

At Last, AI beats professionals in six-player poker

July 11 2019



Noam Brown is a Facebook AI research scientist while finishing his Ph.D. at Carnegie Mellon. Credit: Noam Brown

An artificial intelligence program developed by Carnegie Mellon University in collaboration with Facebook AI has defeated leading

professionals in six-player no-limit Texas hold'em poker, the world's most popular form of poker.

The AI, called Pluribus, defeated [poker](#) professional Darren Elias, who holds the record for most World Poker Tour titles, and Chris "Jesus" Ferguson, winner of six World Series of Poker events. Each pro separately played 5,000 hands of poker against five copies of Pluribus.

In another experiment involving 13 pros, all of whom have won more than \$1 million playing poker, Pluribus played five pros at a time for a total of 10,000 hands and again emerged victorious.

"Pluribus achieved superhuman performance at multi-player poker, which is a recognized milestone in artificial intelligence and in [game theory](#) that has been open for decades," said Tuomas Sandholm, Angel Jordan Professor of Computer Science, who developed Pluribus with Noam Brown, who is finishing his Ph.D. in Carnegie Mellon's Computer Science Department as a research scientist at Facebook AI. "Thus far, superhuman AI milestones in strategic reasoning have been limited to two-party competition. The ability to beat five other players in such a complicated game opens up new opportunities to use AI to solve a wide variety of real-world problems."

A research paper describing this achievement in AI will be published online by the journal *Science* on Thursday, July 11, 2019.

"Playing a six-player game rather than head-to-head requires fundamental changes in how the AI develops its playing strategy," said Brown, who joined Facebook AI last year. "We're elated with its performance and believe some of Pluribus' playing strategies might even change the way pros play the game."

Pluribus' algorithms created some surprising features into its strategy.

For instance, most human players avoid "donk betting"—that is, ending one round with a call but then starting the next round with a bet. It's seen as a weak move that usually doesn't make strategic sense. But Pluribus placed donk bets far more often than the professionals it defeated.

"Its major strength is its ability to use mixed strategies," Elias said last week as he prepared for the 2019 World Series of Poker main event. "That's the same thing that humans try to do. It's a matter of execution for humans—to do this in a perfectly random way and to do so consistently. Most people just can't."

Pluribus registered a solid win with statistical significance, which is particularly impressive given its opposition, Elias said. "The bot wasn't just playing against some middle of the road pros. It was playing some of the best players in the world."

Michael "Gags" Gagliano, who has earned nearly \$2 million in career earnings, also competed against Pluribus.

"It was incredibly fascinating getting to play against the poker bot and seeing some of the strategies it chose" said Gagliano. "There were several plays that humans simply are not making at all, especially relating to its bet sizing. Bots/AI are an important part in the evolution of poker, and it was amazing to have first-hand experience in this large step toward the future."

Sandholm has led a research team studying computer poker for more than 16 years. He and Brown earlier developed [Libratus](#), which two years ago decisively beat four poker pros playing a combined 120,000 hands of heads-up no-limit Texas hold'em, a two-player version of the game.

Games such as chess and Go have long served as milestones for AI

research. In those games, all of the players know the status of the playing board and all of the pieces. But poker is a bigger challenge because it is an incomplete information game; players can't be certain which cards are in play and opponents can and will bluff. That makes it both a tougher AI challenge and more relevant to many real-world problems involving multiple parties and missing information.

All of the AIs that displayed superhuman skills at two-player games did so by approximating what's called a Nash equilibrium. Named for the late Carnegie Mellon alumnus and Nobel laureate John Forbes Nash Jr., a Nash equilibrium is a pair of strategies (one per player) where neither player can benefit from changing strategy as long as the other player's strategy remains the same. Although the AI's strategy guarantees only a result no worse than a tie, the AI emerges victorious if its opponent makes miscalculations and can't maintain the equilibrium.

In a game with more than two players, playing a Nash equilibrium can be a losing strategy. So Pluribus dispenses with theoretical guarantees of success and develops strategies that nevertheless enable it to consistently outplay opponents.

Pluribus first computes a "blueprint" strategy by playing six copies of itself, which is sufficient for the first round of betting. From that point on, Pluribus does a more detailed search of possible moves in a finer-grained abstraction of game. It looks ahead several moves as it does so, but not requiring looking ahead all the way to the end of the game, which would be computationally prohibitive. Limited-lookahead search is a standard approach in perfect-information games, but is extremely challenging in imperfect-information games. A new limited-lookahead search algorithm is the main breakthrough that enabled Pluribus to achieve superhuman multi-player poker.

Specifically, the search is an imperfect-information-game solve of a

limited-lookahead subgame. At the leaves of that subgame, the AI considers five possible continuation strategies each opponent and itself might adopt for the rest of the game. The number of possible continuation strategies is far larger, but the researchers found that their algorithm only needs to consider five continuation strategies per player at each leaf to compute a strong, balanced overall strategy.

Pluribus also seeks to be unpredictable. For instance, betting would make sense if the AI held the best possible hand, but if the AI bets only when it has the best hand, opponents will quickly catch on. So Pluribus calculates how it would act with every possible hand it could hold and then computes a strategy that is balanced across all of those possibilities.

Though poker is an incredibly complicated game, Pluribus made efficient use of computation. AIs that have achieved recent milestones in games have used large numbers of servers and/or farms of GPUs; Libratus used around 15 million core hours to develop its strategies and, during live [game](#) play, used 1,400 CPU cores. Pluribus computed its blueprint [strategy](#) in eight days using only 12,400 core hours and used just 28 cores during live play.

More information: N. Brown et al., "Superhuman AI for multiplayer poker," *Science* (2019). [science.sciencemag.org/lookup/...1126/science.aay2400](https://science.sciencemag.org/lookup/doi/10.1126/science.aay2400)

Provided by Carnegie Mellon University

Citation: At Last, AI beats professionals in six-player poker (2019, July 11) retrieved 17 April 2024 from <https://techxplore.com/news/2019-07-ai-professionals-six-player-poker.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.