Machine Learning in Games

Michael Buro
University of Alberta

BNMI Workshop on
Artificial Stupidity/Artificial Intelligence

Hans Berliner on AI Trends

"I consider the most important trend was that computers got considerably faster in these last 50 years. In this process, we found that many things for which we had at best anthropomorphic solutions, which in many cases failed to capture the real gist of a human's method, could be done by more brute-forcish methods that merely enumerated until a satisfactory solution was found. If this is heresy, so be it."

Outline

- The AI Challenge
- Machine Learning to the rescue
  - Heuristic Evaluation
  - Shaping the Search Tree
- Outlook

The AI Challenge

- Creating machines en par with human experts or better
- Fuzzy understanding of how humans think
- Inferior hardware

Why Games?

- Well defined. Can be tailored to study certain aspects. E.g.
  - Perfect vs. imperfect information
  - Two-player vs. multi-player
  - Turn-based vs. real-time
- Simple rules, yet arbitrarily complex problems
- Popular. Many human experts available
- FUN

Game A.I. Highlights

- 1994 TD-Gammon reaches master level
- 1995 Chinook vs. Lafferty 16.5-15.5
- 1997 Deep Blue vs. Kasparov 3.5-2.5
- 1997 Logistello vs. Murakami 6-0
- 1998 Maven vs. Logan 9-5
- 1998 Mahmoud & Rosenberg vs. GIB 6 IMP

Constructing Game A.I.s

- Extract knowledge from expert players
  - How do they judge positions?
  - How do they determine reasonable moves?
  - How do they look ahead?
- Implement approximations
- Hope they work on today's hardware
- Iterate if not satisfied

Laborious!

Machine Learning to the Rescue

- Tune evaluation function parameters
- Tune look-ahead search parameters
- Find new evaluation features
- Create and analyse new opening lines
- Model opponents
- ...

Combining the Strengths of Human Players and Machines

- Generalization
- Pattern based evaluation
- Planning
- Post mortem analysis
- Fast symbolic computation
- Large and fast memory
- Large-scale numerical optimization

Learned Evaluation Functions

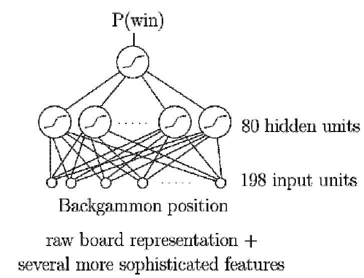

Backgammon (~16K weights)

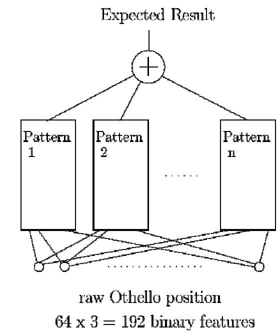

Othello (~ 1 million weights)

Pattern-Based Evaluation

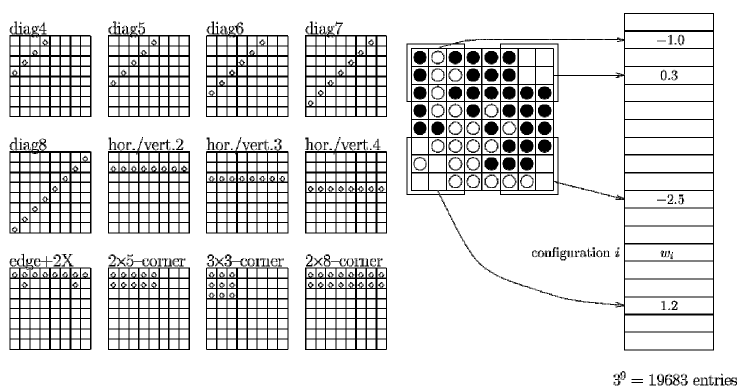

Selective Search

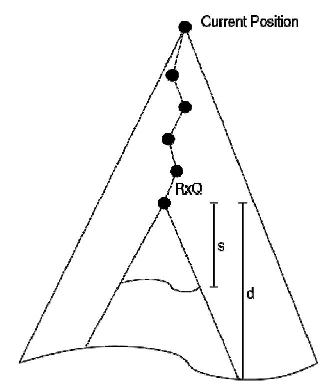

- Mini-Max search is inefficient
- v(d) ~ v(s) + N(0,σ^2)
- Cut if v(s) extreme
- Wins 80% of the games in Othello, 66% in Shogi

Conclusion

- Creating strong AI without understanding every detail is possible
- Human problem solving strategies can be mapped to machines
- ML frees AI programmers from laborious manual tuning

Outlook

- Enhance evaluation and search models in classic games
  - How do humans learn from only a few samples?
  - How do they generate features?
  - How do they create plans?
- AI and ML in commercial games
  - Real-time adversarial planning
  - Opponent modeling

Opening Learning

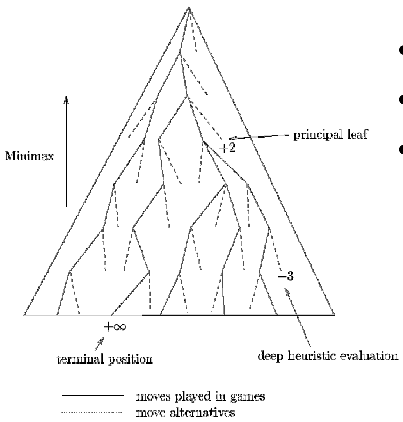

- Avoid repeating losses
- Prepare for opponents
- Explore new openings

Machine Learning in Games

Michael Buro
University of Alberta

BNMI Workshop on
Artificial Stupidity/Artificial Intelligence

Hans Berliner on AI Trends

"I consider the most important trend was that computers got considerably faster in these last 50 years. In this process, we found that many things for which we had at best anthropomorphic solutions, which in many cases failed to capture the real gist of a human's method, could be done by more brute-forcish methods that merely enumerated until a satisfactory solution was found. If this is heresy, so be it."

Outline

- The AI Challenge
- Machine Learning to the rescue
  - Heuristic Evaluation
  - Shaping the Search Tree
- Outlook

The AI Challenge

- Creating machines en par with human experts or better
- Fuzzy understanding of how humans think
- Inferior hardware

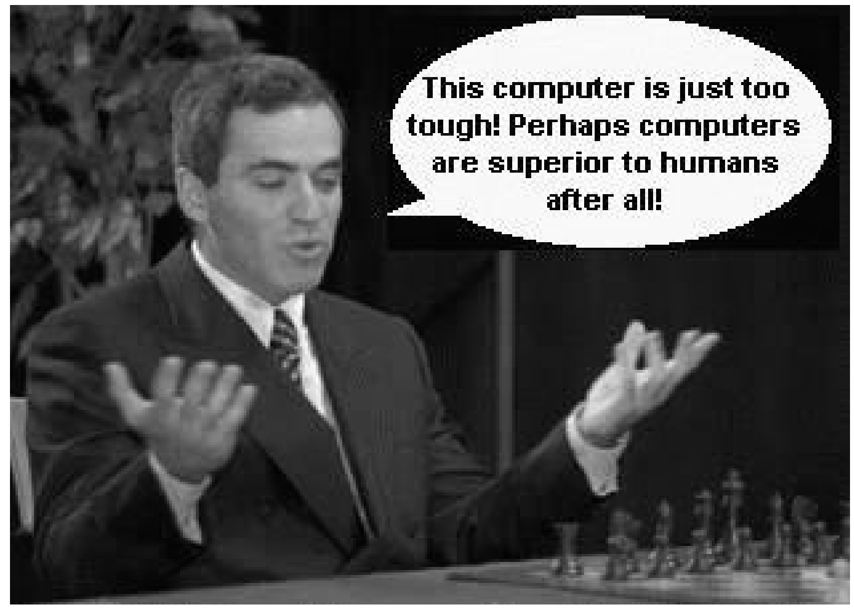

Why Games?

- Well defined. Can be tailored to study certain aspects. E.g.
  - Perfect vs. imperfect information
  - Two-player vs. multi-player
  - Turn-based vs. real-time
- Simple rules, yet arbitrarily complex problems
- Popular. Many human experts available
- FUN

Game A.I. Highlights

- 1994 TD-Gammon reaches master level
- 1995 Chinook vs. Lafferty 16.5-15.5
- 1997 Deep Blue vs. Kasparov 3.5-2.5
- 1997 Logistello vs. Murakami 6-0
- 1998 Maven vs. Logan 9-5
- 1998 Mahmoud & Rosenberg vs. GIB 6 IMP

Constructing Game A.I.s

- Extract knowledge from expert players
  - How do they judge positions?
  - How do they determine reasonable moves?
  - How do they look ahead?
- Implement approximations
- Hope they work on today's hardware
- Iterate if not satisfied

Laborious!

Machine Learning to the Rescue

- Tune evaluation function parameters
- Tune look-ahead search parameters
- Find new evaluation features
- Create and analyse new opening lines
- Model opponents
- ...

Combining the Strengths of Human Players and Machines

- Generalization
- Pattern based evaluation
- Planning
- Post mortem analysis

- Fast symbolic computation
- Large and fast memory
- Large-scale numerical optimization

Learned Evaluation Functions

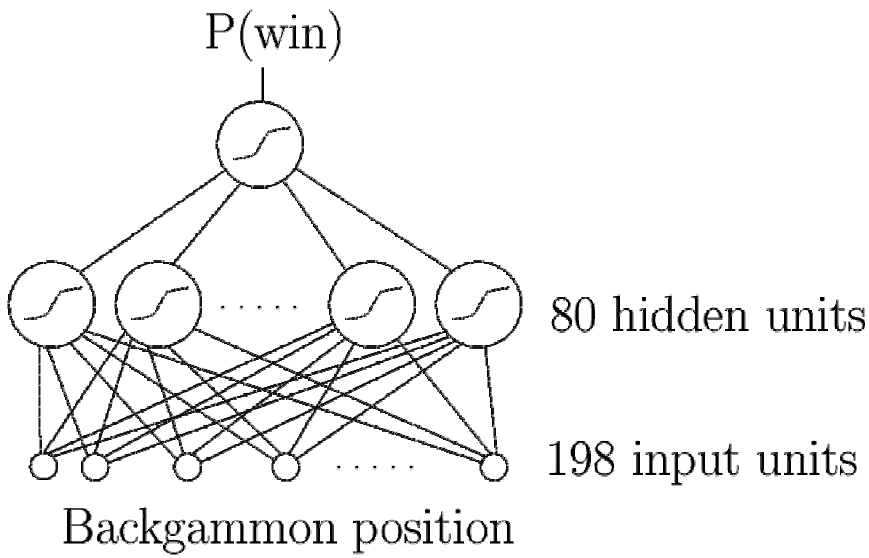

raw board representation +
several more sophisticated features

Backgammon (~16K weights)

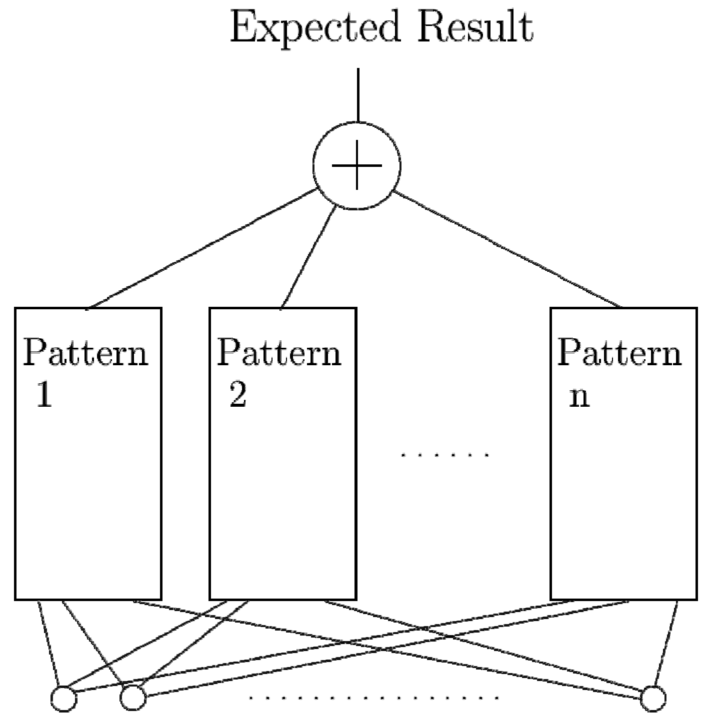

raw Othello position
64 x 3 = 192 binary features

Othello (~ 1 million weights)

Pattern-Based Evaluation

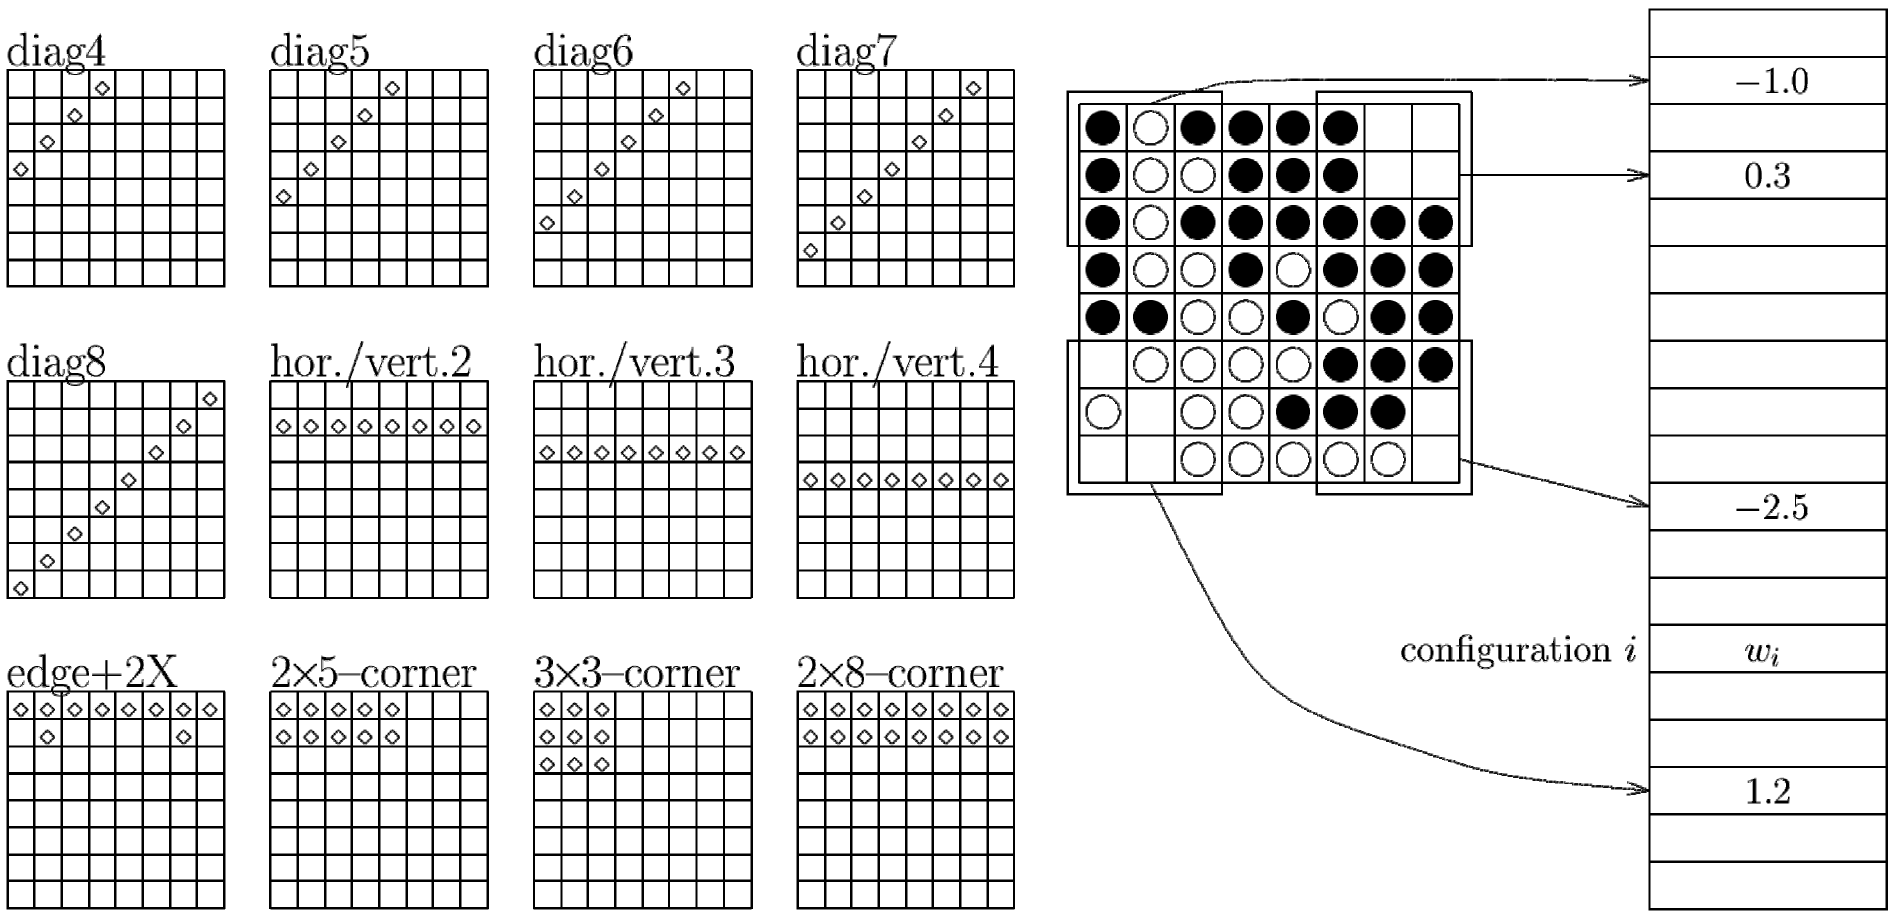

$3^9 = 19683$ entries

Selective Search

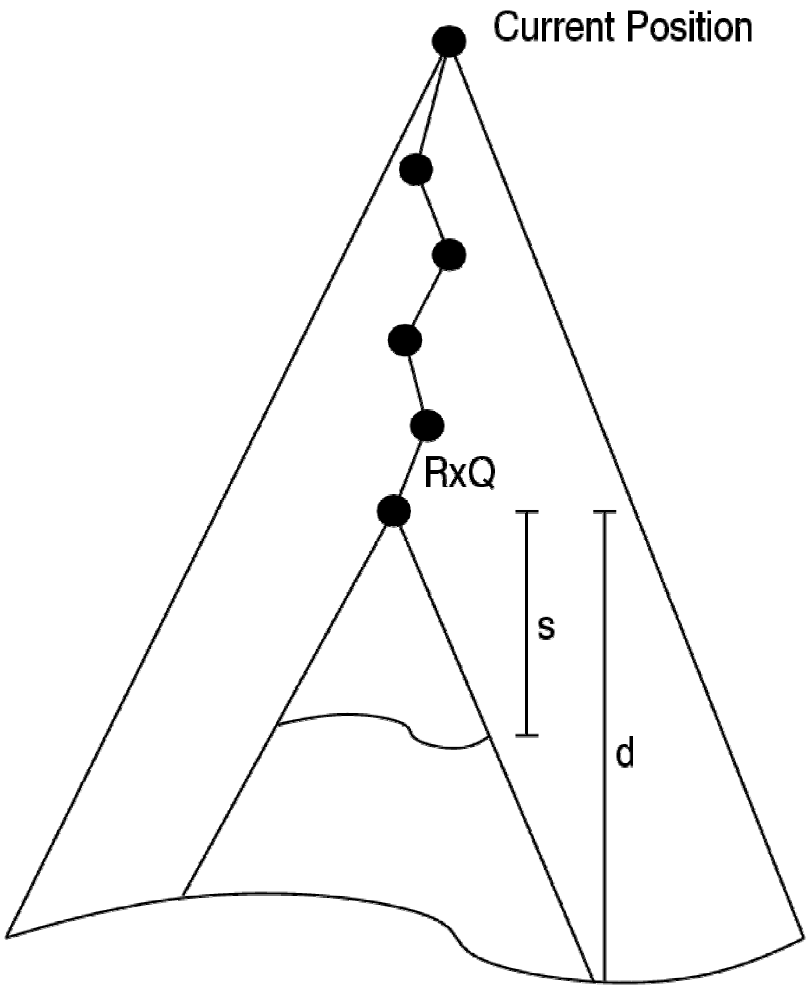

- Mini-Max search is inefficient
- v(d) ~ v(s) + N(0,σ^2)
- Cut if v(s) extreme
- Wins 80% of the games in Othello, 66% in Shogi

Conclusion

- Creating strong AI without understanding every detail is possible
- Human problem solving strategies can be mapped to machines
- ML frees AI programmers from laborious manual tuning

Outlook

- Enhance evaluation and search models in classic games
  - How do humans learn from only a few samples?
  - How do they generate features?
  - How do they create plans?
- AI and ML in commercial games
  - Real-time adversarial planning
  - Opponent modeling

Opening Learning

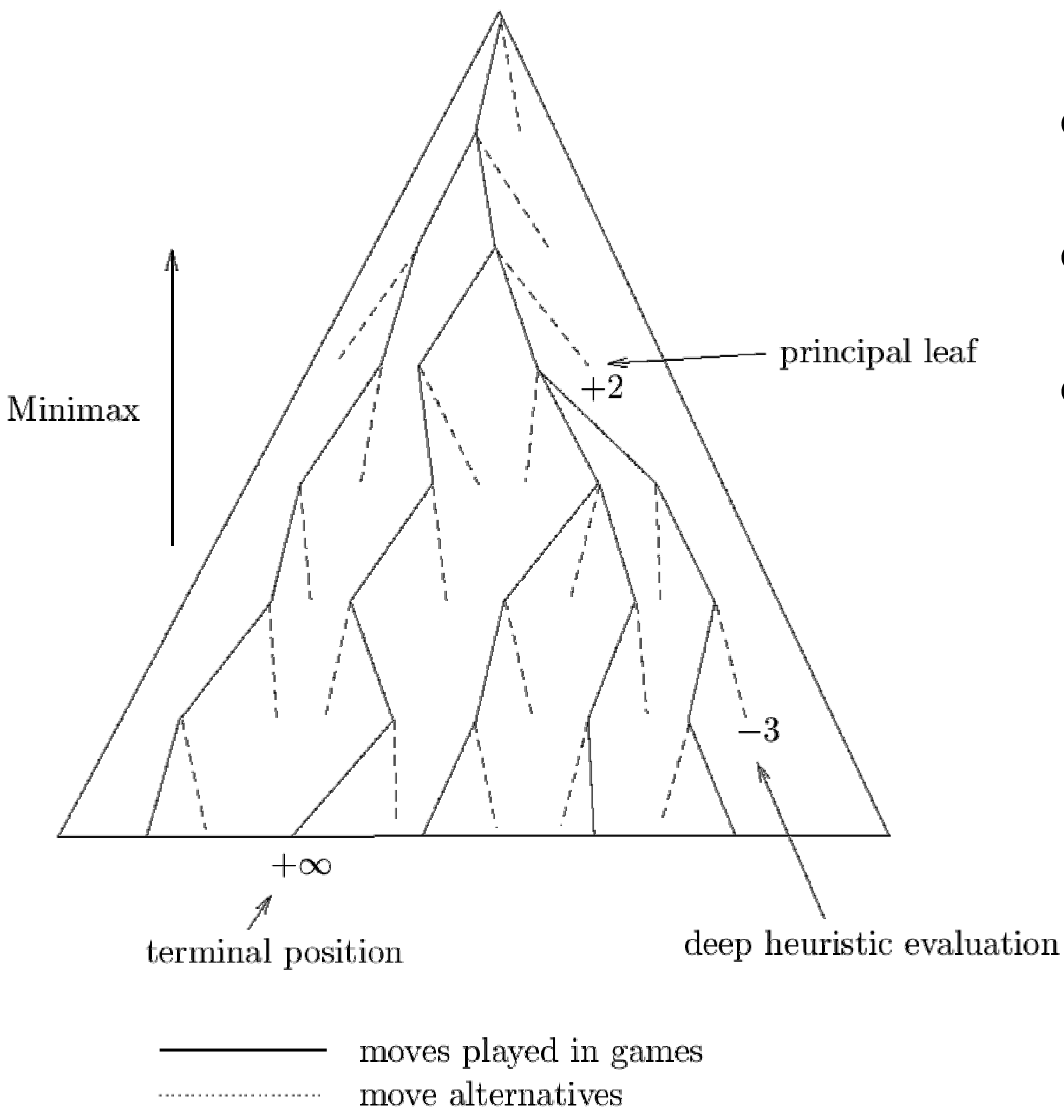

- Avoid repeating losses
- Prepare for opponents
- Explore new openings