

# DeepStack the first computer program to outplay human professionals at heads-up no-limit Texas hold'em poker

2 March 2017



Michael Bowling (center) flanked by co-authors of DeepStack. Credit: Photo by John Ulan for UAlberta.

A team of computing scientists from the University of Alberta's [Computer Poker Research Group](#) is once again capturing the world's collective fascination with artificial intelligence. In a historic result for the flourishing AI research community, the team—which includes researchers from Charles University in Prague and Czech Technical University—has developed an AI system called DeepStack that defeated professional poker players in December 2016. The landmark findings have just been published in *Science*, one of the world's most prestigious peer-reviewed scientific journals.

DeepStack bridges the gap between approaches used for games of perfect information—like those used in checkers, chess, and Go—with those used for imperfect information games, reasoning while it plays using "intuition" honed through deep learning to reassess its strategy with each decision.

"Poker has been a longstanding challenge problem in [artificial intelligence](#)," says Michael Bowling, professor in the University of Alberta's Faculty of Science and principal investigator on the study. "It is the quintessential game of imperfect information in the sense that the players don't have the same information or share the same perspective while they're playing."

Don't let the name fool you: imperfect information games are serious business. These "games" are a general mathematical model that describe how decision-makers interact. Artificial intelligence research has a storied history of using parlour games to study these models, but attention has been focused primarily on perfect information games. "We need new AI techniques that can handle cases where decision-makers have different perspectives," says Bowling, explaining that developing techniques to solve imperfect information games will have applications well beyond the [poker](#) table.

"Think of any real world problem. We all have a slightly different perspective of what's going on, much like each player only knowing their own cards in a game of poker." Immediate applications include making robust medical treatment recommendations, strategic defense planning, and negotiation.

This latest discovery builds on an already impressive body of research findings about artificial intelligence and imperfect information games that stretches back to the creation of the University of Alberta's Computer Poker Research Group in 1996. Bowling, who became the group's principal investigator in 2006, has led the group to several milestones for artificial intelligence. He and his colleagues developed Polaris in 2008, beating top poker players at heads-up limit Texas hold'em

poker. They then went on to solve heads-up limit hold'em with [Cepheus](#), published in 2015 in *Science*.

DeepStack extends the ability to think about each situation during play—which has been famously successful in games like checkers, chess, and Go—to imperfect information games using a technique called continual re-solving. This allows DeepStack to determine the correct strategy for a particular poker situation without thinking about the entire game by using its "intuition" to evaluate how the game might play out in the near future.

"We train our system to learn the value of situations," says Bowling. "Each situation itself is a mini poker game. Instead of solving one big poker game, it solves millions of these little poker games, each one helping the system to refine its intuition of how the game of poker works. And this intuition is the fuel behind how DeepStack plays the full game."

Thinking about each situation as it arises is important for complex problems like heads-up no-limit hold'em, which has vastly more unique situations than there are atoms in the universe, largely due to players' ability to wager different amounts including the dramatic "all-in." Despite the game's complexity, DeepStack takes action at human speed—with an average of only three seconds of "thinking" time—and runs on a simple gaming laptop with an Nvidia graphics processing unit.

To test the approach, DeepStack played against a pool of professional poker players in December, 2016, recruited by the International Federation of Poker. Thirty-three players from 17 countries were recruited, with each asked to play a 3000-hand match over a period of four weeks. DeepStack beat each of the 11 players who finished their match, with only one outside the margin of statistical significance, making it the first computer program to beat professional players in heads-up no-limit Texas hold'em poker.

"DeepStack: Expert-Level Artificial Intelligence in No-Limit Poker" will be published online by the journal *Science* on Thursday, March 2, 2017.

**More information:** "DeepStack: Expert-level artificial intelligence in heads-up no-limit poker," *Science*, [science.sciencemag.org/lookup/...1126/science.aam6960](https://www.sciencemag.org/lookup/doi/10.1126/science.aam6960)

Provided by University of Alberta

APA citation: DeepStack the first computer program to outplay human professionals at heads-up no-limit Texas hold'em poker (2017, March 2) retrieved 23 January 2022 from <https://techxplore.com/news/2017-03-deepstack-outplay-human-professionals-heads-up.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.*