

CHANGING THE GAME

WHY TEACHING AI TO PLAY IS MORE THAN FUN AND GAMES

BY SHANE RICZU, '12 MA

ILLUSTRATION BY PETER AND MARIA HOEY


Artificial Intelligence researchers like teaching computers to play games. It might seem frivolous, but researchers see games the way biologists see fruit flies: a simplified system to test theories, says U of A computer scientist Richard Sutton in a story on *Nature.com*.

Recently, AI beat a world champion at the board game Go. Getting here in 22 years required processing power and new ways to solve problems.

Here are some AI programs that have bested us at our own games.

1994 Chinook vs. Checkers

- ▶ **U of A connection:** Chinook was built here by a team led by current dean of science Jonathan Schaeffer.
- ▶ **Why was it so hard?** Checkers has 500 billion billion possible moves.
- ▶ **The solution:** Use brute-force processing power and improved search algorithms.
- ▶ **What happened?** Chinook won the world championship in 1994 and "solved" checkers in 2007 using up to 200 computers working for a total of 14 years.
- ▶ **So what?** Chinook proved some big problems can be solved with enough time and processing power, making challenges like decoding DNA seem less formidable.




2015 Cepheus vs. Poker

- ▶ **U of A connection:** It was co-developed by computing science professor Michael Bowling.
- ▶ **Why was it so hard?** Cepheus can't see the opponent's cards, so it learned to make the best decision with imperfect information.
- ▶ **The solution:** Cepheus learned by playing six billion hands of heads-up limit hold'em poker per second against itself for two months. With each loss, Cepheus tried a new strategy and learned to avoid losing.
- ▶ **What happened?** Cepheus plays with such defensive prowess, you could play it perfectly for 70 years and never make it go bust.
- ▶ **So what?** This technology helps when information is limited, such as scheduling a limited number of coast guard boats to cover vulnerable stretches of ocean.





1997 Deep Blue vs. Chess

- ▶ **U of A connection:** The team was co-led by Murray Campbell, '79 BSc(Hons), '81 MSc.
- ▶ **Why was it so hard?** Chess has too many permutations to search a database of every winning move.
- ▶ **The solution:** To out-think a human, Deep Blue needed 32 computers to consider 200 million moves per second.
- ▶ **What happened?** Deep Blue beat chess grandmaster Garry Kasparov in 1997 over six games.
- ▶ **So what?** The technology led to Jeopardy-winning computer program (and future AI doctor) Watson and to Blue Gene, a supercomputer that's solving protein folding, the next big challenge in personalized medicine.



2016 AlphaGo vs. Go

- ▶ **U of A connection:** Lead developers were David Silver, '09 PhD, and Aja Huang, who did his post-doc here.
- ▶ **Why was it so hard?** A Go board has more combinations of pieces than there are atoms in the universe.
- ▶ **The solution:** Developers used a process called reinforcement learning to teach AlphaGo by playing itself billions of times.
- ▶ **What happened?** AlphaGo defeated world champion Lee Sedol four games to one.
- ▶ **So what?** This win means AI might be ready to go beyond the game table. Reinforcement learning could help other AI do things that are hard to program, like driving cars and noticing social context in conversation.



What's Next?

Starcraft is a promising candidate to push the boundaries of AI. The video game involves strategy like chess and Go, limits a player's information like poker, and relies on bluffing to get an opponent to make a mistake. So far, AI has been steamrolled by the top human players, but U of A researcher David Churchill has been hosting an AI-only Starcraft tournament for the last five years—and AI is improving.

If all of this makes you nervous about an impending AI takeover, perhaps there's small comfort in this thought: they might still let us give them cute human names and play games with them. Which, of course, we will lose.