


The game of Go

Martin Müller, University of Alberta


Go

- ❖ Classic two-player board game
- ❖ Invented in China thousands of years ago
- ❖ Simple rules, complex strategy
- ❖ Played by millions, hundreds of professional players
- ❖ Until 2016: computers weaker than best humans






Go Rules

- ❖ Start with empty board
- ❖ Move: place stone of your own color
- ❖ Black goes first
- ❖ Surround empty points
- ❖ Capture opponent stones






Blocks, Liberties and Capture

- ❖ *Block*: Connected stones of the same color
- ❖ *Liberty* of block: adjacent empty point
- ❖ *Capture*: occupy last liberty



Ko rules: Avoiding Infinite Loops

- ❖ *Loop*: repeat board position by capture and re-capture
- ❖ Go rules *forbid* full-board repetition
- ❖ In example: Black must play elsewhere




White captures:
OK

Black re-capture:
illegal repetition

End of Game

- ❖ How to win: control more than half the board
- ❖ Count stones + territory
- ❖ Komi: compensate for first player advantage
- ❖ Example: AlphaGo trained with 7.5 komi for White



Final score,
9x9 board

Search - Game Tree Search

- ❖ All possible move sequences
- ❖ Combined in a tree structure
- ❖ Root is the current game position
- ❖ Leaf node is end of game
- ❖ Search used to find good move sequences
- ❖ Minimax principle




Image Source:
<http://web.emn.fr>

Go: Massive Size of Game Tree




Image Source: DeepMind

AlphaGo



- ❖ Program by DeepMind
- ❖ Based in London, UK and **Edmonton (2017-23)**
- ❖ Expertise in **Reinforcement Learning** and **search**
- ❖ 2014-17: worked on Go program for about 3 years, at first in secret


AlphaGo Matches

- ❖ Fall 2015 - **beat pro player**, European champion Fan Hui by 5:0 (kept secret)
- ❖ January 2016 article in Nature, announced win vs Fan Hui
- ❖ March 2016 **match vs top player Lee Sedol**, AlphaGo wins 4:1
- ❖ January 2017, wins fast games online 60:0 against many top players
- ❖ May 2017 match vs **world #1 Ke Jie** Wins 3:0 then retires




AlphaGo Zero

- ❖ **AlphaGo Zero:**
October 2017 article in Nature
“Mastering the game of Go without human knowledge”
- ❖ New simplified architecture
- ❖ **Learn entirely from self play**
- ❖ Only human knowledge: rules of game
- ❖ **Stronger** than previous AlphaGo



AlphaZero

- ❖ 2017 preprint, 2018 final version in Science
- ❖ “A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play”
- ❖ Further generalized, simplified architecture
- ❖ Wins test matches against top programs in all three games - chess, shogi, Go



Is the Game of Go Solved Now?

- ❖ **No!**
- ❖ AlphaZero is incredibly strong but...
 - ❖ ... it is all based on learned heuristics
- ❖ AlphaZero still makes mistakes
- ❖ The largest Go board that is solved is 5x6

