Analyzing the Impact of Knowledge and Search in Monte Carlo Tree Search in Go

Farhad Haqiqat and Martin Müller
University of Alberta
Edmonton, Canada

Contents

Motivation

- MCTS works extremely well in Go
- Combined with strong knowledge it works even better
- Why?
  - Many empirical results
  - Little in-depth analysis and understanding
- Detailed experiments to study relation between knowledge and search in MCTS
- Most of our work is with “old-fashioned” programs, without deep networks

Goals of this Research

- Examine relation between knowledge and search in Go programs
    - How do these two impact each other?
- Evaluation tools:
    - move prediction in master games
    - play against another programs
    - How do these two relate to each other?
- Evaluate the impact of knowledge strength on performance
- How does longer and deeper search improve the strength of a MCTS program, in the presence of knowledge?

Knowledge in Go Engines

- Playout policies
- Simple Features
- Small and large-scale patterns
- Neural Networks

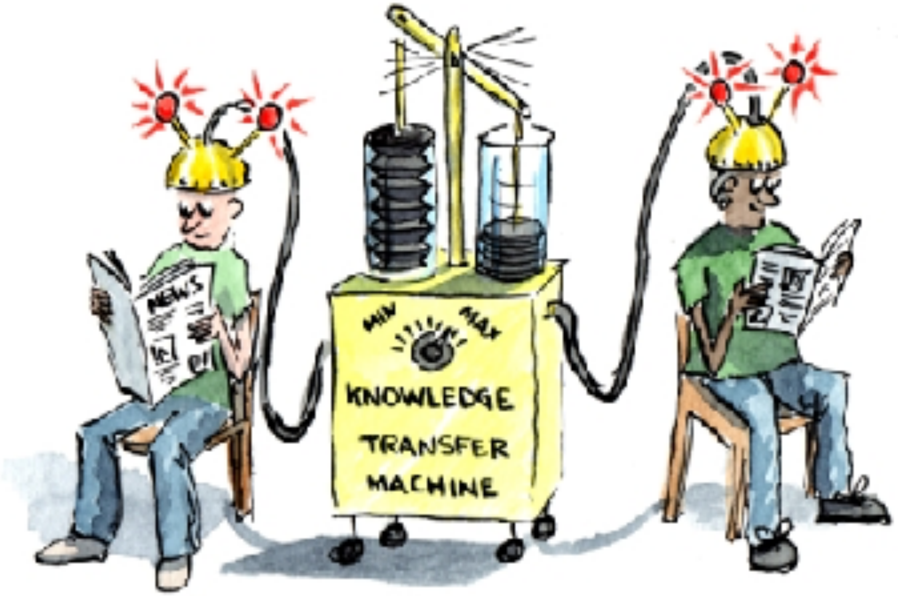

https://researchleap.com/wp-content/uploads/2016/08/image.png

Simple Features

Based on known properties of the game:

- Passing
- Distance to other stones
- Capturing
- Number of liberties
- Patterns

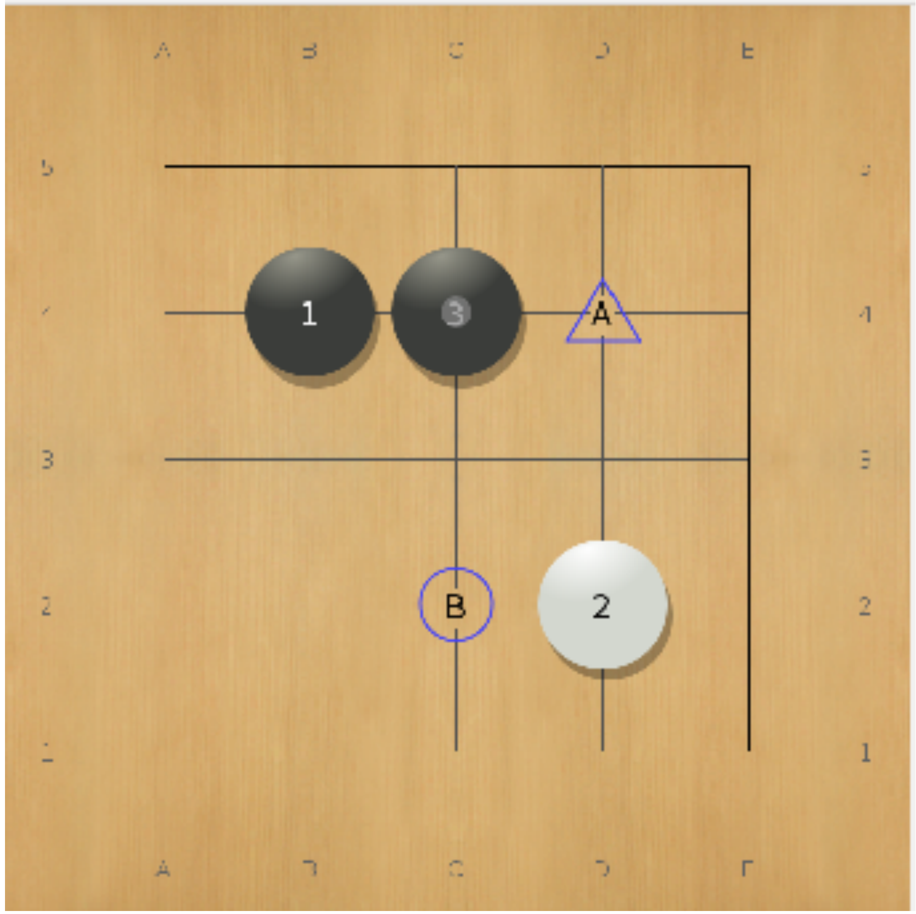

Patterns

- Square shape pattern
- Diamond shape pattern

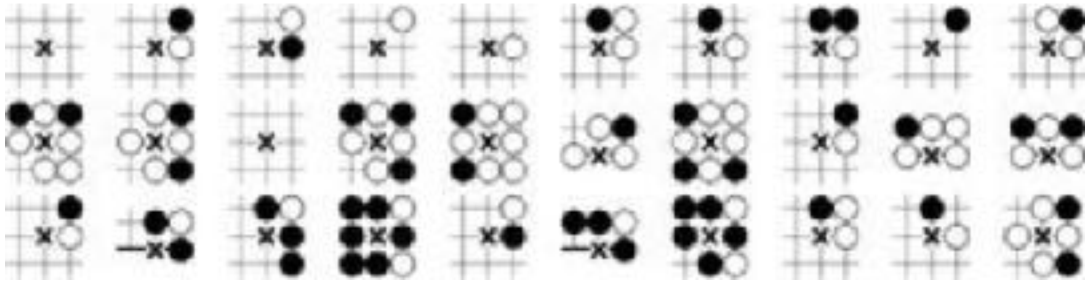

https://www.sciencedaily.com/images/2012/04/120416100437_1_540x360.jpg

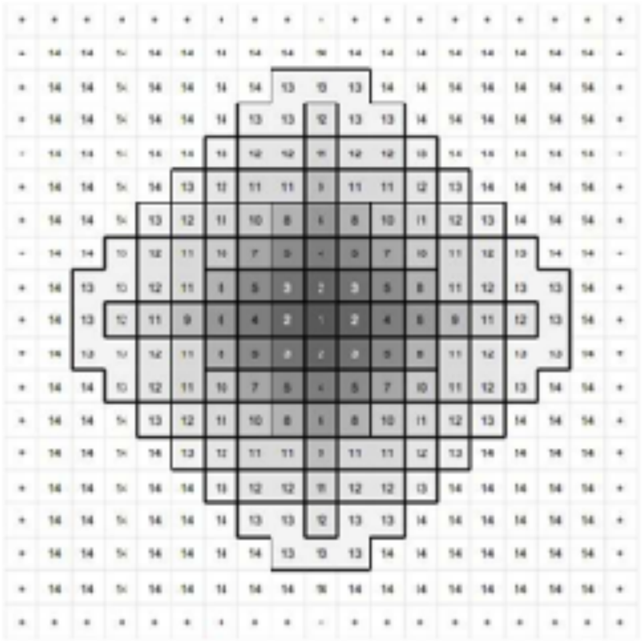

D. Stern, R. Herbrich, and T. Graepel. Bayesian pattern ranking for move prediction in the game of Go. In Proceedings of the 23rd international conference on Machine learning, pages 873–880. ACM, 2006

Monte Carlo Tree Search (MCTS)

- Selection
- Playout
- Expansion
- Update

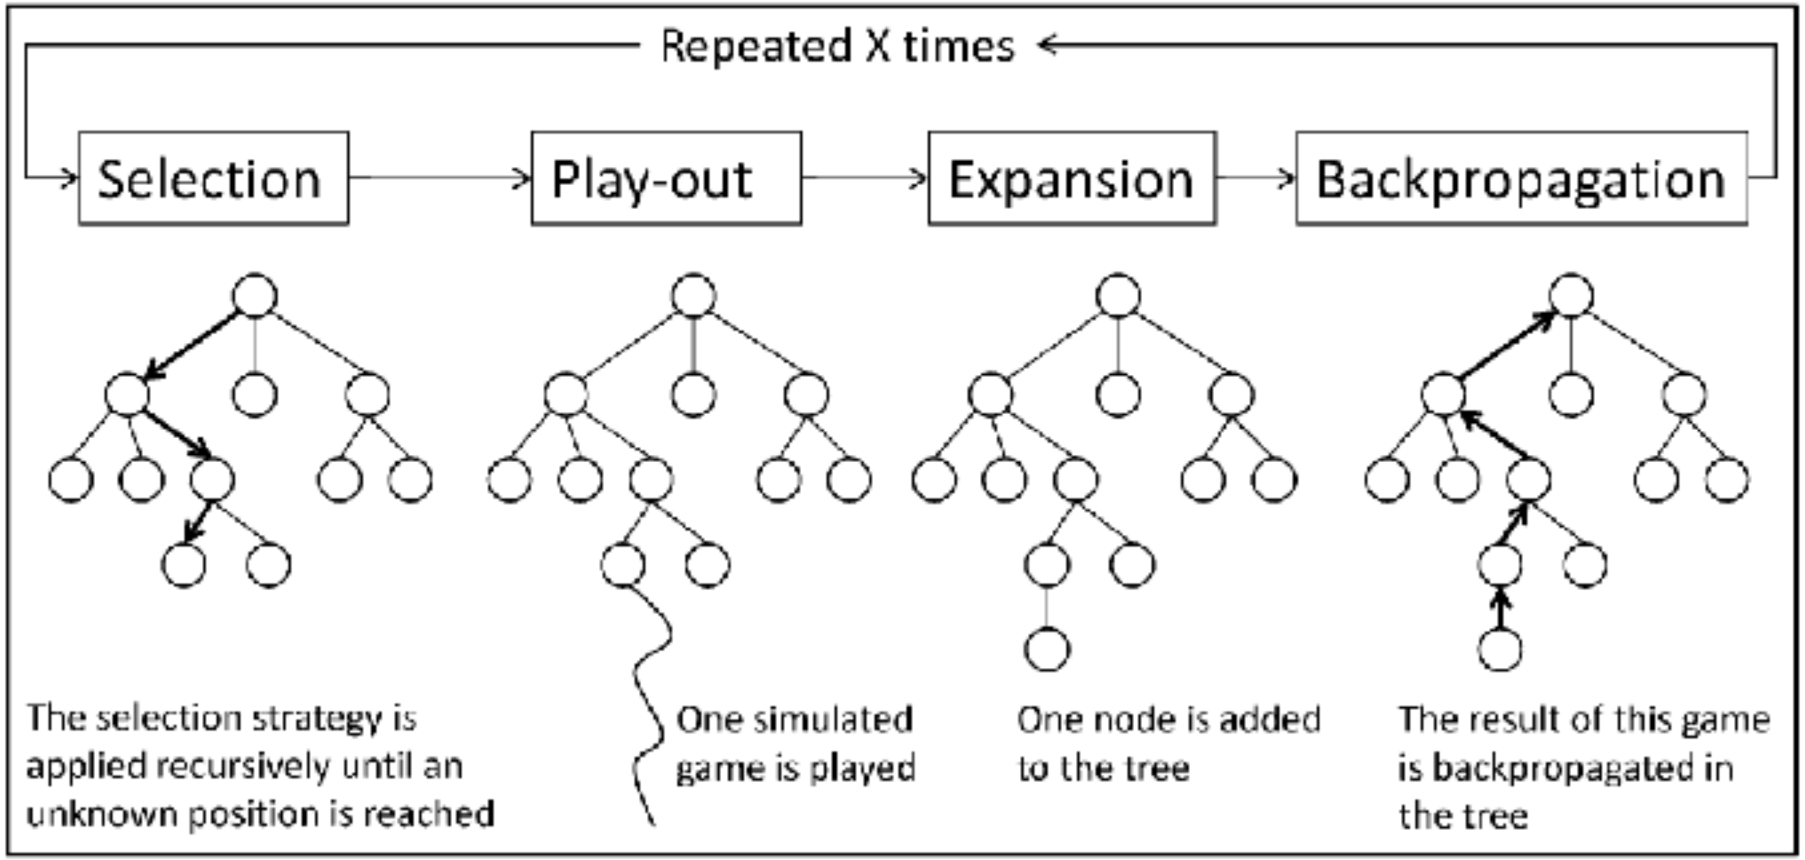

Types of Knowledge in Fuego

- Additive Knowledge
  - Diamond shape patterns
  - Evaluation term added to UCT formula
- Simple Features Knowledge
  - Initialization of nodes in search tree
  - Not scaled, can have negative values
- Small patterns
  - Used in playout policy

Fuego-Based Players

Players	Search Type
Playout policy-only	No search
Simple feature-only	No search
No Knowledge	MCTS
No Additive	MCTS
Default Fuego	MCTS

Evaluation Methods for Game Engines

- Move prediction
- Playing strength

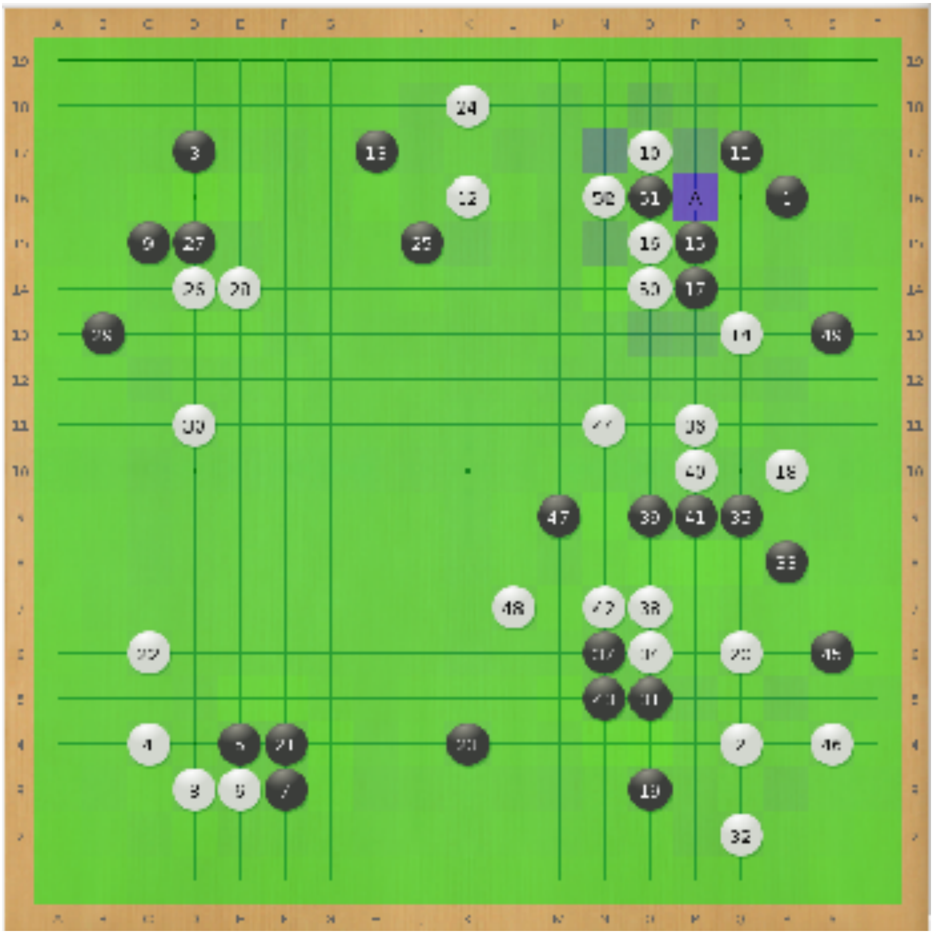

Move Prediction Task

- 4621 games played by professional players
  - All positions of all games
  - 19x19 board, no handicap

Baseline Experiments - Move Prediction

- Horizontal lines: prediction from feature knowledge much stronger than from playout policy
- Bottom diagonal line: More simulations help the no-knowledge MCTS
- Top two lines: *strange result*!
  deeper search hurts prediction for MCTS with knowledge
- Growing gap between No Additive and default player

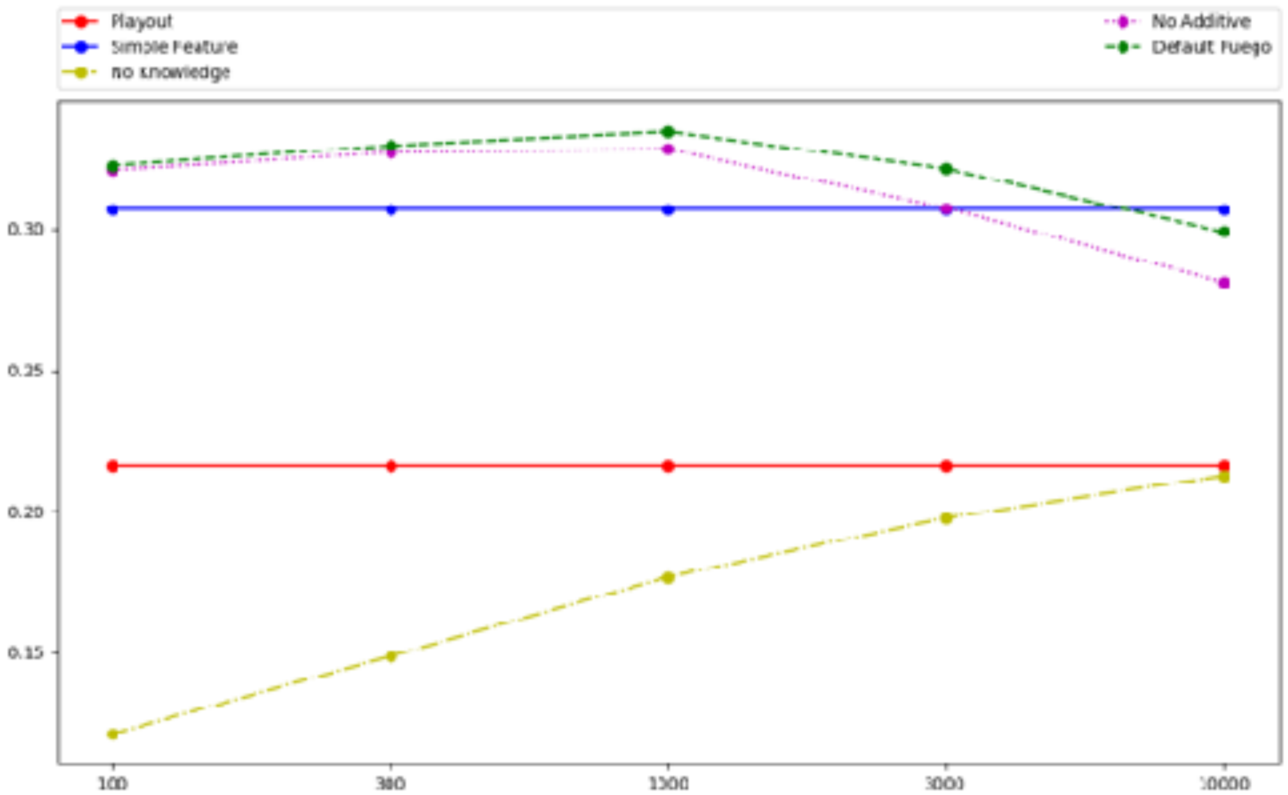

X-Axis: number of simulations in MCTS players
Y-Axis: prediction rate

Baseline Experiments - Playing Strength

- Green:
  Additive knowledge has minimal impact
- Orange & Blue:
  - Knowledge very significant, still increases with more simulations
  - More simulation let’s knowledge players inspect good moves deeply
- Red:
  No-knowledge search eventually beats no-search feature knowledge

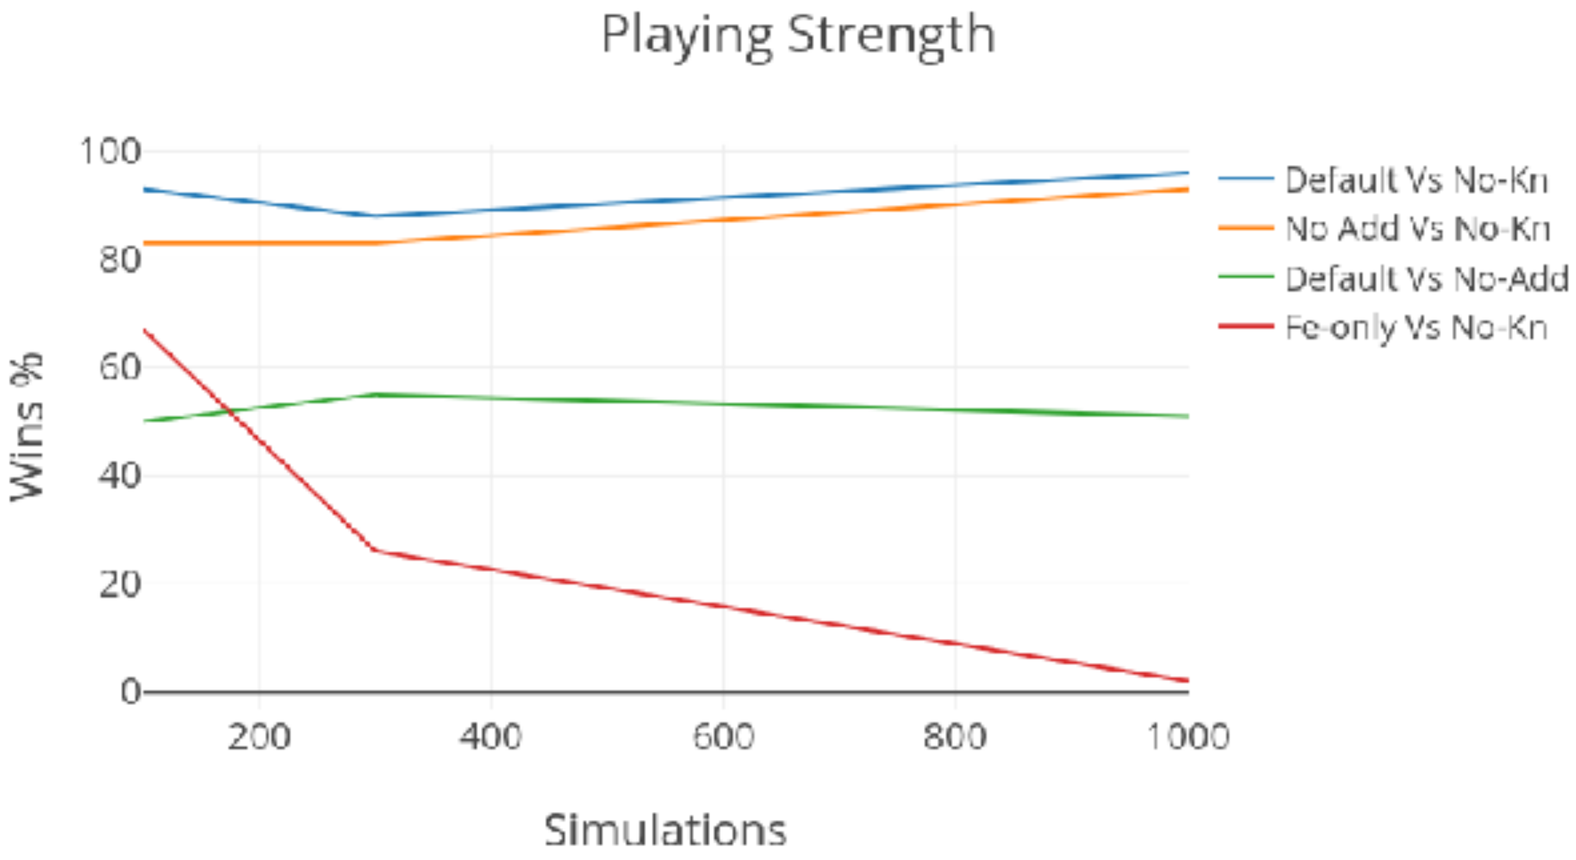

Experiments - Playing Strength

- Same player with more simulations almost always wins

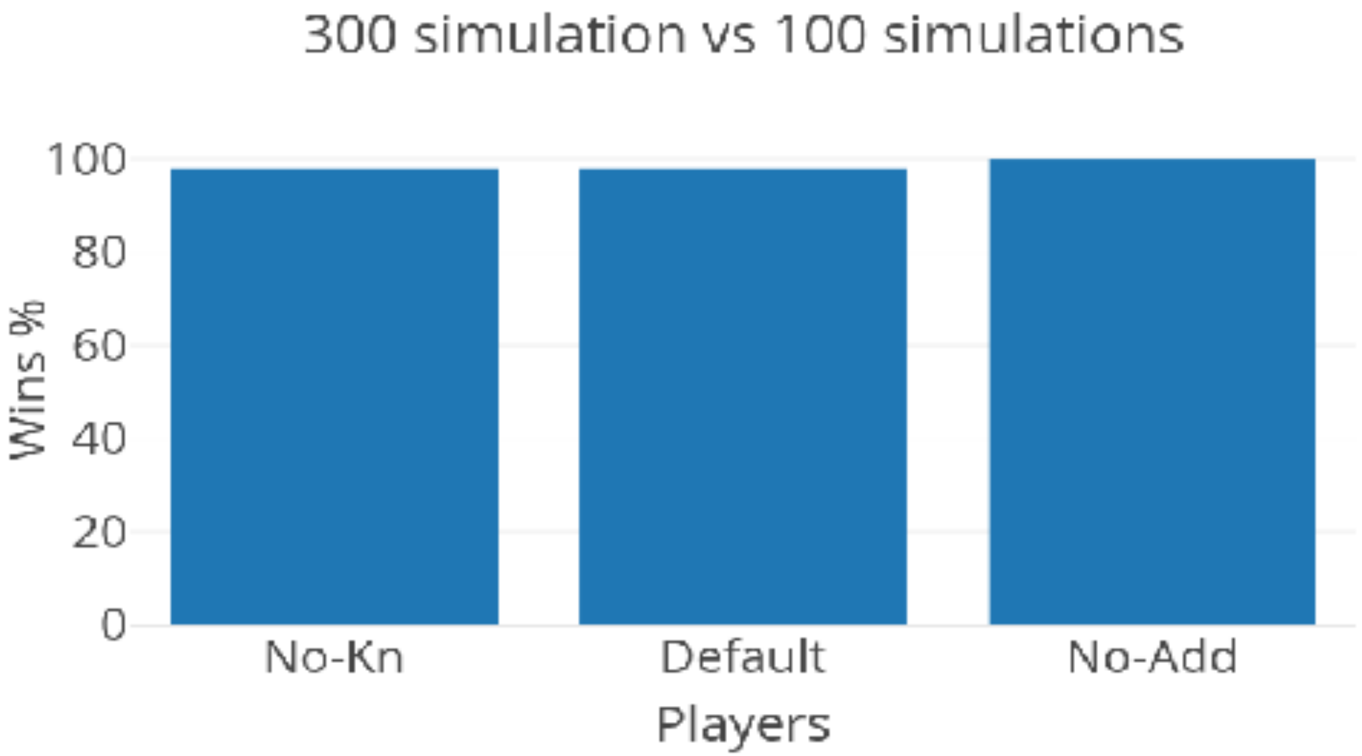

Explaining Strange Move Prediction Results

- Why does more search not help move prediction rate of knowledge-based players?
- Approach:
  - Divide games into 6 phases
  - Ignored very late endgame, moves 300+, due to limited sample size

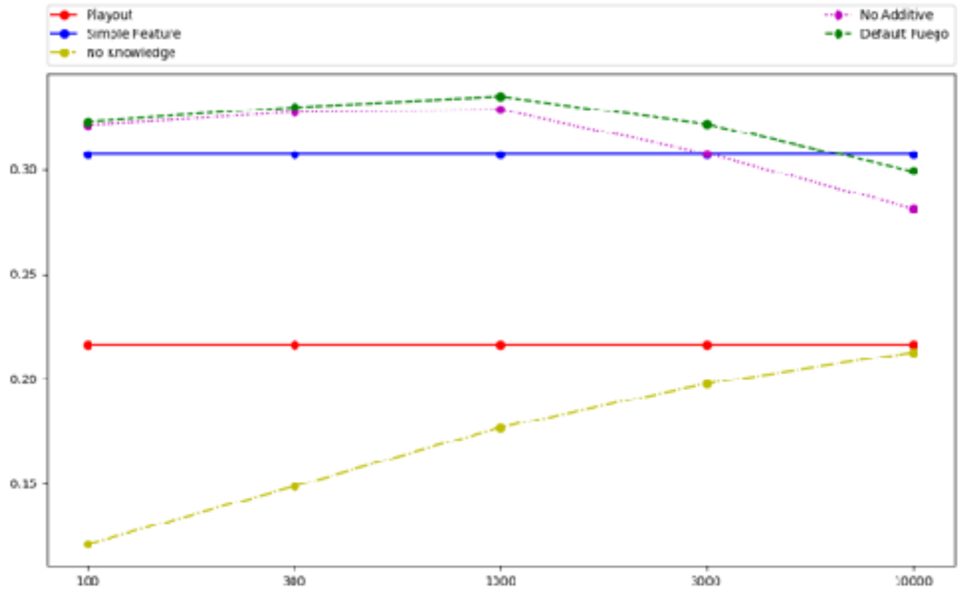

Move Prediction Rate with 100 and 1000 Simulations

- Default Fuego:
  Blue (100 sim) vs Red (1000 sim)
- In early game phases, more search helps prediction
- In endgame it reverses
- Reason:
- Fuego maximizes winning probability, not score
- Professional players don’t like to lose points in endgame

- No-knowledge player (beige vs grey):
- does not reverse
- search benefit is largest in middle game

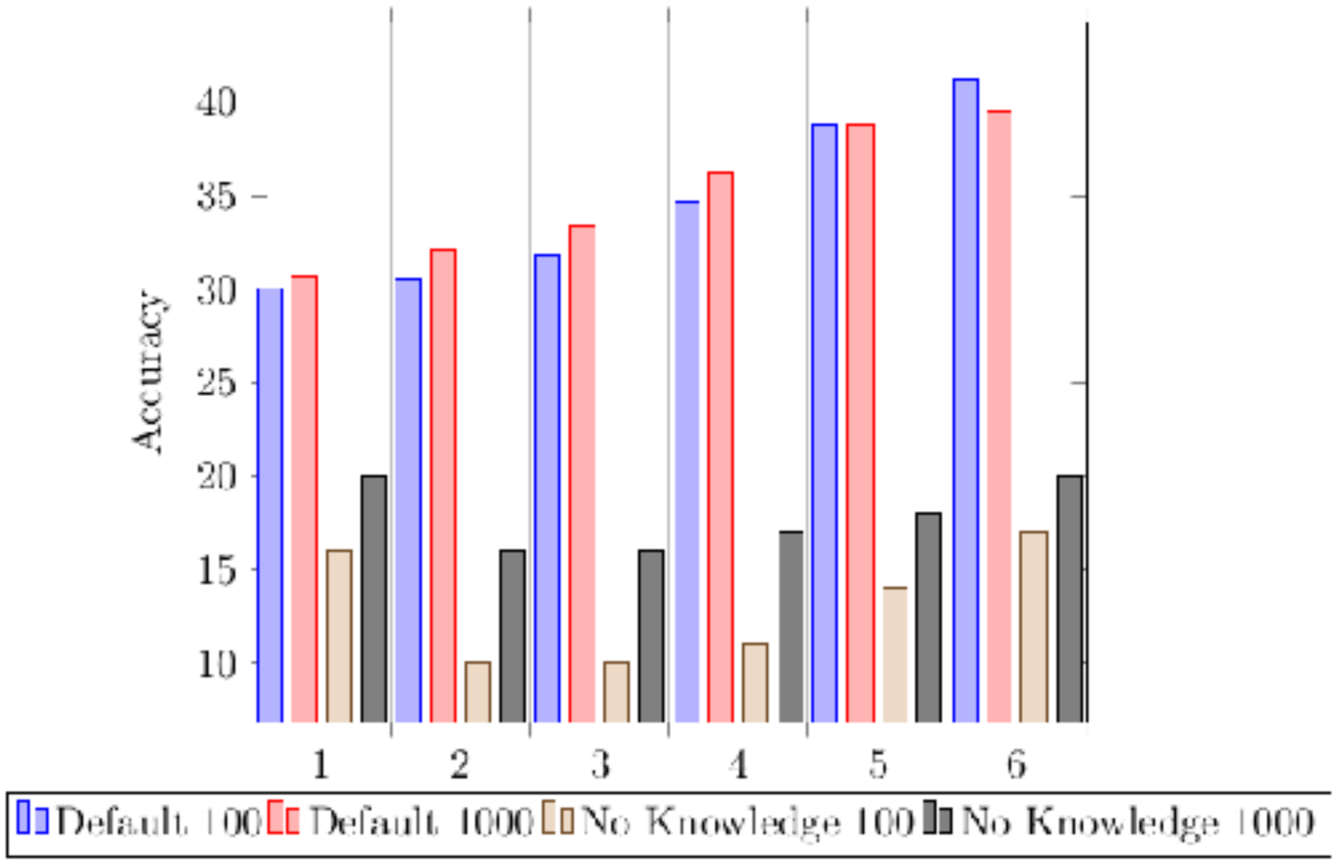

Analyzing Feature Frequencies

- Study moves by different players in terms of their *simple features*
- Express the difference between players in these terms
- Frequency: count features present for each move chosen by a player

Master Move Features

- Understand the types of moves professionals play, and the differences to the programs
- Compare:
    - All moves played by professional players
    - Moves by professional players than have less than 1% of total simulations

Experiments - Feature Frequency of Master Moves

- Highlights only here, more in the paper

- Professional players play close to last move of own or opponent - more than 80% of the time

- “Tenuki” moves by professional players are not found by Fuego

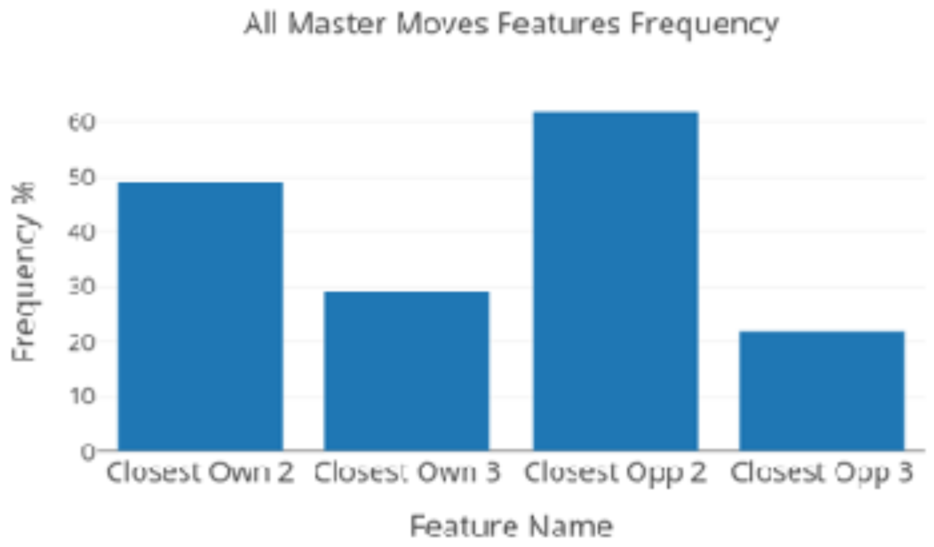

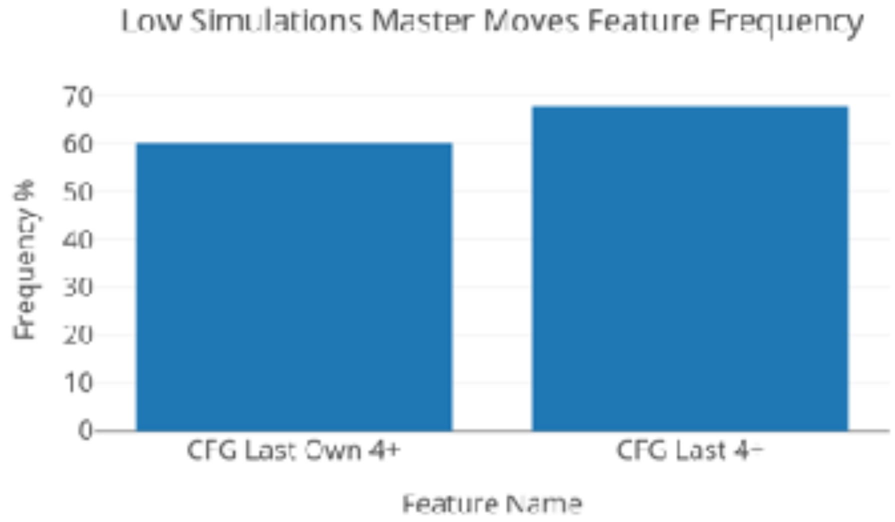

Are programs significantly different in which Master Moves they predict?

- Features of professional moves predicted:
  - correctly by player A
  - not predicted by player B
- Are there types of moves that one player misses systematically?
- Short answer: no.

Impact of Additive Term

- Compare feature frequencies:
  - All moves by default Fuego
  - All moves by No Additive player
  - Both with 3000 simulations

Experiments - Impact of Additive Term

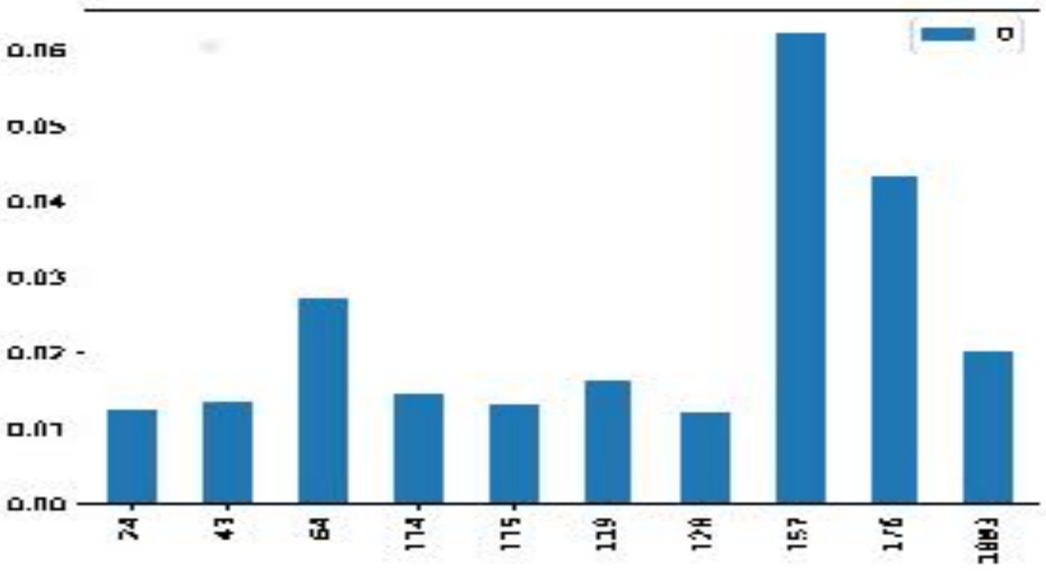

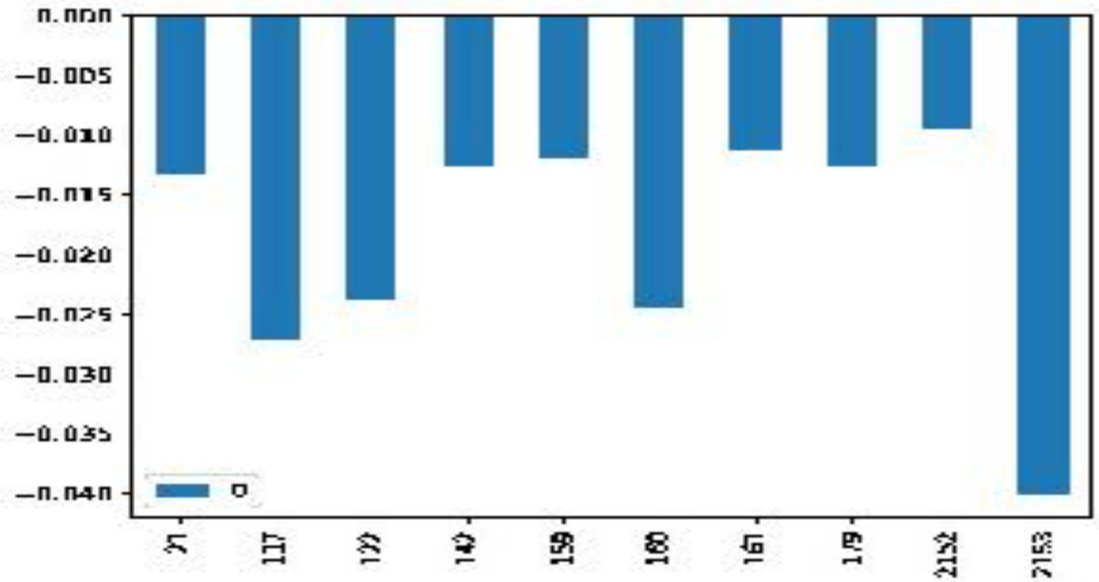

64: PLAYOUT POLICY 3X3 PATTERN
157: DIST CLOSEST OWN STONE 2
176: DIST CLOSEST OPP STONE 2

117: CFG DISTANCE LAST 4+
122: CFG DISTANCE LAST OWN 4+
160: DIST CLOSEST OWN STONE 5
2153: EMPTY 3X3 PATTERN

- Additive knowledge encourages playing close to previous stones
- The No Additive player plays more often in empty areas of the board (feature 2153, 3×3 empty pattern)

Impact of Knowledge

- Compare feature frequencies:
  - All moves by default Fuego
  - All moves by No Knowledge player
  - Both with 3000 simulations

Experiments - Impact of Knowledge: 3000 Simulations

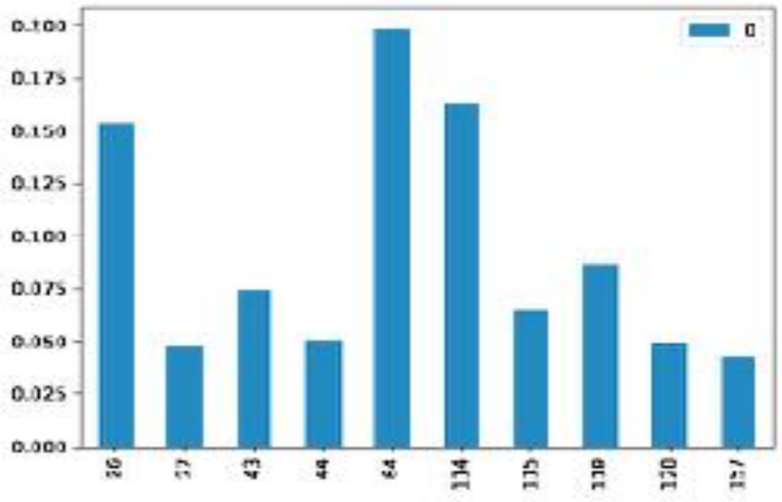

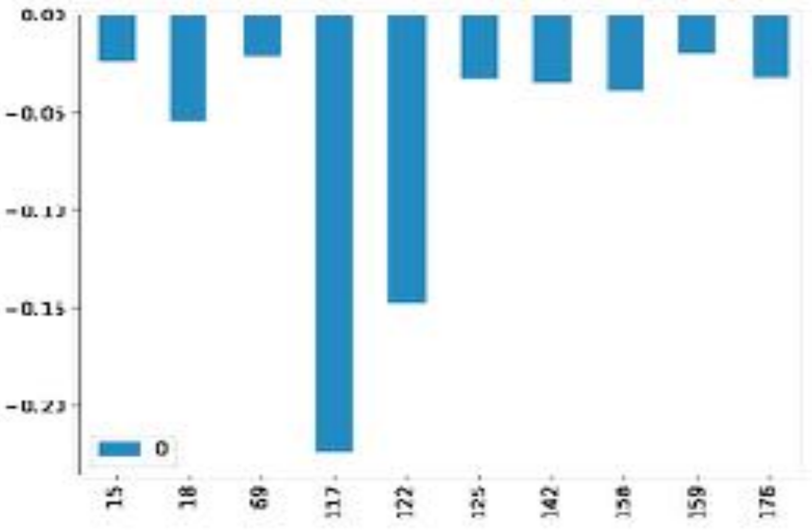

26: DIST PREV 2

64: PLAYOUT POLICY 3X3 PATTERN
114: CFG DISTANCE LAST 1

117: CFG DISTANCE LAST 4+
122: CFG DISTANCE LAST OWN 4+

- No Knowledge plays tenuki moves way more often
- Feature knowledge encourages local response to same moves

Why do Programs Ignore some Master Moves?

- Compare feature frequency
- Moves by default Fuego with 3000 simulations
- Moves by professionals
- Restrict to positions where professional move receives less than 1% of total number of simulations in Fuego

Professional Moves with Low Simulations

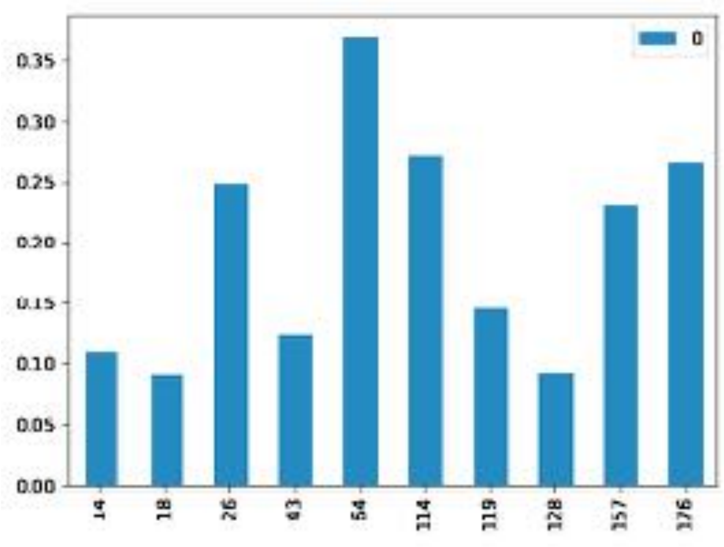

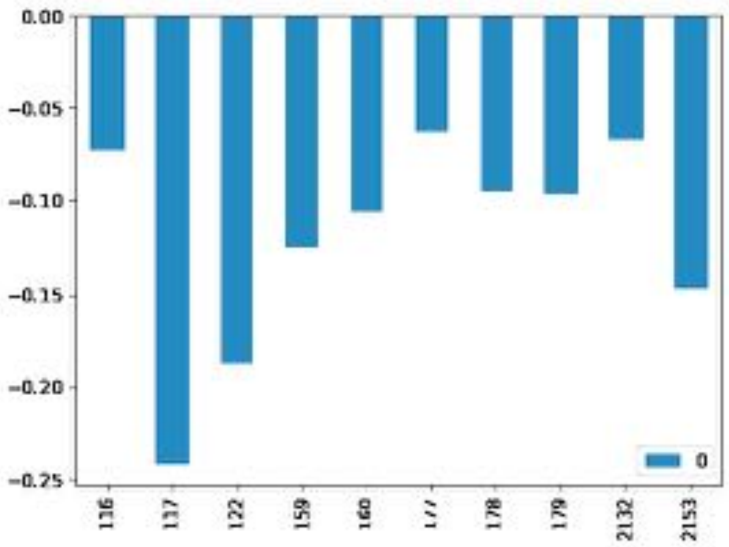

26: DIST PREV 2
64: PLAYOUT POLICY 3X3 PATTERN
114: CFG DISTANCE LAST 1
119: CFG DISTANCE LAST OWN 1
157: DIST CLOSEST OWN STONE 2
176: DIST CLOSEST OPP STONE 2

117: CFG DISTANCE LAST 4+
122: CFG DISTANCE LAST OWN 4+
159: DIST CLOSEST OWN STONE 4
160: DIST CLOSEST OWN STONE 5
178: DIST CLOSEST OPP STONE 4
179: DIST CLOSEST OPP STONE 5
2153: EMPTY 3X3 PATTERN

- Distance 1 or 2 up to 25% more often in Fuego
- Distance 4 or more up to 24% more in master moves

Move Selection Analysis

- Impact of knowledge initialization on number of simulations
- Initial weight of features on moves chosen by Fuego
- Initial weight of features on moves played by professionals
- Maximum weight in that position of the game
- Percent of simulations received by each move

Move Selection Analysis - Fuego Move

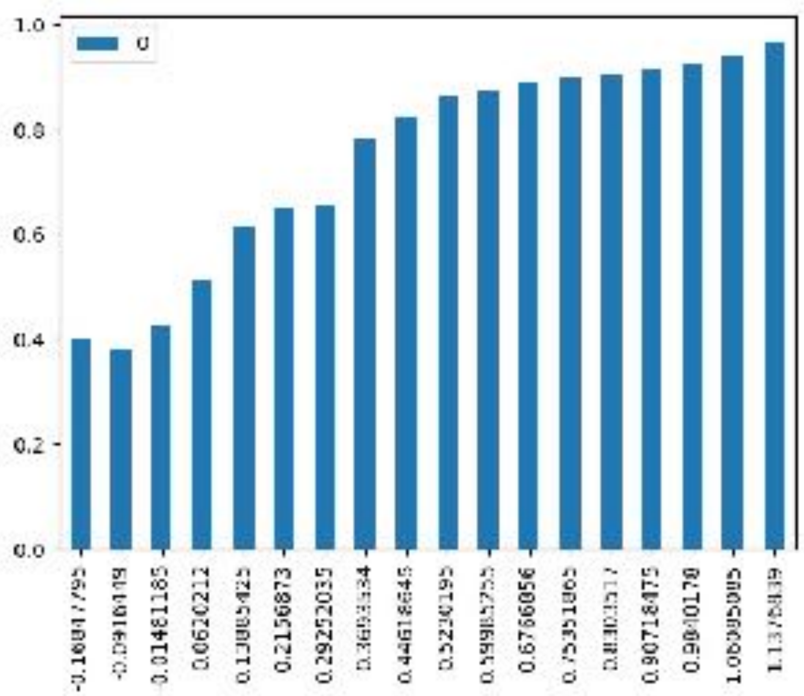

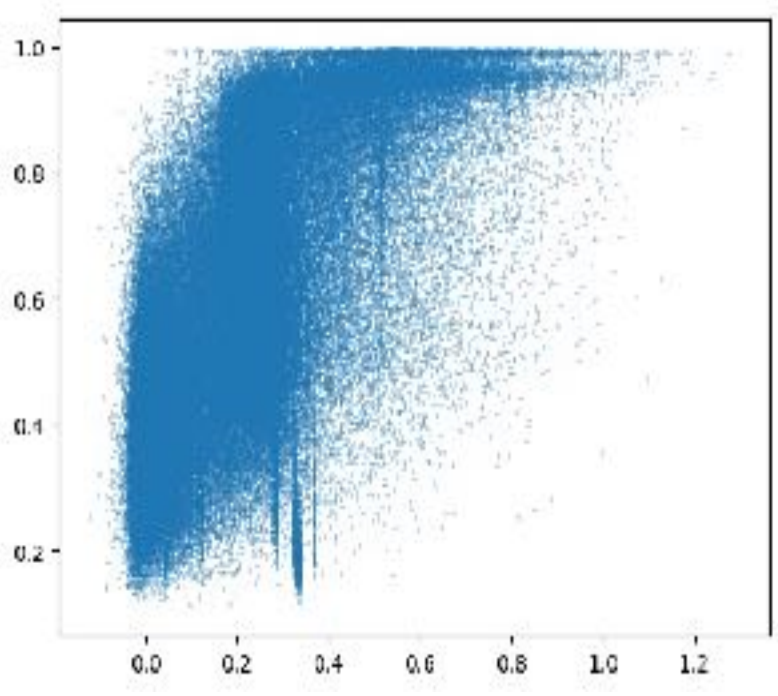

Initial Weight

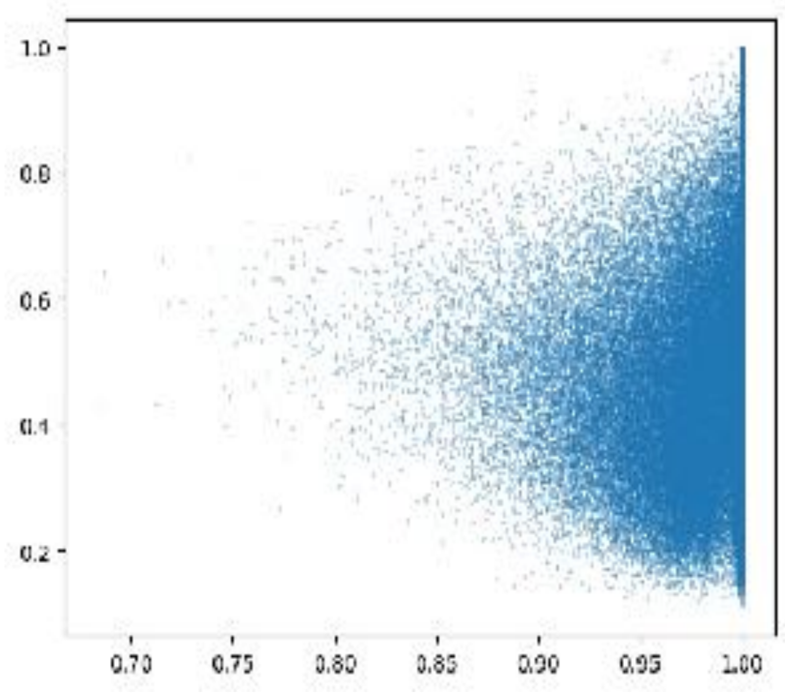

Sigmoid of initial weight to Sigmoid of max weight

- Most Fuego moves have weight very close to maximum
- Majority of all simulations assigned to them

Move Selection Analysis - Professional Move

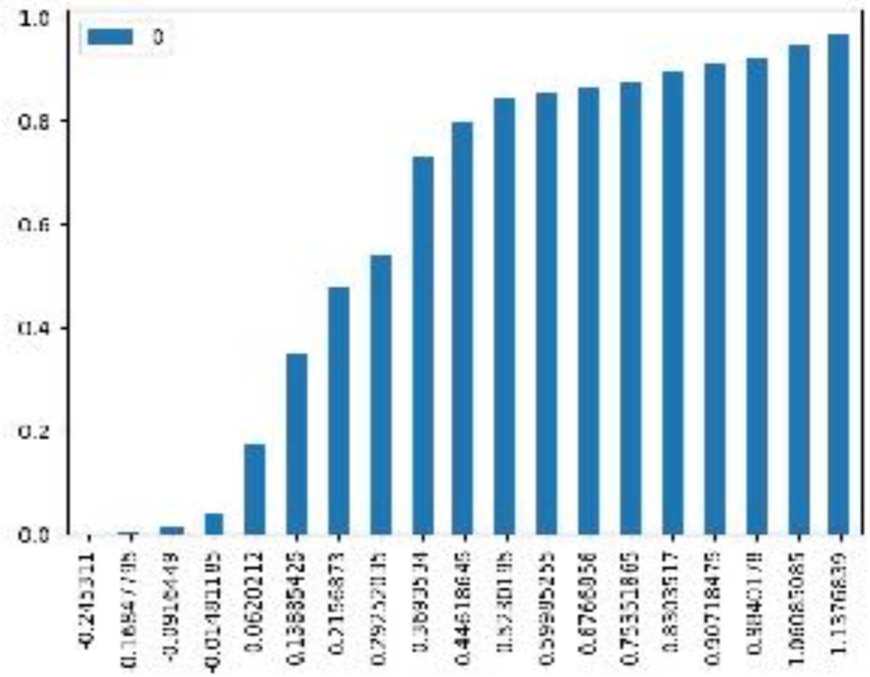

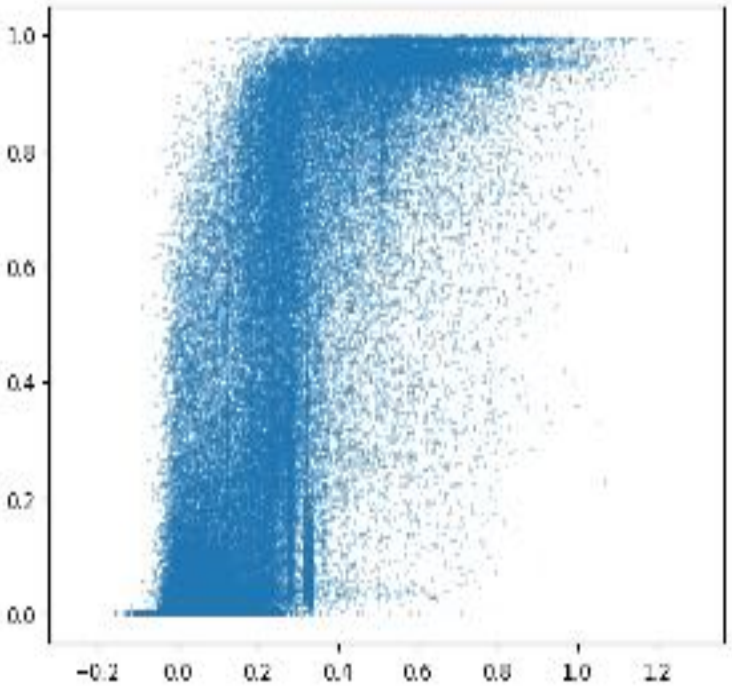

Initial Weight

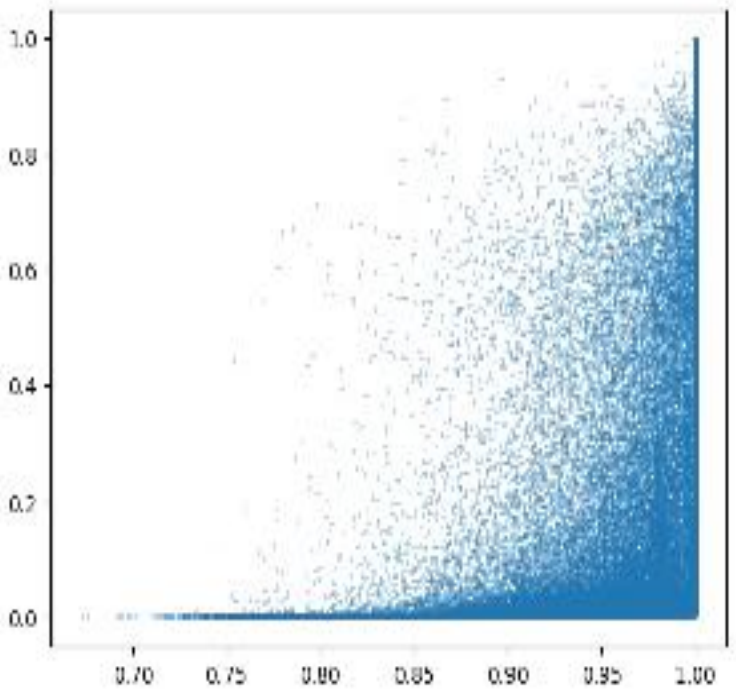

Sigmoid of initial weight to Sigmoid of max weight

- Professional players most of the time either get majority of simulation or nothing
- Higher evaluation needed for professional players move to receive majority of simulations

Professional Move vs Fuego Move

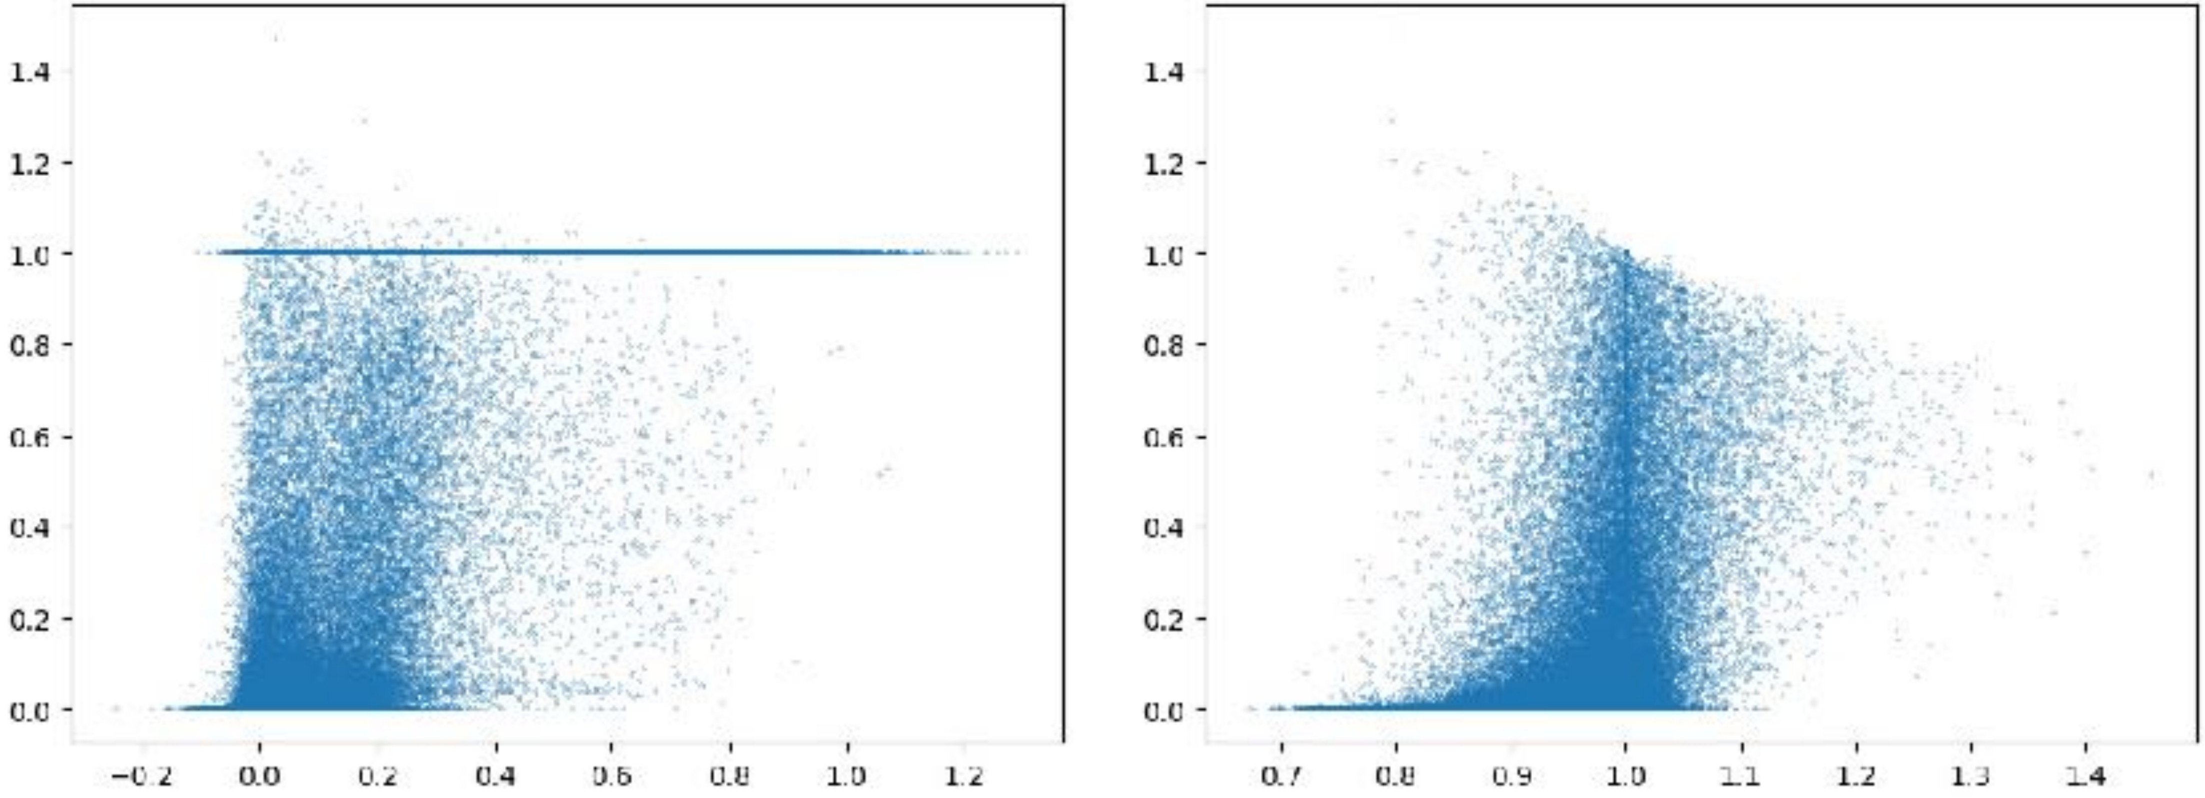

- Same as Fuego Move
- If they differ master moves most of the time has less than 20% of Fuego move simulations
- 7% of Professional moves have higher weight and less simulations

Extra: some Leela Zero Experiments

- Leela Zero
- Strongest open source program
- Super-human strength
- Re-implementation of AlphaGo Zero
- Super-strong knowledge in deep neural net trained by self-play

Leela Zero - Move Prediction Rate per Game Phase

- Deep nets have much higher prediction rate than simple features (about 50% vs 35-40%)
- Small amounts of search boost prediction rate, then it drops with more search, even below “raw net” rate
- Does Leela Zero find better moves than human masters?
- Steady increase from opening to endgame
    - Why?

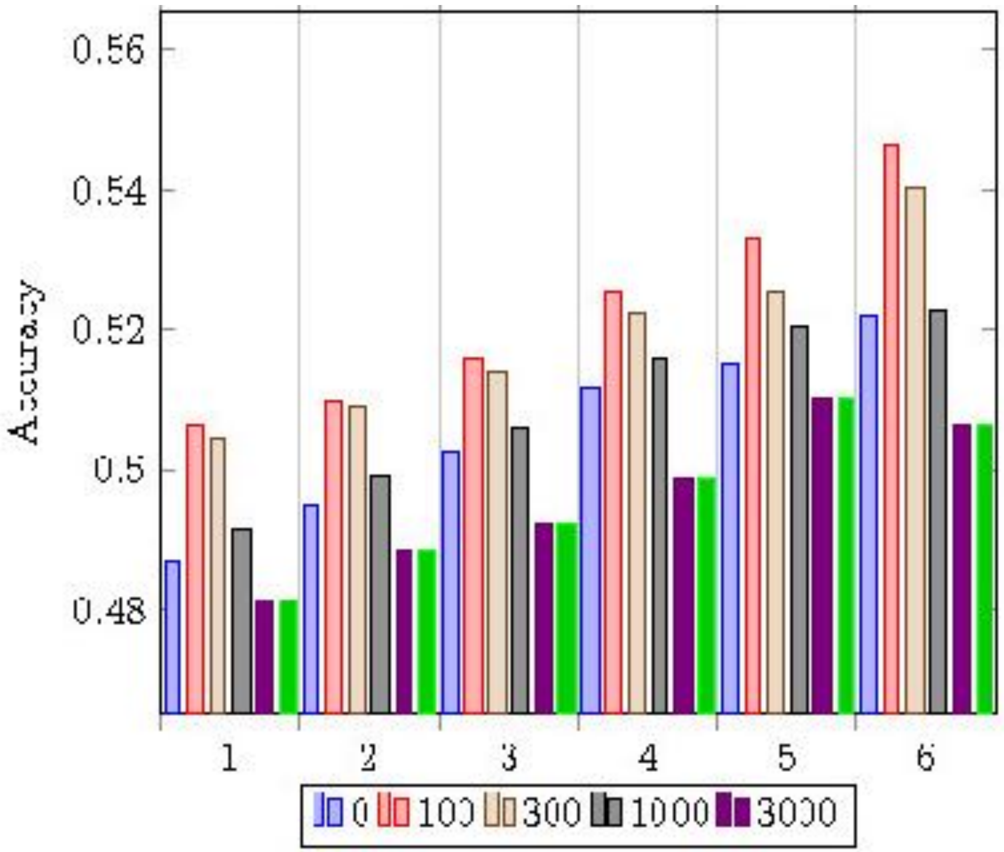

Figure 1: Move prediction accuracy per game phase with varying simulation number in LeelaZ. Each group has 50 moves.

Leela Zero - Feature Frequency of Master Moves

25: DIST PREV 1
64: PLAYOUT POLICY 3X3 PATTERN
114: CFG DISTANCE LAST 1
115: CFG DISTANCE LAST 2
117: CFG DISTANCE LAST 4+
122: CFG DISTANCE LAST OWN 4+
157: DIST CLOSEST OWN STONE 2
158: DIST CLOSEST OWN STONE 3
176: DIST CLOSEST OPP STONE 2
177: DIST CLOSEST OPP STONE 3

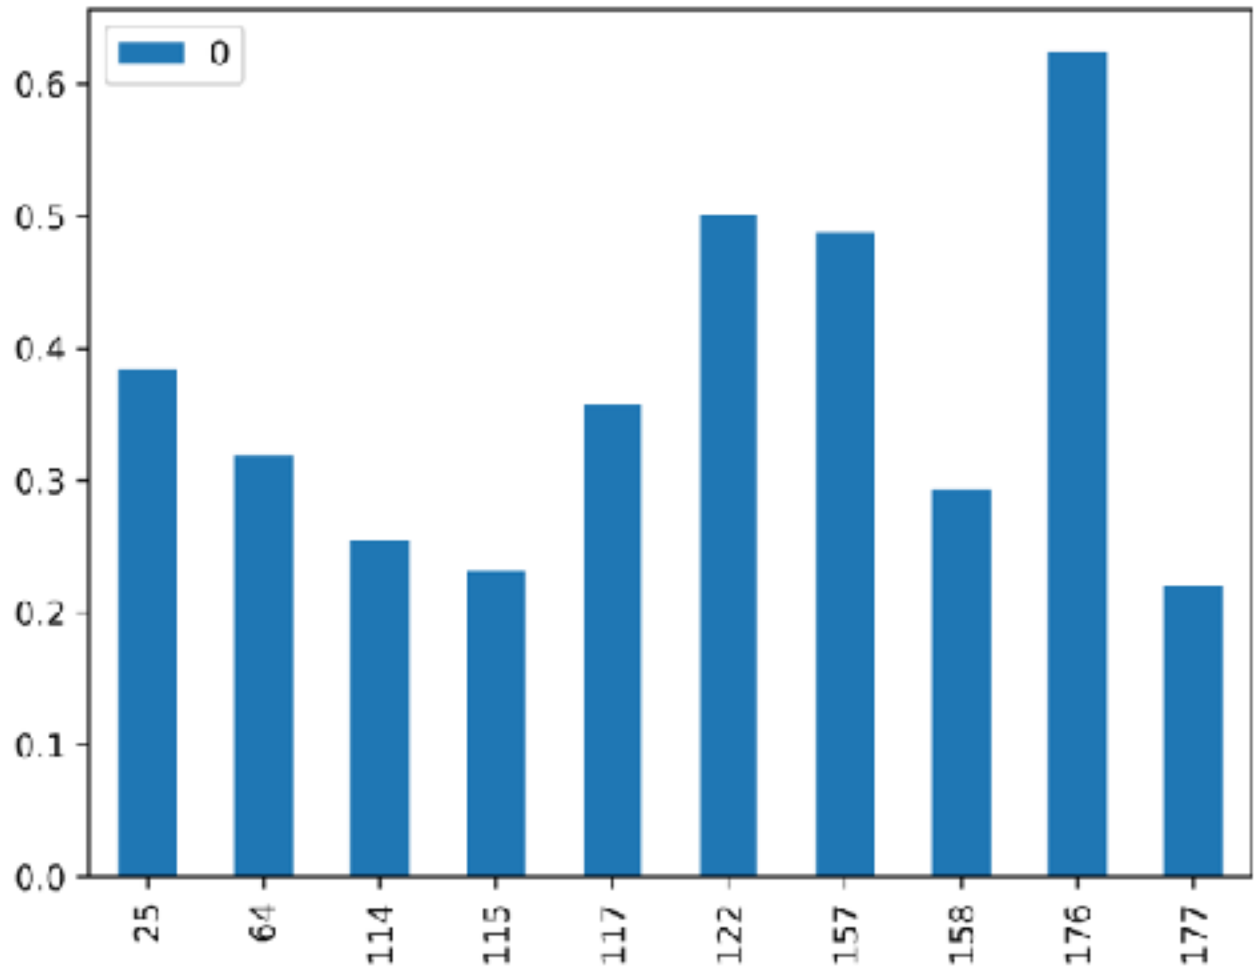

Feature Frequency of Leela Zero with 1000 Simulations

25: DIST PREV 1
64: PLAYOUT POLICY 3X3 PATTERN
114: CFG DISTANCE LAST 1
115: CFG DISTANCE LAST 2
117: CFG DISTANCE LAST 4+
122: CFG DISTANCE LAST OWN 4+
157: DIST CLOSEST OWN STONE 2
158: DIST CLOSEST OWN STONE 3
176: DIST CLOSEST OPP STONE 2
177: DIST CLOSEST OPP STONE 3

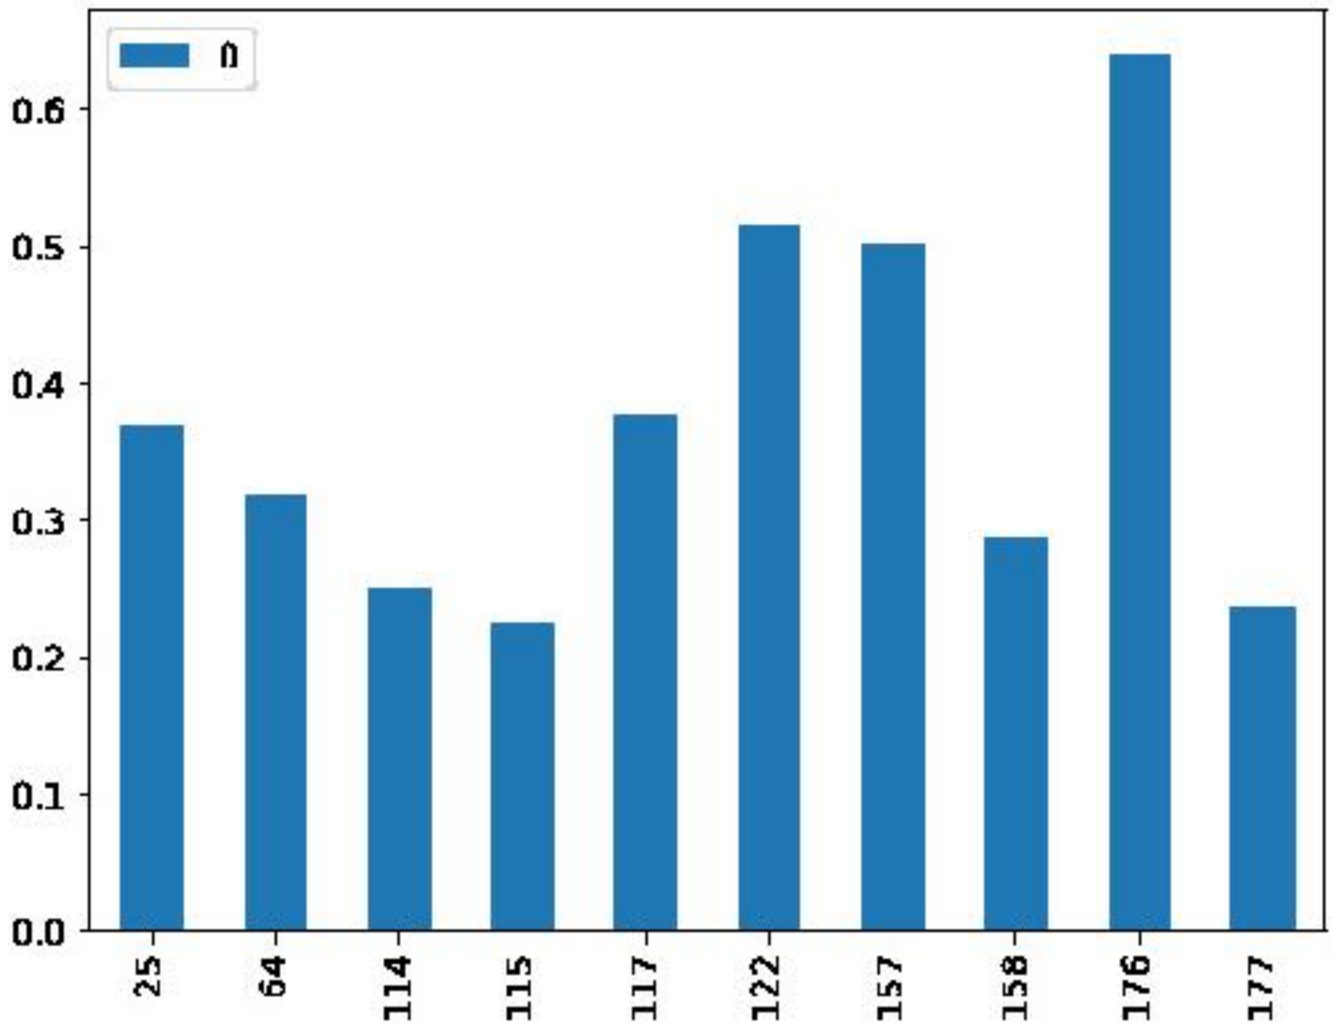

Frequency difference - Leela Zero (1000 sims) vs Human Master

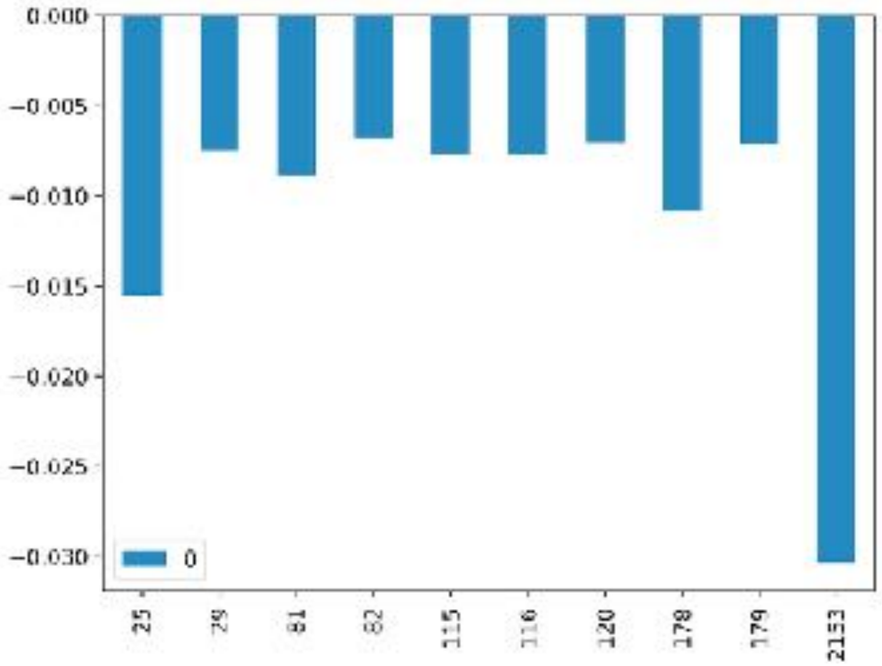

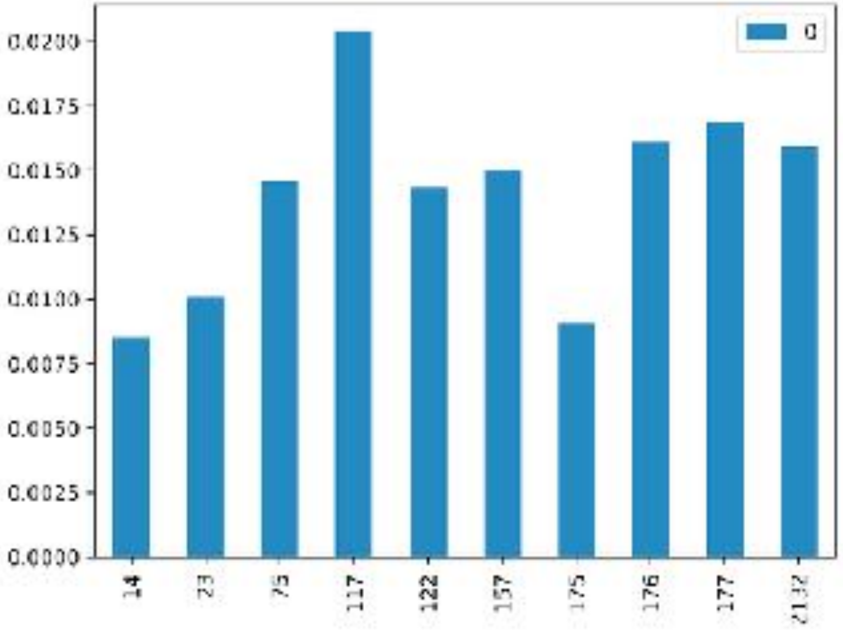

2153 is empty 3x3 pattern

Experiments - Non-Master Vs Master in 1000 simulations

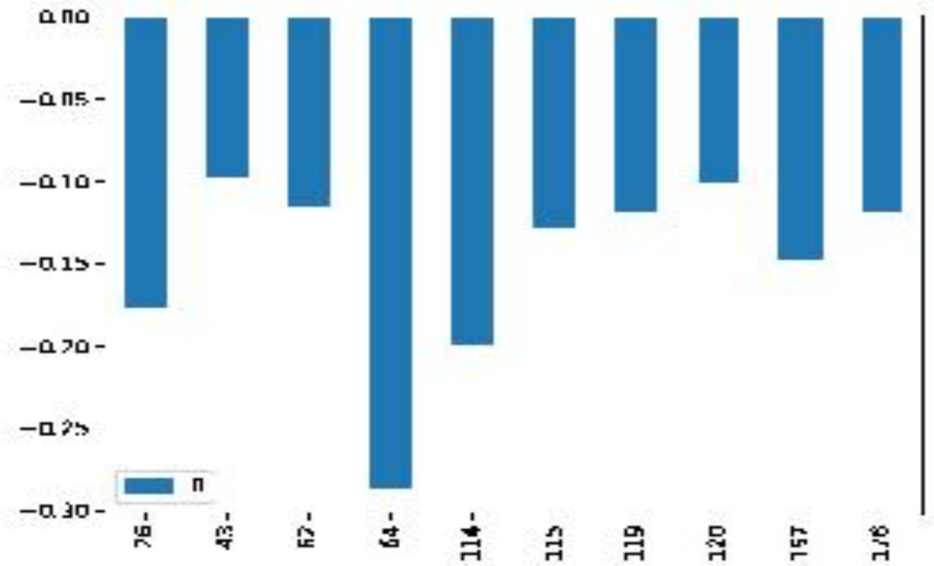

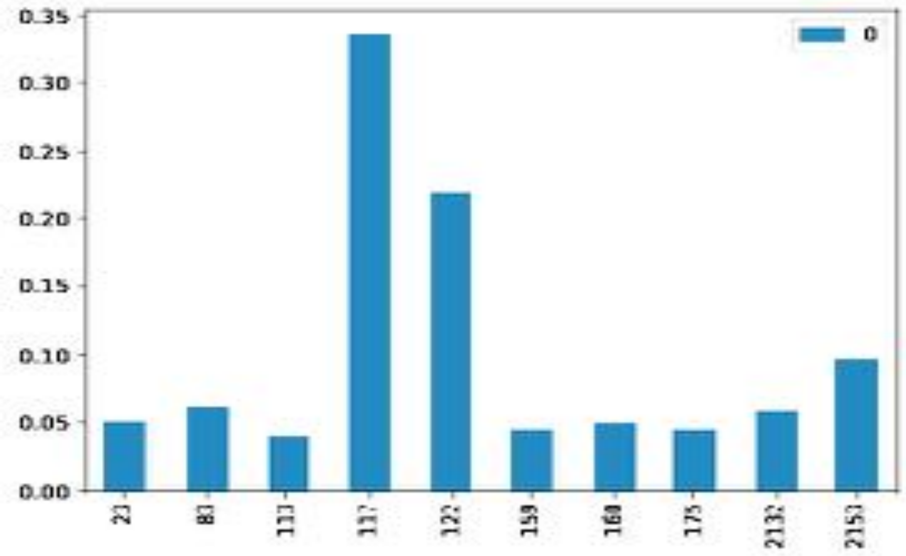

26: DIST PREV 2
43: DIST PREV OWN 2
62: ATARI DEFEND
64: PLAYOUT POLICY 3X3 PATTERN
114: CFG DISTANCE LAST 1
115: CFG DISTANCE LAST 2
119: CFG DISTANCE LAST OWN 1
120: CFG DISTANCE LAST OWN 2
157: DIST CLOSEST OWN STONE 2
176: DIST CLOSEST OPP STONE 2

117: CFG DISTANCE LAST 4+
122: CFG DISTANCE LAST OWN 4+
2153: EMPTY 3X3 PATTERN

Test set = moves played by Leela Zero in master games
plots = nonmaster frequency - master frequency
Feature 64 happens more often in master moves also found by Leela

Conclusions

Evaluation Methods:

- Relation of move prediction to playing strength is complex.
  - Early+middle game prediction is better than full-game prediction
- Better knowledge scales well with more search

Feature Frequencies in different players:

- Many “Tenuki” moves by professional players initially not found by Fuego
  - Up to 24% of master moves not found by Fuego are at distance 4 or more
  - More search can find some
- Additive knowledge likes playing close to existing stones
- Features knowledge likes local responses to previous move

Knowledge Initialization:

- Most Fuego moves have weight very close to maximum, and get most of the simulations
- Professional moves usually get either the majority of simulations, or nothing