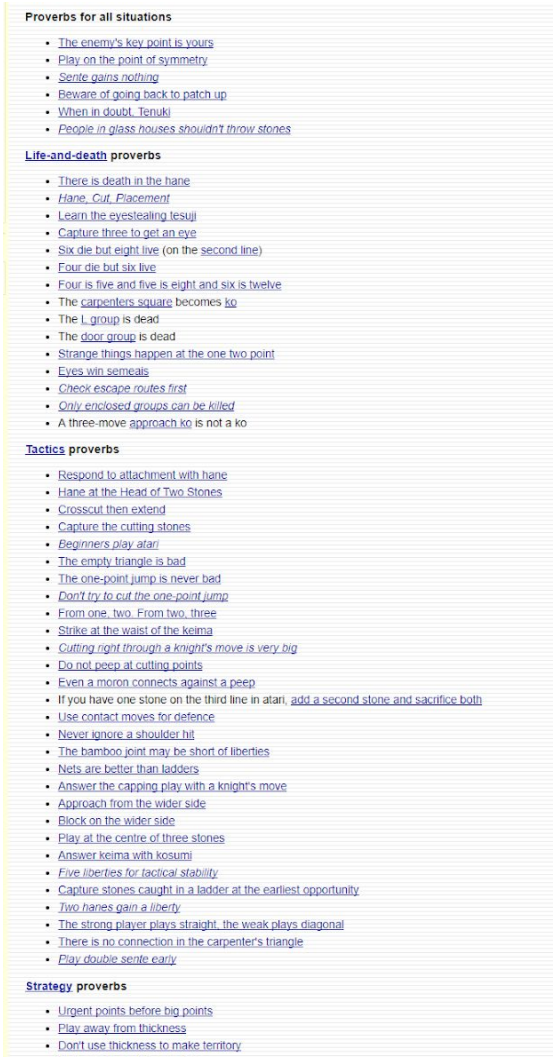

Rules

Demo in Drago

- Counting
- Capturing
- Suicide
- Ko
- Example for complex patterns: Ladder
- Handicap

Complexity

- Local vs. global
- Highest state space complexity of any information complete game

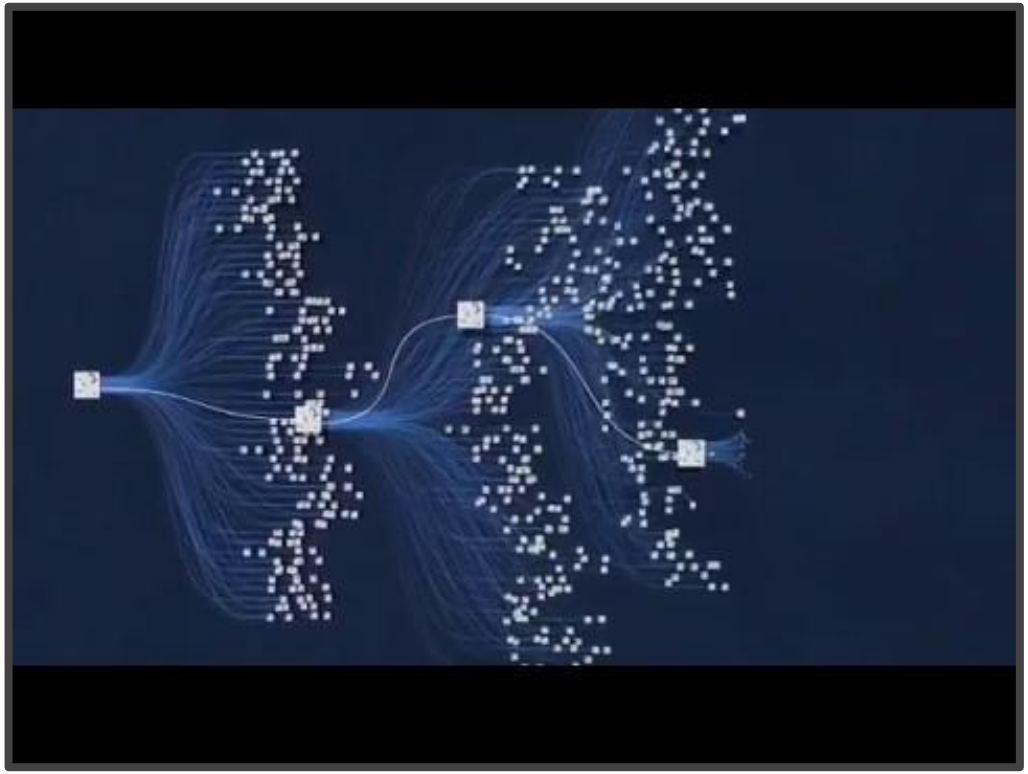

Chess vs. Go search space - animation from AlphaGo movie

Computer Go Timeline: **Early decades**

- **1968**: 1st Go program by A. L. Zobrist, thesis on pattern recognition, influence function to estimate territory and Zobrist hashing to detect ko
- **1981**: Jonathan K Millen, “Wally” on a 15x15 board, 1 KB RAM
- **1994**: The winning program of the World Computer Go Championship, “Go Intellect”, loses consistently against youth players despite 15-stone handicap. Weaknesses were easy to exploit.

Many programs were developed. Screenshot from Sensei’s Library GoPlayingPrograms page:

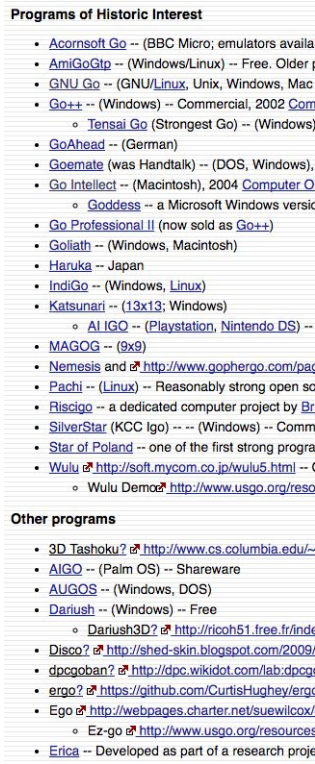

Programs of Historic Interest

- Acornsoft Go -- (BBC Micro; emulators availa
- AmiGoGtp -- (Windows/Linux) -- Free. Older
- GNU Go -- (GNU/Linux, Unix, Windows, Mac
- Go++ -- (Windows) -- Commercial, 2002 Com
  - Tensai Go (Strongest Go) -- (Windows)
- GoAhead -- (German)
- Goemate (was Handtalk) -- (DOS, Windows),
- Go Intellect -- (Macintosh), 2004 Computer O
  - Goddess -- a Microsoft Windows versio
- Go Professional II (now sold as Go++)
- Goliath -- (Windows, Macintosh)
- Haruka -- Japan
- IndiGo -- (Windows, Linux)
- Katsunari -- (13x13; Windows)
  - AI IGO -- (Playstation, Nintendo DS) --
- MAGOG -- (9x9)
- Nemesis and http://www.gophergo.com/pa
- Pachi -- (Linux) -- Reasonably strong open so
- Riscigo -- a dedicated computer project by Br
- SilverStar (KCC Igo) -- -- (Windows) -- Comm
- Star of Poland -- one of the first strong progra
- Wulu http://soft.mycom.co.jp/wulu5.html --
  - Wulu Demo http://www.usgo.org/reso

Other programs

- 3D Tashoku? http://www.cs.columbia.edu/~
- AIGO -- (Palm OS) -- Shareware
- AUGOS -- (Windows, DOS)
- Dariush -- (Windows) -- Free
  - Dariush3D? http://ricoh51.free.fr/inde
- Disco? http://shed-skin.blogspot.com/2009/
- dpcgoban? http://dpc.wikidot.com/lab:dpcg
- ergo? https://github.com/CurtisHughey/ergo
- Ego http://webpages.charter.net/suewilcox/
  - Ez-go http://www.usgo.org/resources
- Erica -- Developed as part of a research proje

Computer Go Timeline: **21st century**

Developments in Monte Carlo tree search and machine learning…

- **2008:** “MoGo” beat Kim MyungWan (8p) with a 9-stone handicap, on 19×19 game, ran on an 800-node supercomputer
- **2009:** first programs holding dan-level ranks on KGS Go Server
- **2013**: “Crazy Stone” beat Yoshio Ishida (9p) with a 4-stone handicap
- **2014**: “Crazy Stone” lost an even game against Franz-Jozef Dickhut (6d), wins one match though

Range	Stage
30–20k	Beginner
19–10k	Casual player
9–1k	Intermediate amateur
1–7d	Advanced amateur
1–9p	Professional Player

en.wikipedia.org/wiki/Go_ranks_and_ratings

Computer Go Timeline: **Deep Learning**

- **2015**: “AlphaGo” by Google DeepMind beat Fan Hui (2p), the European Go champion, in all 5 games
- ★ **2016** ★ “AlphaGo” beat Lee Sedol (9p) in 4 of 5 games
- **2017**: “AlphaGo” beat Ke Jie (9p), in 3 of 3 games
- **October 2017**: “AlphaGo Zero”, no human input, beating its previous Ke Jie version in 89% of games

The basic principles of AlphaGo were published. Other teams produced high-level programs inspired by it. By 2017 others were also capable of defeating high level professionals some of the time.

Computer Go Timeline: **Handicap curve**

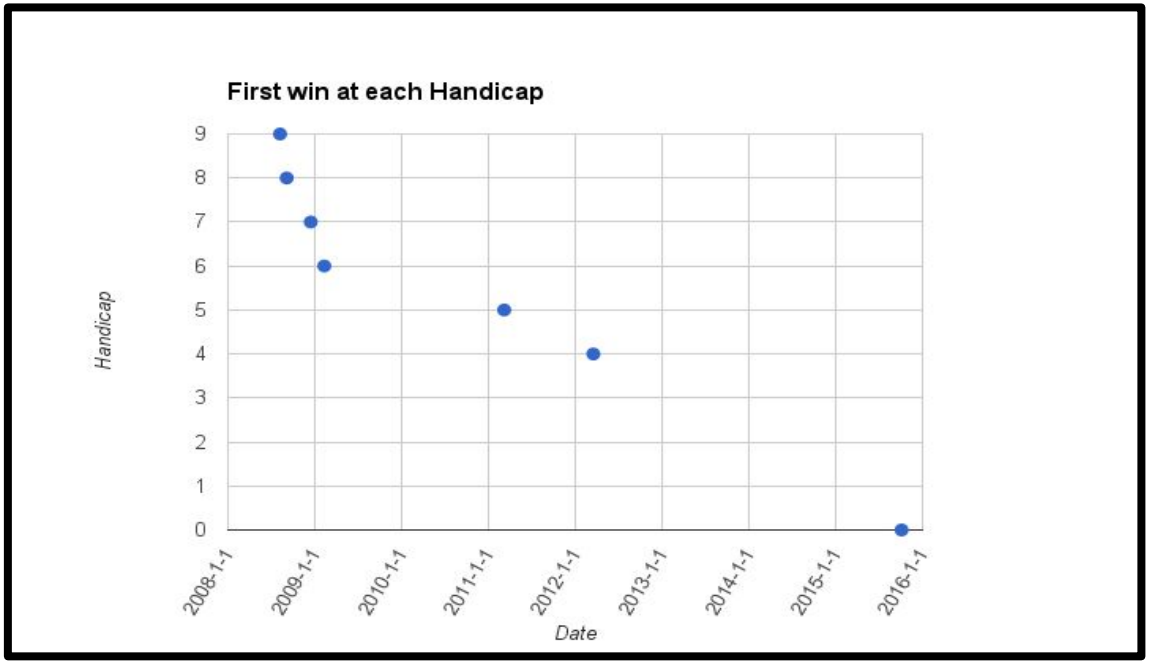

Handicap	Date	Computer	Human
9	2008-8-7	MoGo	Myungwan Kim 8p
8	2008-9-4	CrazyStone	Kaori Aoba 4p
7	2008-12-14	CrazyStone	Kaori Aoba 4p
6	2009-2-9	MoGo	Li-Chen Chien 1p
5	2011-3-8	Zen	Kozo Hayashi 6p
4	2012-3-17	Zen	Takemiya Masaki 9p
0	2015-10-5	AlphaGo	Fan Hui 2p
0	2016-03-09	AlphaGo	Lee Sedol 9p

AlphaGo movie

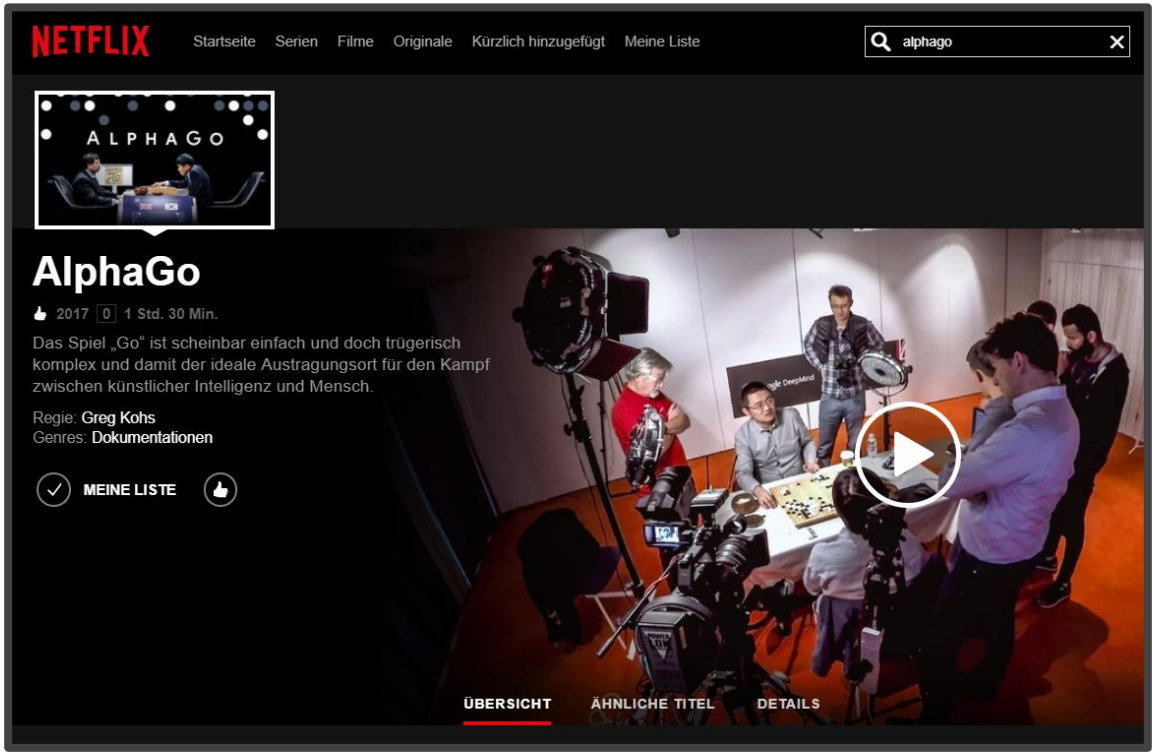

Our Project

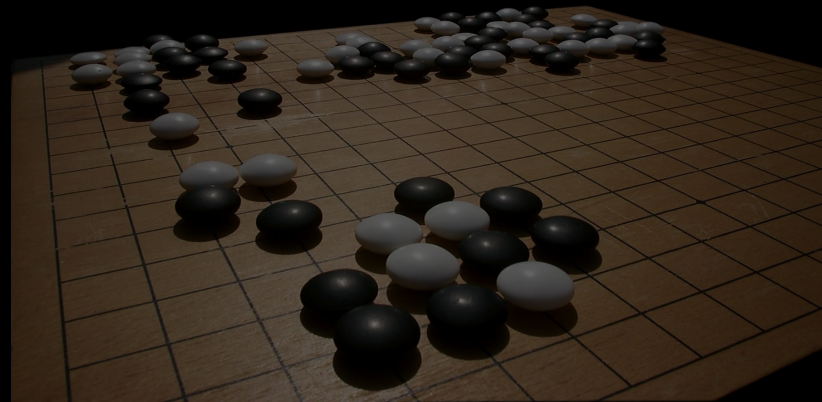

Our Project

Play Go using simple neural networks

Our Project

Play Go

using simple neural networks

Playing Go on a computer

- Game engine
- Command line interface
- GUI
- Communication Protocol (GTP)
- Implement a random agent

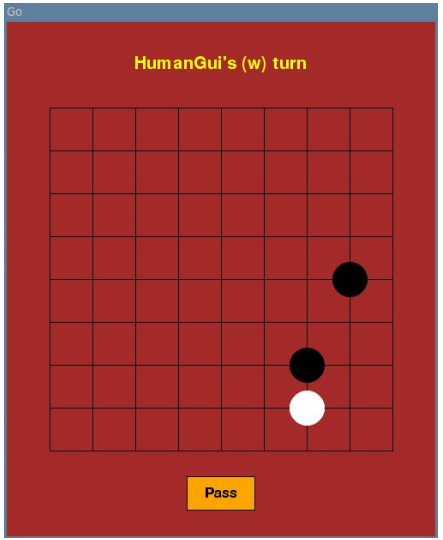

Neural Networks

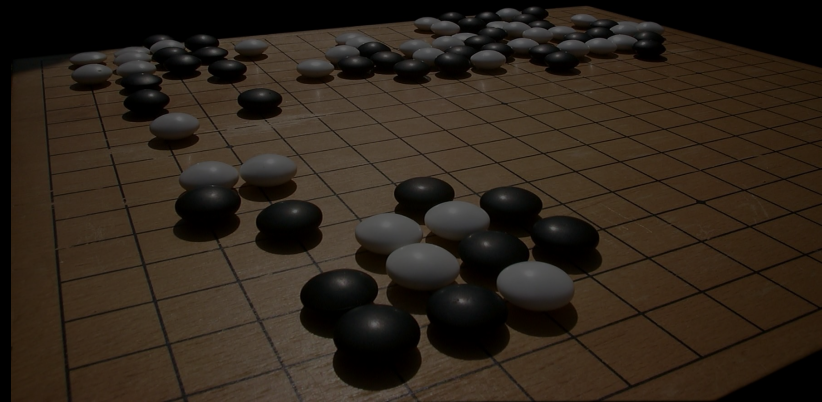

Our Go Neural Networks: **Architecture**

Goal: Find **good** approximation of (f: input → output) through **learning**

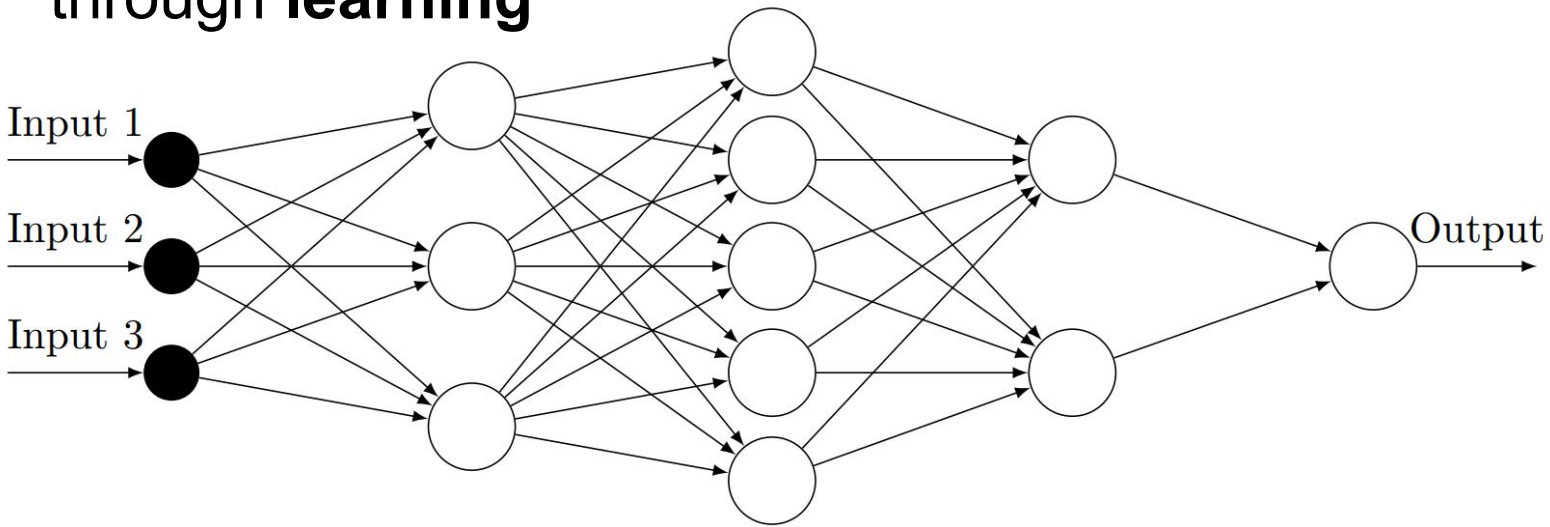

Our Go Neural Networks: **Gradient Descent**

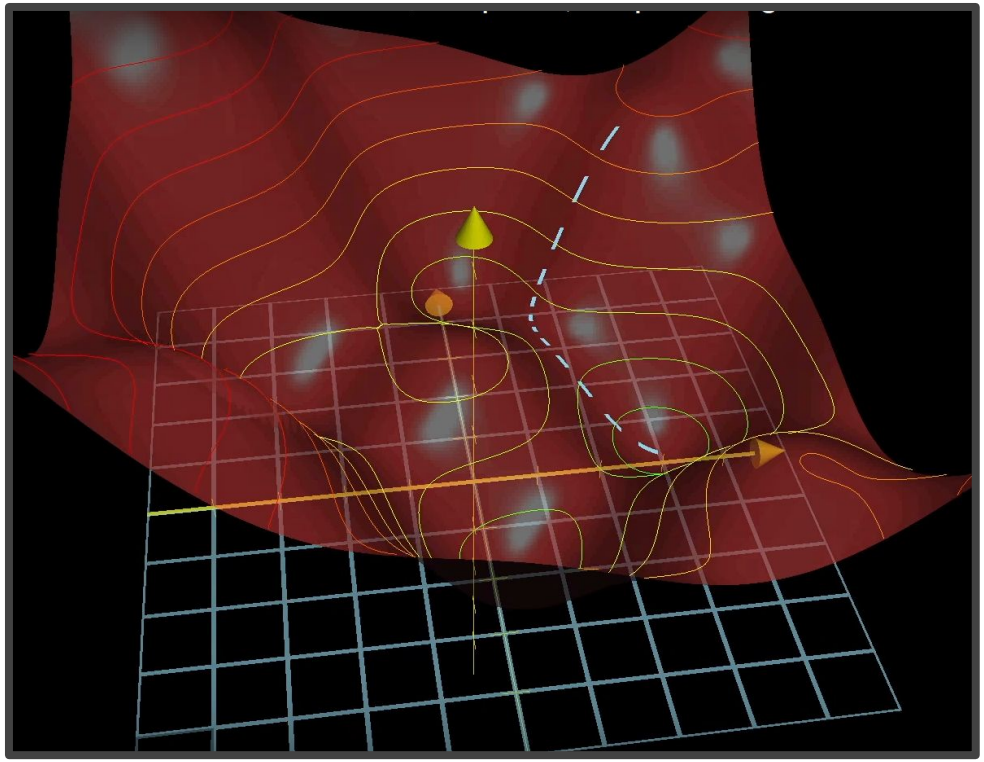

https://www.youtube.com Channel: 3blueonebrown

Our Go Neural Networks: **Input**

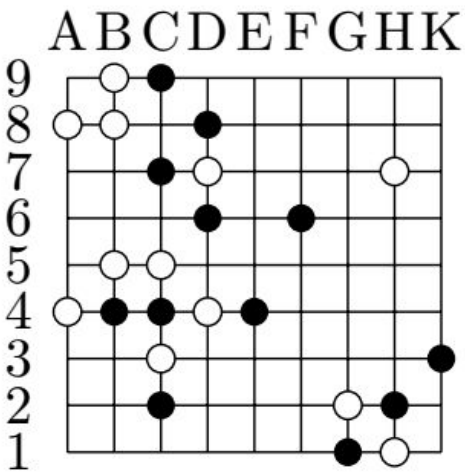

$\simeq$

$$\begin{pmatrix}
0 & 1 & -1 & 0 & 0 & 0 & 0 & 0 & 0 \\
1 & 1 & 0 & -1 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & -1 & 1 & 0 & 0 & 0 & 1 & 0 \\
0 & 0 & 0 & -1 & 0 & -1 & 0 & 0 & 0 \\
0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\
1 & -1 & -1 & 1 & -1 & 0 & 0 & 0 & 0 \\
0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & -1 \\
0 & 0 & -1 & 0 & 0 & 0 & 1 & -1 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & -1 & 1 & 0
\end{pmatrix}$$

Our Go Neural Networks: **Input**

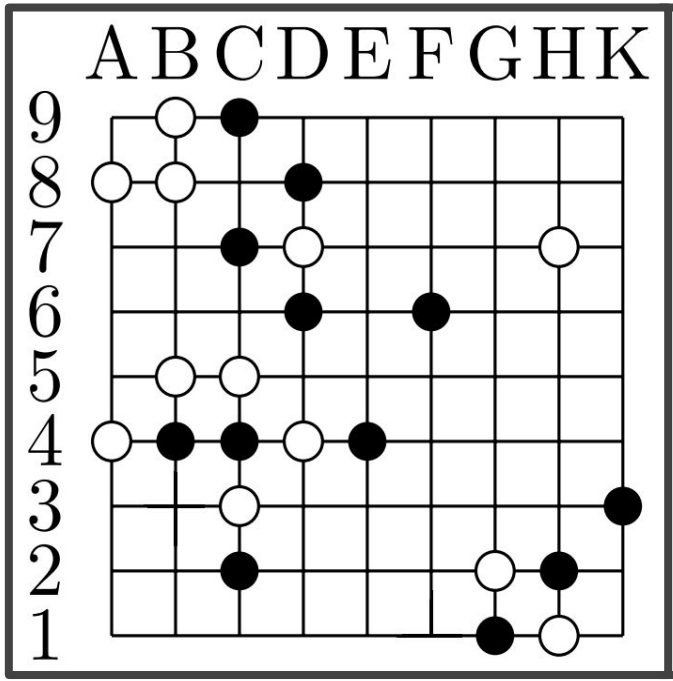

$$\begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \end{bmatrix}$$

Our Go Neural Networks: **Output**

Two types of possible outputs:

- **Policy net**:
  relative distribution of moves on this particular board (“heatmap”)
- **Value net:**
  (win, loss)-ratio given the current board and the player whose turn it is

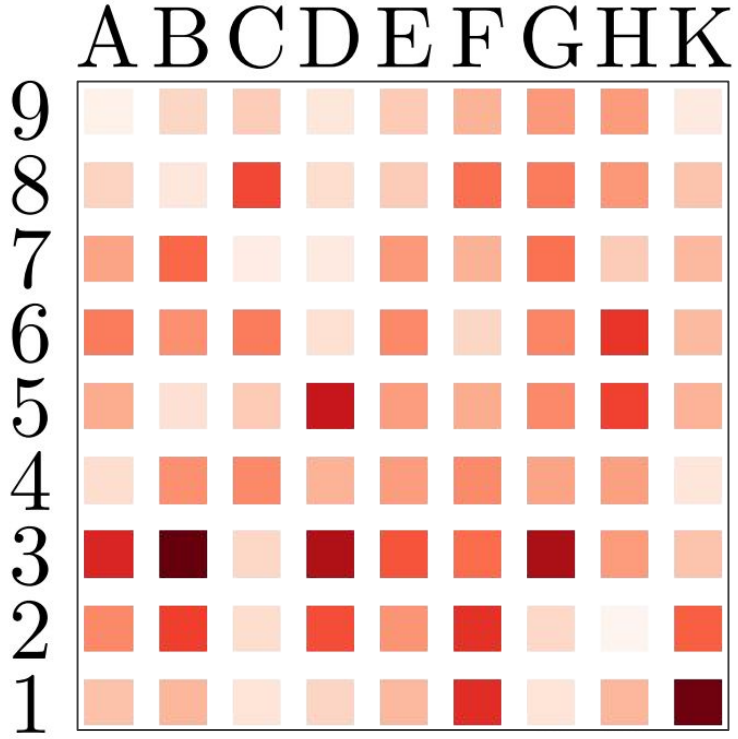

Results

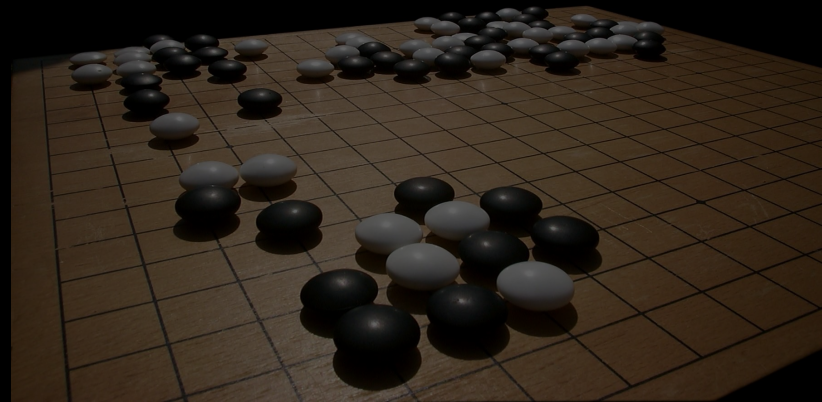

Performance

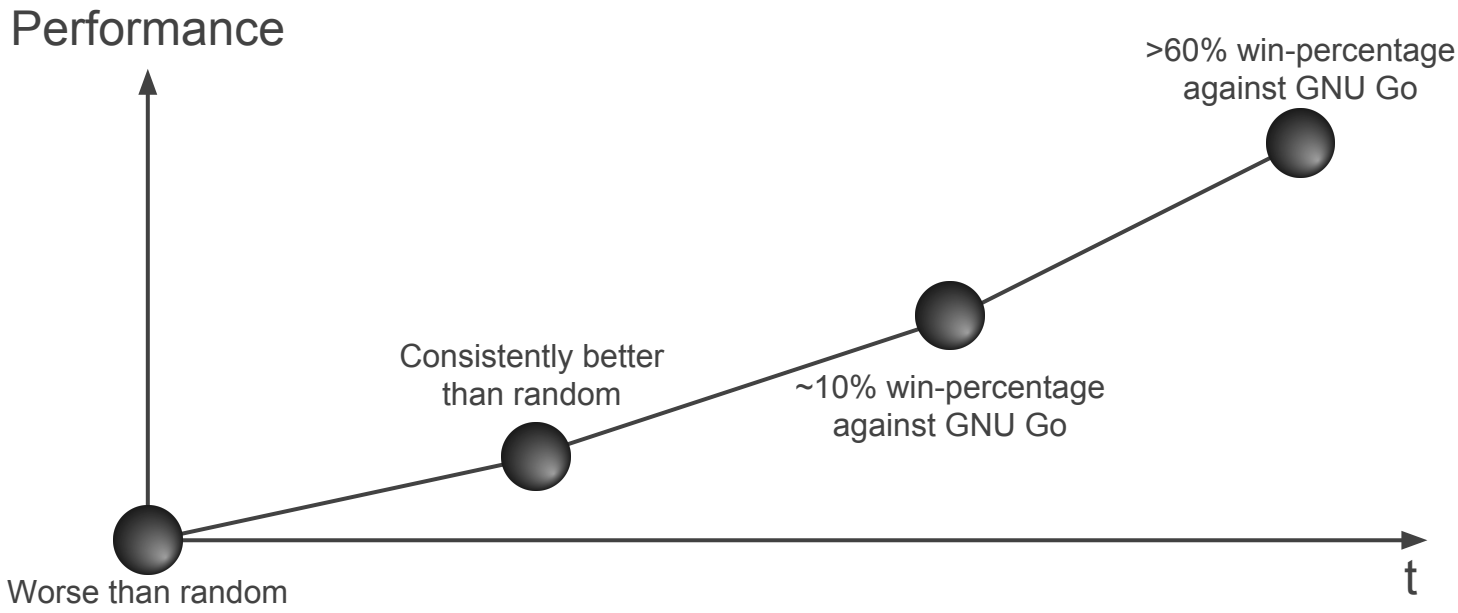

Performance

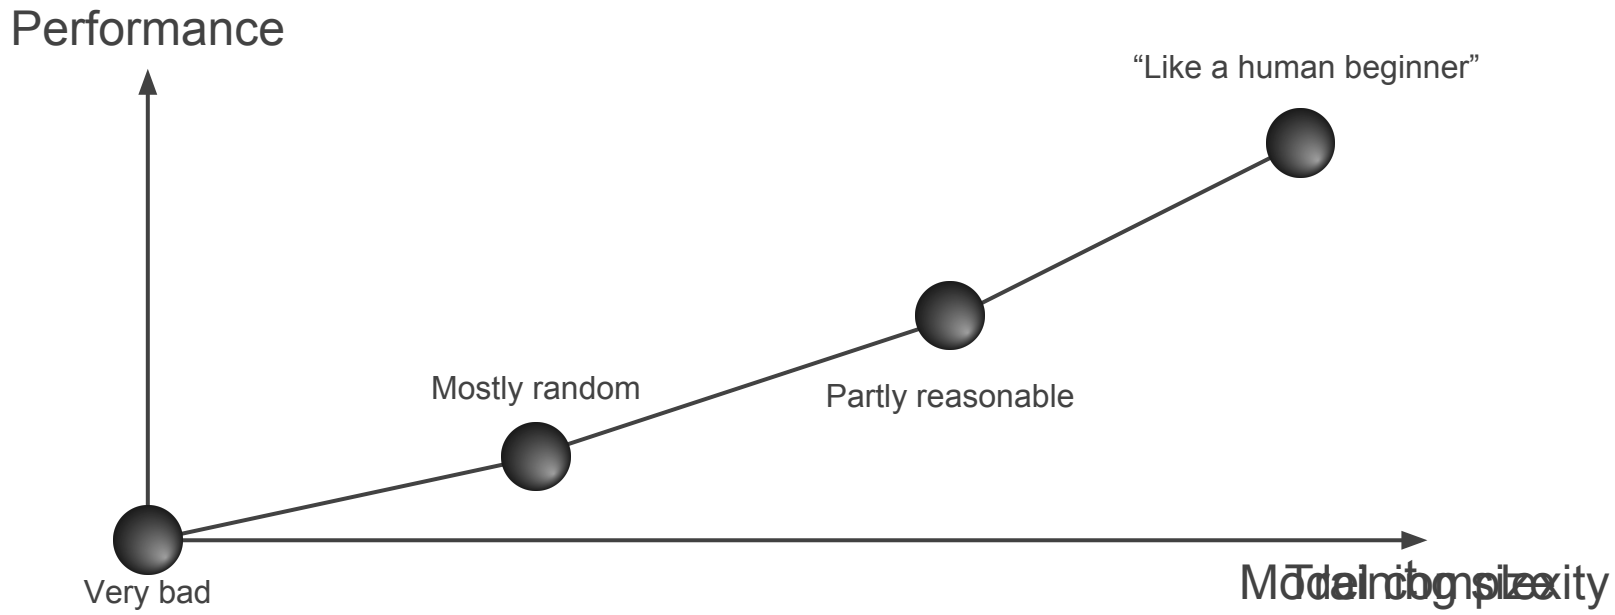

Insights: **Programming**

python™

K Keras

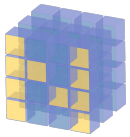

NumPy

theano

PYTORCH

git

- Python
- Implementing neural networks from scratch
- Deep learning libraries
- Software engineering
- Git

Insights: **Go & Computer-Go**

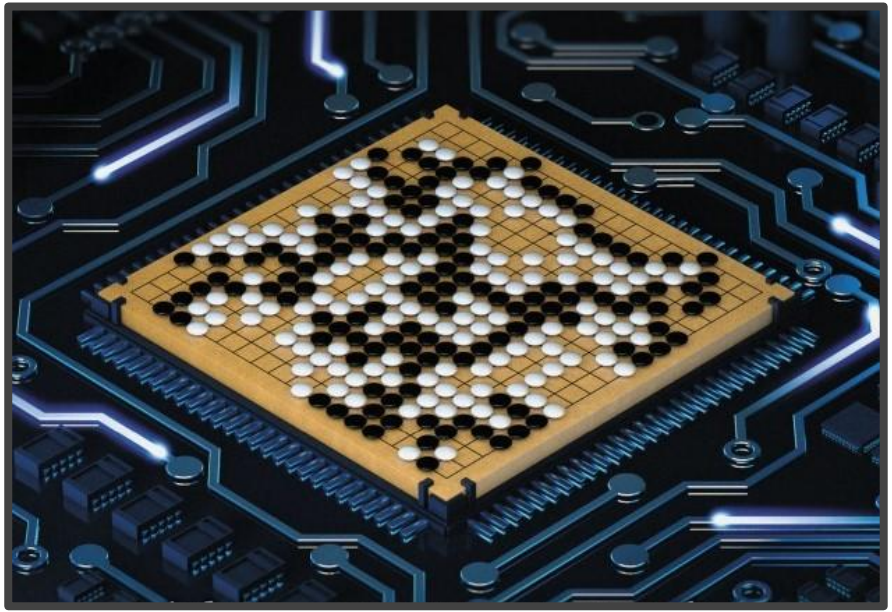

- Simple rules, but complex gameplay
- Challenging machine learning environment
- Rewarding project

Image: https://deepmind.com/research/alphago/

Learning from AI

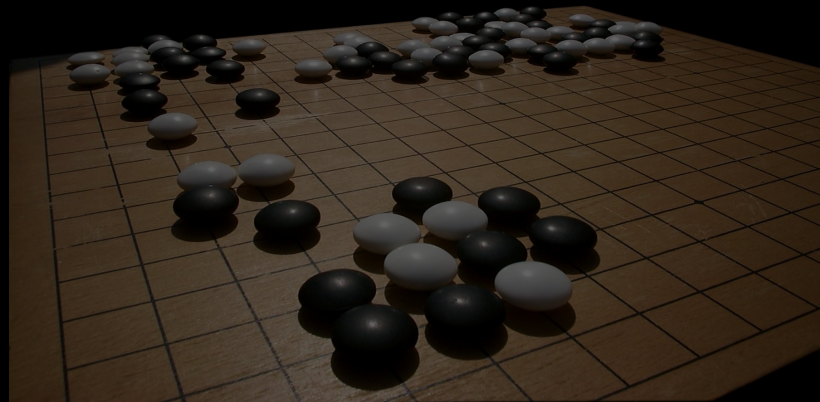

AlphaGo Zero

"WE'VE REMOVED THE CONSTRAINTS OF HUMAN KNOWLEDGE." - AlphaGo Zero's lead programmer, David Silver.

www.theverge.com/2017/10/18/16495548/deepmind-ai-go-alphago-zero-self-taught
www.theguardian.com/science/2017/oct/18/its-able-to-create-knowledge-itself-google-unveils-ai-learns-all-on-its-own

AlphaGo Zero: **From Random to Champion**

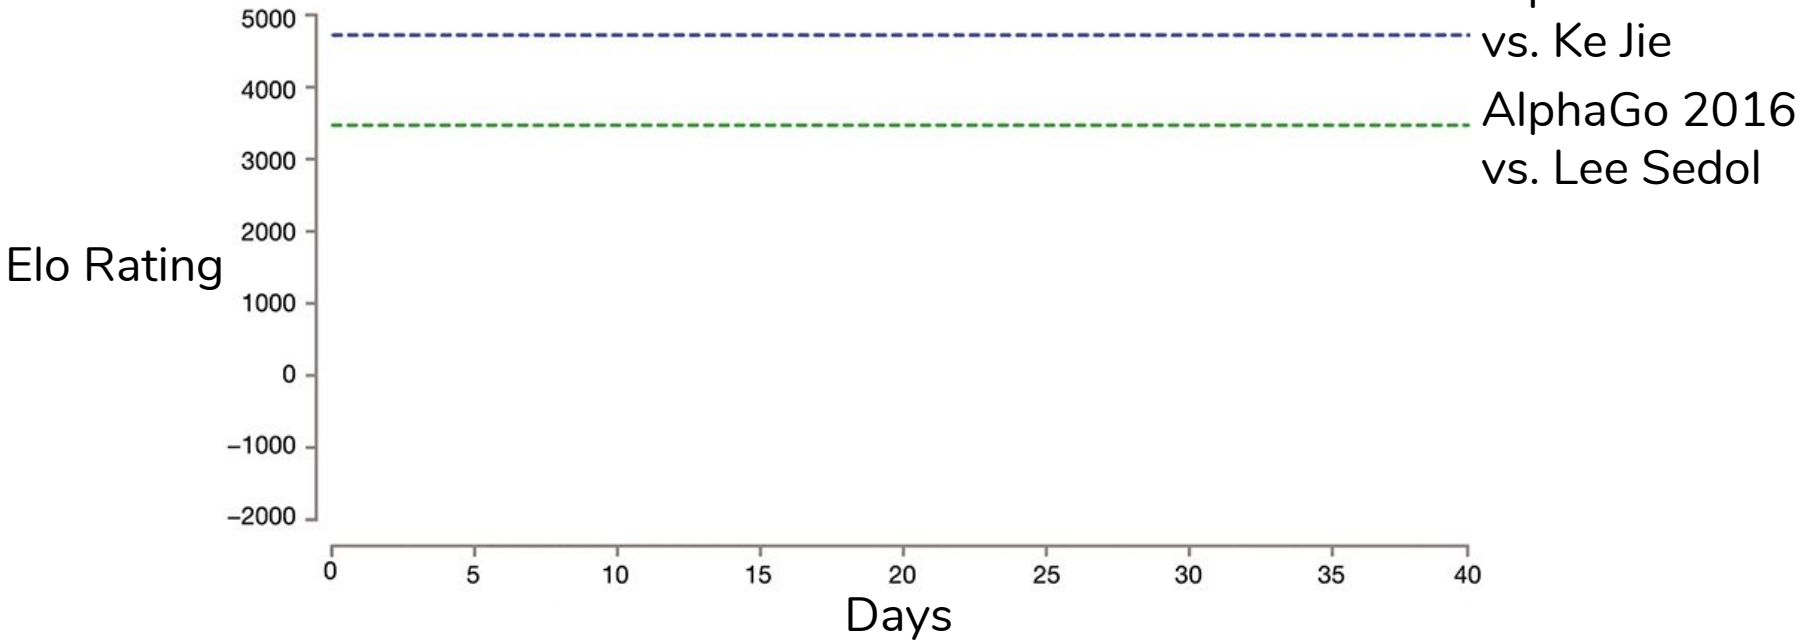

https://deepmind.com/blog/alphago-zero-learning-scratch/
https://en.chessbase.com/post/the-future-is-here-alphazero-learns-chess

Alpha Zero: **Redefining Chess**

MIT Technology Review

Alpha Zero's "Alien" Chess Shows the Power, and the Peculiarity, of AI

The Guardian International edition

AlphaZero AI beats champion chess program after teaching itself in four hours

THE TIMES

Google's DeepMind robot becomes world-beating chess grandmaster in four hours

www.thetimes.co.uk/article/google-s-deepmind-alphazero-becomes-world-beating-chess-grandmaster-in-four-hours-hcppp9vr2
www.technologyreview.com/s/609736/alpha-zeros-alien-chess-shows-the-power-and-the-peculiarity-of-ai/
www.theguardian.com/technology/2017/dec/07/alphazero-google-deepmind-ai-beats-champion-program-teaching-itself-to-play-four-hours

Creating Knowledge

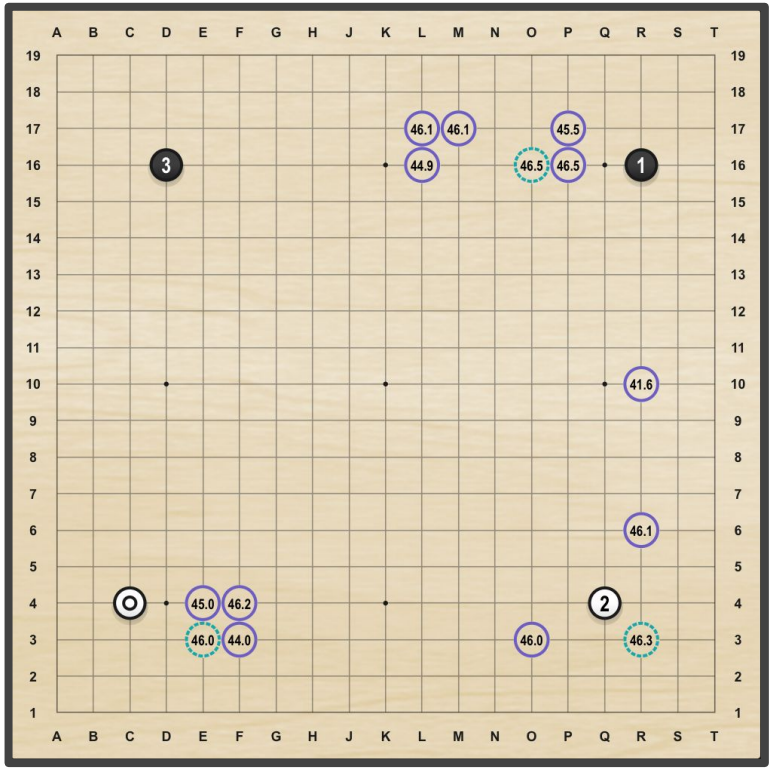

The Guardian International edition

'It's able to create knowledge itself'

KEY

Moves AlphaGo would play

Moves human players would play

https://alphagoteach.deepmind.com

Thank you: **xiè xie - arigatou - gamsahada**

A Machine Learning playing GO

TUM Data Innovation Lab

Sponsored by: Lehrstuhl für Geometrie und Visualisierung

Mentor: M.Sc. Bernhard Werner

TUM-DI-LAB Coordinator & Project Manager: Dr. Ricardo Acevedo Cabra

Head of TUM-DI-LAB: Prof. Dr. Massimo Fornasier

Nathanael Bosch, Benjamin Degenhart, Karthikeya Kaushik, Yushan Liu, Benedikt Mairhörmann, Stefan Peidli, Faruk Toy, Patrick Wilson