

Heuristic Search in the Age of Deep Learning



Martin Müller


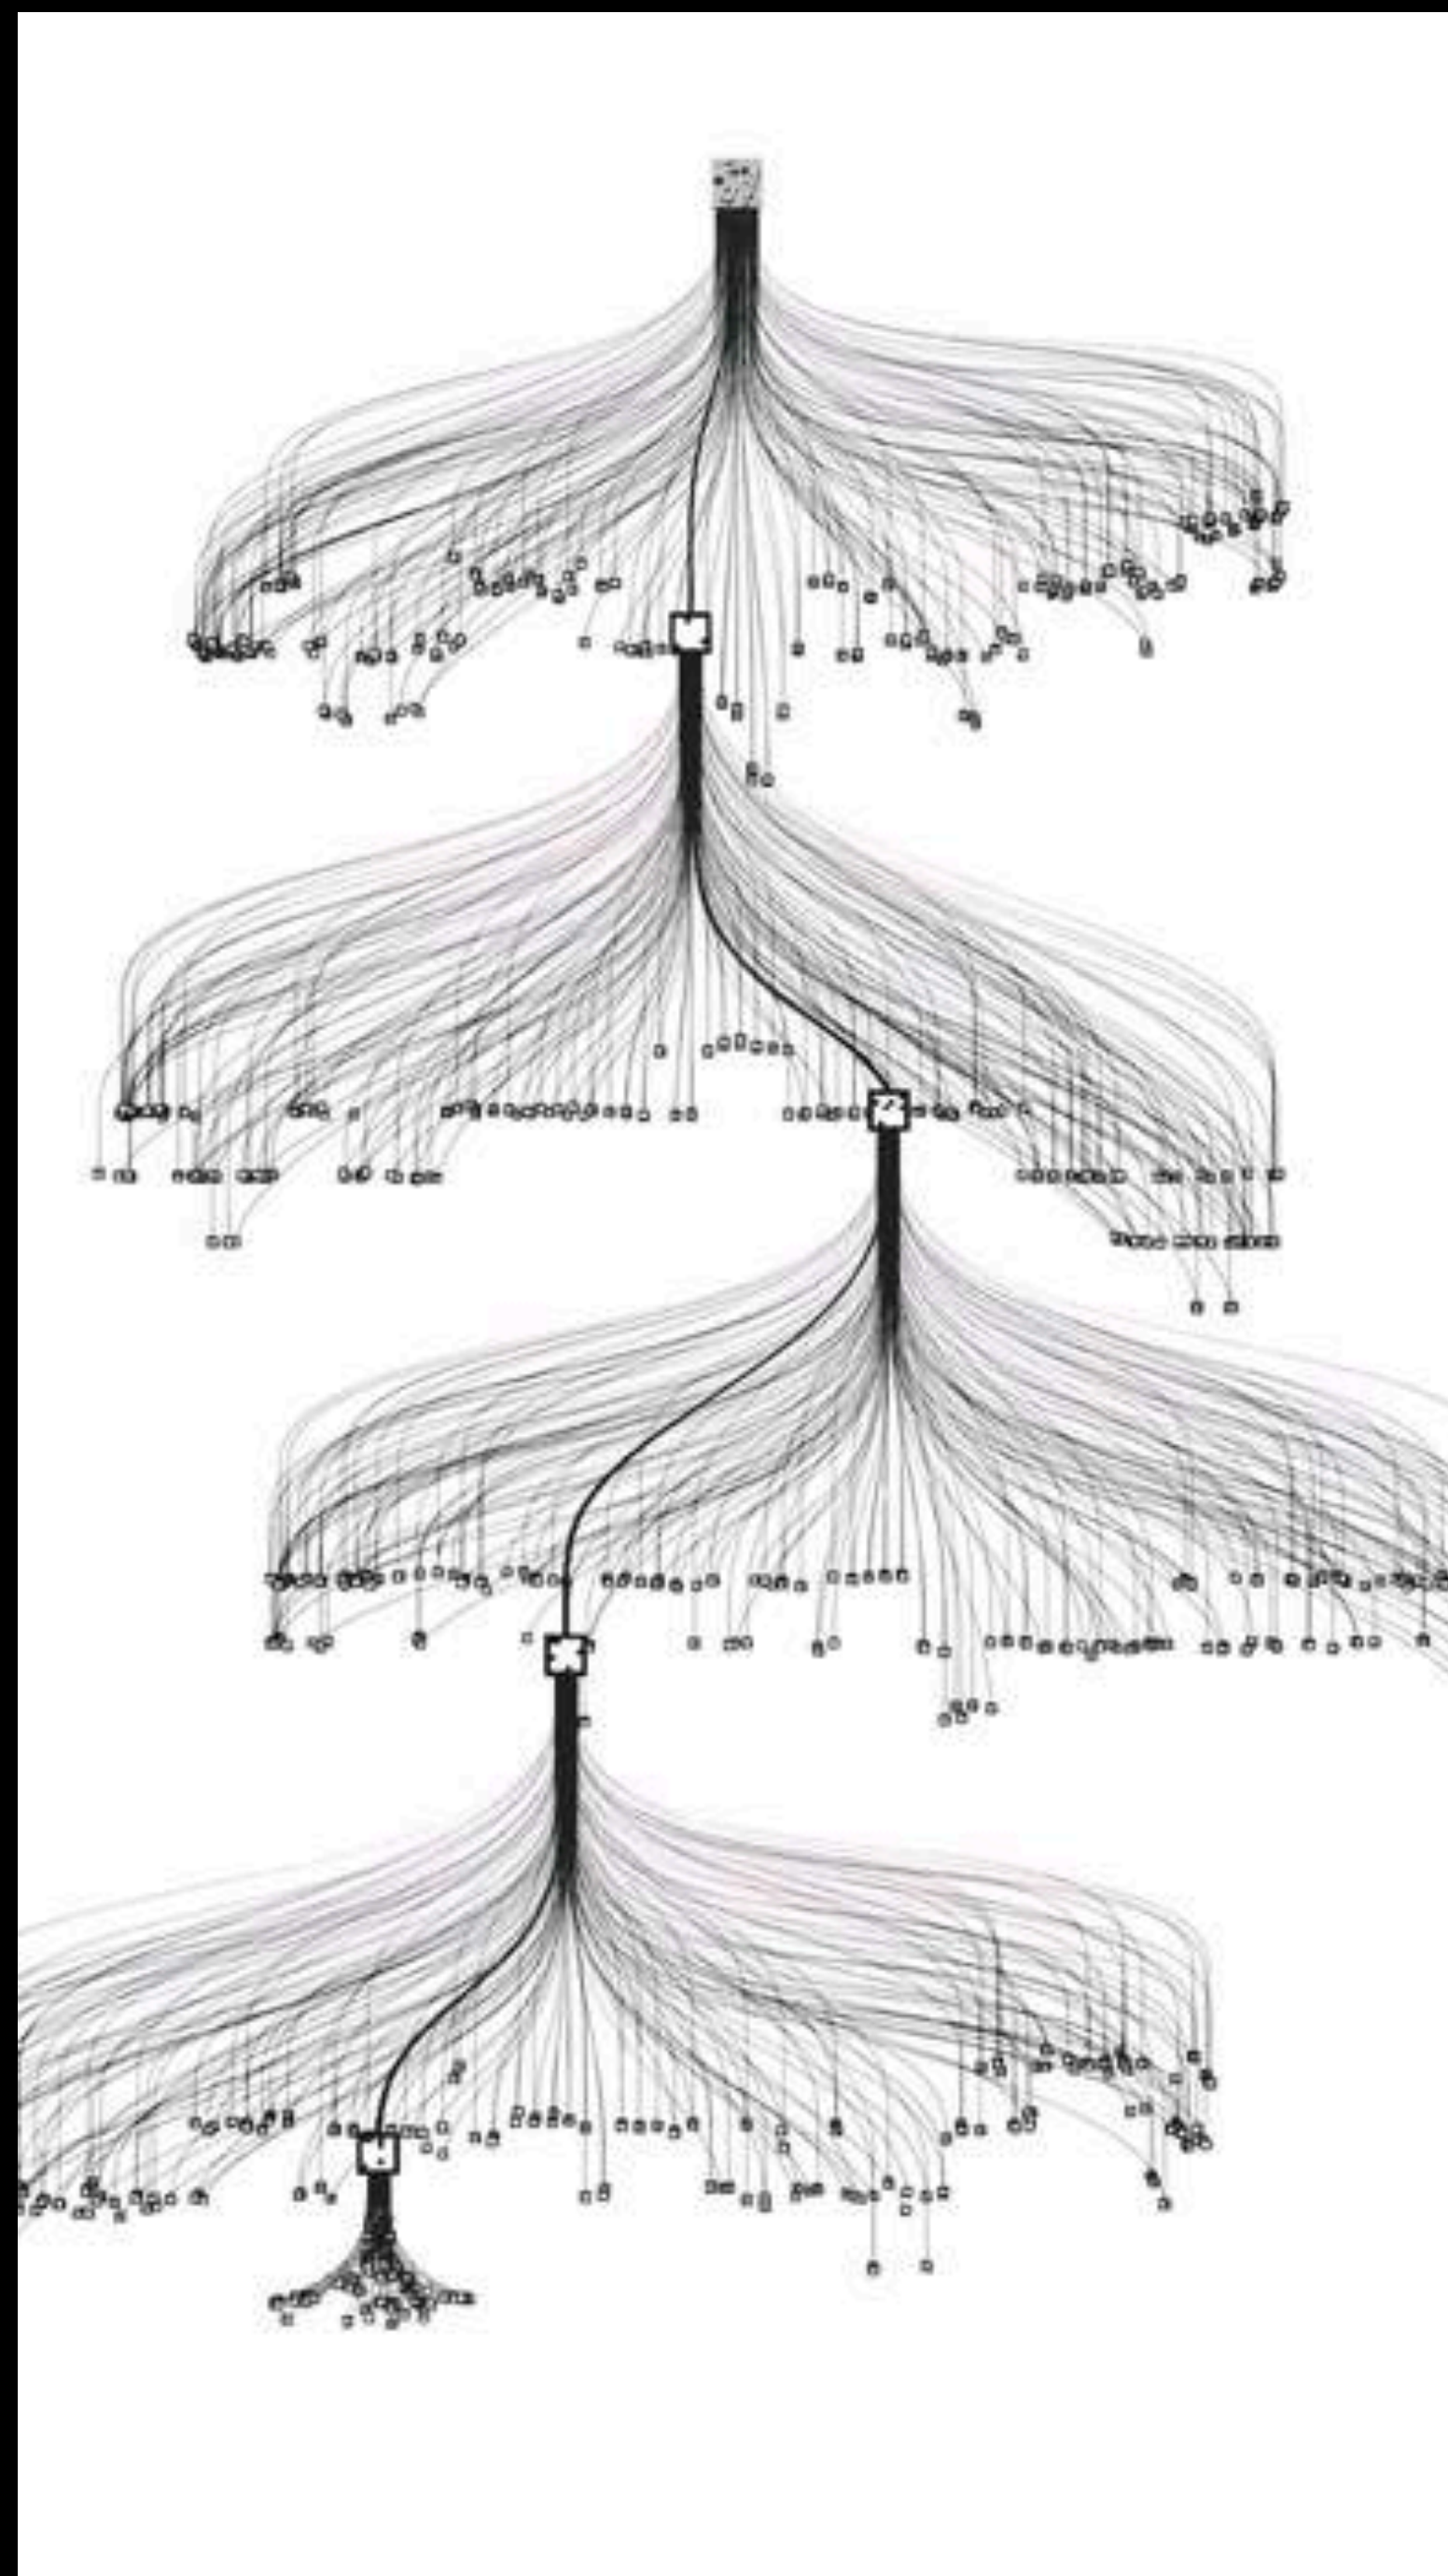


UNIVERSITY OF
ALBERTA

EDMONTON · ALBERTA · CANADA

Contents

- 一. Why Heuristic Search?
- 二. Past Successes
- 三. Modern Successes
- 四. Future



Why Heuristic Search?

Make complex decisions:

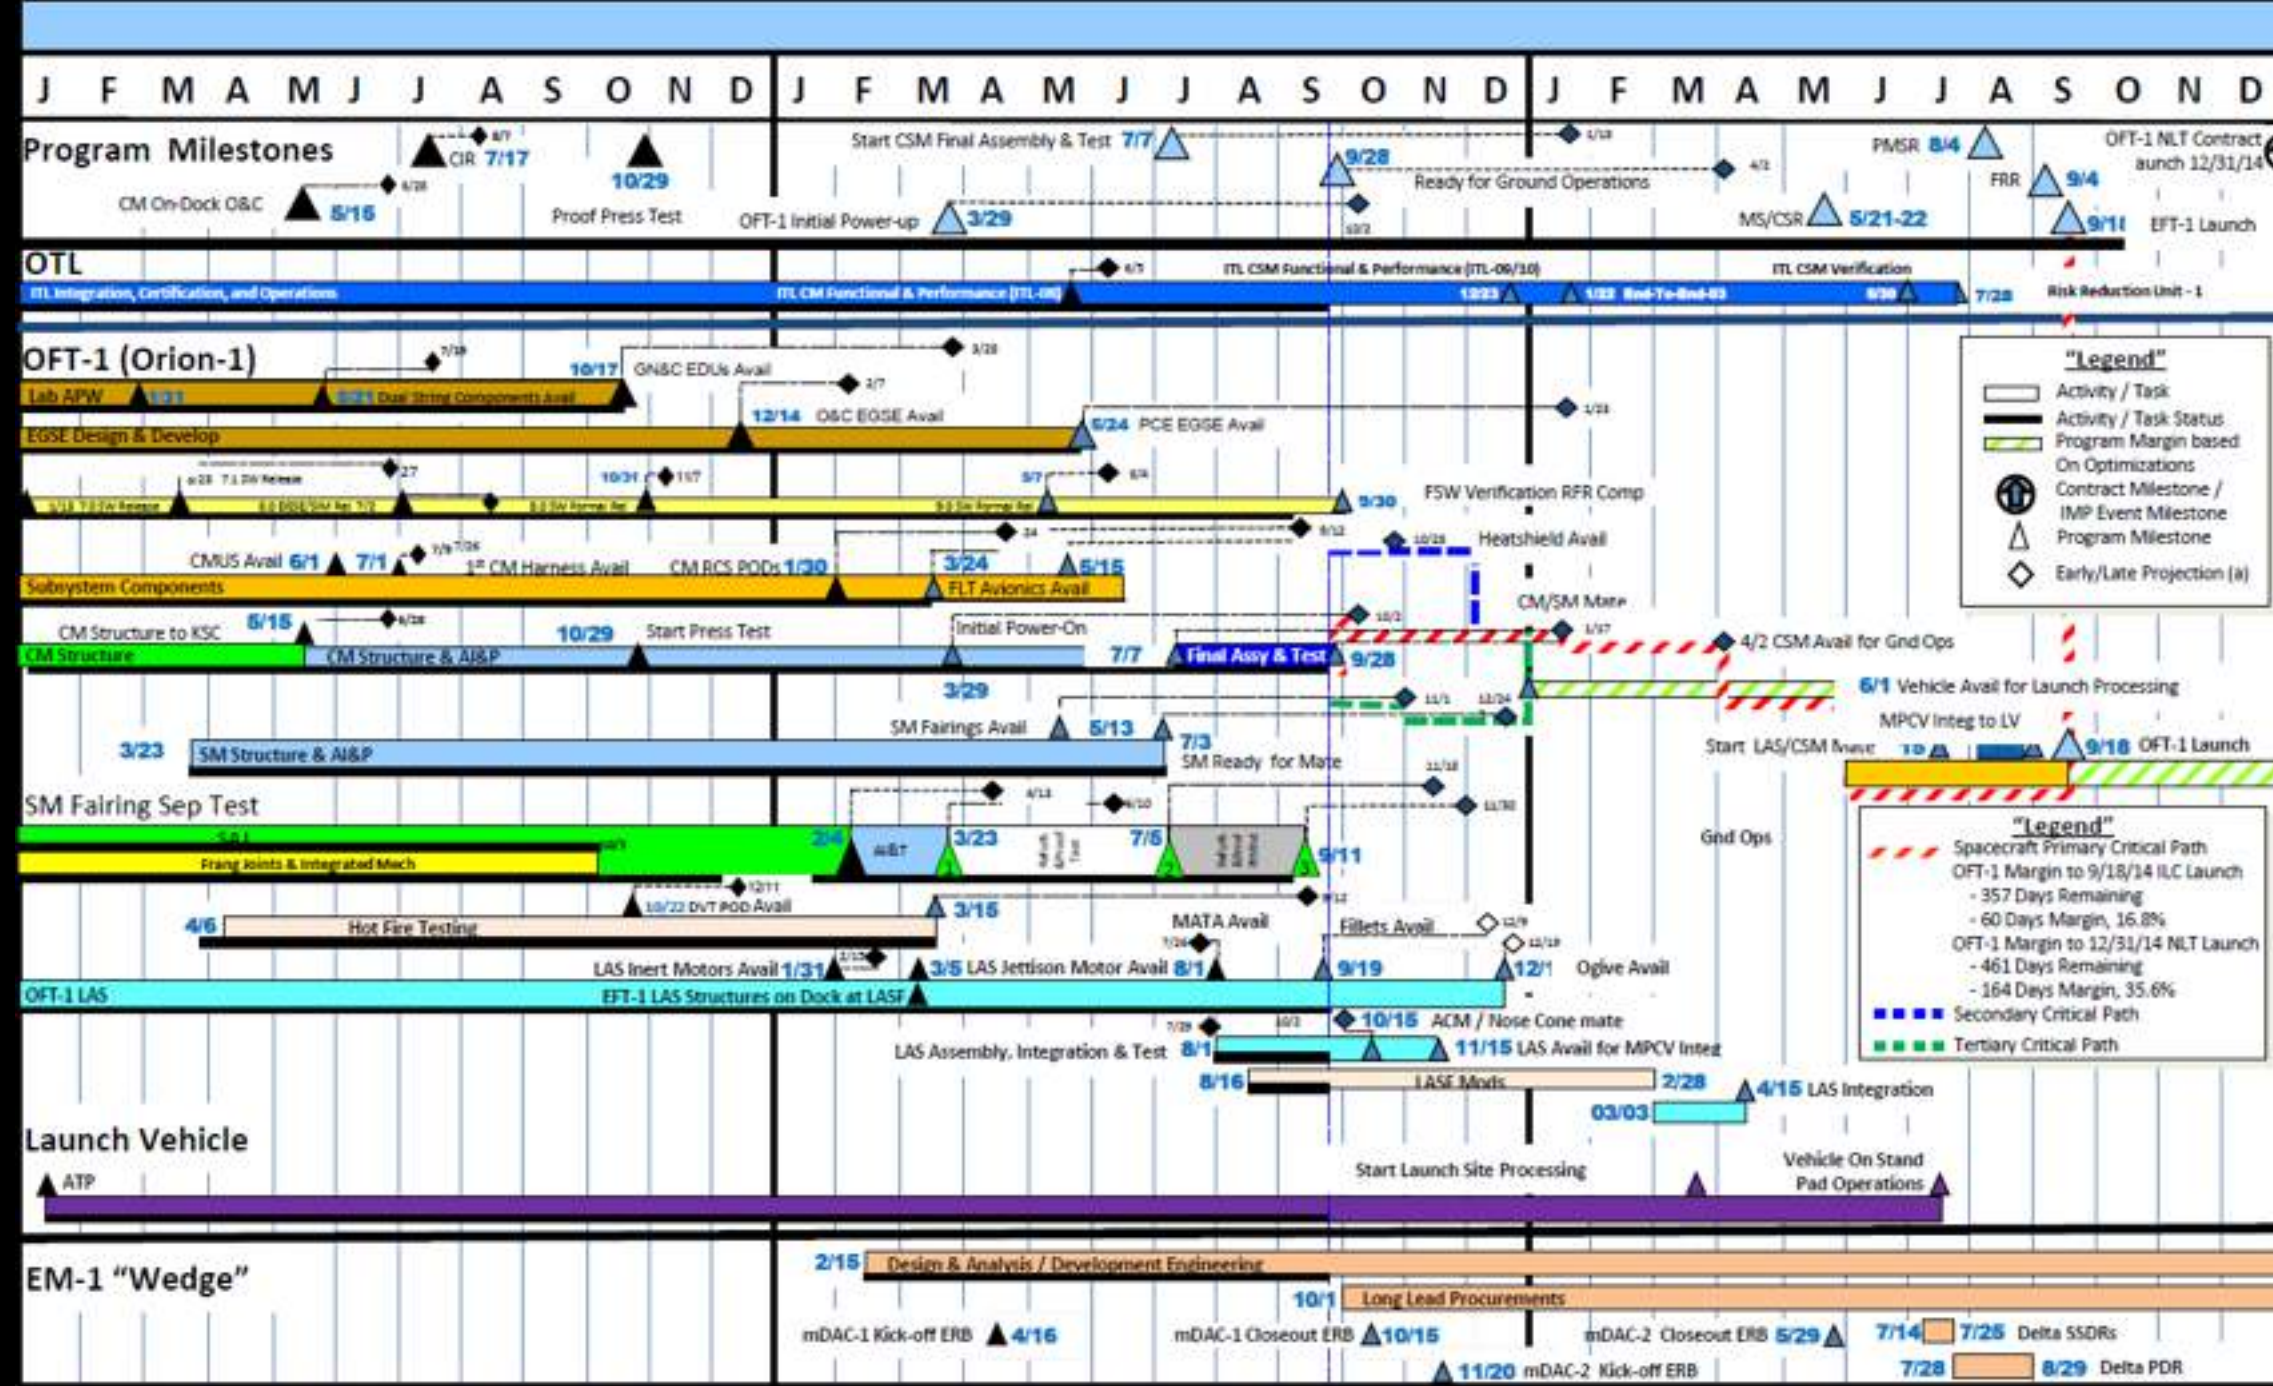
- Look ahead
- Evaluate different options
- Simulate future scenarios



Examples



- Prediction
- Planning
- **Games**



Why Games?

- Simple clear rules
- Try many new ideas
- Use results in real world
- Games are good business!



Past Successes

- 一. Search: minimax, alphabeta
- 二. Knowledge: handcoded rules
- 三. Perfect endgame databases
- 四. Case studies: **chess, checkers**
- 五. **TD-Gammon**: a glimpse of the future



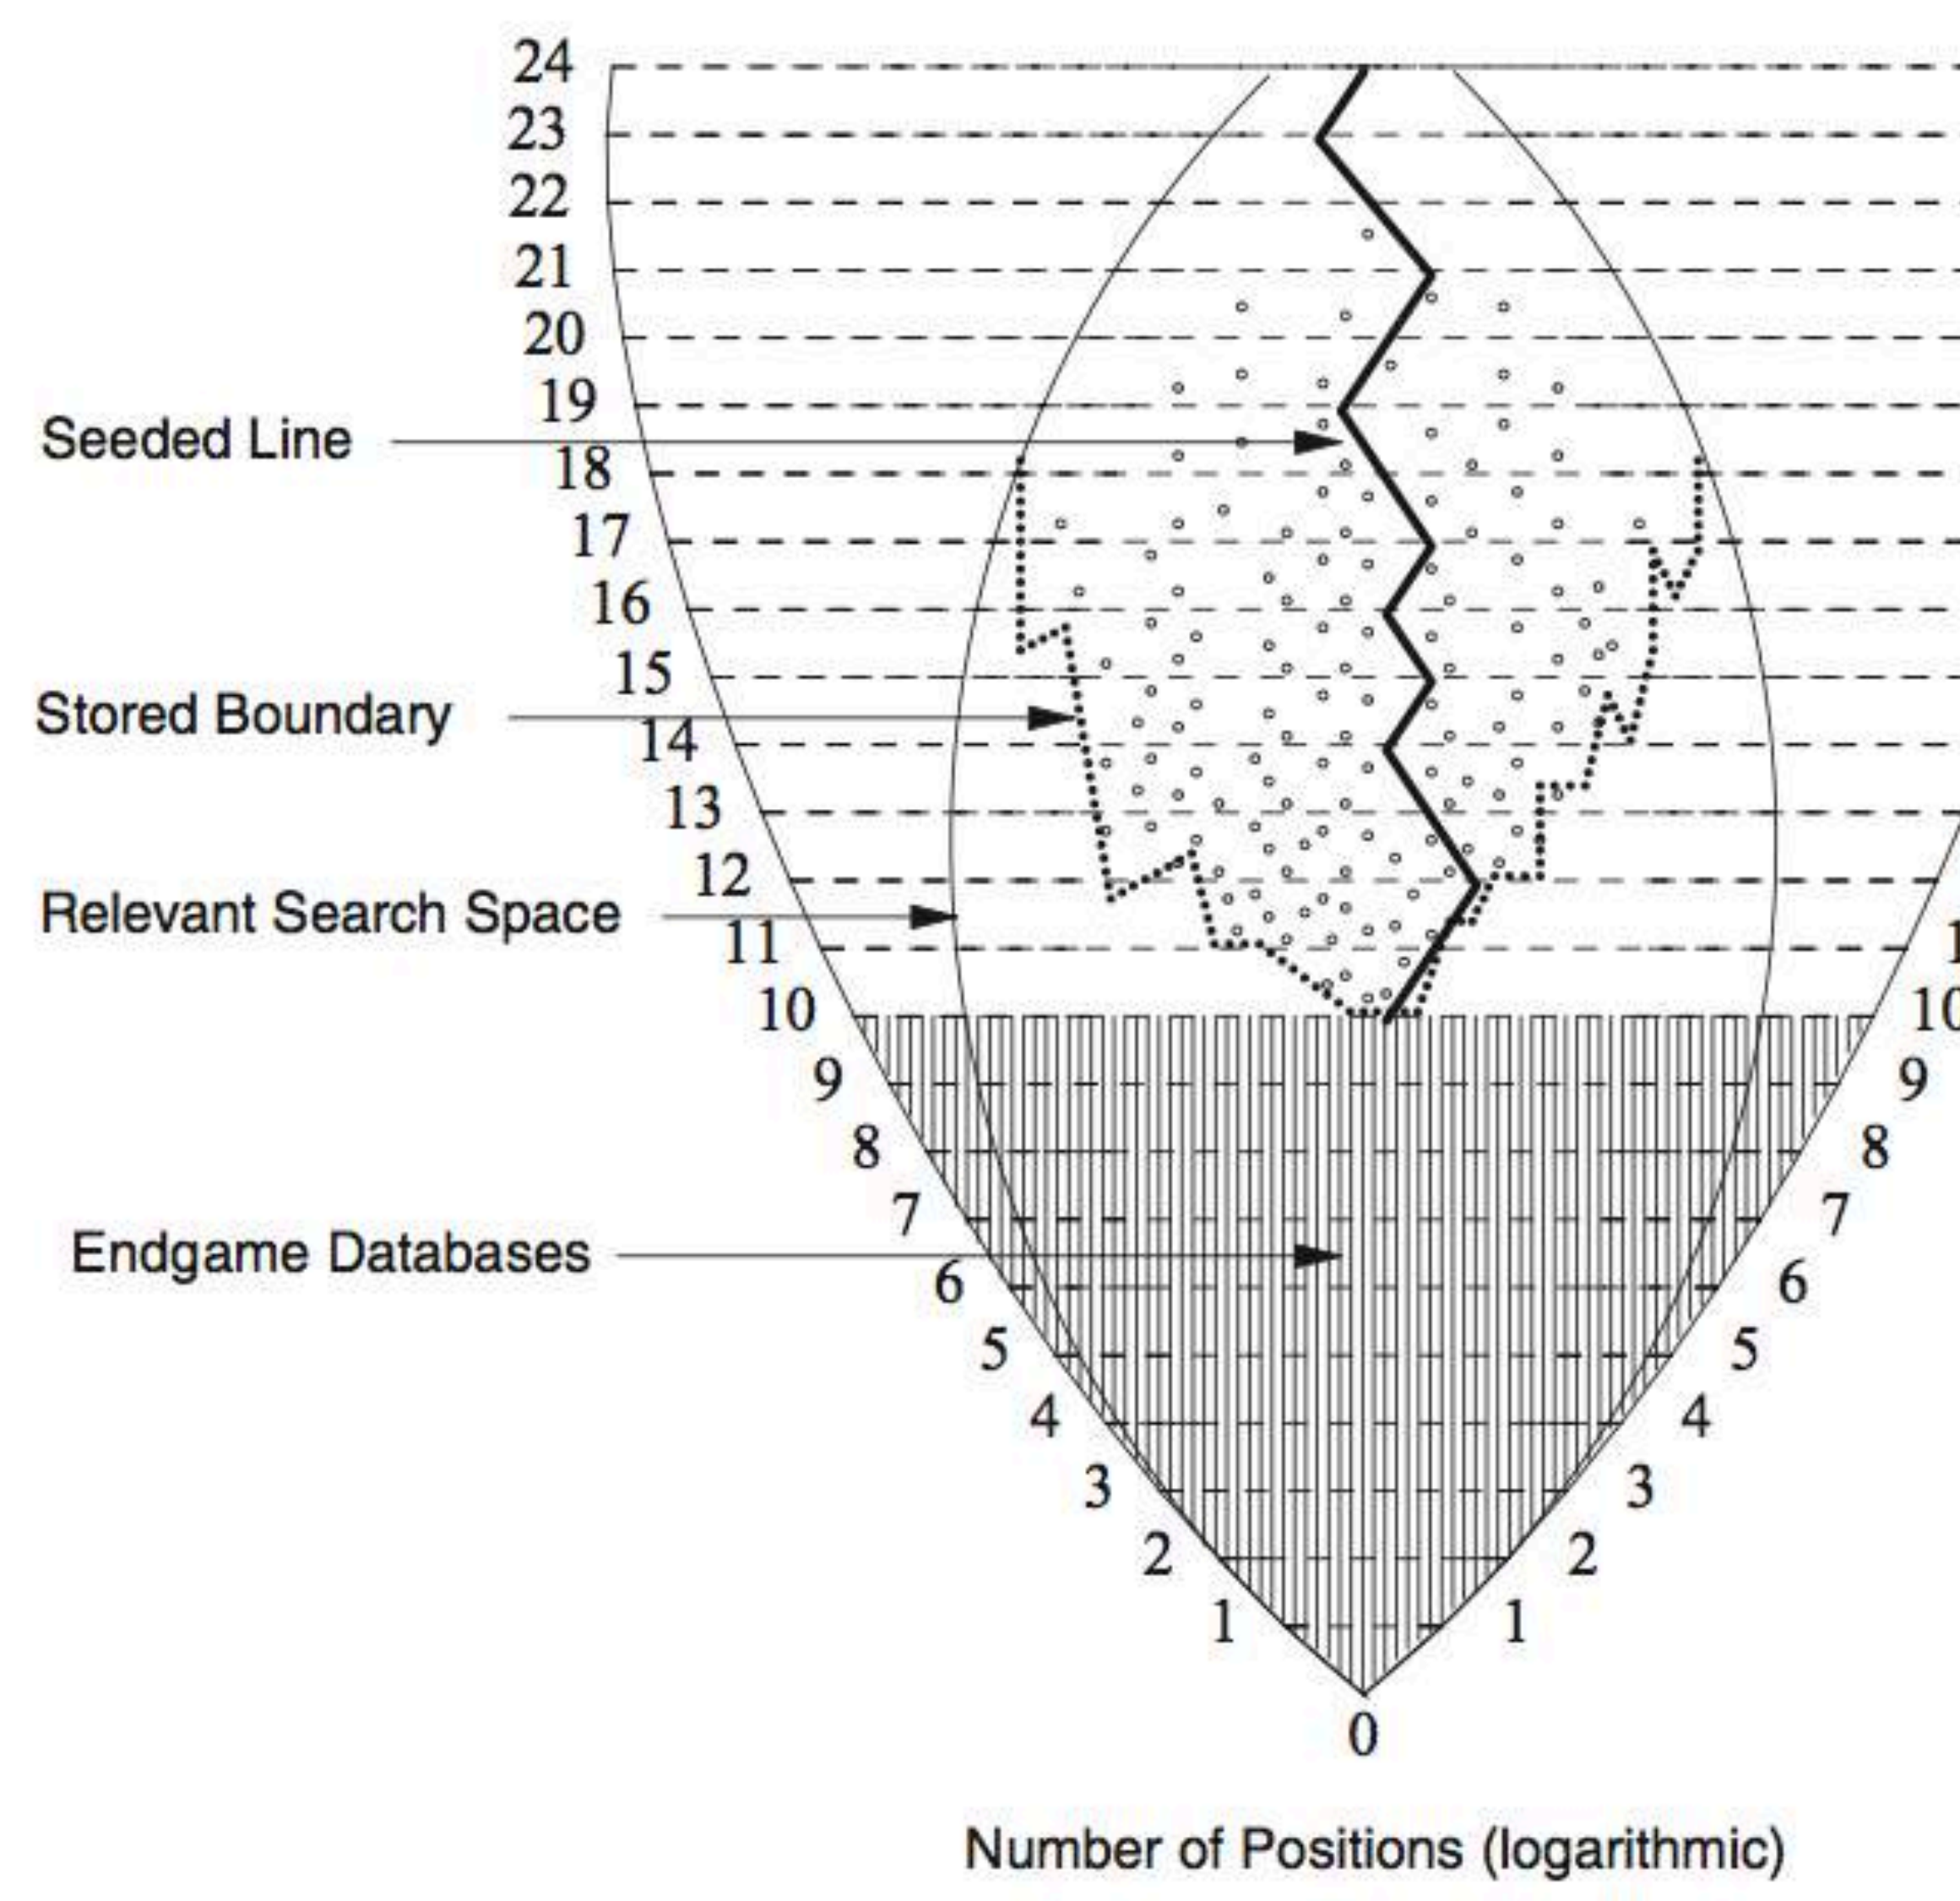


Chess

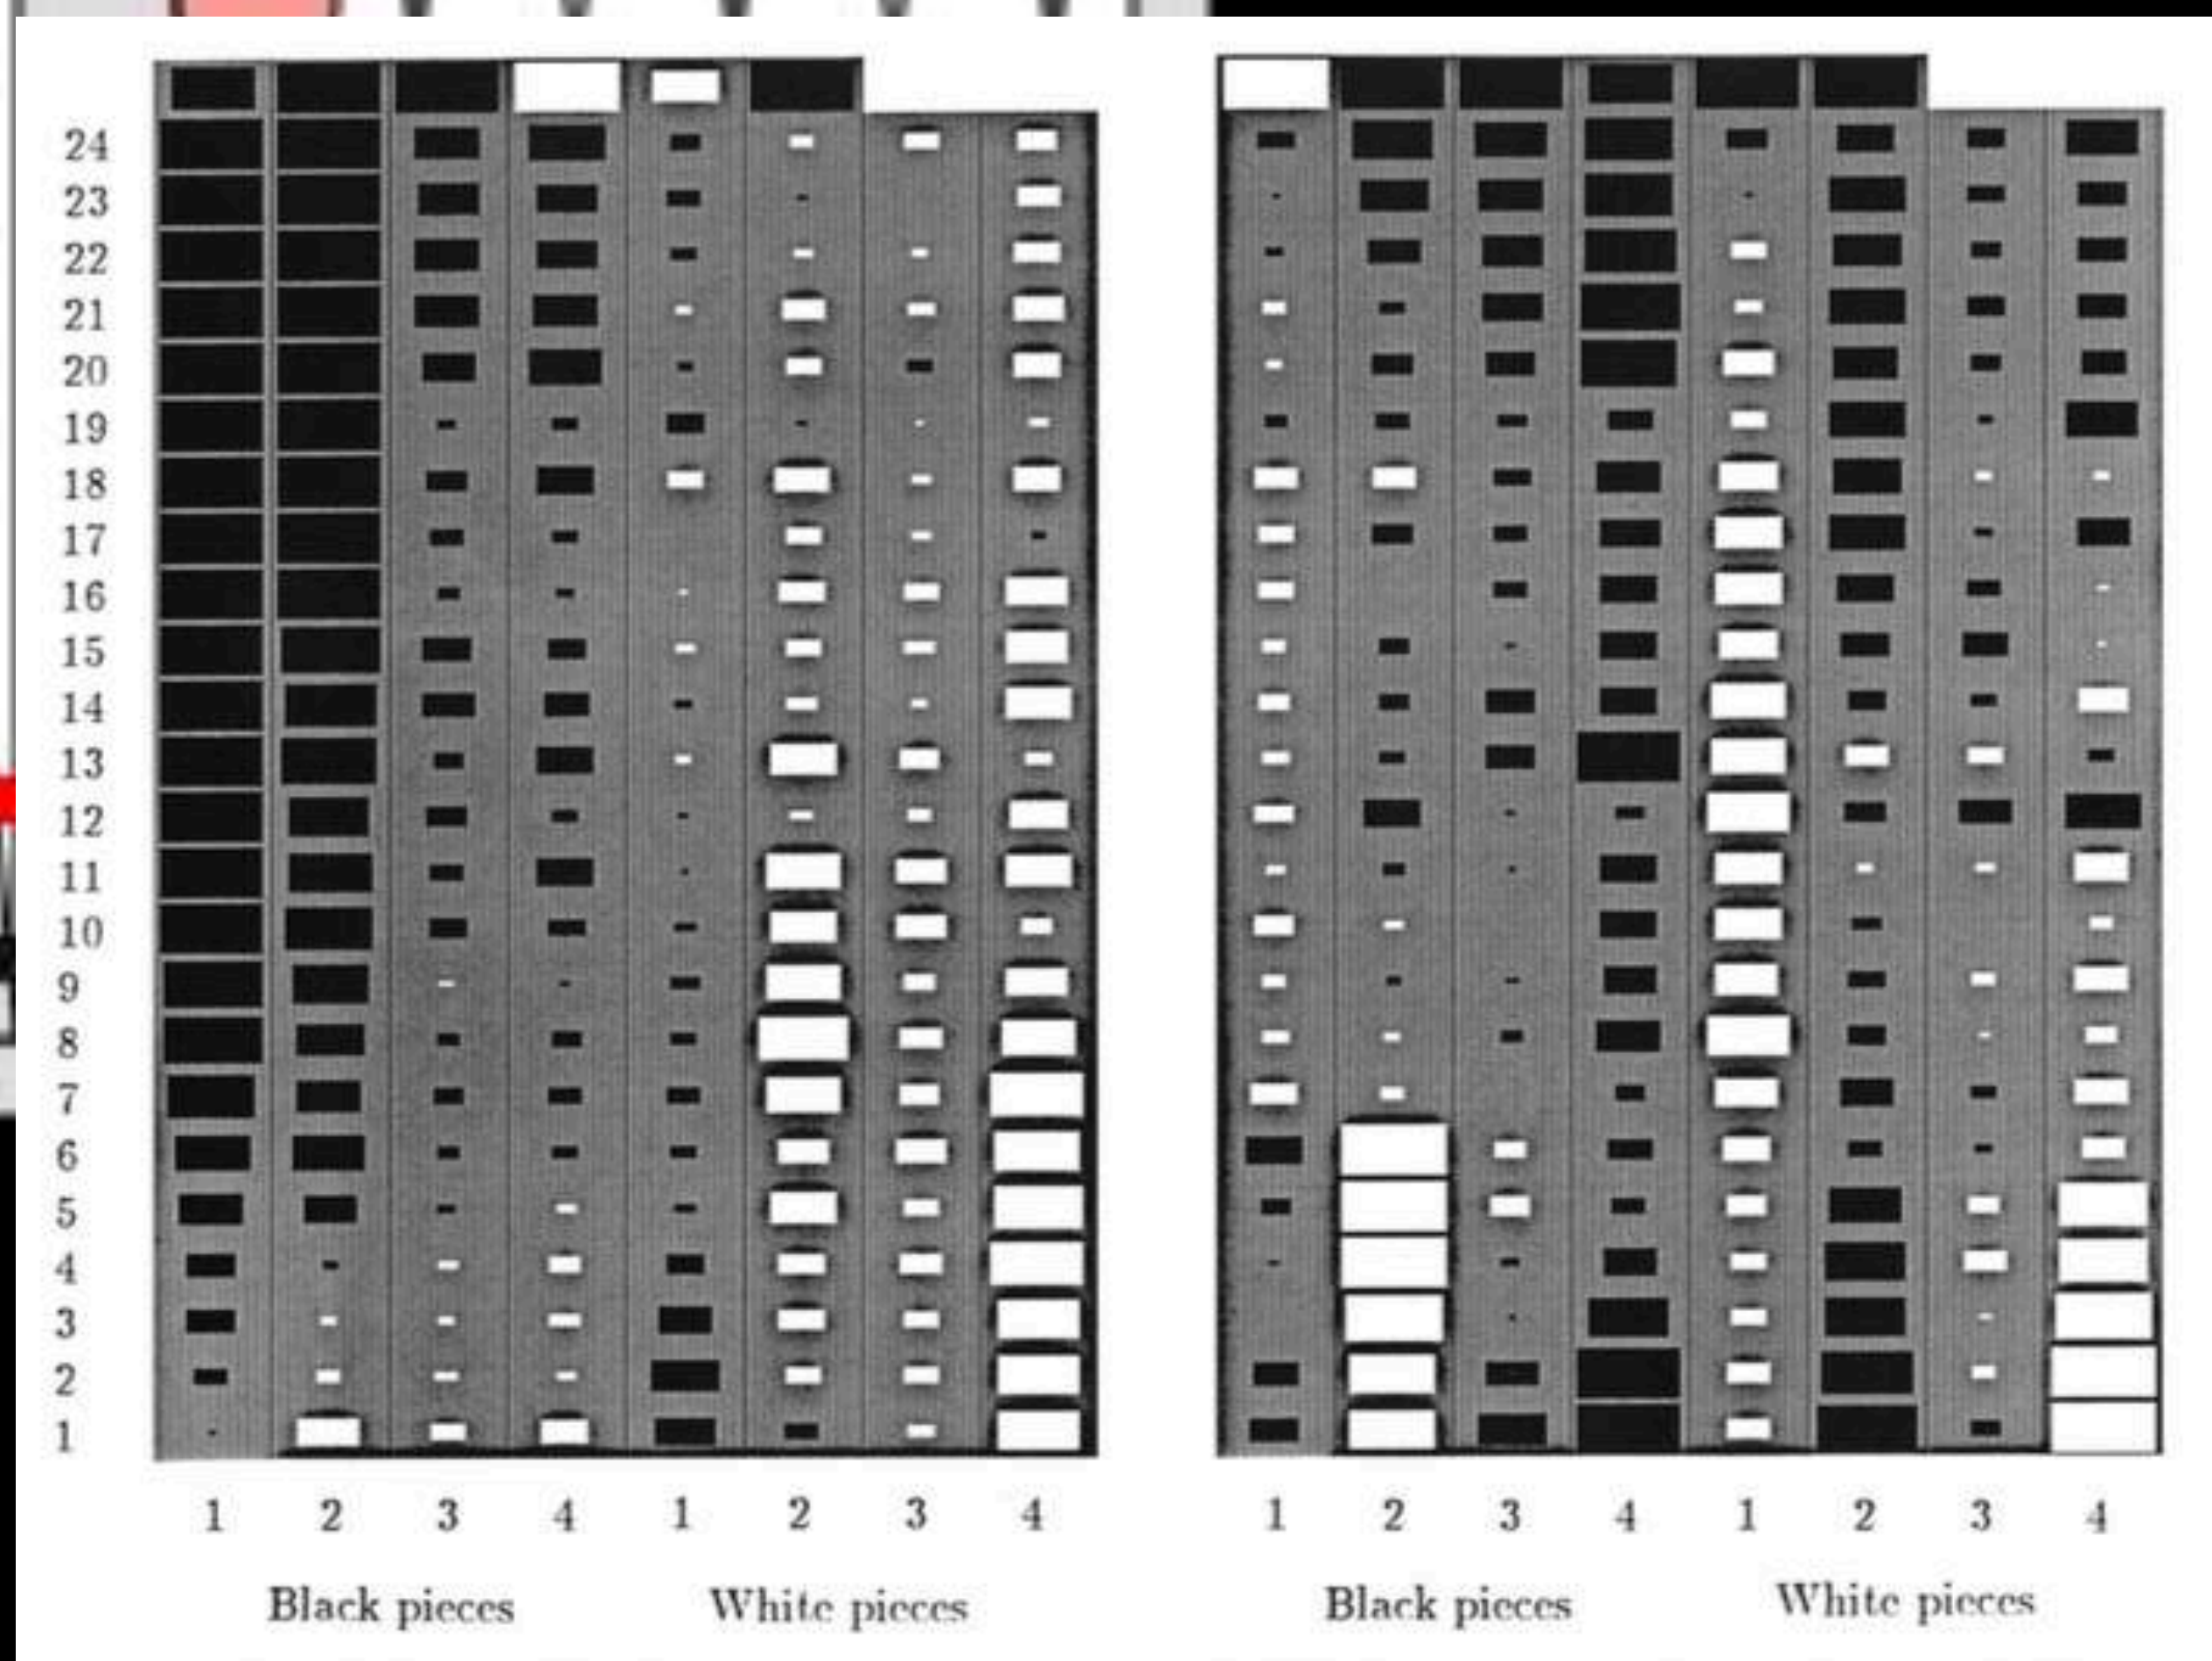

- 20 years ago, IBM's Deep Blue beat Kasparov
- Deep search, parallel computing
- Evaluation with hand-designed features
- Some automated tuning

Checkers

- Solved ten years ago, Jonathan Schaeffer
- Exact, proven result
- Massive search, endgame databases
- Hand-built evaluation



Backgammon - TD-Gammon

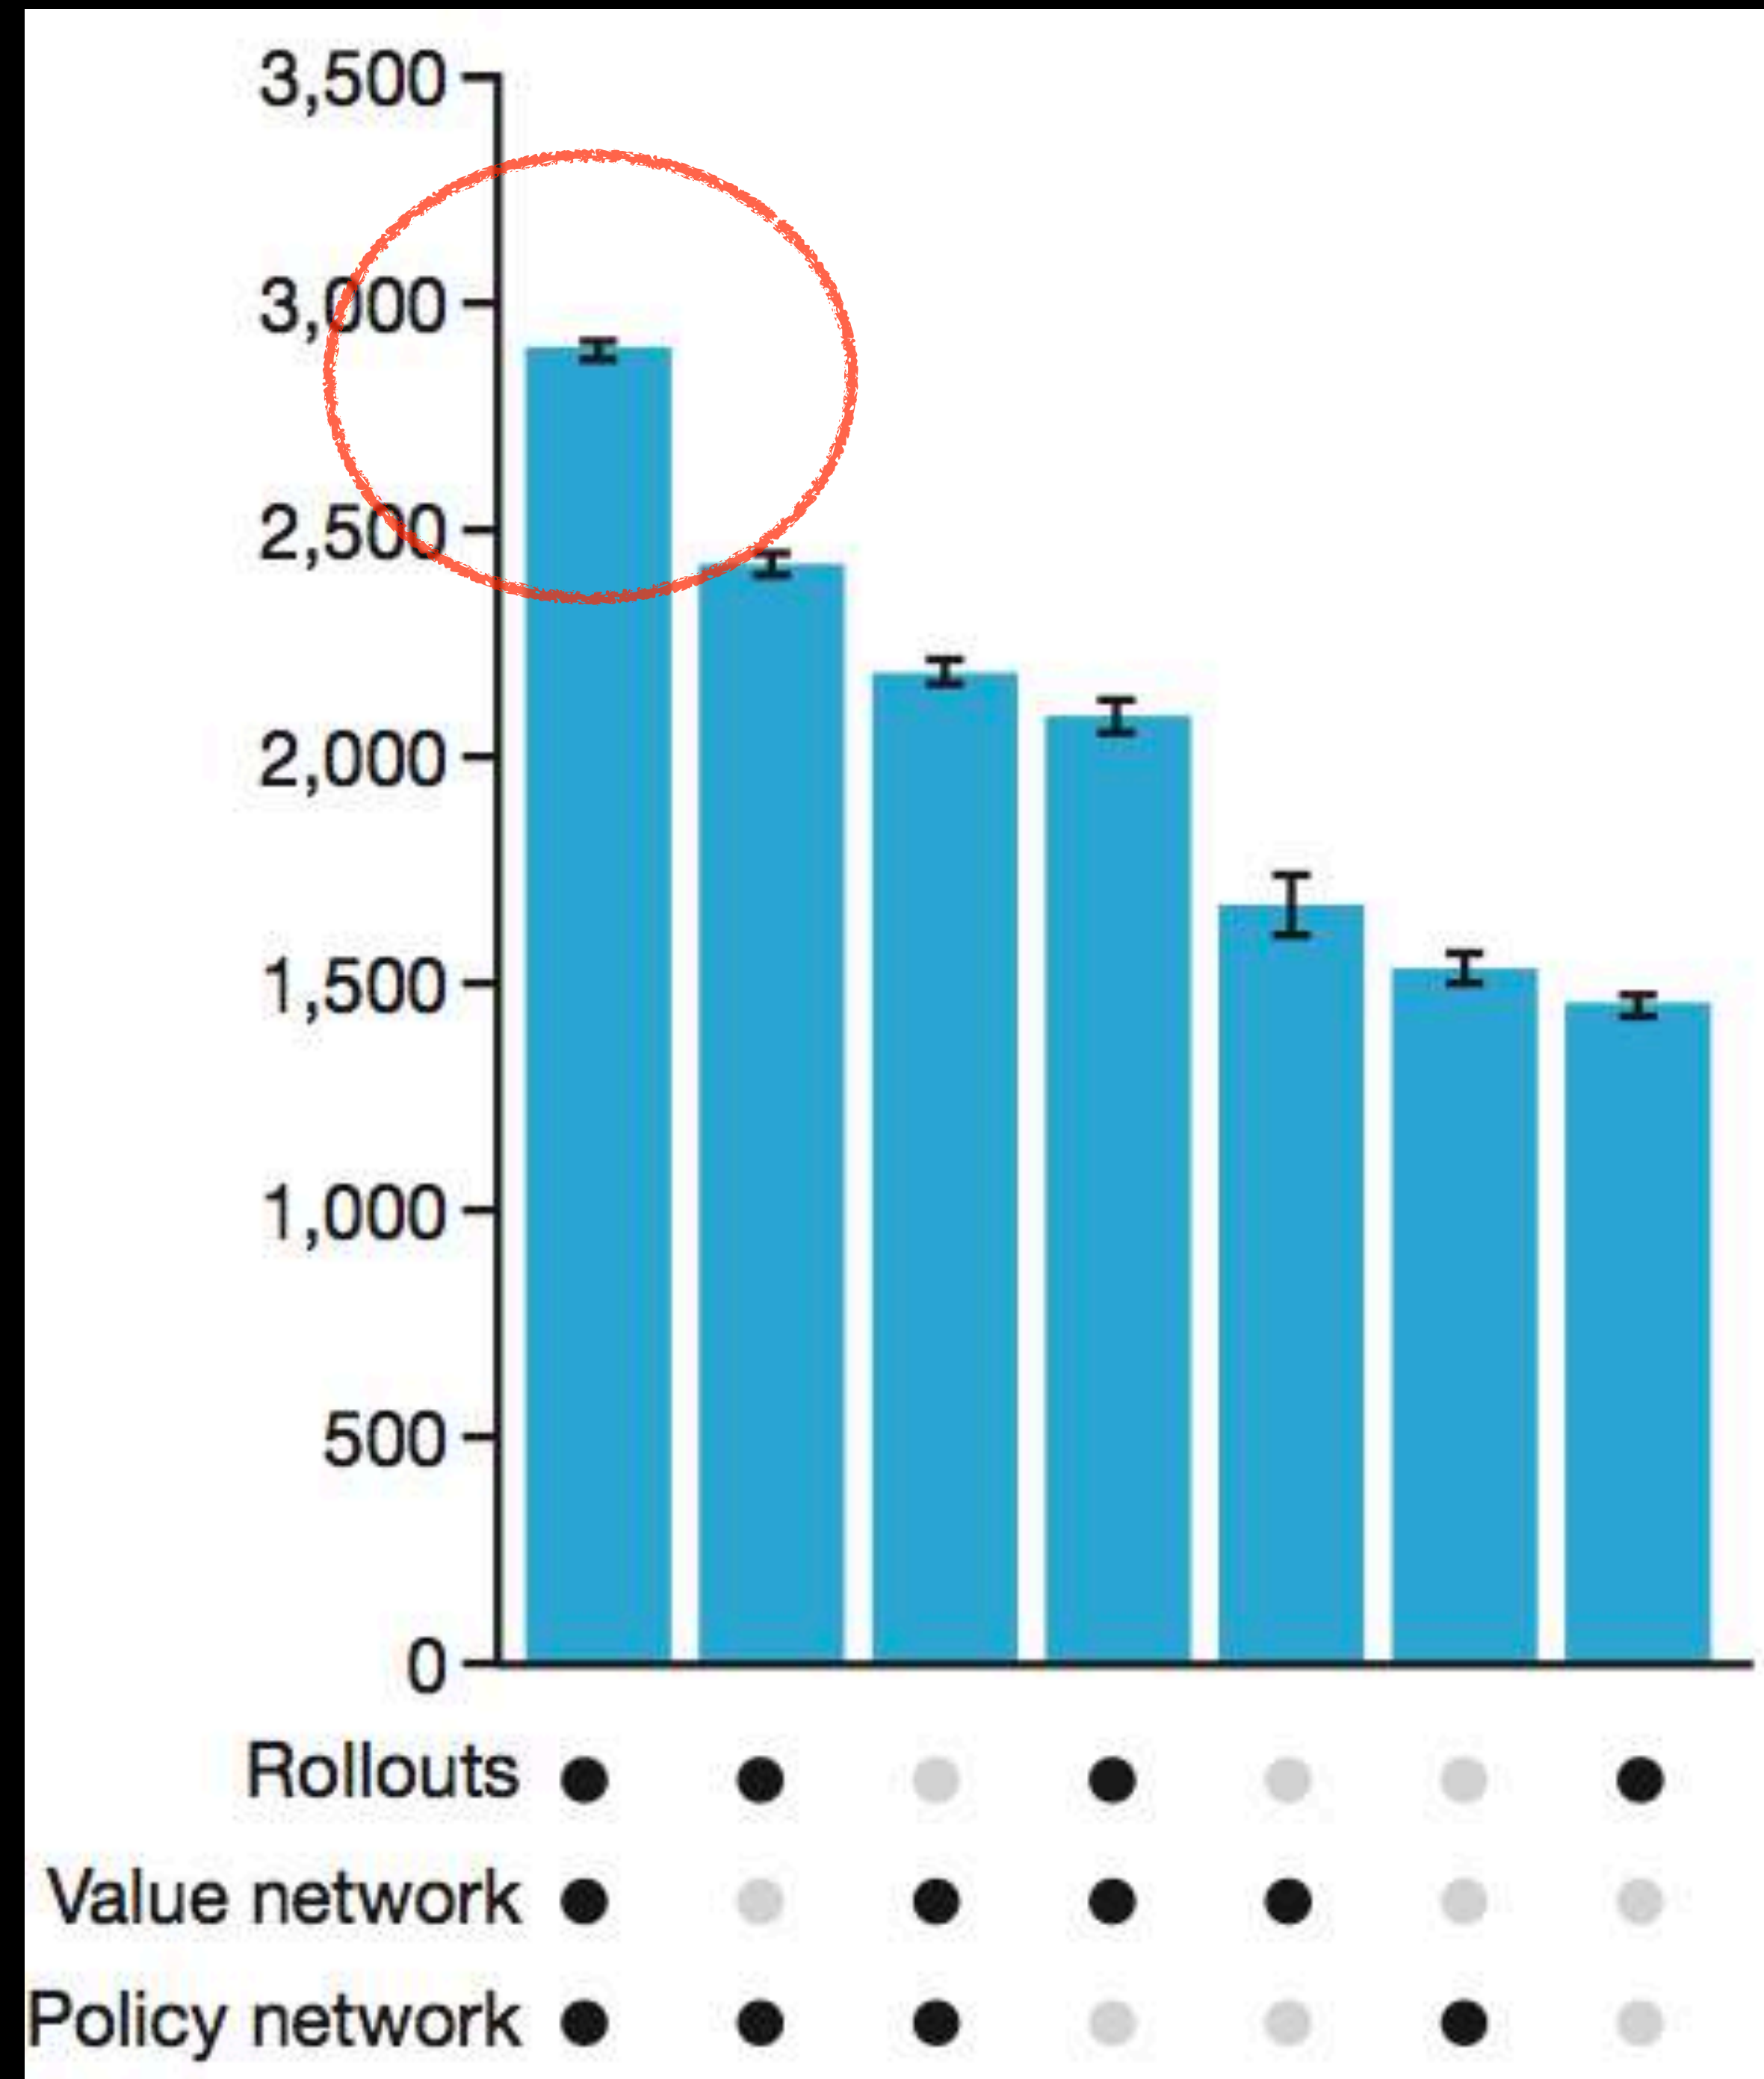


- Gerry Tesauro, IBM
- 25 years ago
- Early success of
 - Neural networks
 - Reinforcement learning

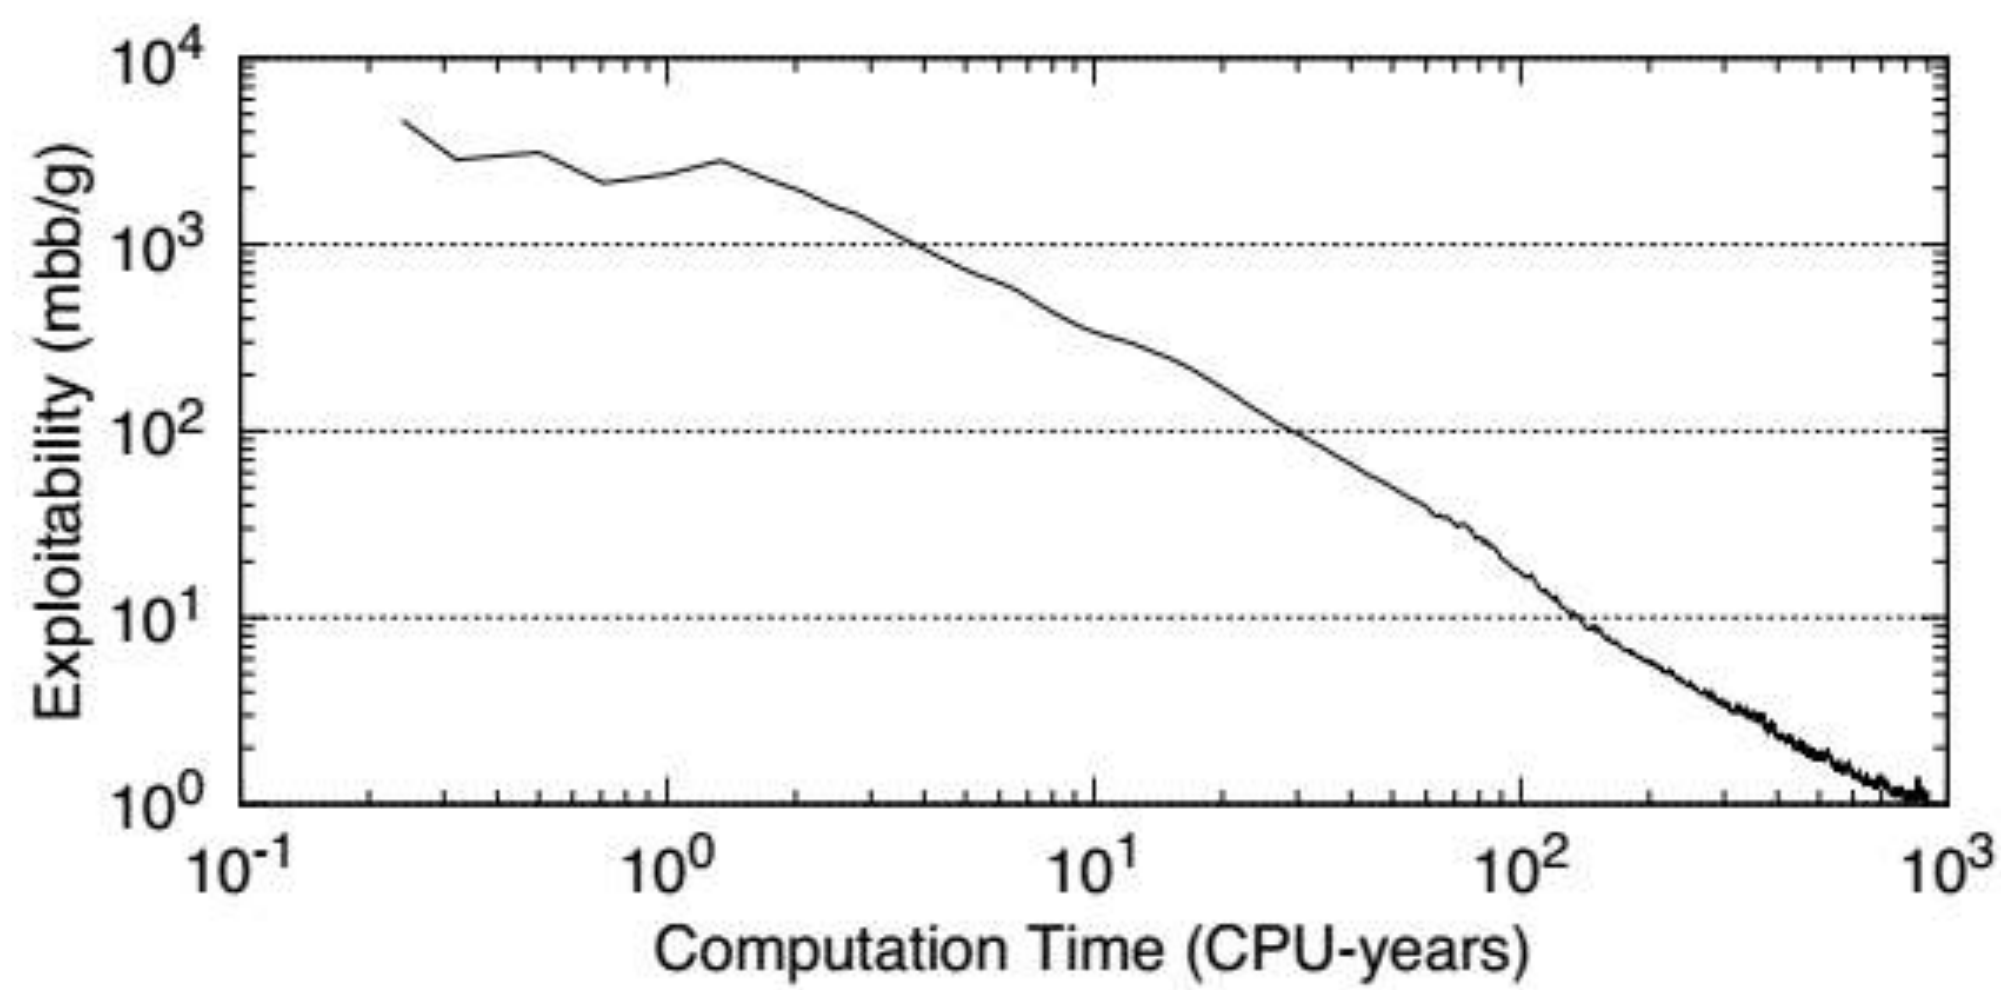
Modern Successes

- 一. Search: Monte Carlo Tree Search
- 二. Knowledge: large scale machine learning, deep networks
- 三. Massively parallel, use GPU/TPU
- 四. Case studies: **Go, poker**

Go (Weiqi)



- AlphaGo
- Search, Knowledge, Simulations
- Supervised and reinforcement learning
- Combination much stronger than each alone



Poker

- Bowling et al 2017, **DeepStack**
- Heads-up No-Limit Texas

Hold'em Poker

- Incomplete information, huge state space

- First application of heuristic search to such problems


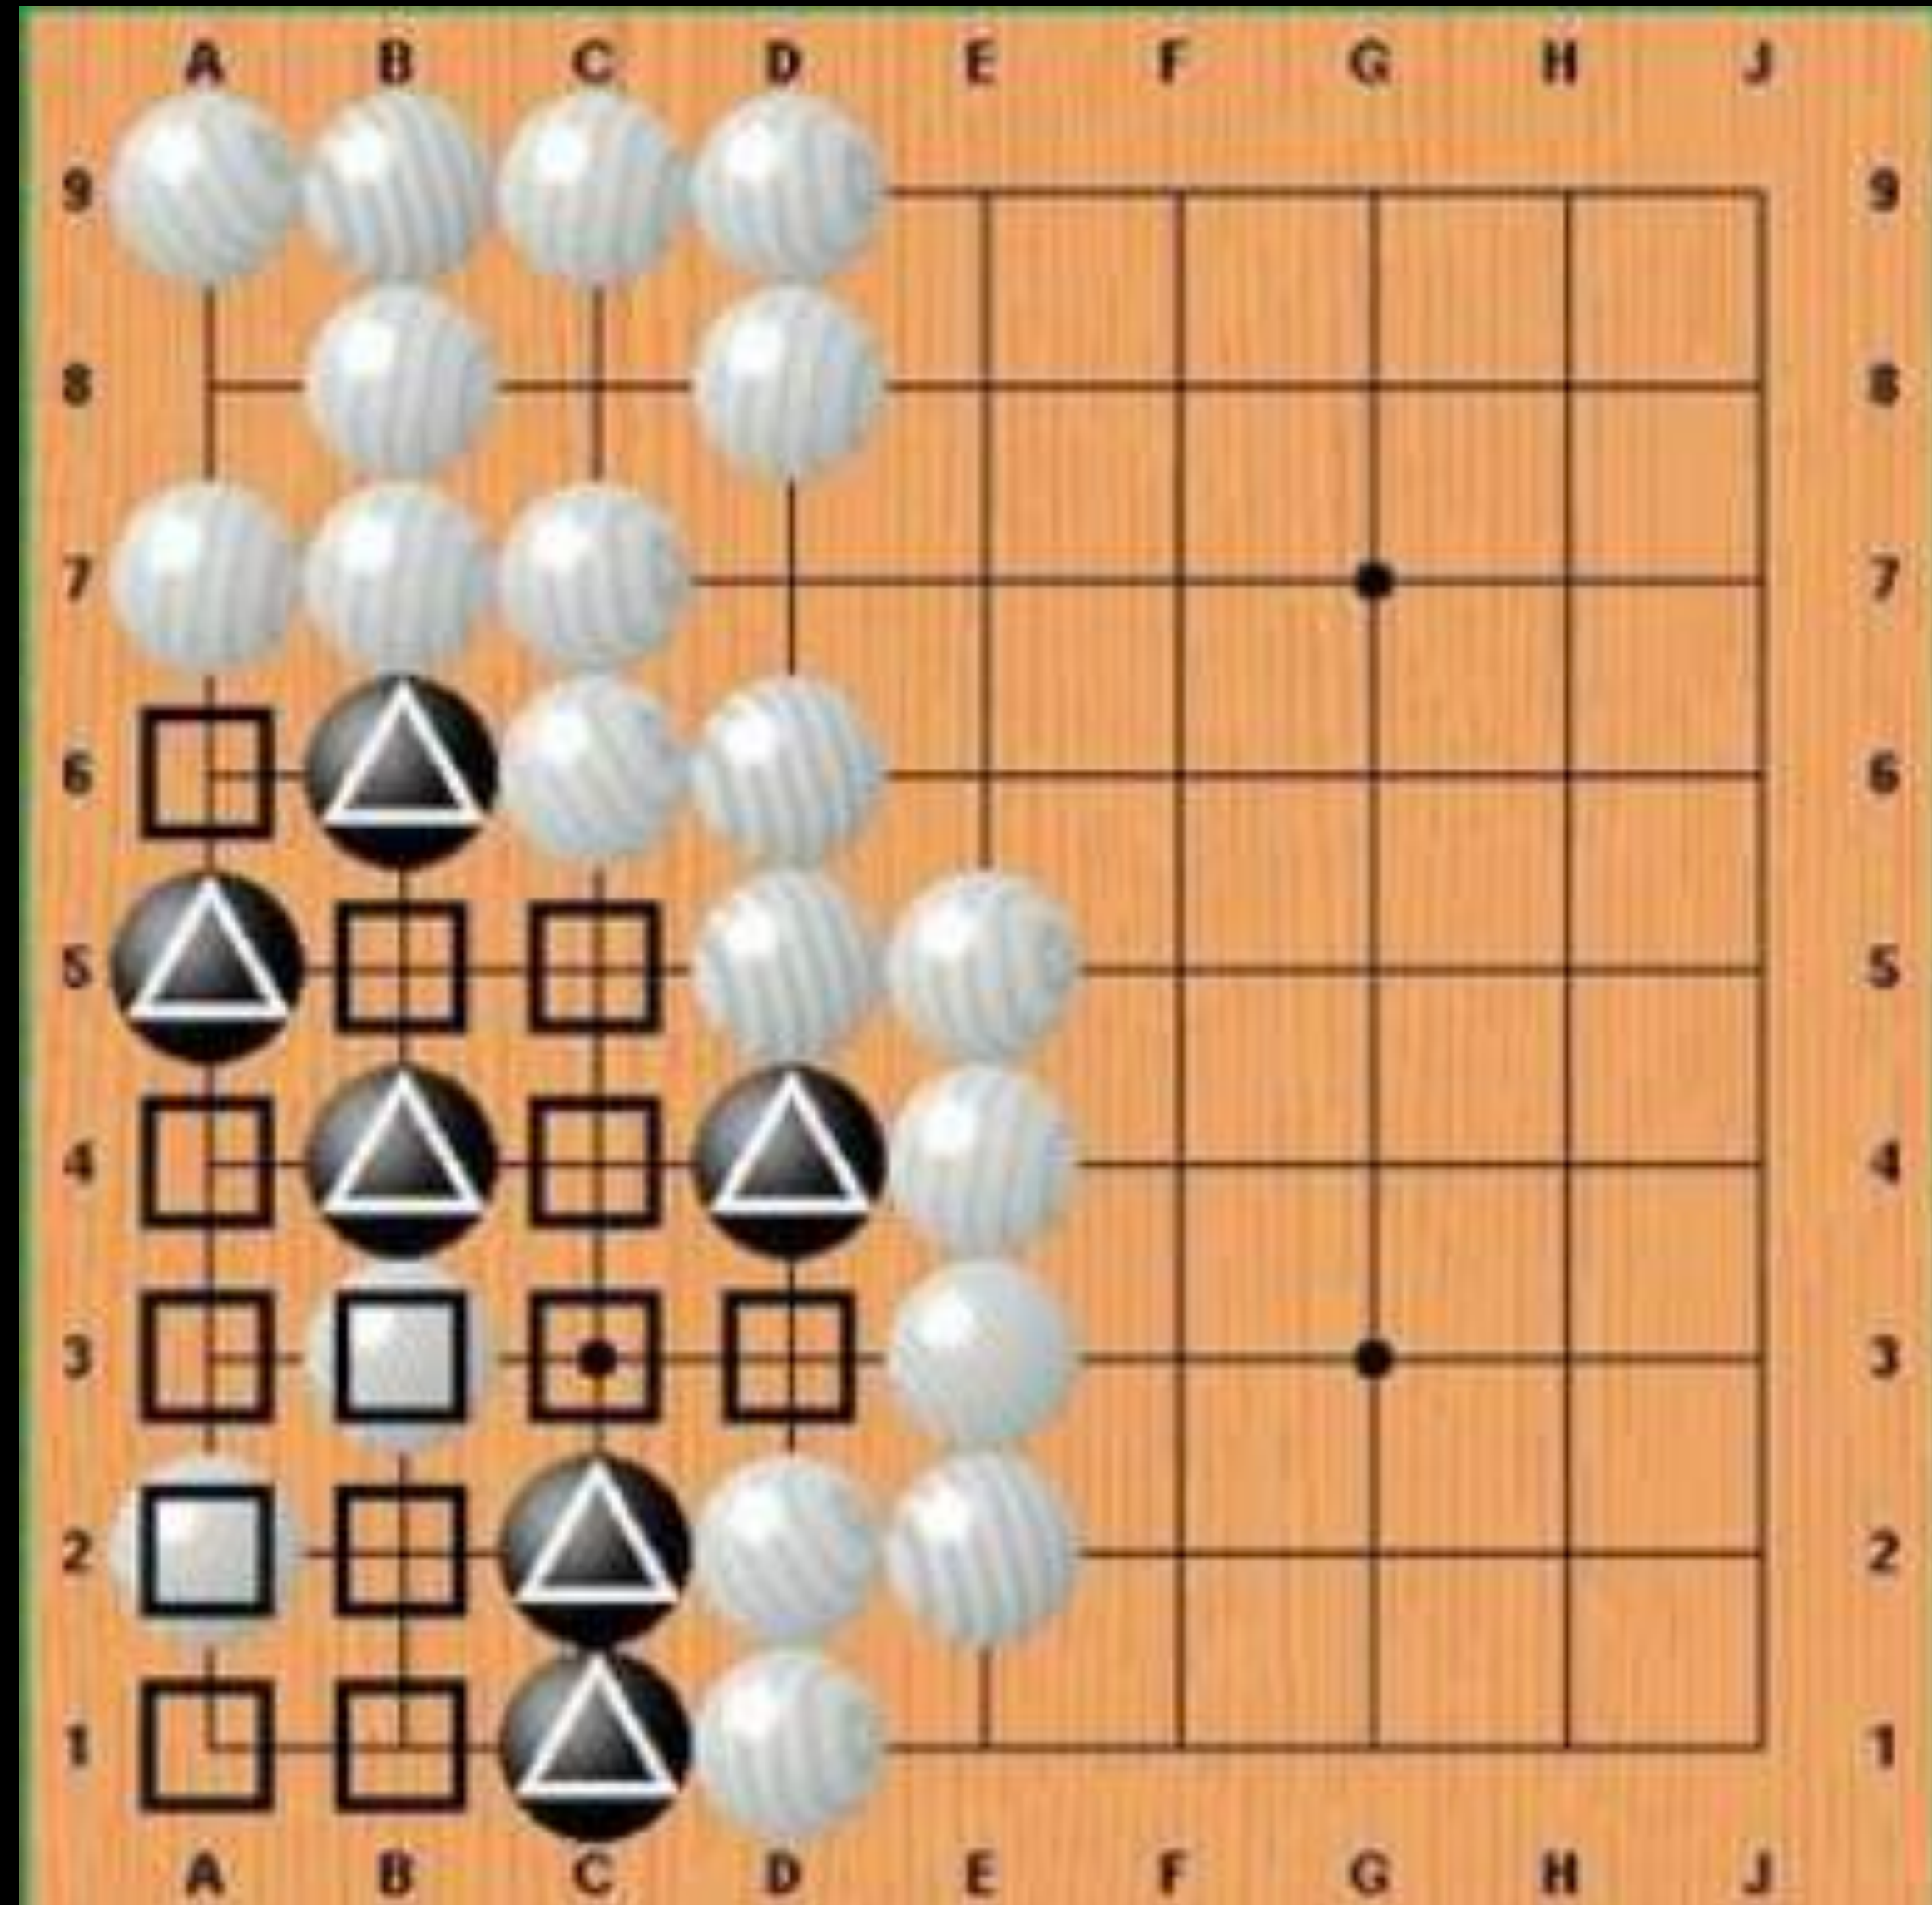


Fig. 3. Deep counterfactual value network. The inputs to the network are player ranges, which are first processed into hand clusters. The output layer is postprocessed to guarantee the values satisfy the zero-sum constraint of counterfactual values.


Future Challenges

- Search: combine exact and heuristic methods
- Knowledge: learn models from data
- Application: **your** decision-making problem



Challenge: Combine Exact and Heuristic Methods

- Deep Learning is very powerful
- Heuristic, not exact
- Many critical applications require certainty, exact methods
- Challenge: how to combine?



Challenge: Learn Models from Data

- Games:
 - Exact rules known
 - Can simulate billions of steps
- Real world: **not known**



Summary - Heuristic Search

- Core technology for complex decision-making
- Advances in deep learning greatly expand its power