

How poker and other games help artificial intelligence evolve

April 18 2018, by Alexandra Zabjek



Besting human champions at complex games like checkers, poker and Go is an ideal challenge for artificial intelligence—and a stepping stone to wide-ranging applications beyond the game room, say UAlberta experts. Credit: University of Alberta

When he was growing up in Ohio, his parents were avid card players, dealing out hands of everything from euchre to gin rummy. Meanwhile, he and his friends would tear up board games lying around the family home and combine the pieces to make their own games, with new

challenges and new markers for victory.

Bowling has come far from his days of playing with colourful cards and plastic dice. He has three degrees in computing science and is now a professor at the University of Alberta.

But, in his heart, Bowling still loves playing games. His research in [artificial intelligence](#)—and how it intersects with games and machine learning—has put him at the forefront of the rapidly evolving field.

"Games are such a beautiful microcosm of interesting decision-making problems," he says.

"For (a moment), you can keep it as a self-contained thing, where your artificial intelligence agents look at the world of the game and try to solve it. In the same way, it's fun for humans to play that game for a minute and simplify their worlds to that."

Bowling heads the university's Computer Poker Research Group, where he and a team solved a problem last year once thought unsolvable: the group developed an algorithm that beats professional poker players in no-limit Texas hold 'em poker.

It's called DeepStack. Its success is considered another milestone in artificial intelligence research, where complex algorithms often start with scientists who use card games, [board games](#) and video games as their testing grounds.

From checkers to the human genome

In 1989, computing scientist Jonathan Schaeffer had the idea of writing a computer program, Chinook, to win the World Checkers/Draughts Championship.

While people in the checkers-playing community were intrigued by the novelty of a computer competing in their tournaments, they were surprised when, in just its second tournament, Chinook earned the right to play for the World Championship.

The governing bodies of world checkers at first resisted the idea of a computer playing for the world championship, but the game's top federations eventually relented—somewhat—and formed a new event: the Man vs. Machine World Championship.

In 1992, Chinook lost to Marion Tinsley, who had retired the previous year as the world champion of checkers. But in a rematch in 1994, the program prevailed. Chinook became the first program to win a human world championship in any game, a feat recognized by the Guinness Book of World Records.

That battle of man versus machine has played out again and again in the decades since: in 1997, IBM's Deep Blue—which was co-authored by U of A alumnus Murray Campbell and two others—beat world chess champion Garry Kasparov; in 2011, IBM's Watson took home a \$1-million Jeopardy! championship; and last year, AlphaGo became the first computer program to beat a human in the complex board game, Go.

Schaeffer, now dean of the U of A's Faculty of Science, has been researching artificial intelligence since 1979.

"I do AI research. And one of the tests that I use, the most popular with students, happens to be games. And if the ideas are good, they can be applied to other applications."

Schaeffer clearly relishes the fun of a game, the competition and the broad appeal of working on something that so many can relate to. But he is also clear that game research can have impacts far beyond a deck of

cards or a playing board.

He points to something ubiquitous in modern-day life: the GPS system. And he connects the technology behind it to the Rubik's Cube.

"The Rubik's Cube is jumbled up. It's your 'start' and you're trying to get to a position where everything is in place. What's the route to go from one point to another? Think of each move (in the game) as part of your trip, and you want to get from start to end as fast as possible—which is exactly what the GPS does.

"I'd like to think some of the ideas we've generated here are used in GPS systems. We don't know that for sure. Our work is public and companies do not have to disclose what they put inside their products."



Jonathan Schaeffer developed Chinook, a checkers-playing program that defeated former world champion Marion Tinsley in 1994—the first program to win a human world championship in any game. Credit: Richard Siemens

Schaeffer's work on Chinook involved a massive amount of computing power that had to organize data and compress it down into something very small. The program also had to find information quickly, among the trillions of data points that were compressed.

Soon after Chinook's [world championship](#), a biologist came into Schaeffer's office. He was working on the human genome project, which also involved an enormous amount of data. He also needed to compress that data, and identify elements quickly.

A year later, a company called BioTools formed to support the [human genome project](#). Some of the ideas used in their products came from the checkers research.

"It was one of those pleasant surprises, completely unexpected, coming out of left field."

Building the 'intuitive' machine

The recent work from the U of A to master no-limit Texas hold 'em poker landed in the March 2017 edition of the prestigious journal *Science*. What made the feat groundbreaking is that the game is an "imperfect information" scenario—unlike chess or checkers, where both players can see everything on the table, players in poker keep some cards to themselves.

The challenge for Bowling and his team was to develop an "intuition" for DeepStack.

"Our algorithms have to think deeply about what COULD the other agent know about my cards right now and what could I know about their cards. The reasoning has to encompass beliefs and not just what you can see on the table."

To develop that gut instinct, that "intuition," Bowling and his team had to run DeepStack through tens of millions of poker scenarios. The program began to recognize some situations as good and some as bad, some as less good or less bad.

This built up DeepStack's general experience, and it began to recognize new situations as being various degrees of good or bad—and allowed it come up with the best plays based on that.

The team proved DeepStack's competence by having it play against some of the best poker players in the world. They played so many hands that the findings were statistically significant.

When asked whether DeepStack will essentially kill poker—after all, who wants to play against a machine if it's guaranteed the machine will win?—Bowling points to chess.

"That was said of chess and now there's maybe more grandmasters than at any other time ... it seems to have accelerated the skill level," he says, noting that amateurs can play against the best computers to hone their skills and there are competitions where chess players team up with computers to enhance the competition.

But he also acknowledges that poker is a bit different. People play poker for money. It's a gambling activity, and it makes no sense to play against an entity that is statistically likely to win.

"Maybe that means that poker starts to become about something different than just a pure gambling activity ... I hope maybe it will move the [game](#) away from the gambling piece. Let's try to figure out who the best players in the world are and maybe we could start to highlight this for what it is, a skill-based activity."

Bowling first started his work on DeepStack with the approach that most "real-world" problems are, in fact, imperfect information scenarios.

But his work isn't just driven by how his AI breakthroughs might be applied beyond the realm of games.

"If we're going to push forward to see more capable artificial intelligence, it might be a distraction to say, 'Here's an application that you can make a lot of money off of.' Is that really the path to get us to better, more capable artificial intelligence?"

"There should be people who push on the longer-term path, who push it forward to the next thing without being certain how it's going to be a money-making activity."

Provided by University of Alberta

Citation: How poker and other games help artificial intelligence evolve (2018, April 18) retrieved 19 April 2024 from

<https://techxplore.com/news/2018-04-poker-games-artificial-intelligence-evolve.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.