

Maia Chess

A Human-Like Neural Network Chess Engine

Ashton Anderson

Reid McIlroy-Young

University of Toronto



Reid McIlroy-Young
University of Toronto



Siddhartha Sen
Microsoft Research

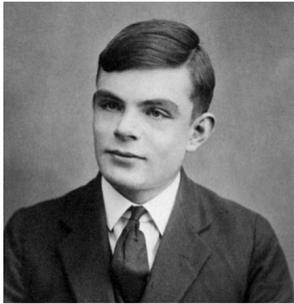


Jon Kleinberg
Cornell University



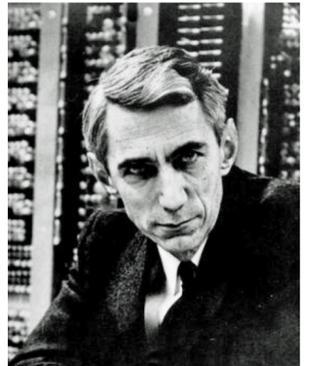
Ashton Anderson
University of Toronto

Chess and AI: A Long History



Alan Turing created the first chess-playing algorithm in 1948, before computers could run it

Claude Shannon described minimax in chess in 1949



Long-standing AI problem: “Chess is the drosophila of artificial intelligence” — John McCarthy, 1967

“**However**, computer chess has developed much as genetics might have if the geneticists had concentrated ... on breeding racing Drosophila. We would have some science, but mainly we would have very fast fruit flies.” — John McCarthy

Chess and AI: A Long History

Deep Blue defeats Kasparov in 1997



AlphaZero defeats Stockfish in 2017

Chess AI and People



+0.4 Stockfish 11+ WASMX
Depth 22 CLOUD +

+0.4 5... d5 6. cxd5 ♖b4 7. d3 ♜fxd5 8. ♜xd5 ♜xd5 9. ♜f3 ♜d6 1...

+0.5 5... d6 6. d3 g6 7. ♜ge2 ♜g7 8. O-O O-O 9. a3 ♜e6 10. ♜d5

+0.8 5... h5 6. h3 d6 7. d3 ♜e6 8. ♜f3 d5 9. ♜g5 dxc4 10. ♜xe6

“Here’s what I would do” —3400-rated
calculation beast 🤖

“But what should I do?” —1100-rated human 😞

Chess AI and People

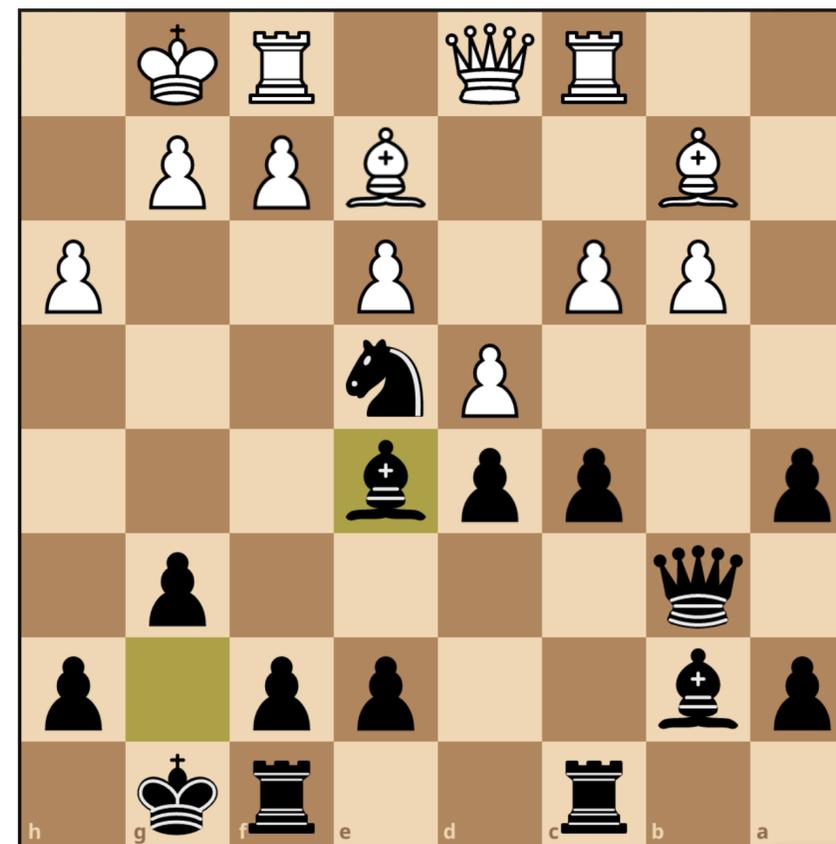
Stockfish level 4



Ashton

Chess AI and People

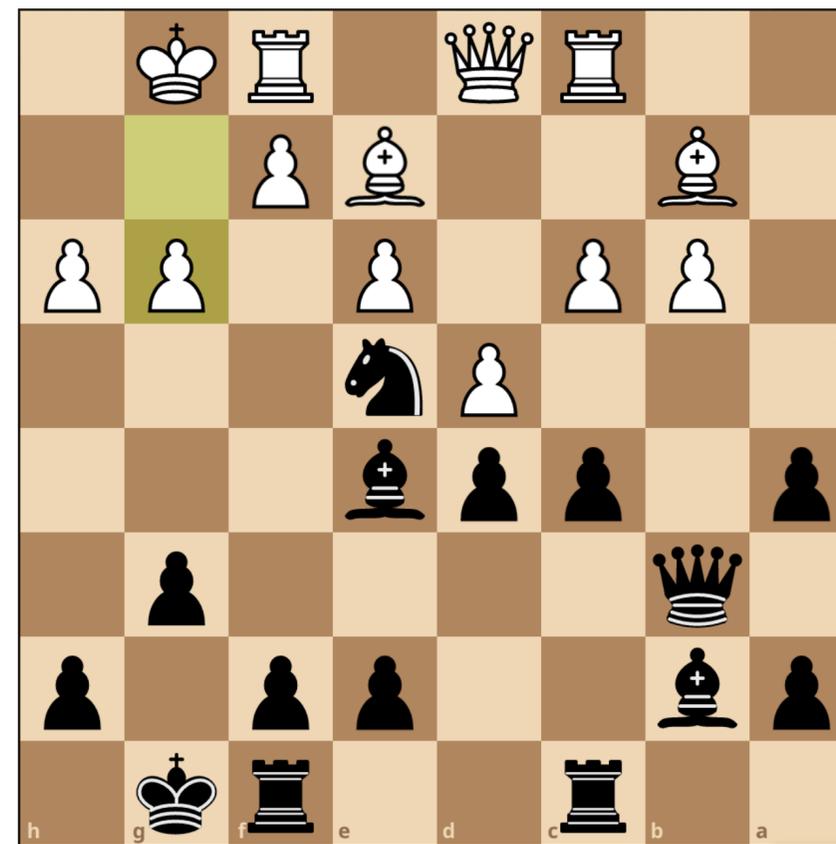
Stockfish level 4



Ashton

Chess AI and People

Stockfish level 4



Ashton

Chess AI and People

How do we bridge the gap between artificial and human intelligence?

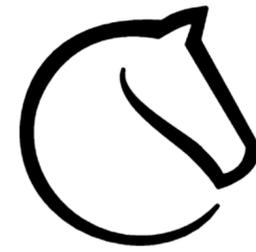
Chess as a model system:

- 1) AI reached superhuman performance at least 15 years ago
- 2) There is voluminous data on human activity in chess
- 3) Wide range of human skill levels

How can we algorithmically capture
human style in chess?

**Predict the next move a human, at a specific skill level,
will make in a real game**

Data



lichess.org
game database

The largest openly available dataset of human games
1.7B games and counting
Players, moves, move times, etc.

Background: Attenuated Engines

Depth-limited Stockfish



Limit search depth to simulate fallible thinking

Early Leela models



Take Leela versions early in their self-play training to simulate not-fully-evolved understanding

Do Attenuated Engines Perform Well?

Attenuated engines match aggregate human performance (**rating**)

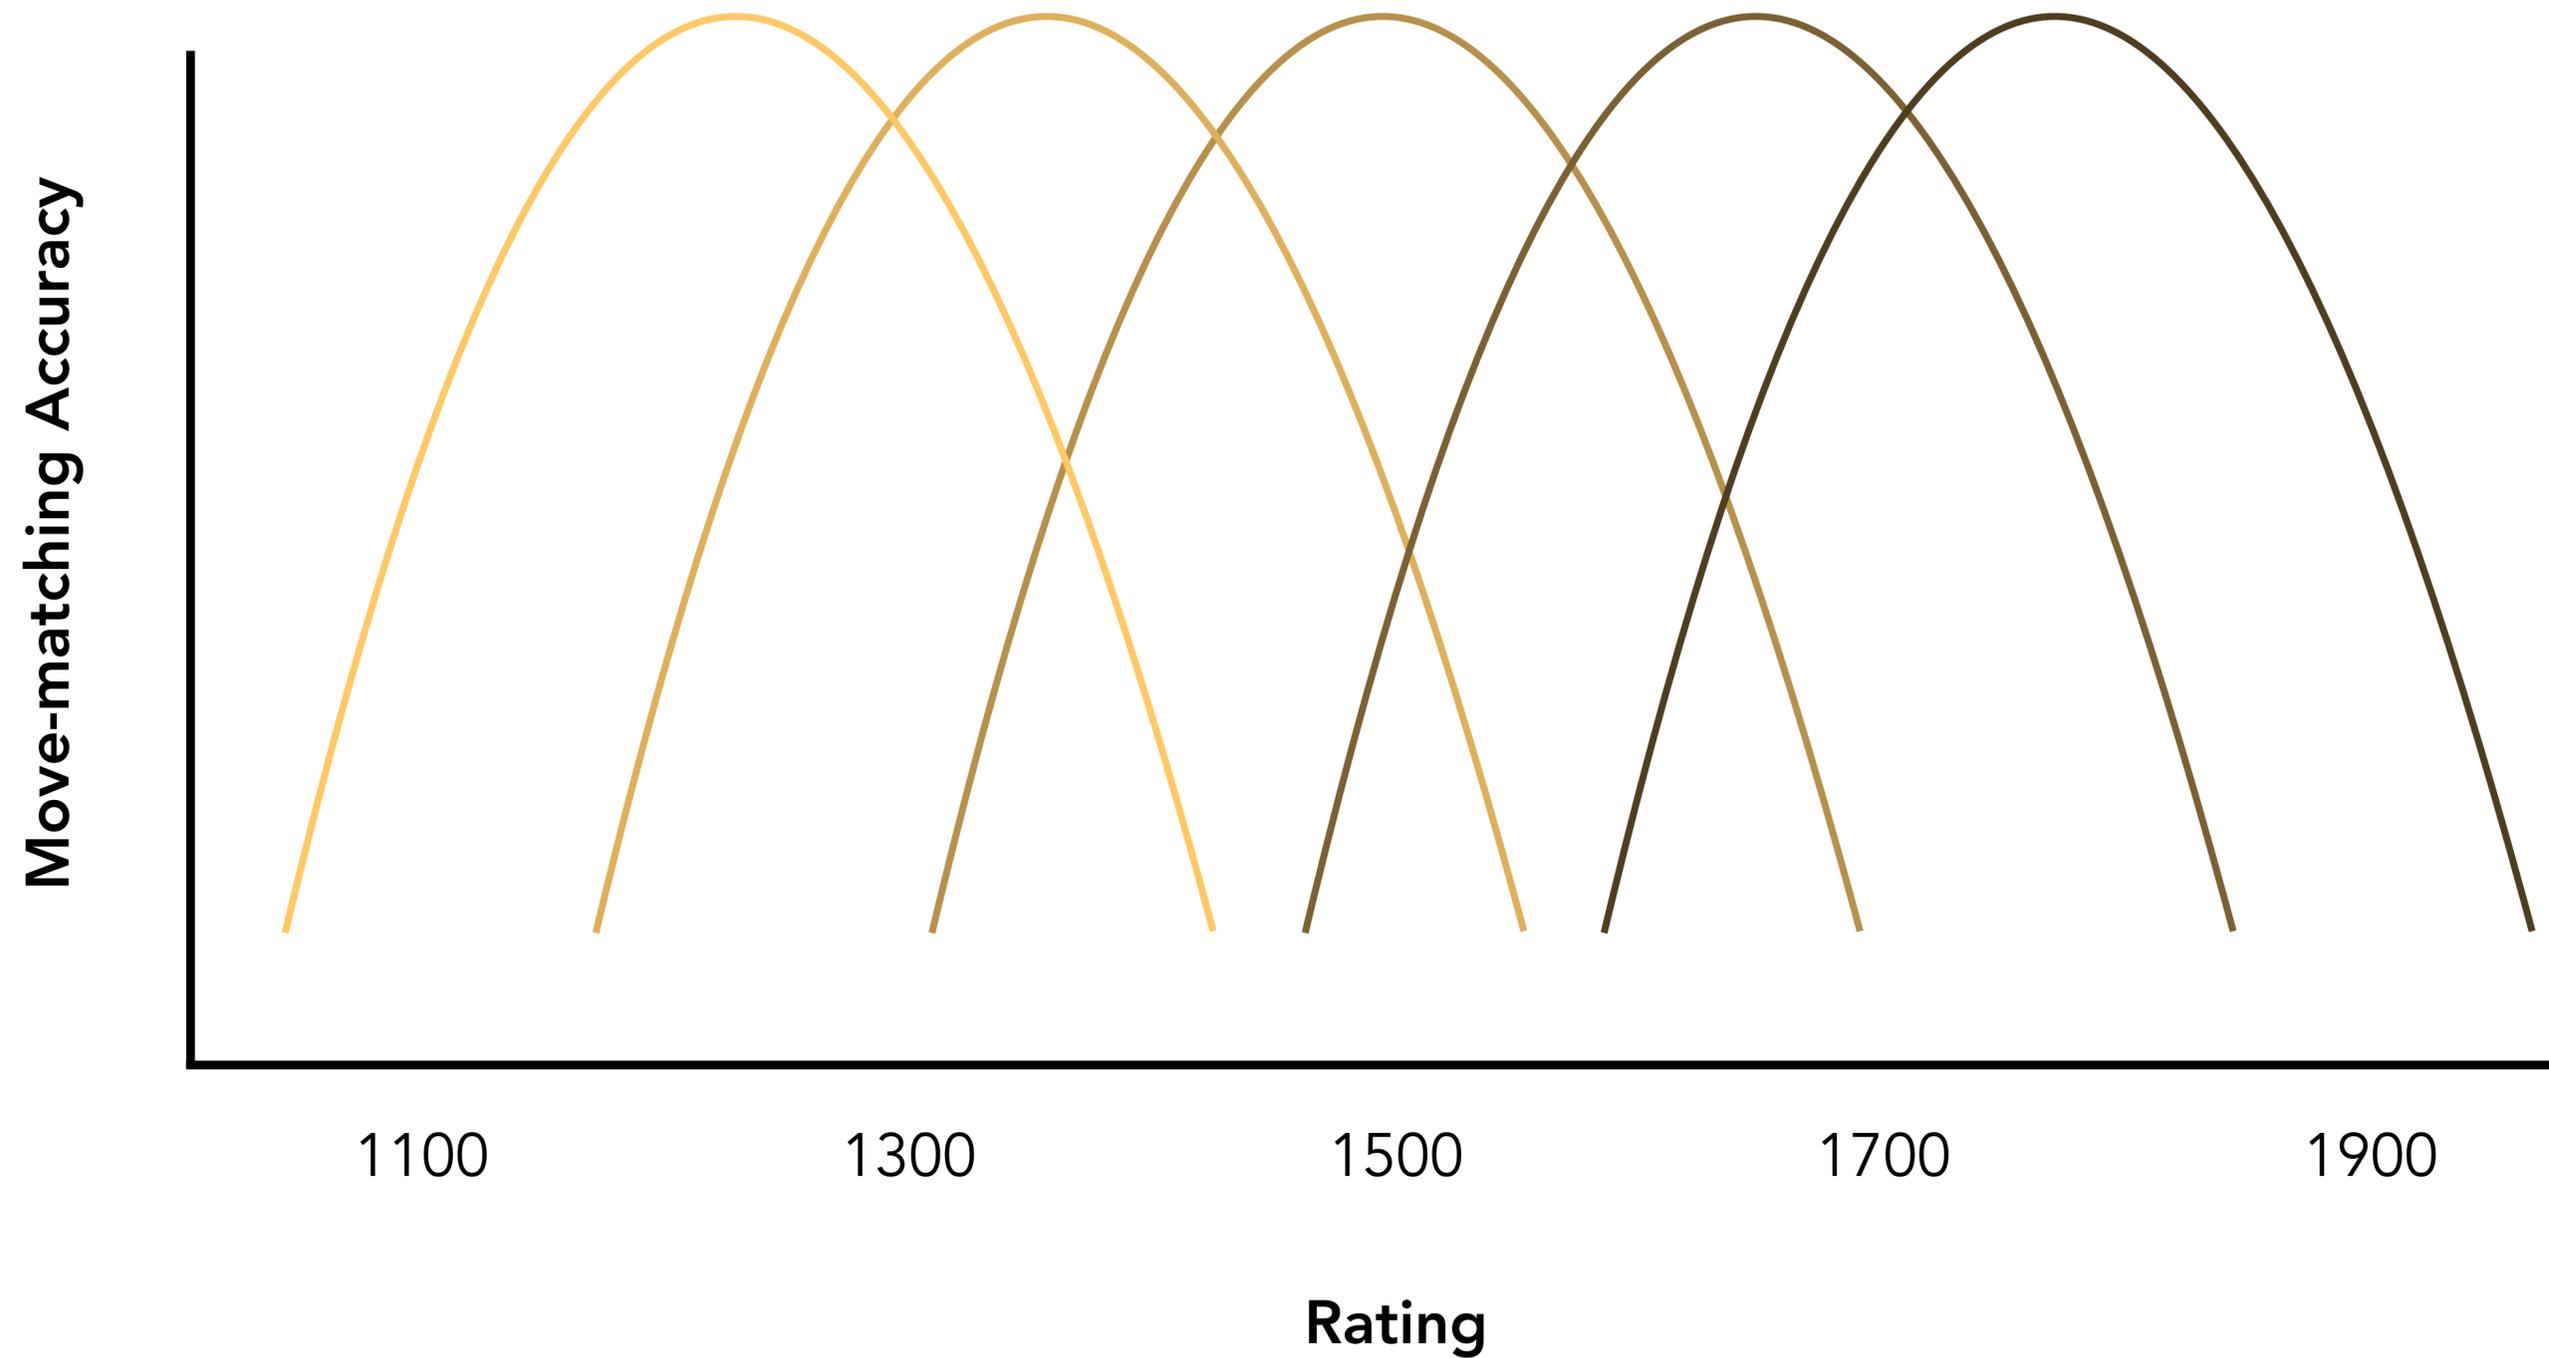
Do they match granular human decision-making (**moves**)?

Evaluation Set

From December 2019:

1. Create bins for each range of 100 rating points
2. Divide games into bins by rating of both players
3. Select 10,000 games from each bin between 1100 and 1900
4. Metric: **move-matching accuracy**: % of positions for which model's move matches the human move played in the game.

What Does It Mean to Perform Well?



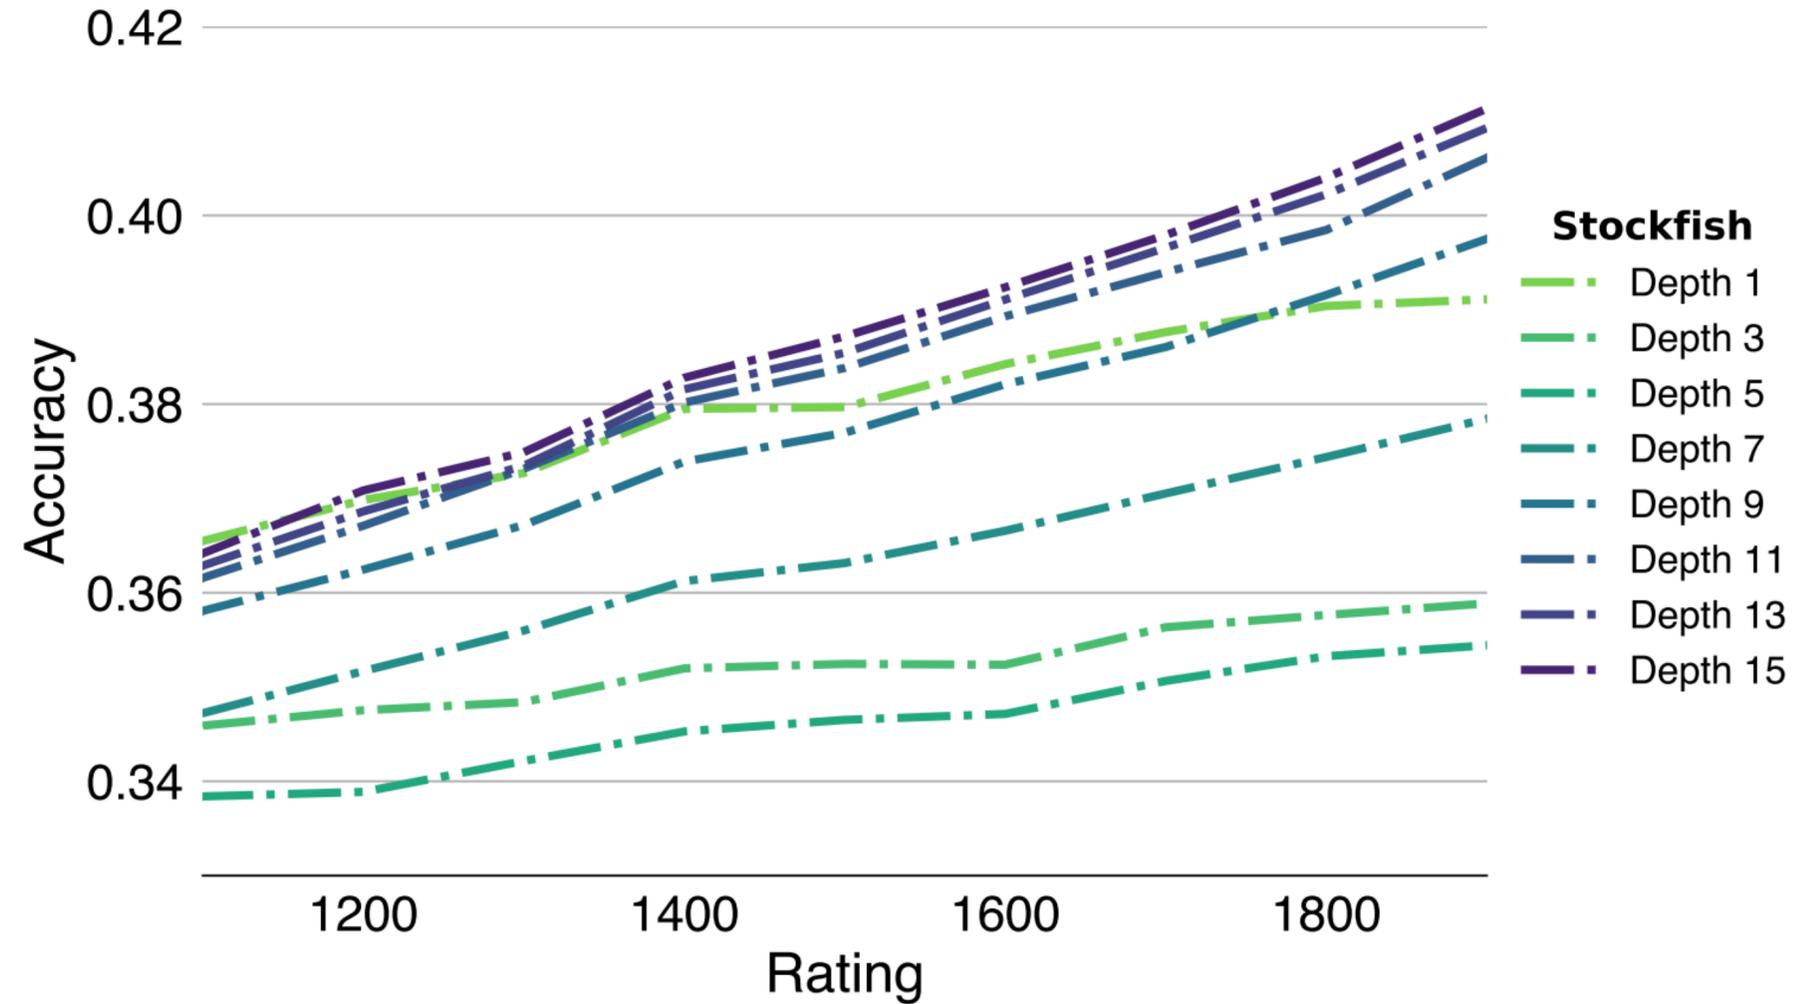
Stockfish

Attenuated Stockfish: limited to various depths

Move-matching accuracy increases with rating of players being predicted

Relatively low accuracy, no skill-level targeting

Interesting non-monotonicity: d1 and d15 both more accurate than d5

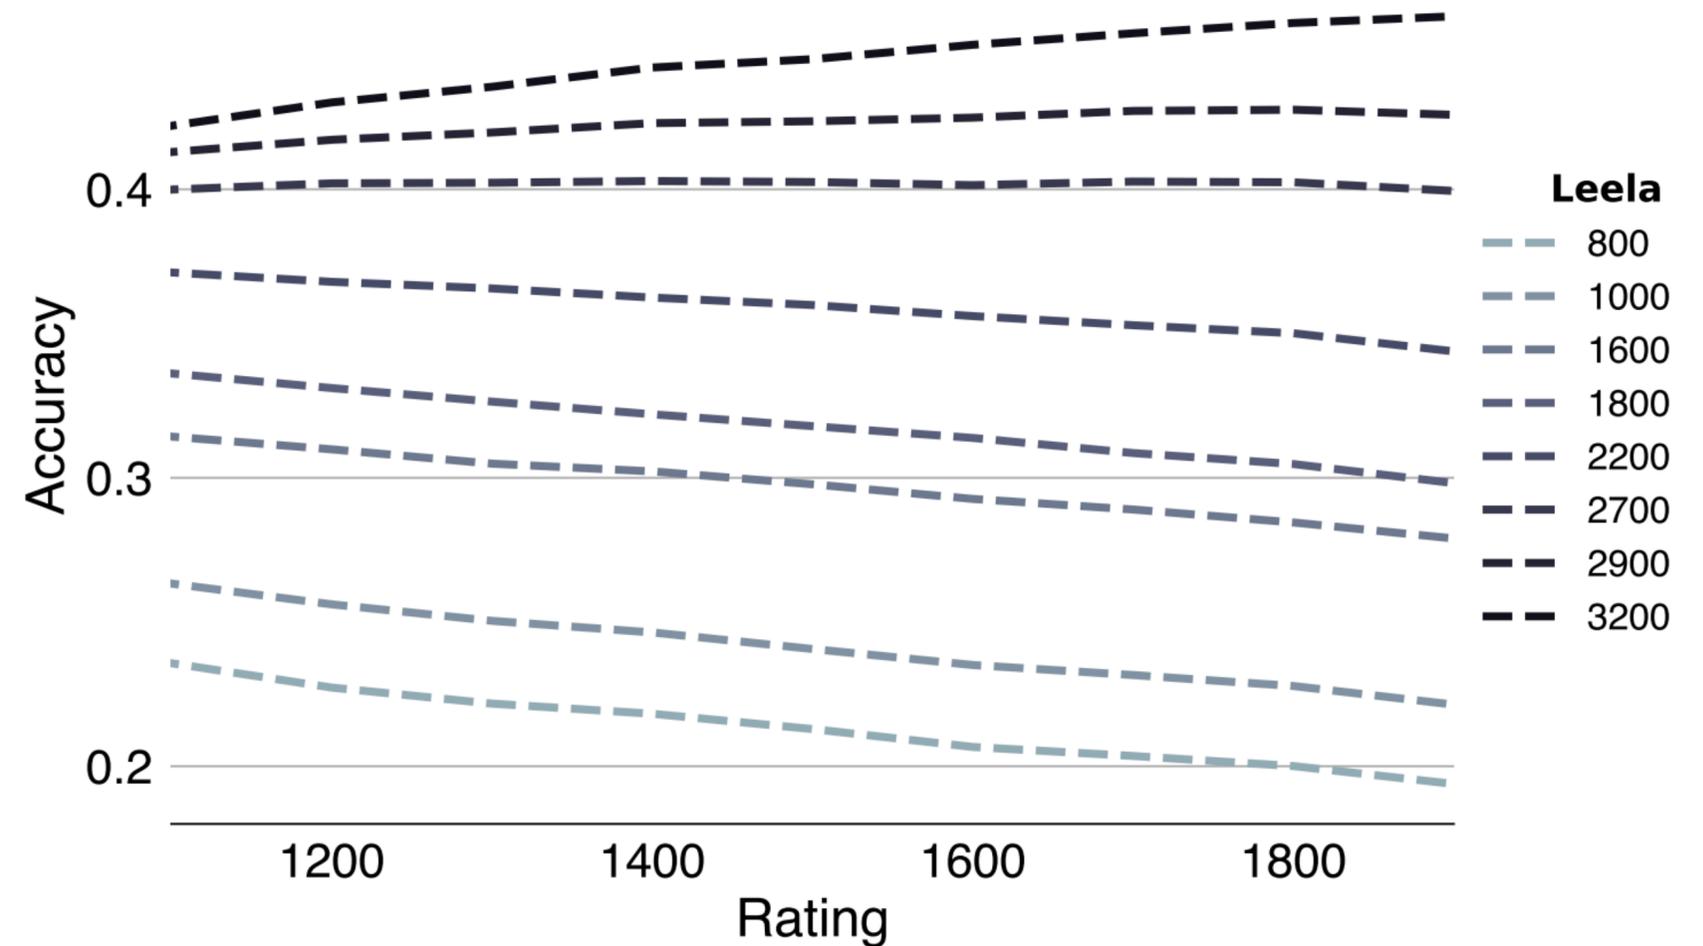


Leela

Attenuated Leela: versions at different points in training

Move-matching accuracy flat across ratings of players being predicted

Somewhat higher accuracy, no skill-level targeting

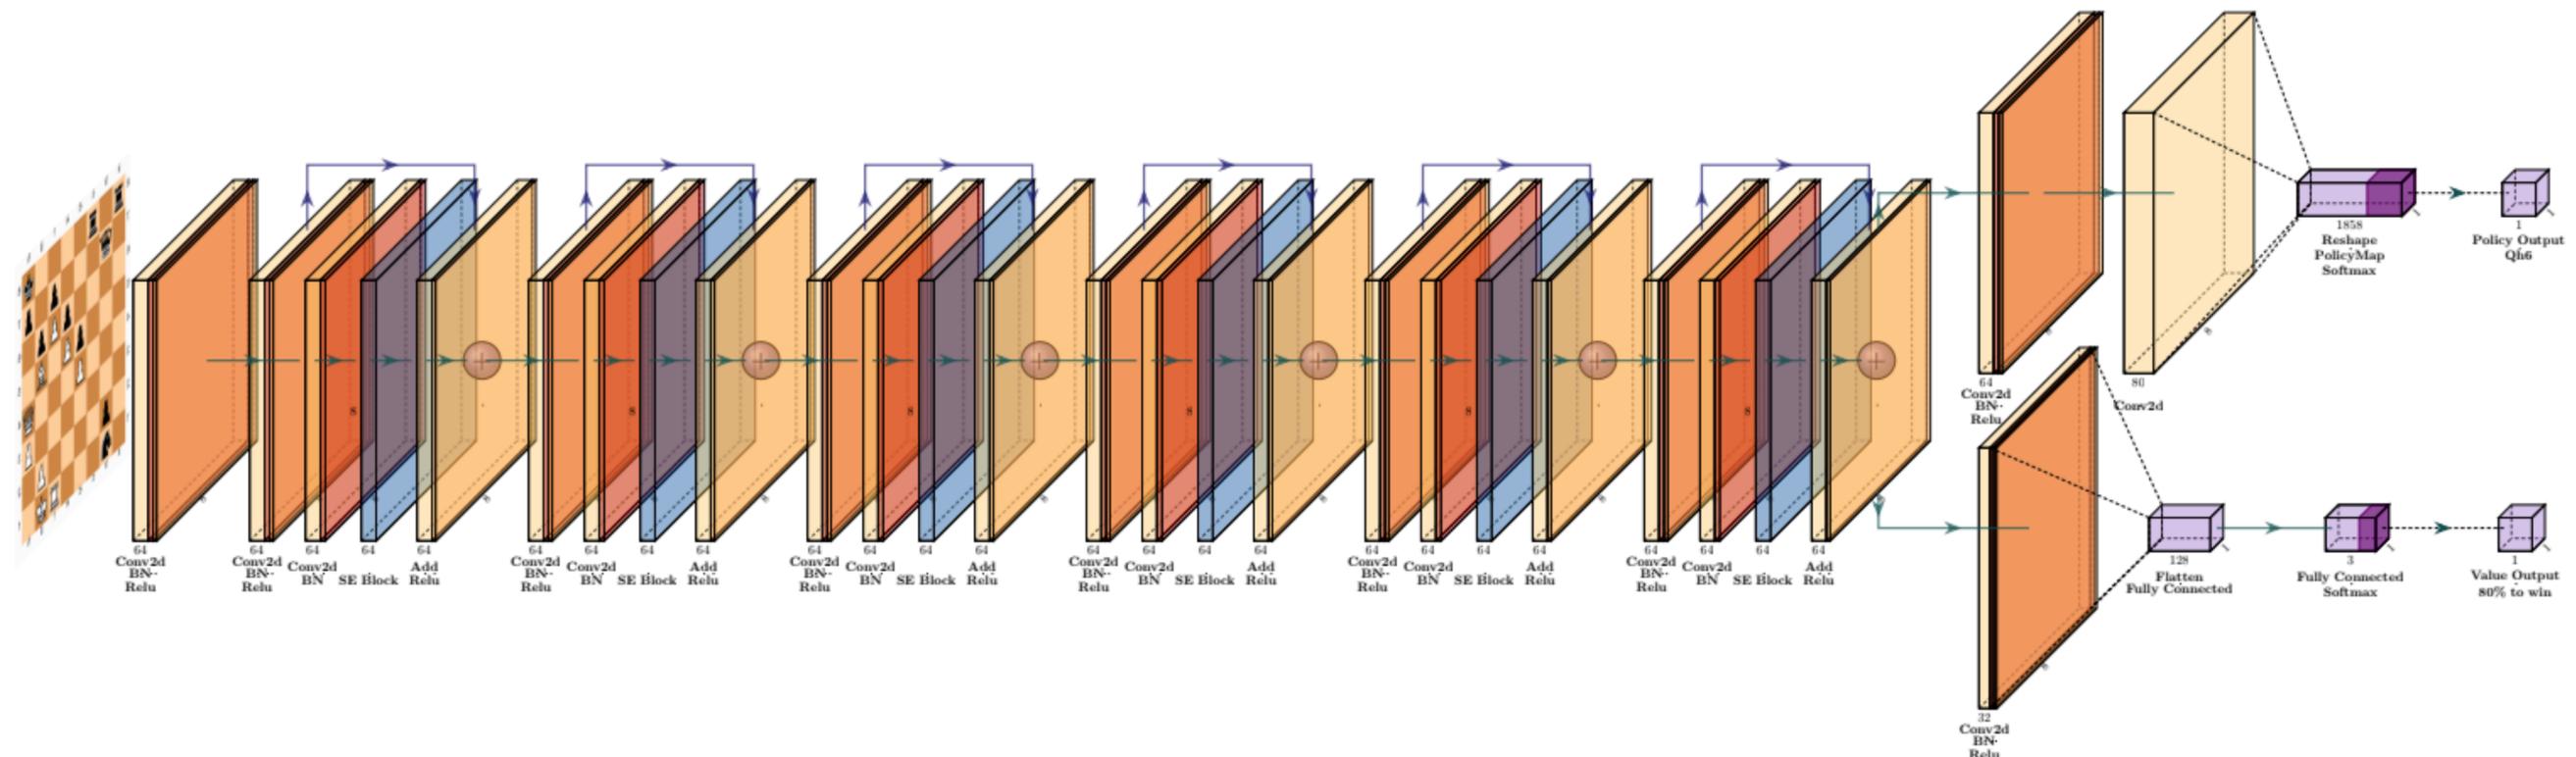


Maia

Key idea:

Learn from human play instead of self-play

Goal: predict the next move a human will play



Maia–Leela Comparison

Leela

Maia

Data	Self-play games	Human games
Goal	Optimal play	Human-like play
Method	Reinforcement learning	Classification
Tree search	Monte-Carlo tree search	No tree search
Policy	Optimal moves	Human moves
Value	Probability of winning given optimal play	Probability of human winning

Maia Training

Train a Maia for each rating level between 1100 and 1900

Maia 1100 12M games between 1100s

Maia 1200 12M games between 1200s

Maia 1300 12M games between 1300s

Maia 1400 12M games between 1400s

Maia 1500 12M games between 1500s

Maia 1600 12M games between 1600s

Maia 1700 12M games between 1700s

Maia 1800 12M games between 1800s

Maia 1900 12M games between 1900s

Maia Move-Matching Performance

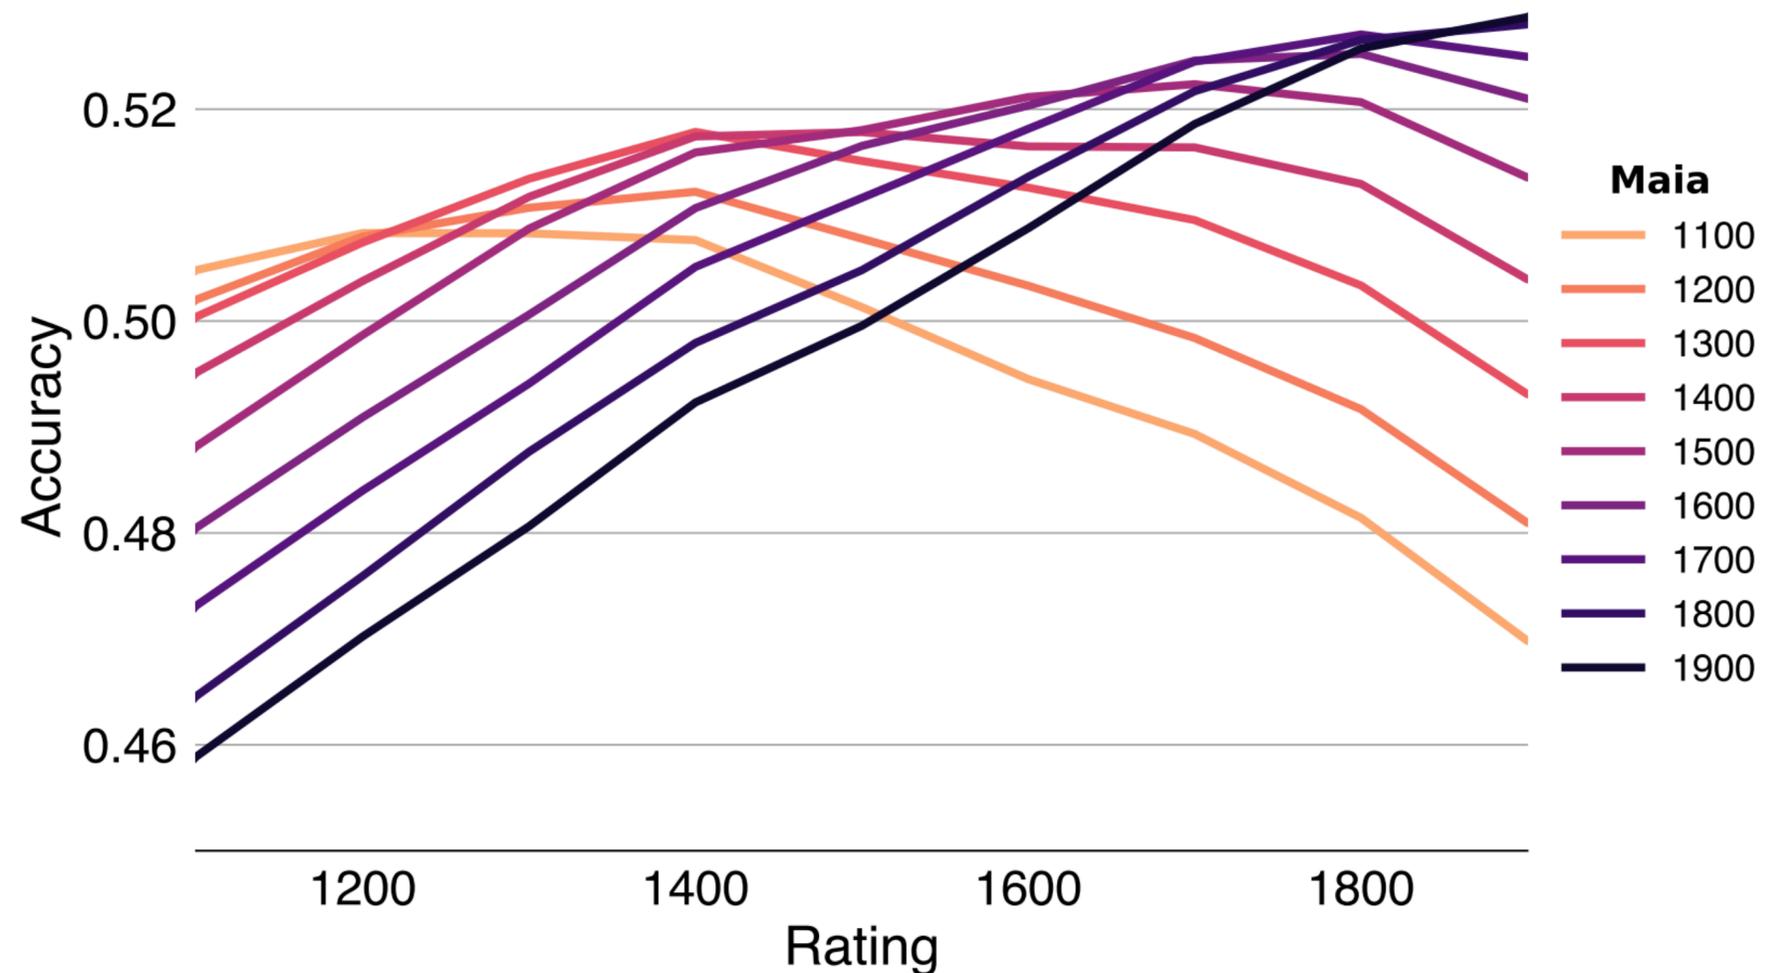
Maia: versions trained on different rating levels

High accuracy: best performance >50%

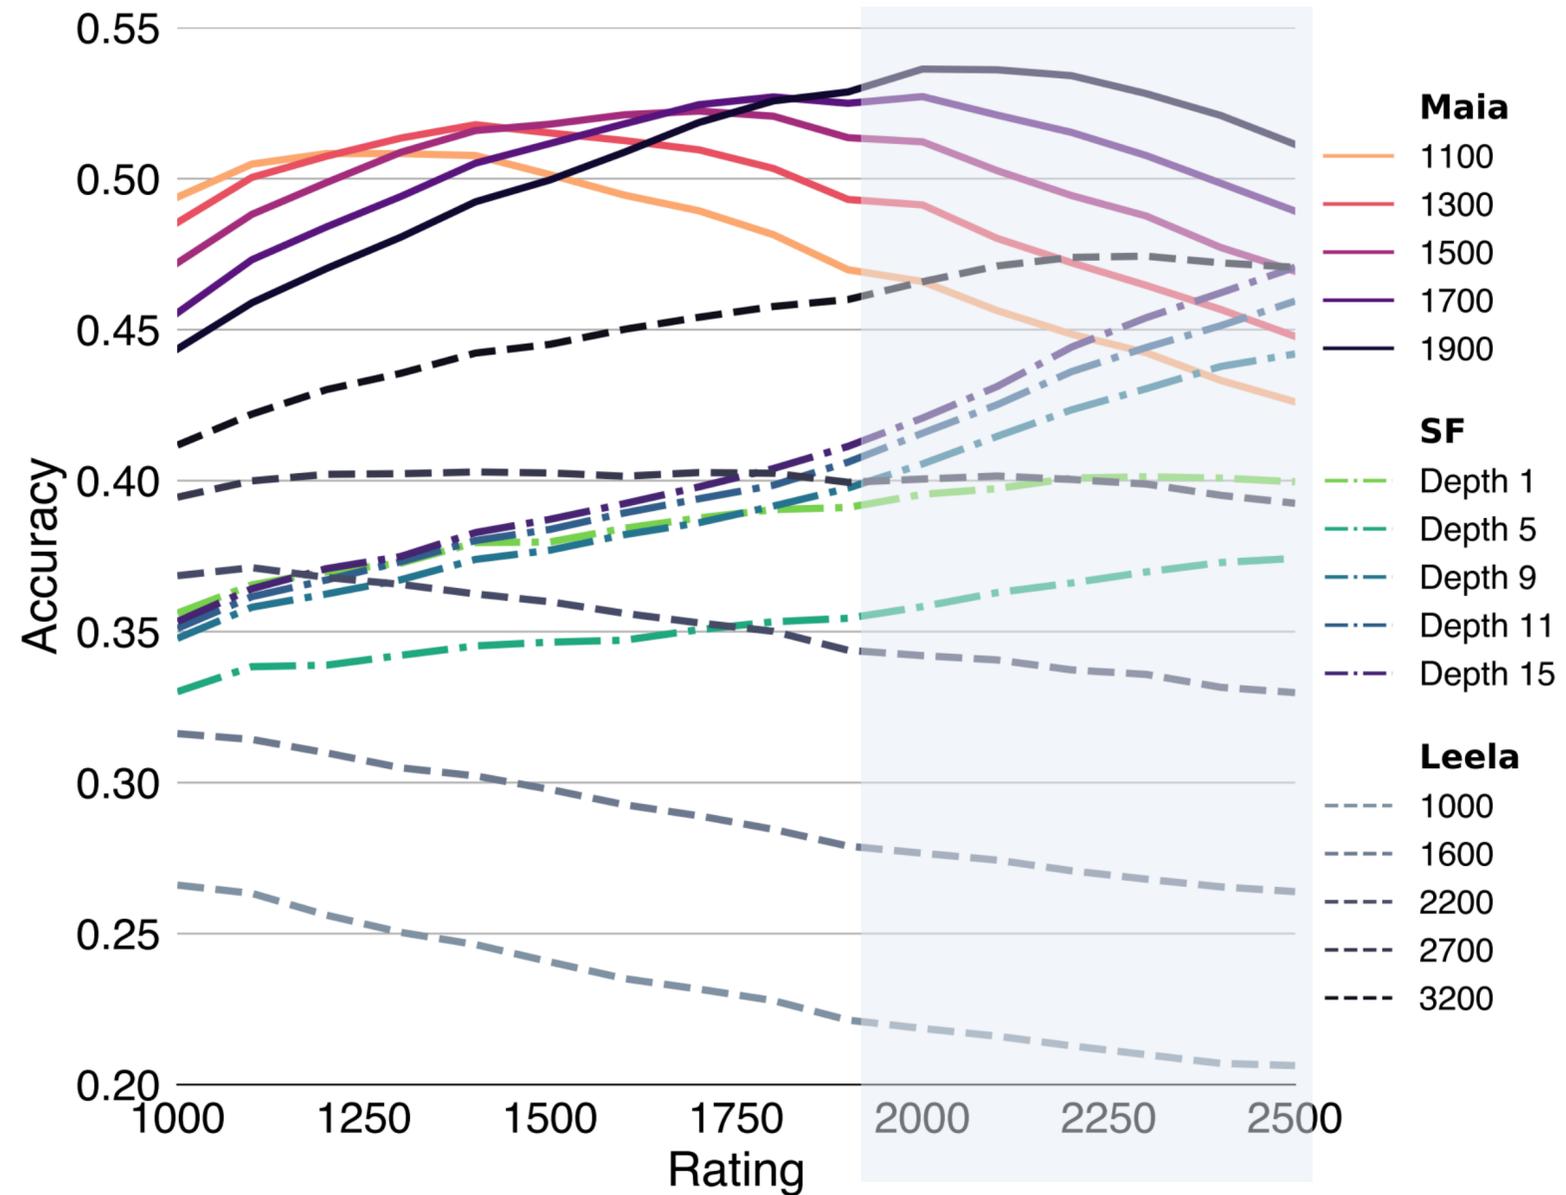
Move-matching accuracy maximised at target rating level

Maia's lowest accuracy, Maia 1900 tested on 1100-rated players (46%), is still higher than any Stockfish or Leela model we tested

High accuracy and skill-level targeting



Maia Move-Matching Performance



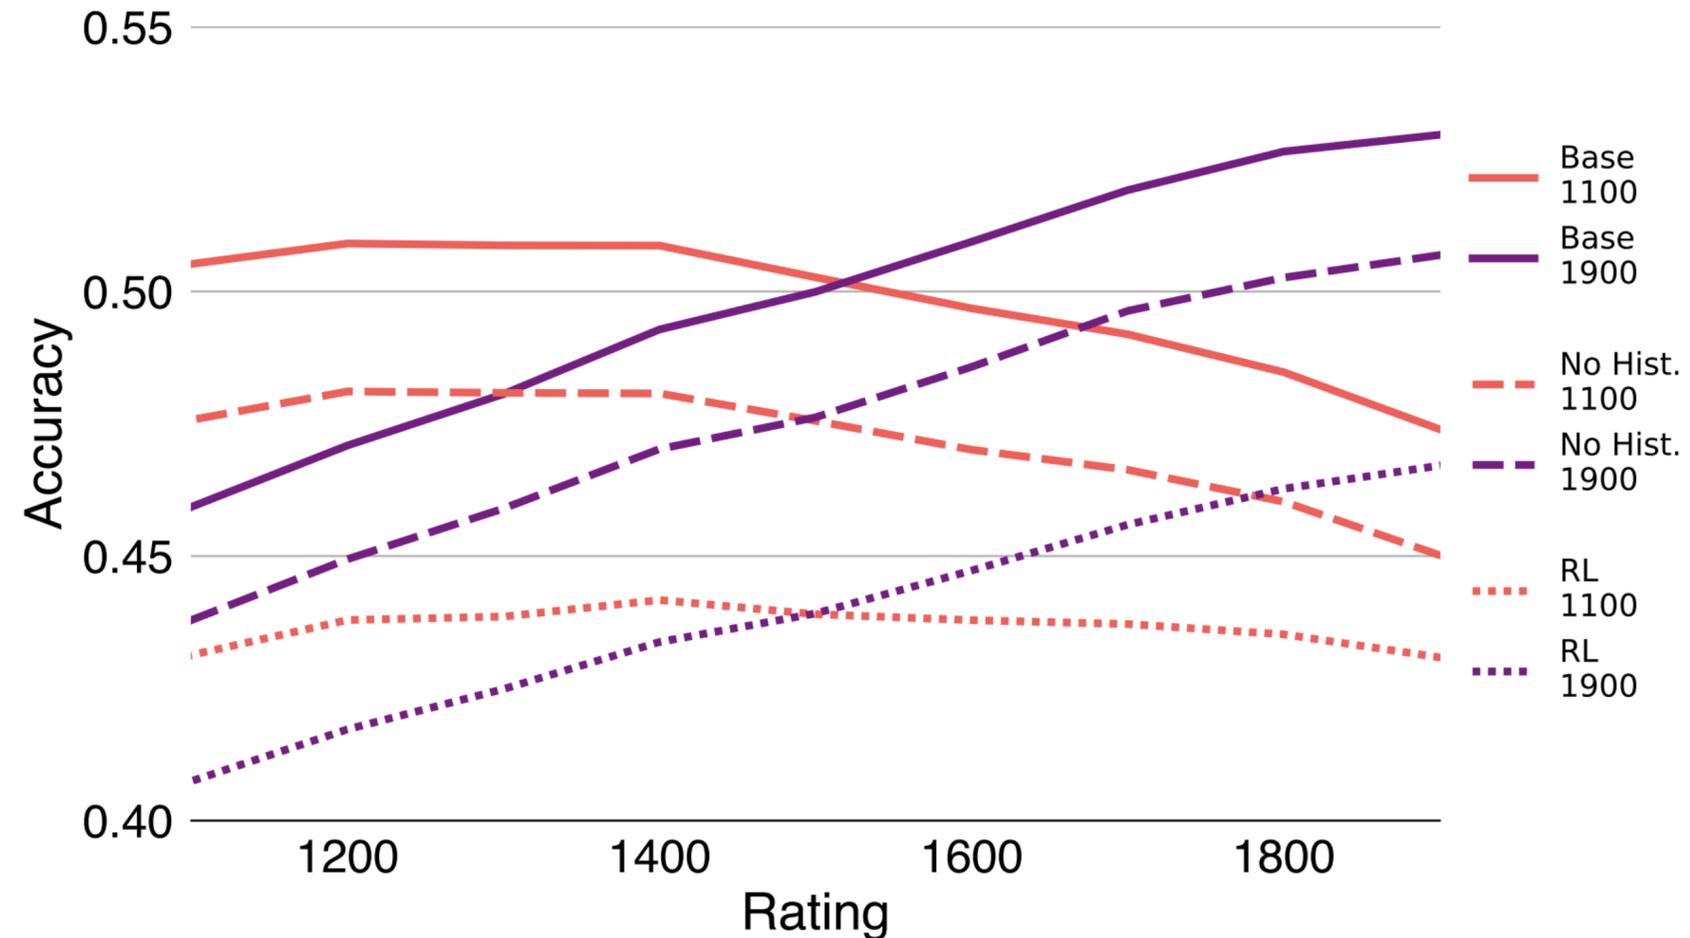
Shaded region: rating levels not explicitly targeted (2000+)

Maia: Architecture Decisions

Base Maia vs. a version with no history and a version with Leela-style rollouts (calculation)

Including calculation hurts move-matching performance by ~5–8% (moves become too strong)

Dropping history hurts performance by 2–3%



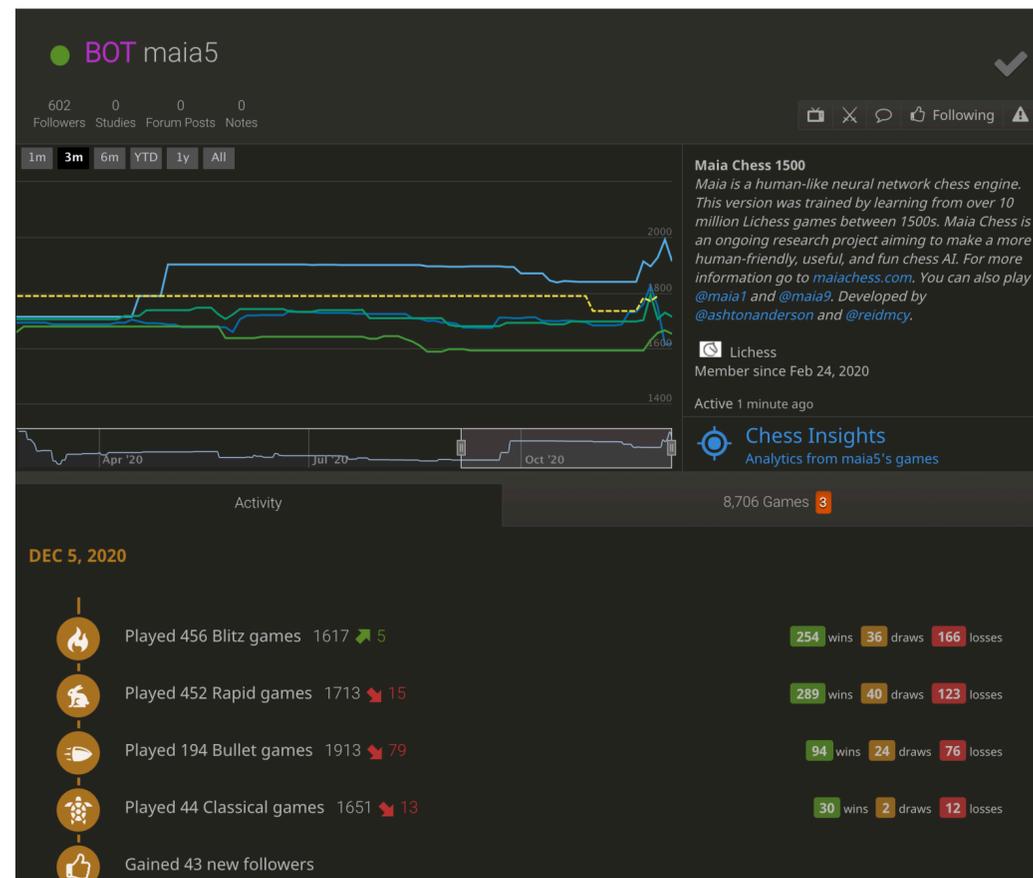
Maia Bots

@maia1: Maia 1100



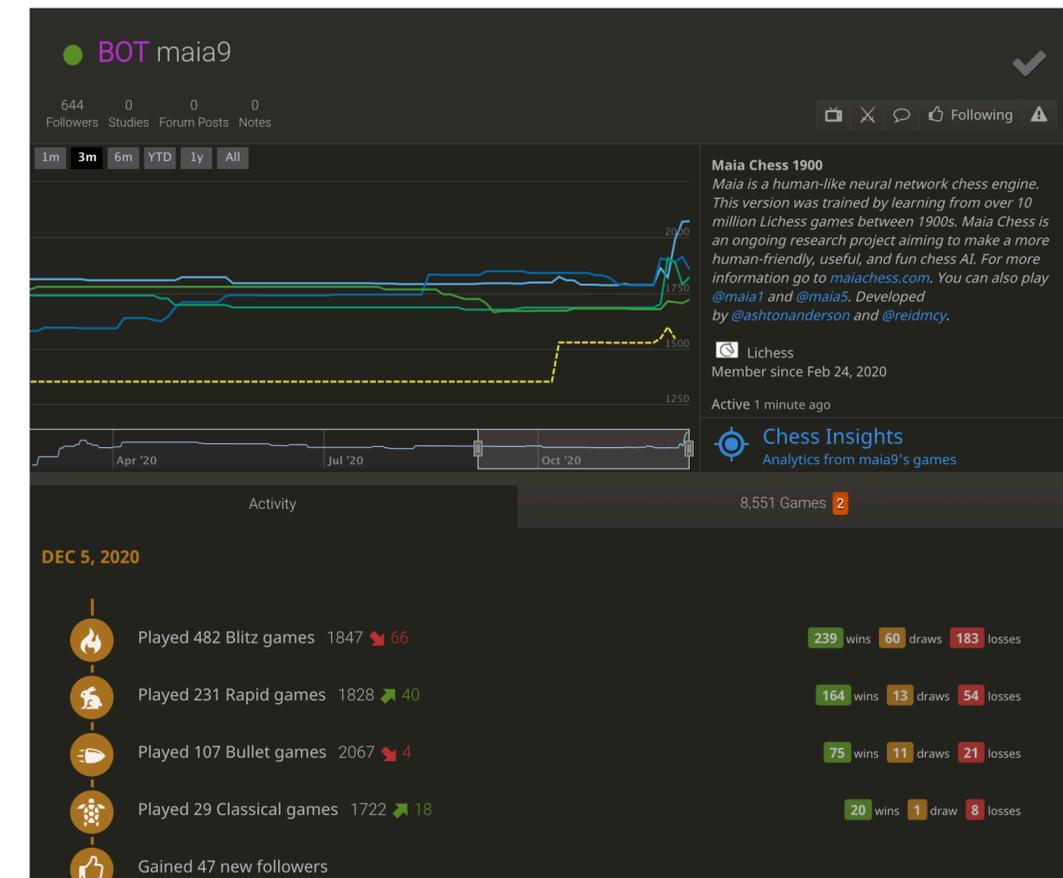
21,000+ games

@maia5: Maia 1500



8,000+ games

@maia9: Maia 1900



8,000+ games

Playing Maia 1100 is like playing a committee of 1100-rated players, etc.

Maia Bots

7,000+ unique human opponents in 4 days

Max # of games started in an hour: 1,491

"I've been playing engines for over 15 years. This is the first time I'm playing an engine that feels like a person. When it makes mistakes, they are natural mistakes, like missing something in the midst of tricky tactics."

Top opponent by # of games: 149

"this is amazing to practice against...Definitely a 9.9/10 from me"



agadmator vs Maia || A Human-like Neural Network...

agadmator's Chess Chann... ✓
115K views • 3 days ago

Meet Maia Chess || A Human-like Neural Network Chess...

agadmator's Chess Chann... ✓
103K views • 3 days ago

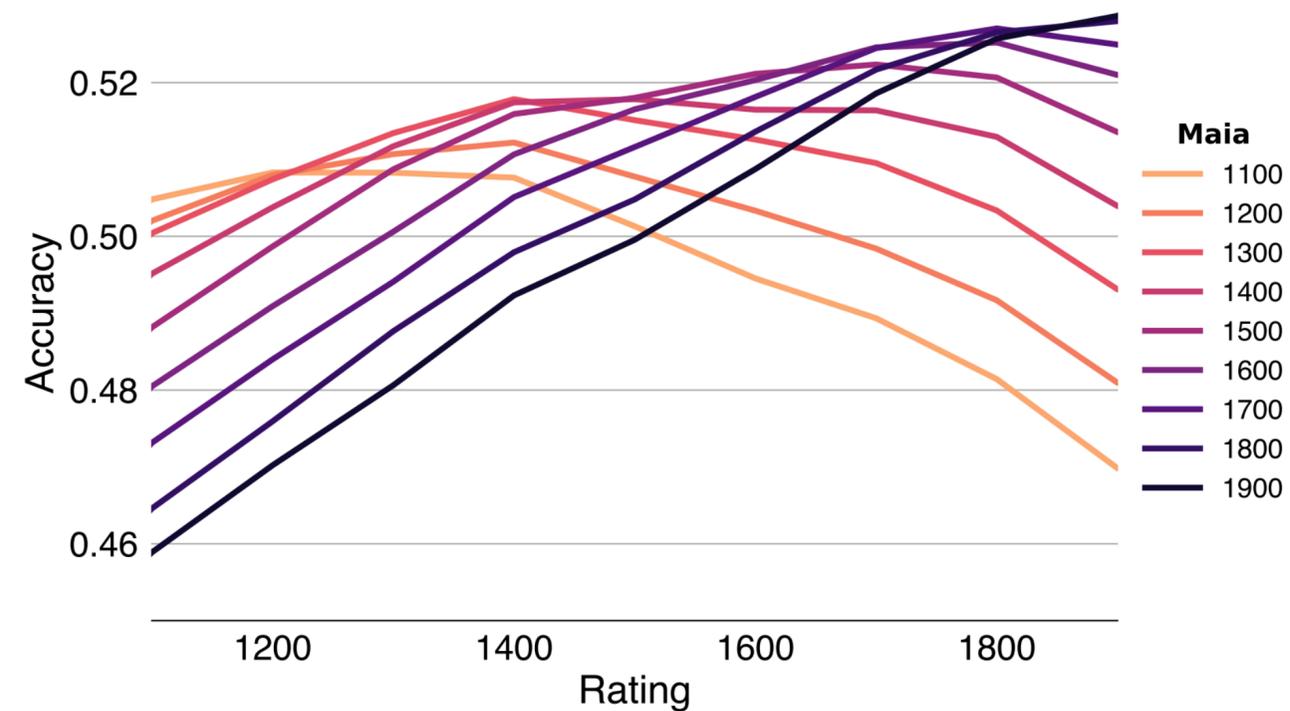


Maia Chess: A human-like neural network chess engine

27K views • 3 days ago

Maia: A Framework to Understand Human Play

Since we can **predict human play at different levels**, there is a reliable, predictable, and maybe even **algorithmically teachable difference** between one human skill level and the next



Maia: A Framework to Understand Human Play

In this position, Maia levels 1100–1400 correctly predict White will play the tempting but wrong move b6 (the move played in the game)

Maia levels 1500–1900 predict that, on average, players rated 1500 and above will play the correct bxa6, forcing the Queenside open to decisive effect

Model	Move	Agreed	↗ Human move / correct prediction	↗ Incorrect Predictions	↗ Stockfish
Stockfish	bxa6	✗			
Maia 1100	b6	✓			
Maia 1200	b6	✓			
Maia 1300	b6	✓			
Maia 1400	b6	✓			
Maia 1500	bxa6	✗			
Maia 1600	bxa6	✗			
Maia 1700	bxa6	✗			
Maia 1800	bxa6	✗			
Maia 1900	bxa6	✗			

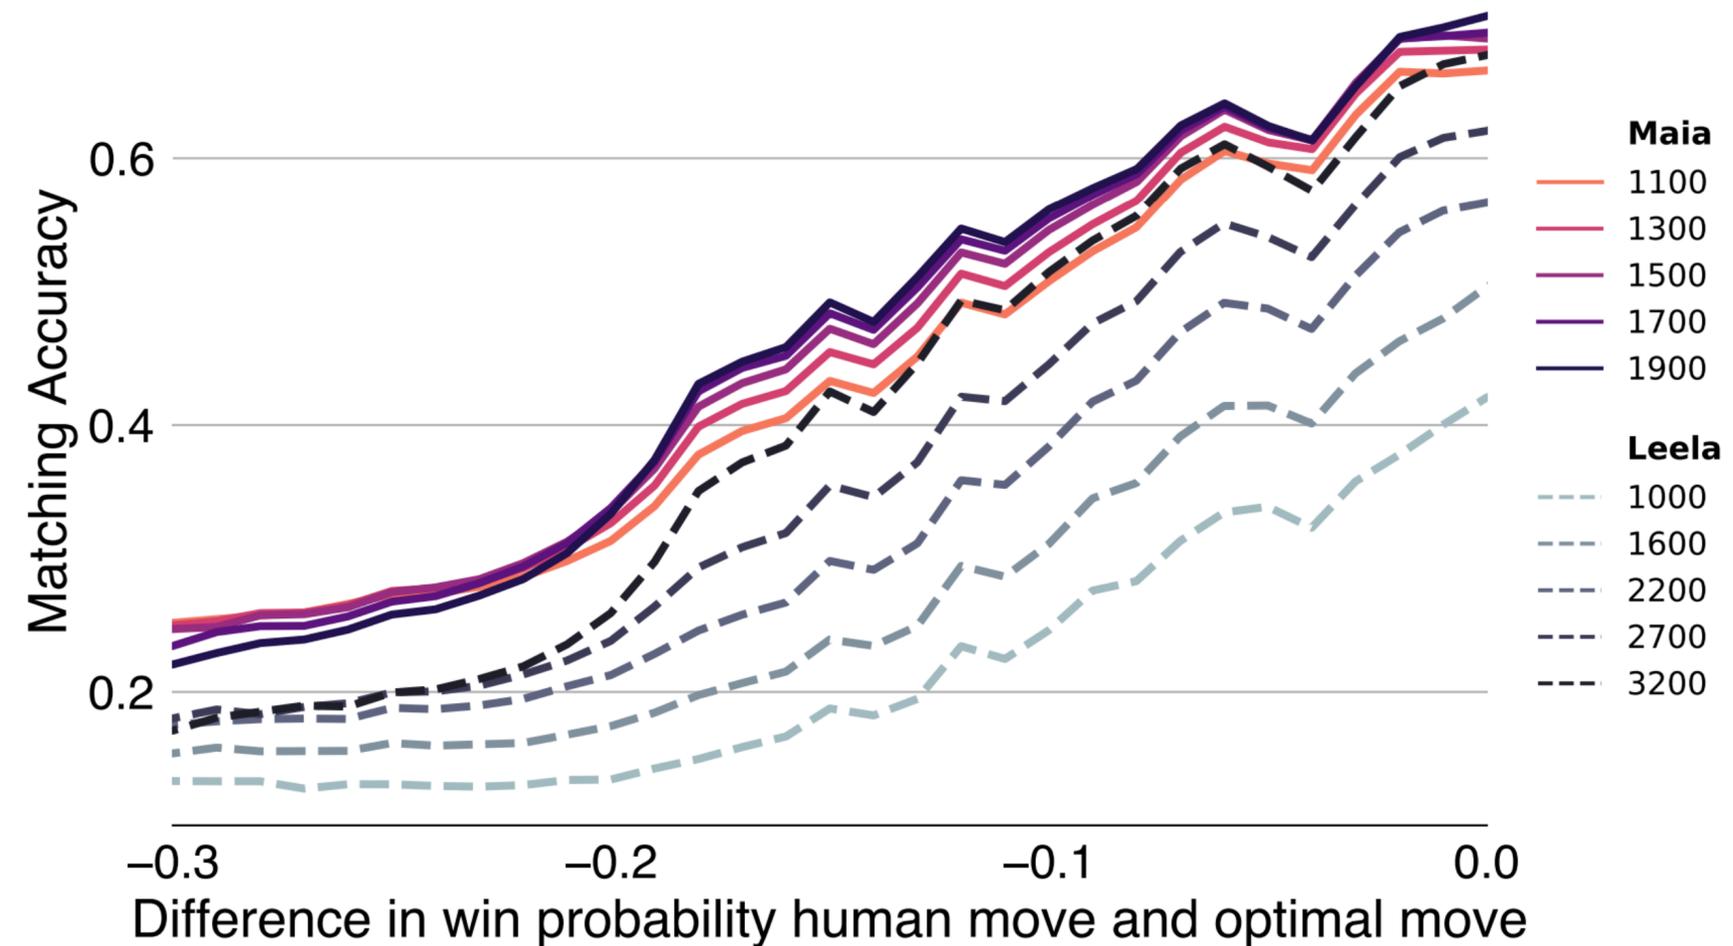
The chessboard diagram shows a position where White is to move. The board is oriented with White on top. The pieces are: White King on c8, White Rook on f8, White Rook on g8, White Queen on c7, White Bishop on d7, White Pawn on b6, White Bishop on f6, White Pawn on a5, White Pawn on c5, White Pawn on d5, White Pawn on e5, White Knight on f5, White Pawn on g5, White Pawn on a3, White Pawn on c3, White Knight on e3, White Queen on c2, White Bishop on e2, White Pawn on d2, White Pawn on e2, White Pawn on f2, White Pawn on g2, White Pawn on h2, White Rook on a1, White Pawn on b1, White Pawn on c1, White Pawn on d1, White Pawn on e1, White Rook on f1, White King on g1. Black pieces are: Black Pawn on b7, Black King on c7, Black Bishop on d7, Black Pawn on e6, Black Pawn on f6, Black Pawn on h6, Black Pawn on a5, Black Pawn on c5, Black Pawn on d5, Black Pawn on e5, Black Knight on f5, Black Pawn on g5, Black Pawn on a3, Black Pawn on c3, Black Knight on e3, Black Pawn on a2, Black Pawn on c2, Black Knight on e2, Black Pawn on d2, Black Pawn on e2, Black Pawn on f2, Black Pawn on g2, Black Pawn on h2, Black Rook on a1, Black Pawn on b1, Black Pawn on c1, Black Pawn on d1, Black Pawn on e1, Black Rook on f1, Black King on g1.

Maia: A Framework to Understand Human Play

Maia predicts mistakes surprisingly well

When players make even the most glaring blunders, Maia predicts the exact move >25% of the time

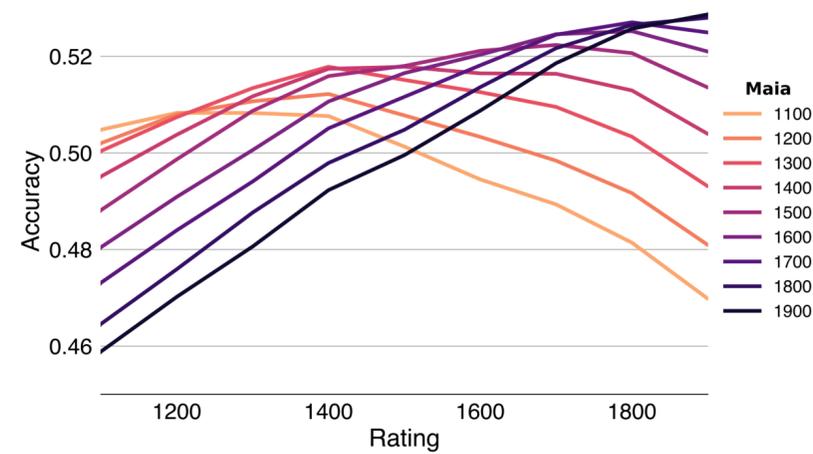
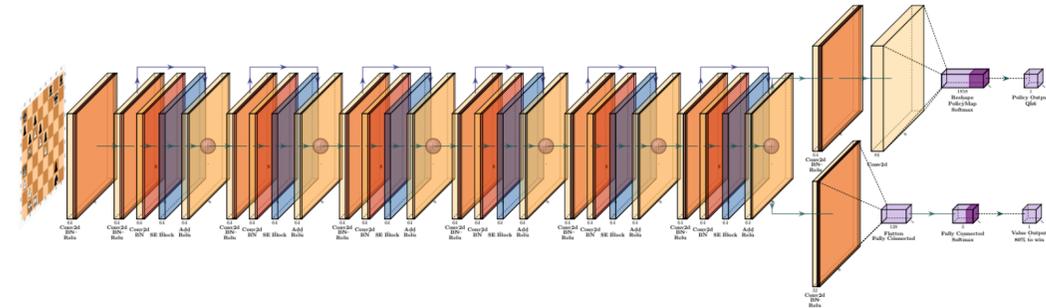
Performance increase relative to other models is larger for blunders



Horrific blunders ← → Correct moves

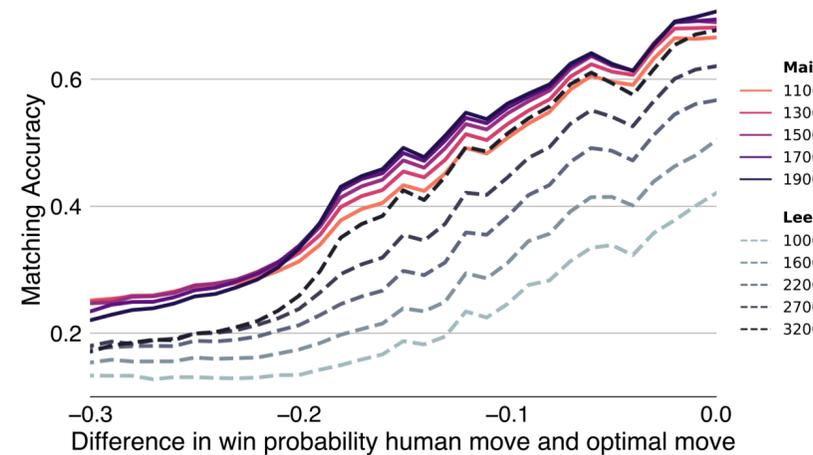
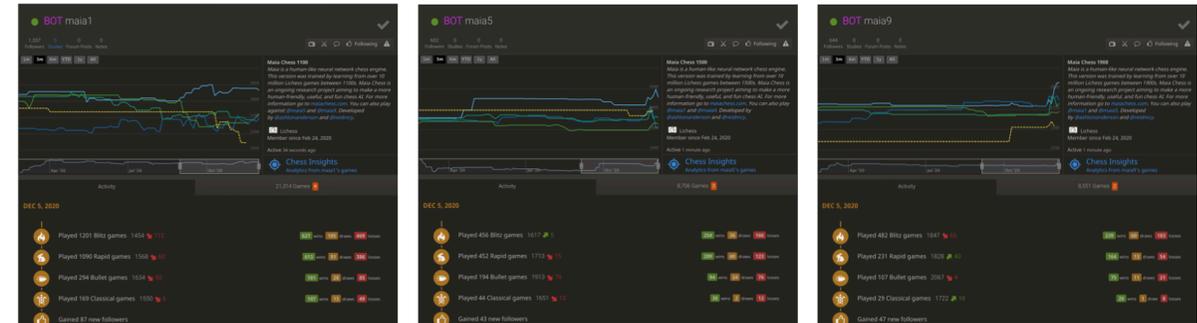
Summary

Maia is a human-like neural network chess engine



Maia captures human style in chess at targeted skill levels

You can play Maia online now!



Next: algorithmic and data-driven improvement tools

Thanks!

“Aligning Superhuman AI with Human Behavior: Chess as a Model System”
Knowledge Discovery and Data Mining (KDD), 2020.

<http://www.cs.toronto.edu/~ashton/pubs/maia-kdd2020.pdf>

maiachess.com

@maia1 @maia5 @maia9 on Lichess



Reid McIlroy-Young
University of Toronto



Siddhartha Sen
Microsoft Research



Jon Kleinberg
Cornell University



Ashton Anderson
University of Toronto