Teaching AI to Play Chess Like People

Austin A. Robinson
University of Minnesota - Morris
Computer Science Senior Seminar
Spring 2021

"If you program a machine, you know what it's capable of. If the machine is programming itself, who knows what it might do?"

— Garry Kasparov, Deep Thinking: Where Machine Intelligence Ends and Human Creativity Begins

- Computer's have fundamentally changed chess
- Computers can serve a different purpose

https://www.pri.org/stories/2018-01-05/garry-kasparov-and-game-artificial-intelligence

Outline

- Background
- AlphaZero
- Maia
    - Development
    - Results
    - Comparison
- Application
- Conclusion

Background: Chess Terminology

- **Elo rating**
- Blunder
- Time Control

https://ratings.fide.com/profile/1503014

Background: Chess Terminology

- Elo rating
- **Blunder**
- Time Control

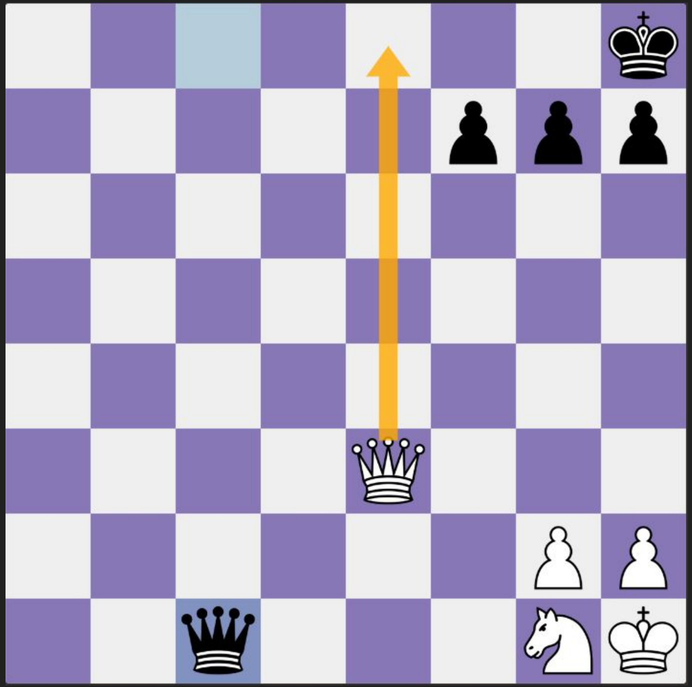

Background: Chess Terminology

- Elo rating
- Blunder
- **Time Control**

https://upload.wikimedia.org/wikipedia/commons/d/d3/Schachuhr_mechanisch.jpg

Background: Deep Learning

- **Machine Learning**
    - “The hierarchy of concepts enables the computer to learn complicated concepts by building them out of simpler ones. If we draw a graph showing how these concepts are built on top of each other, the graph is deep, with many layers. For this reason,we call this approach to AI deep learning."
      - *Deep Learning*, by Ian Goodfellow, Yoshua Bengio and Aaron Courville
- Deep Learning is allowing computers to learn from experience
- Reinforcement learning

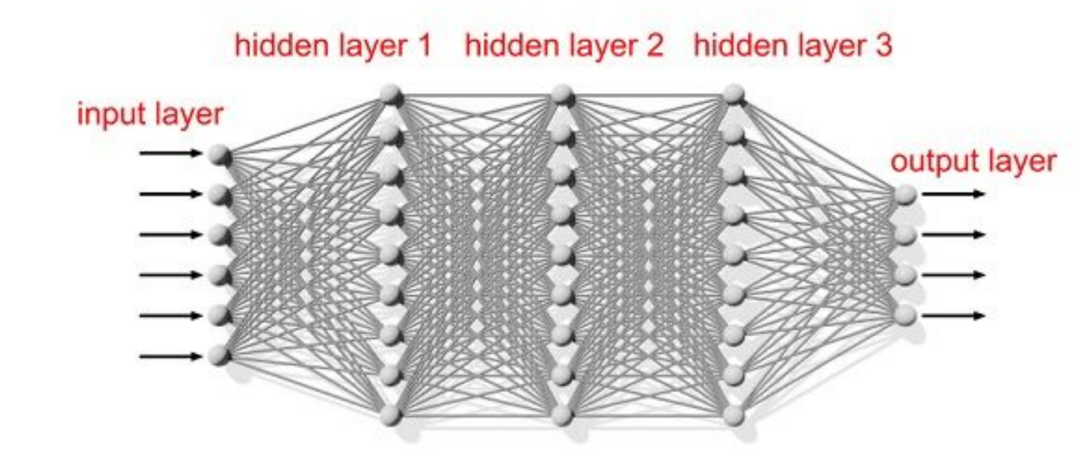

https://insidebigdata.com/wp-content/uploads/2018/10/Neural-Network-diagram.jpg

Background: Deep Learning

- Machine Learning
  - “The hierarchy of concepts enables the computer to learn complicated concepts by building them out of simpler ones. If we draw a graph showing how these concepts are built on top of each other, the graph is deep, with many layers. For this reason,we call this approach to AI deep learning."
    - *Deep Learning*, by Ian Goodfellow, Yoshua Bengio and Aaron Courville
- **Deep Learning is allowing computers to learn from experience**
- Reinforcement learning

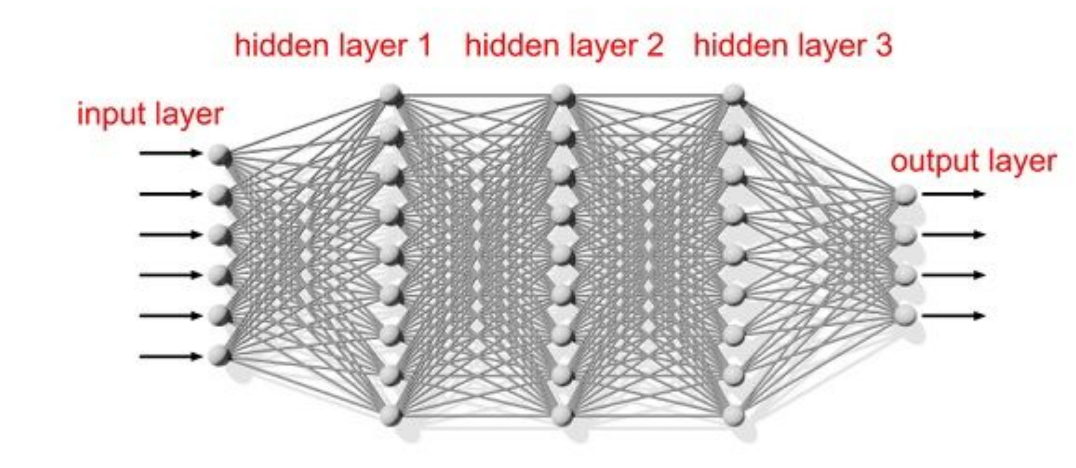

https://insidebigdata.com/wp-content/uploads/2018/10/Neural-Network-diagram.jpg

Background: Deep Learning

- Machine Learning
    - “The hierarchy of concepts enables the computer to learn complicated concepts by building them out of simpler ones. If we draw a graph showing how these concepts are built on top of each other, the graph is deep, with many layers. For this reason,we call this approach to AI deep learning."
      - *Deep Learning*, by Ian Goodfellow, Yoshua Bengio and Aaron Courville
- Deep Learning is allowing computers to learn from experience
- **Reinforcement learning**

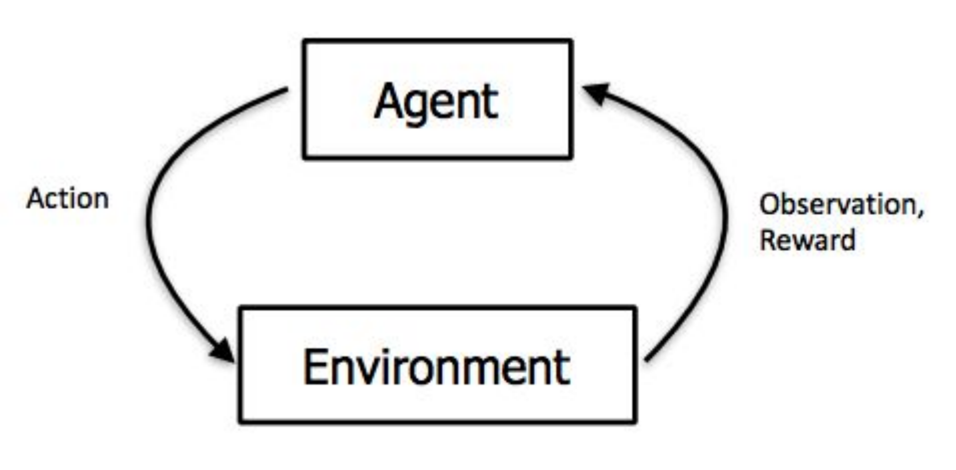

https://miro.medium.com/max/478/1*QVsnwatDVz8wcqJUsLJejw.png

Background: Monte Carlo Tree Search(MCTS)

- **MCTS is a tree search that also implements machine learning principles of reinforcement learning**
- 4 primary steps
  - Simulation

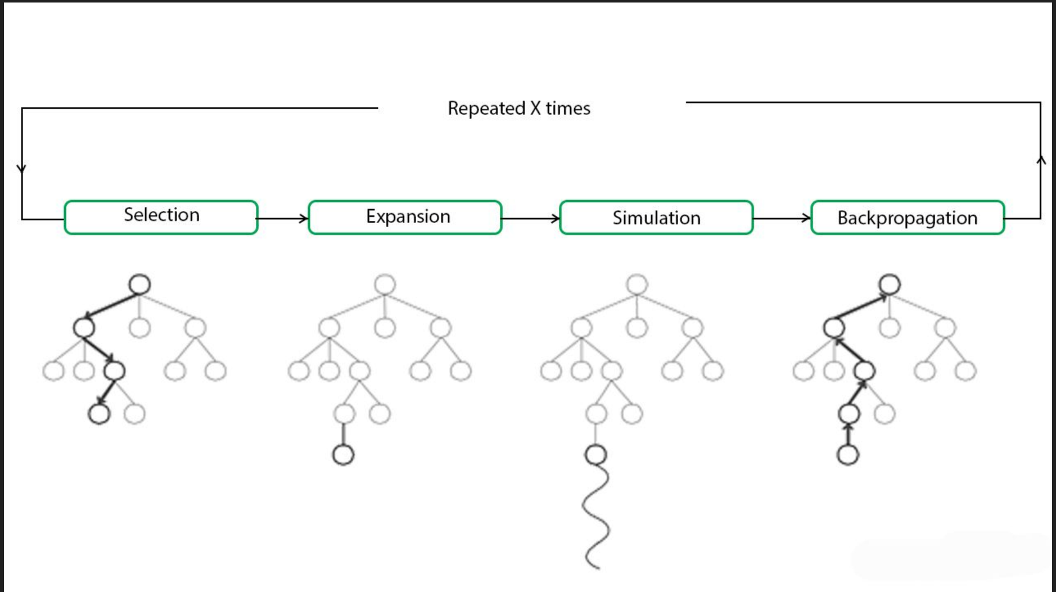

[3]

Background: Monte Carlo Tree Search(MCTS)

- MCTS is a tree search that also implements machine learning principles of reinforcement learning
- **4 primary steps**
    - **Simulation**

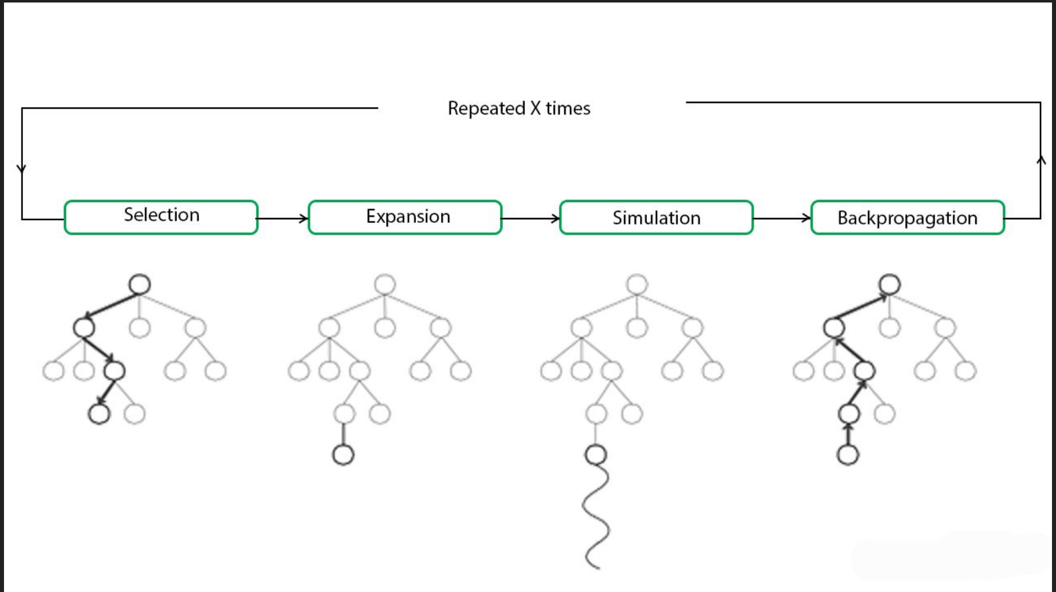

[3]

Background: Chess Engines

- **Software that is used to generate and analyse positions**
- Stockfish
  - More traditional chess engine
  - 3564
- AlphaZero and Leela
  - Deep Learning
  - Monte Carlo Tree Search
  - 3463

Background: Chess Engines

- Software that is used to generate and analyse positions
- **Stockfish**
  - **More traditional chess engine**
  - **3564**
- AlphaZero and Leela
  - Deep Learning
  - Monte Carlo Tree Search
  - 3463

https://stockfishchess.org/

Background: Chess Engines

- Software that is used to generate and analyse positions
- Stockfish
  - More traditional chess engine
  - 3564
- **AlphaZero and Leela**
  - **Deep Learning**
  - **Monte Carlo Tree Search**
  - **3463**

https://lczero.org/blog/

AlphaZero

AlphaZero: Architecture

- **Inputs**
- Output
    - Probabilities
    - Expected outcome
    - Values
- Parameters
- Obtaining Outputs
    - Reinforcement through self play

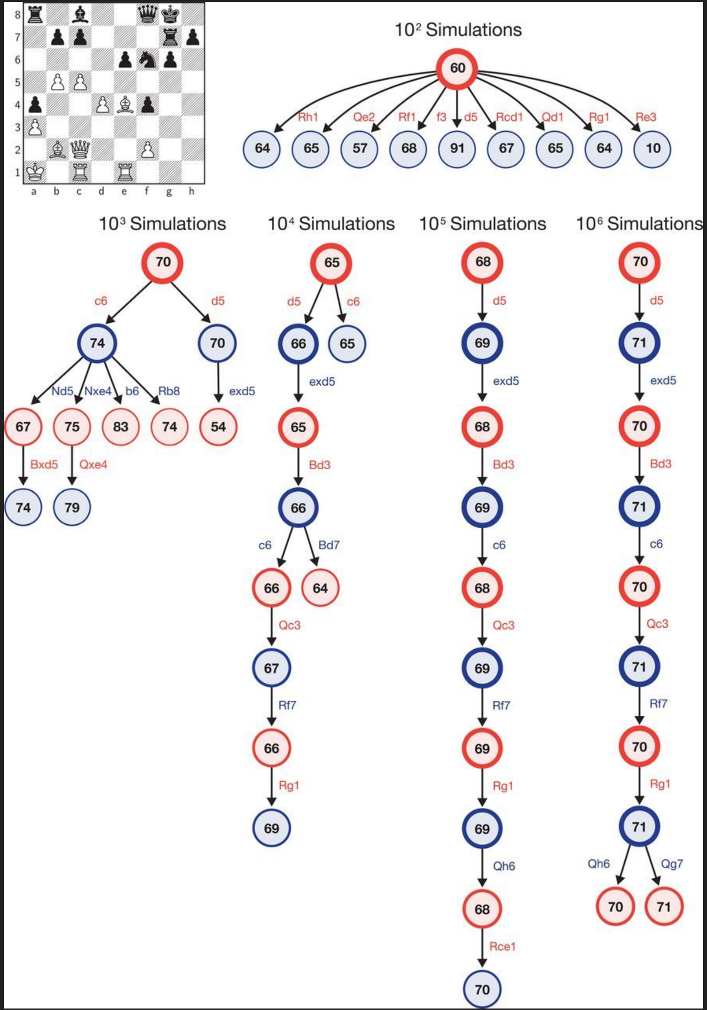

[2]

AlphaZero: Architecture

- Inputs
- **Output**
  - **Probabilities**
  - Expected outcomes
  - Values
- Parameters
- Obtaining Outputs
  - Reinforcement through self play

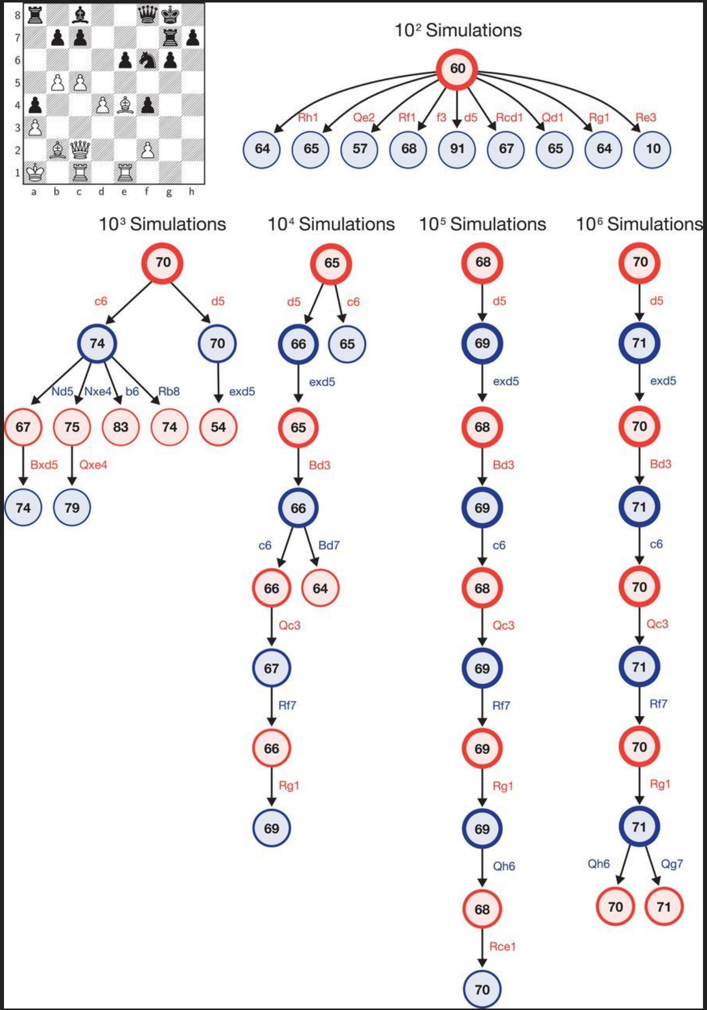

[2]

AlphaZero: Architecture

- Inputs
- **Output**
  - Probabilities
  - **Expected outcomes**
  - Values
- Parameters
- Obtaining Outputs
  - Reinforcement through self play

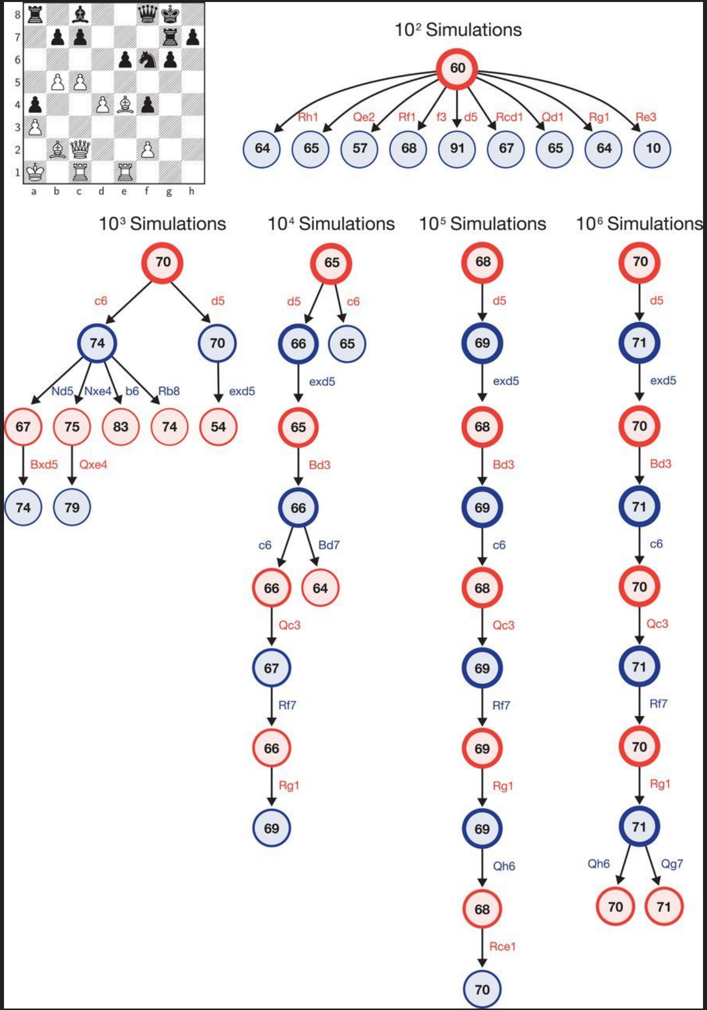

[2]

AlphaZero: Architecture

- Inputs
- **Output**
  - Probabilities
  - Expected outcomes
  - **Values**
- Parameters
- Obtaining Outputs
  - Reinforcement through self play

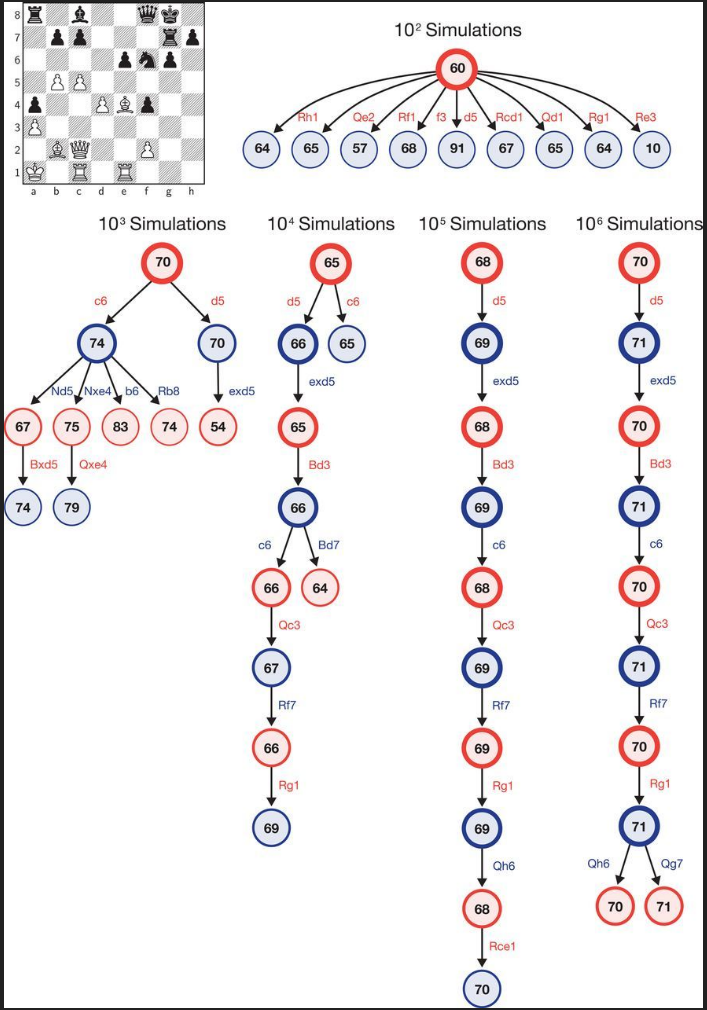

[2]

AlphaZero: Architecture

- Inputs
- Output
  - Probabilities
  - Expected outcomes
  - Values
- **Parameters**
- Obtaining Outputs
  - Reinforcement through self play

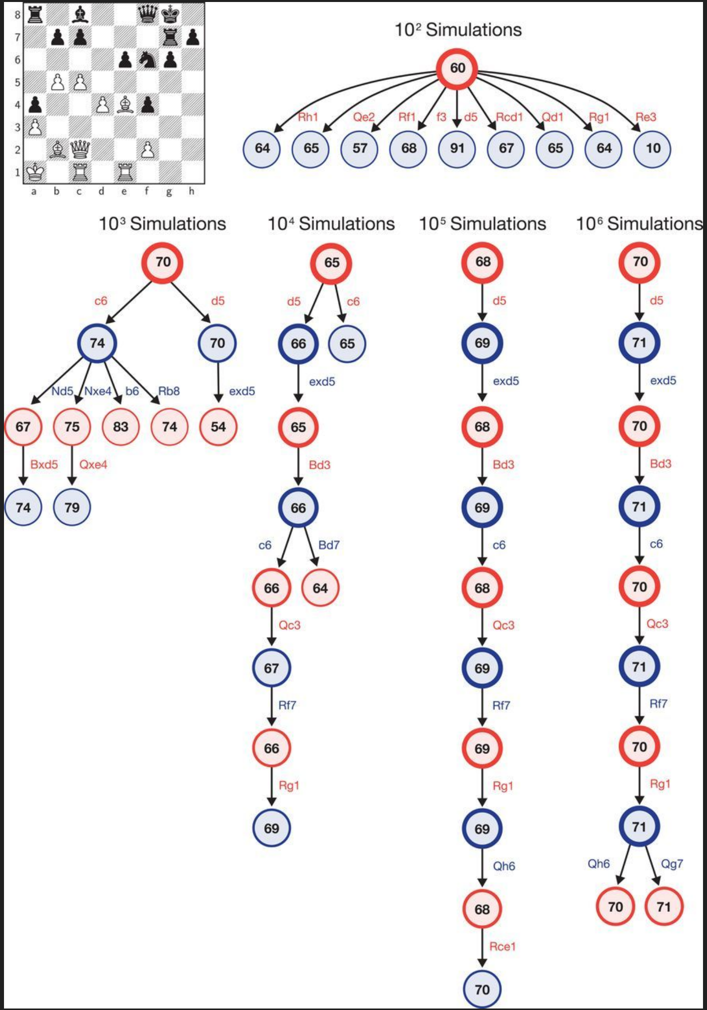

[2]

AlphaZero: Architecture

- Inputs
- Output
    - Probabilities
    - Expected outcomes
    - Values
- Parameters
- **Obtaining Outputs**
    - **Reinforcement through self play**

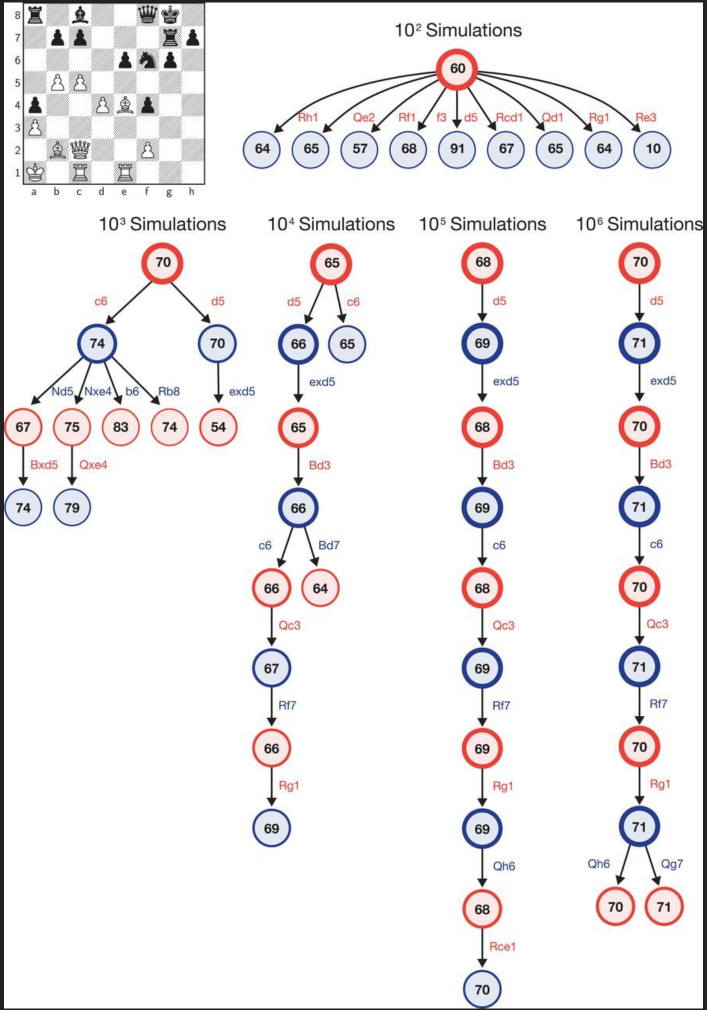

[2]

AlphaZero: Training

- **9 hours of training time**
  - **Reached an Elo rating of ~3200**
- Completely random at beginning of the training
- Learned from win, losses and draws

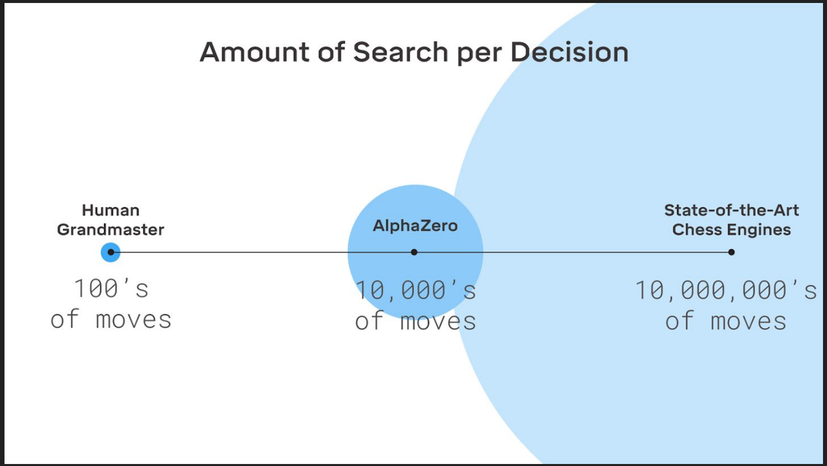

https://deepmind.com/blog/article/alphazero-shedding-new-light-grand-games-chess-shogi-and-go

AlphaZero: Training

- 9 hours of training time
  - Reached an Elo rating of ~3200
- **Completely random at beginning of the training**
- Learned from win, losses and draws

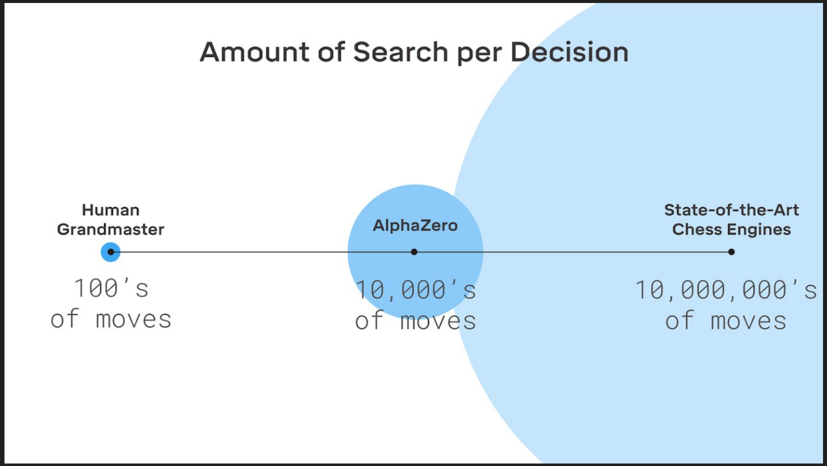

https://deepmind.com/blog/article/alphazero-shedding-new-light-grand-games-chess-shogi-and-go

AlphaZero: Training

- 9 hours of training time
  - Reached an Elo rating of ~3200
- Completely random at beginning of the training
- **Learned from win, losses and draws**

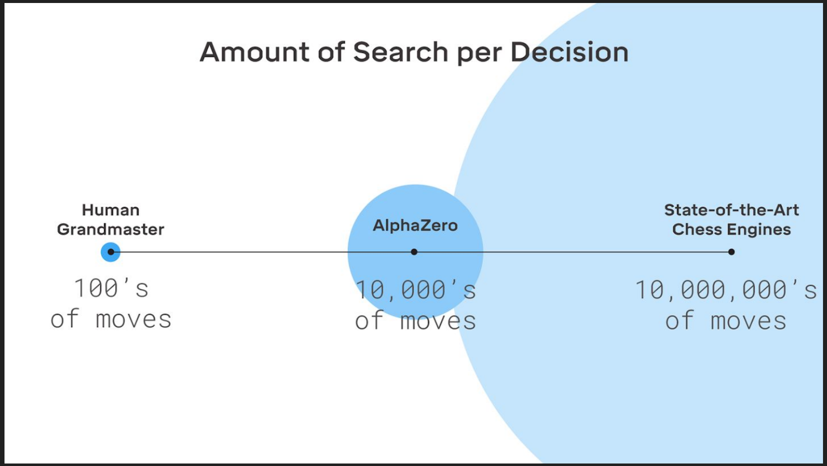

https://deepmind.com/blog/article/alphazero-shedding-new-light-grand-games-chess-shogi-and-go

AlphaZero: Results

- **Defeated Stockfish (won 155 games, drew 839 games and lost 6 games)**
  - **Leela was also able to defeat the same version of StockFish**
- Strategies Learned by AlphaZero
  - Common human strategies
  - Unique to AlphaZero

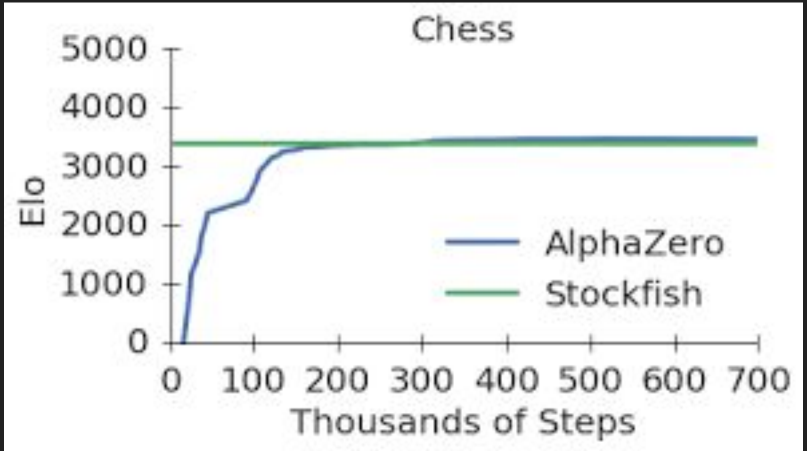

[3]

AlphaZero: Results

- Defeated Stockfish (won 155 games, drew 839 games and lost 6 games)
  - Leela was also able to defeat the same version of StockFish
- **Strategies Learned by AlphaZero**
  - **Common human strategies**
  - **Unique to AlphaZero**

ichess.net

Maia

Maia: Development

- **The Goal of Maia is to play the most like a human**
- Maia utilizes a large amount of code from Leela
- Data Sets
- Move Prediction
- Models of Maia

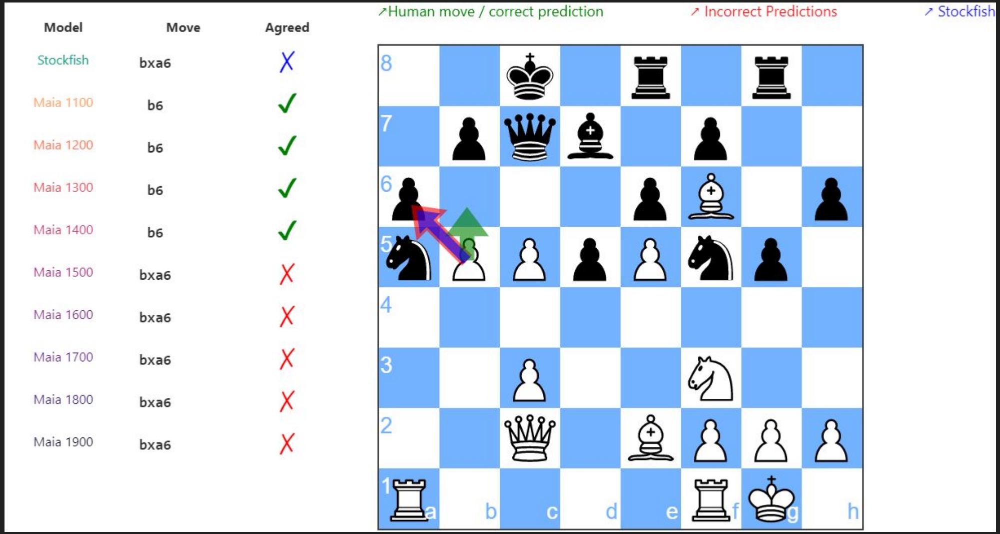

Model	Move	Agreed
Stockfish	bxa6	X
Maia 1100	b6	✓
Maia 1200	b6	✓
Maia 1300	b6	✓
Maia 1400	b6	✓
Maia 1500	bxa6	X
Maia 1600	bxa6	X
Maia 1700	bxa6	X
Maia 1800	bxa6	X
Maia 1900	bxa6	X

[4]

Maia: Development

- The Goal of Maia is to play the most like a human
- **Maia utilizes a large amount of code from Leela**
- Data Sets
- Move Prediction
- Models of Maia

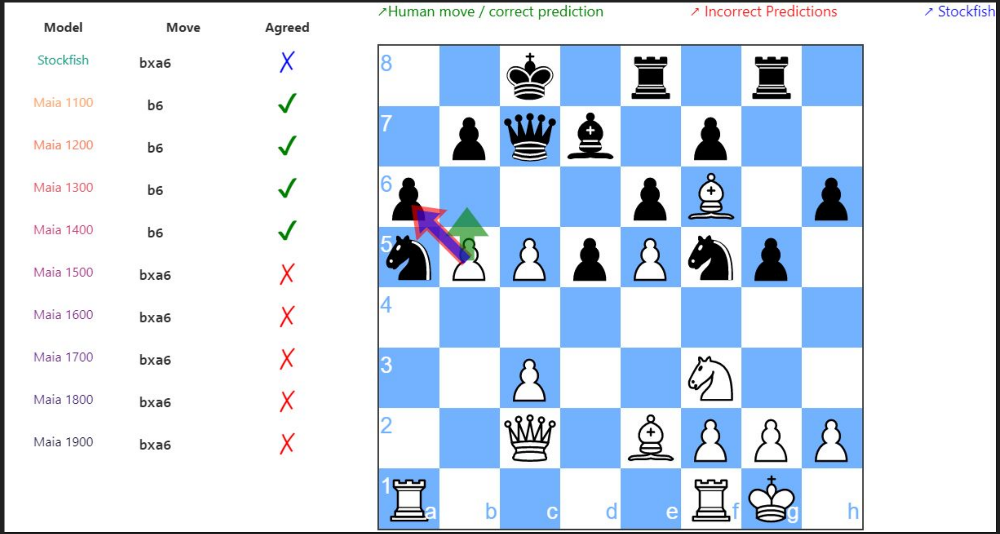

[4]

Maia: Development

- The Goal of Maia is to play the most like a human
- Maia utilizes a large amount of code from Leela
- **Data Sets**
- Move Prediction
- Models of Maia

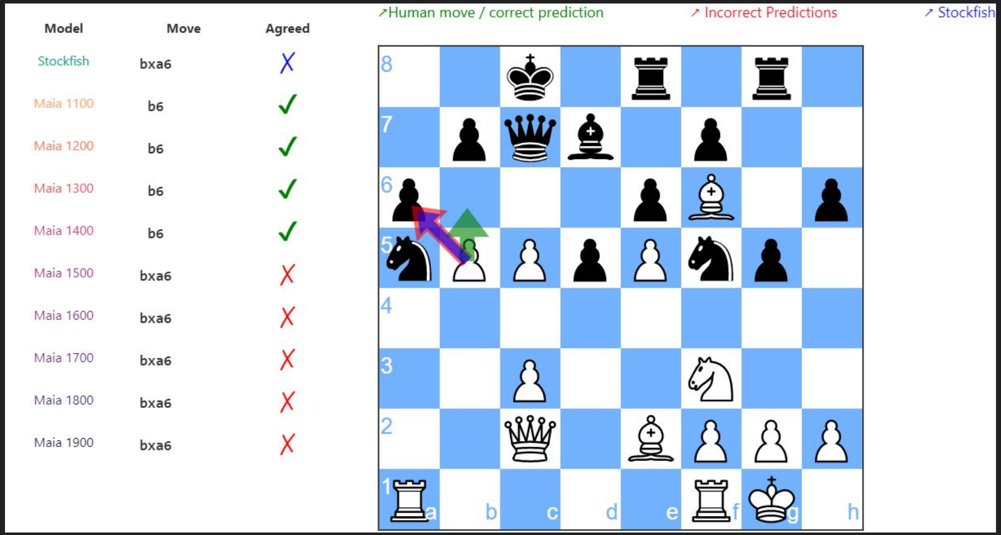

Model	Move	Agreed
Stockfish	bxa6	X
Maia 1100	b6	✓
Maia 1200	b6	✓
Maia 1300	b6	✓
Maia 1400	b6	✓
Maia 1500	bxa6	X
Maia 1600	bxa6	X
Maia 1700	bxa6	X
Maia 1800	bxa6	X
Maia 1900	bxa6	X

[4]

Maia: Development

- The Goal of Maia is to play the most like a human
- Maia utilizes a large amount of code from Leela
- Data Sets
- **Move Prediction**
- Models of Maia

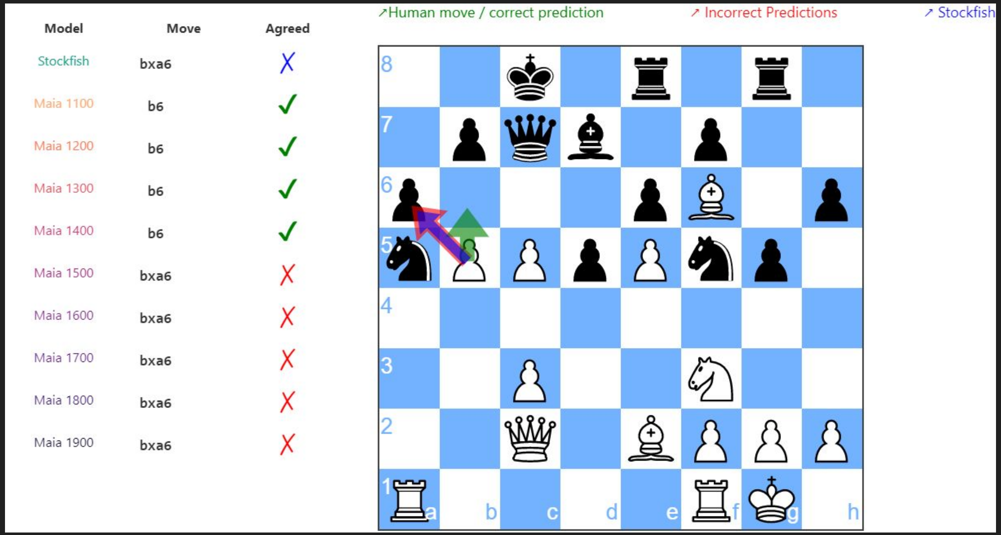

Model	Move	Agreed
Stockfish	bxa6	X
Maia 1100	b6	✓
Maia 1200	b6	✓
Maia 1300	b6	✓
Maia 1400	b6	✓
Maia 1500	bxa6	X
Maia 1600	bxa6	X
Maia 1700	bxa6	X
Maia 1800	bxa6	X
Maia 1900	bxa6	X

[4]

Maia: Development

- The Goal of Maia is to play the most like a human
- Maia utilizes a large amount of code from Leela
- Data Sets
- Move Prediction
- **Models of Maia**

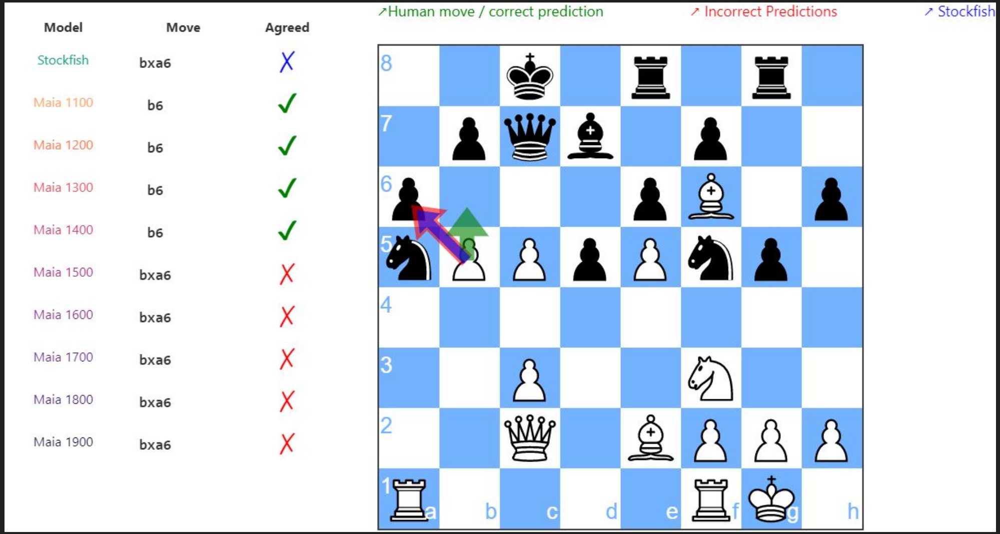

Model	Move	Agreed
Stockfish	bxa6	X
Maia 1100	b6	✓
Maia 1200	b6	✓
Maia 1300	b6	✓
Maia 1400	b6	✓
Maia 1500	bxa6	X
Maia 1600	bxa6	X
Maia 1700	bxa6	X
Maia 1800	bxa6	X
Maia 1900	bxa6	X

[4]

Maia: Results

- **Overall results should an accuracy around 50%**
- Maia was able to show that different ELO ranges have unique play styles
- Everyone plays Chess differently; Maia is the average player for an ELO range
- Stockfish
- Leela

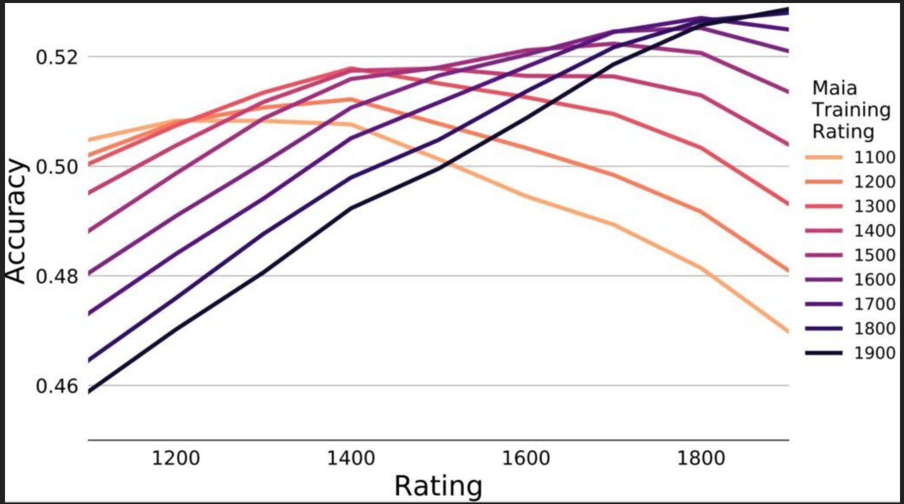

[1]

Maia: Results

- Overall results should an accuracy around 50%
- **Maia was able to show that different ELO ranges have unique play styles**
- Everyone plays Chess differently; Maia is the average player for an ELO range
- Stockfish
- Leela

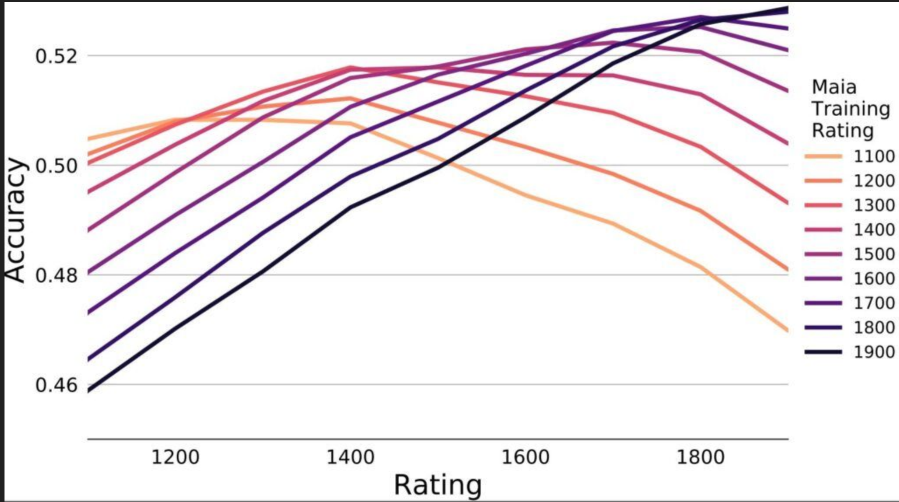

[1]

Maia: Results

- Overall results should an accuracy around 50%
- Maia was able to show that different ELO ranges have unique play styles
- **Everyone plays Chess differently; Maia is the average player for an ELO range**
- Stockfish
- Leela

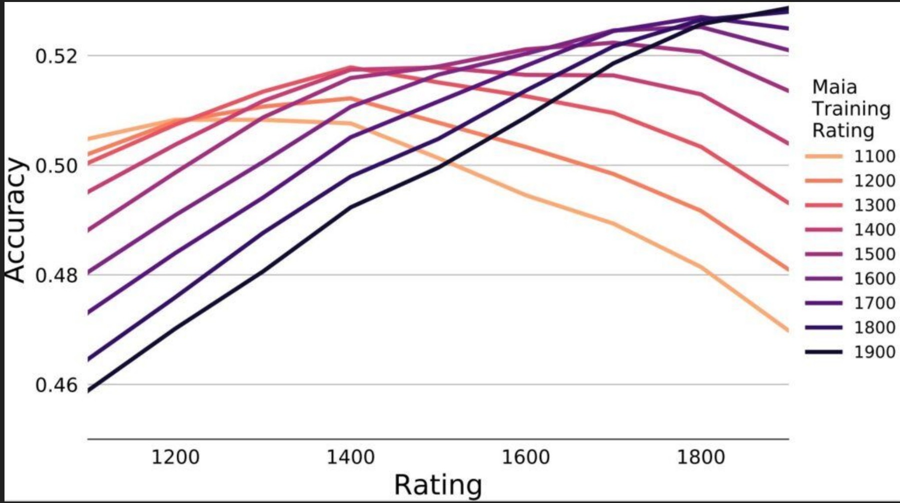

[1]

Maia: Results

- Overall results should an accuracy around 50%
- Maia was able to show that different ELO ranges have unique play styles
- Everyone plays Chess differently; Maia is the average player for an ELO range
- **Stockfish**
- Leela

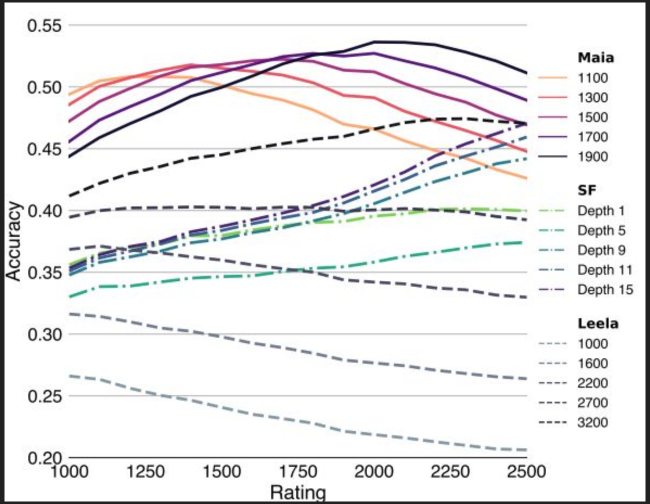

[1]

Maia: Results

- Overall results should an accuracy around 50%
- Maia was able to show that different ELO ranges have unique play styles
- Everyone plays Chess differently; Maia is the average player for an ELO range
- Stockfish
- **Leela**

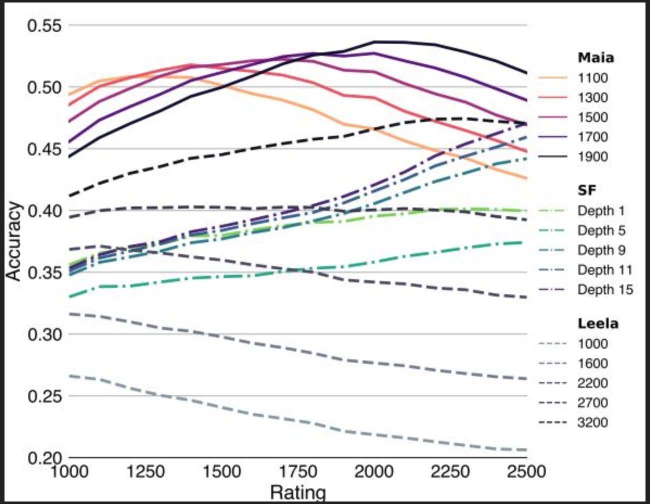

[1]

Applications

- **In Chess**
  - **Hundreds of Thousand of daily online chess players**
  - Median Elo Rating 1500
  - Training
- Human and AI interactions
  - Harvard two-player Atari study
  - Through Deep Learning were able to produce AI that helped improve human performance

Applications

- **In Chess**
  - Hundreds of Thousand of daily online chess players
  - **Median Elo Rating(Lichess) 1500**
  - Training
- Human and AI interactions
  - Harvard two-player Atari study
  - Through Deep Learning were able to produce AI that helped improve human performance

Applications

- **In Chess**
  - Hundreds of Thousand of daily online chess players
  - Median Elo Rating 1500
  - **Training**
- Human and AI interactions
  - Harvard two-player Atari study
  - Through Deep Learning were able to produce AI that helped improve human performance

Applications

- In Chess
  - Hundreds of Thousand of daily online chess players
  - Median Elo Rating 1500
  - Training
- **Human and AI interactions**
  - **Harvard two-player Atari study**
  - **Through Deep Learning were able to produce AI that helped improve human performance**

Conclusion

- Maia shows it possible to capture the play style of people
- Development of more human like chess engines could lead to better training for people

Questions

Sources

[1] Aligning Superhuman AI with Human Behavior: Chess as a Model System
-Young, Sen, Kleinberg, Anderson

[2] Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm - Silver

[3] https://www.geeksforgeeks.org/ml-monte-carlo-tree-search-mcts/

[4] https://www.microsoft.com/en-us/research/blog/the-human-side-of-ai-for-chess/?OCID=msr_blog_MaiaChess_tw