

## AI beats professionals at six-player Texas Hold 'Em poker

11 July 2019

By Donna Lu

Artificial intelligence has finally cracked the biggest challenge in poker: beating top professionals in six-player, no-limit Texas Hold 'Em, the most popular variant of the game.

Over 20,000 hands of online poker, the AI beat 15 of the world's top poker players, each of whom has won more than \$1 million playing the game professionally.

The AI, called Pluribus, was tested in 10,000 games against five human players, as well as in 10,000 rounds where five copies of Pluribus played against one professional – and did better than the pros in both.

Pluribus was developed by Noam Brown of Facebook AI Research and Tuomas Sandholm at Carnegie Mellon University in Pennsylvania. It is an improvement on their previous poker-playing AI, called Libratus, which in 2017 outplayed professionals at Heads-Up Texas Hold 'Em, a variant of the game that pits two players head to head.

Part of what makes poker so difficult for AI to master is the huge number of possible actions to make, says Tristan Cazenave at the Paris Dauphine University. There are more possibilities than there are atoms in the universe.

It also involves hidden information, in which a player has access only to the cards that they see – meaning that an AI has to take into account how it would act with different cards so it isn't obvious when it has a good hand.

"If you look at real-world interactions, most of them involve hidden information, multiple participants or both," says Brown. Pluribus's approach could be applied to situations in cybersecurity, or in having self-driving cars navigate traffic, he says.

Pluribus learned to master the game by playing against five copies of itself, an approach that has been used by other AIs to master games such as Go, Dota 2 and StarCraft II. It started as a poker novice with no knowledge of the game, learning the rules over trillions of hands and improving its strategy by reviewing the decisions it made every round.

### Plays like a bot

In games against five human professionals, Pluribus won by an average of 48 milli-big blinds per game – a measure of how many big blinds were won on average per thousand hands of poker.

Each human player was given an alias for the duration of the tournament, to deter people who knew each other from potentially teaming up against Pluribus.

"We made no effort to hide who the bot was," says Brown, partially because its play style was obvious – Pluribus plays the first few actions in a round instantaneously because it has already prepared its strategy for those moves, while a human player typically takes a few seconds to decide.

Knowing which player was Pluribus meant the human player could attempt to trick the AI, says Jason Les, a professional poker player who was involved in the tournament. He played in the rounds that pitted five humans against Pluribus, playing an estimated 2000 hands over 12 days.

"You really want to push the AI, try everything you can to find a weakness," says Les. "Obviously we weren't able to."

Les also played against Libratus in 2017. "I was pretty amazed that they had made so much progress in just a couple of years," he says. "What was particularly impressive about this

challenge was that the AI played faster and on much less computing power.”

To reduce the number of potential choices that Pluribus needed to consider, the AI grouped similar hands – for example, a king-high flush and queen-high flush – and only considered a few different sizes of bets for a given hand.

“At the end of the day, betting \$150 is a lot like betting \$151,” says Brown. Instead of treating those bets separately, Pluribus groups them together and treats them identically.

### **No guarantees**

“We actually use very few computing resources to produce this AI,” says Brown. Training Pluribus required less than 512 gigabytes of memory, which would cost less than \$150 using cloud computing services.

While Pluribus played better than human poker players, according to a game theory

principle called the Nash equilibrium there was no theoretical guarantee it would always win, says Cazenave.

A Nash equilibrium occurs in non-cooperative games where each player has a list of strategies and no player can improve on their performance by changing to a different strategy. While a Nash equilibrium strategy is unbeatable in Heads Up Texas Hold 'Em, we still have no way of finding one for the six-player variant of the game.

“This is actually why the AI community finds this so surprising,” says Brown. “A lot of people didn’t think that this would be possible – to beat top humans using these techniques.”

Cazenave says that similar approaches could be used to develop AIs that can play other complex multiplayer games such as mahjong and bridge.

Read more: <https://www.newscientist.com/article/2209631-ai-beats-professionals-at-six-player-texas-hold-em-poker/>