

How rival bots battled their way to poker supremacy

Artificial-intelligence programs use game-theory strategies and deep learning to defeat human professionals in two-player hold 'em — a game with 'incomplete' information.

BY ELIZABETH GIBNEY

A complex variant of poker is the latest game to be mastered by artificial intelligence (AI). And it has been conquered not once, but twice — by two bots developed by rival research teams.

The algorithms play a 'no limits' two-player version of Texas Hold 'Em. And both have in recent months hit a crucial AI milestone: they have beaten human professional players.

The game first fell in December to DeepStack, developed by computer scientists at the University of Alberta in Edmonton, Canada, with collaborators from Charles University and the Czech Technical University in Prague. A month later, Libratus, developed by a team at Carnegie Mellon University (CMU) in Pittsburgh, Pennsylvania, achieved the feat.

Over the past decade, the groups have pushed each other to make ever-better bots, and now the team behind DeepStack has formally published details of its AI in *Science* (M. Moravčík *et al.* *Science* <http://doi.org/b2jd>; 2017).

Nature looks at how the two AIs stack up, what the accomplishments could mean for online casinos and what's left for AI to conquer.

Why do AI researchers care about poker?

AI has cracked several board games, including chess and the complex strategy game Go. But poker has a key difference from board games that adds complexity: players must work out their strategies without being able to see all of the information about the game on the table. They must consider what cards their opponents might have and what the opponents might guess about their hand based on previous betting.

Games that have such 'imperfect information' mirror real-life scenarios, such as auctions and financial negotiations, and poker has become an AI test bed for these situations.

Algorithms have already cracked simpler forms of poker: the Alberta team essentially solved a limited version of two-player hold 'em poker in 2015. The form played by DeepStack and Libratus is still a two-player game, but there are no limits on how much an individual player can bet or raise — which makes it considerably more complex for an AI to navigate.



JUICE/ALAMY

Top professional poker players have been beaten by artificial-intelligence bots at no-limits hold 'em.

How did the human-versus-AI games unfold?

Over 4 weeks beginning in November last year, DeepStack beat 10 of 11 professional players by a statistically significant margin, playing 3,000 hands against each.

Then, in January, Libratus beat 4 better professionals who are considered specialists at the game, over a total of around 120,000 hands.

What are the mathematics of the algorithms?

Game theory. Both AIs aim to find a strategy that is guaranteed not to lose, regardless of how an opponent plays. One-on-one poker is a zero-sum game — meaning that one player's loss is always the opponent's gain — so game theory says that such a strategy always exists. Whereas a human player might exploit a weak opponent's errors to win big, an AI with this strategy isn't concerned by margins — it plays only to win.

Previous poker-playing algorithms have tried to work out strategies ahead of time, computing massive 'game trees' that outline solutions for all the different ways that a game could unfold. But the number of possibilities is so huge — 10^{160} — that mapping all of them is impossible. So researchers settled for solving fewer possibilities. In a game, an algorithm compares a live situation to those that it has previously calculated.

It finds the closest one and 'translates' the corresponding action to the table.

Now, however, both DeepStack and Libratus have found ways to compute solutions in real time — as is done by AIs that play chess and Go.

How do the bots' approaches compare?

Instead of trying to work out a whole game tree ahead of time, DeepStack recalculates only a short tree of possibilities at each point in a game. The developers created this approach using deep learning, a technique that uses brain-inspired architectures known as neural networks (and that helped a computer to beat one of the world's best players at Go).

By playing itself in more than 11 million game situations, and learning from each one, DeepStack gained an 'intuition' about the likelihood of winning from a given point in the game. This allows it to calculate fewer possibilities in a relatively short time — about 5 seconds — and make real-time decisions.

The Libratus team has yet to publish its method, so it's not as clear how the program works. Early in a hand, it seems to use previously calculated possibilities and the 'translation' approach, although it refines the strategy as the game gives up more information. But for the

rest of each hand, as the possible outcomes narrow, the algorithm also computes solutions in real time. And Libratus has a learning element — a self-improvement module that analyses its playing strategy to learn how an opponent exploited its weaknesses.

The two methods require substantially different computing power: DeepStack trained using 175 core years, the equivalent of running a processing unit for 175 years. During games, it can run off a single laptop. Libratus, by contrast, uses a supercomputer before and during the match, and the equivalent of around 2,900 core years.

Can they bluff?

Yes. People often see bluffing as something human, but to a computer it has nothing to do with reading an opponent, and everything to do with the mathematics of the game. Bluffing is merely a strategy to ensure that a player's betting pattern never reveals to an opponent the cards that they have.

OK, so which result was more impressive?

It depends on whom you ask. Experts could quibble over the intricacies of both methods. But both AIs played enough hands to generate statistically significant wins — and both did so against professional players.

Libratus played more hands, but DeepStack didn't need to because its team used a sophisticated statistical method that enabled it to prove a significant result from fewer games. Libratus beat much better professionals than did DeepStack, but on average, DeepStack won by a bigger margin.

Will the two AIs now face off?

Maybe. A sticking point is likely to be the difference in computing power, and so the speed of play. This could make it difficult to find rules to which both sides can agree.

The match would carry a big caveat: the winner might not be the better bot. Both are trying to play the perfect game, but the strategy closest to that ideal doesn't always come out in head-to-head play. One program could accidentally hit on a hole in the opponent's strategy, but that wouldn't mean that the strategy has more or bigger holes overall.

Does this mean the end of online poker?

No. Many online casinos forbid the use of bots in matches, although top players have started to train against machines.

What's left for AI tackle?

The natural next target is multiplayer poker. This could mean almost starting from scratch, because zero-sum game theory does not apply: in three-player poker, for instance, a bad move by one opponent can indirectly hinder, rather than always advantage, another player. More-complicated games such as bridge are also unsolved. ■

PUBLISHING

Psychologists push for open data

Journal board member is asked to resign in peer-review spat.

BY GAUTAM NAIK

A consulting editor on the board of a journal published by the prestigious American Psychological Association (APA) has been asked to resign in a controversy over data sharing in peer review. Gert Storms — who says he won't step down — is one of a few hundred scientists who have vowed from the start of this year to begin rejecting papers if authors won't publicly share the underlying data, or explain why they can't.

The idea, called the Peer Reviewers' Openness Initiative, was launched by psychologists hoping to increase transparency in a field beset by reports of fraud and dubious research practices. And the APA, which does not ask that data be made available to peer reviewers or shared openly online, seems set to become an early testing ground for the initiative's influence. With Storms' situation still unresolved, the society's council of editors will discuss whether it should change its policies at a meeting in late March.

Storms, a psychologist at the Catholic University of Leuven in Belgium and a consulting editor for the APA's *Journal of Experimental Psychology: Learning, Memory, and Cognition*, accepted an invitation last year to review a study for the journal, and pointed out his new open-data policy. The journal's editor, Robert Greene, wrote back to say that Storms' stance set "a terrible precedent" because it was unfair to the author of the paper and opposed the APA's policies and the guidelines followed by other reviewers. "Given that your policy conflicts with that of the journal, I think that it's best that you step down from the editorial board," he wrote.

Storms refused, writing that he would continue to do what he thought was necessary to "prevent sloppy science". And he forwarded his correspondence to other editors. Two of them, Robert Hartsuiker and Marc Brysbaert, both psychologists at Ghent University in Belgium, wrote to Greene saying that they would quit if Storms was forced to resign. "The policy of asking people to leave rather than inviting a discussion and getting critical voices — I found that quite inappropriate," said Hartsuiker.

Greene, a psychologist at Case Western

Reserve University in Cleveland, Ohio, notes that Storms' stance is inconsistent with APA rules and that all submissions to the society's journals should be treated in the same way. "At this point we're just letting things lie and I am not removing [Storms] from the editorial board. We'll see what happens at the council of editors," he adds. The APA's journals publisher, Rose Sokol-Change, declined to comment on Storms' case. "While we support open sharing of data when it can be ethically shared, we leave the decision of whether to do so to the author," she said.

PSYCHOLOGY'S DATA STRIFE

The conflict marks the latest effort by some psychologists to change their discipline's policies on data sharing. The APA, like many publishers in the field, asks authors to make their data available to others after publication. But as far back as 2006, a study found that 73% of psychologists were unwilling or unable to do so, even though they had agreed to share (J. M. Wicherts *et al. Am. Psych.* **61**, 726–728; 2006).

Calls for change gathered force after 2011, when the full extent of years of research fraud by Dutch psychologist Diederik Stapel came to light. An investigative committee noted that Stapel often refused to share his scientific data with colleagues, including his co-authors.

Psychologists have since taken the lead on efforts to revisit earlier work, questioning statistical data and showing that some textbook findings in the field are hard to replicate. Simine Vazire, a psychologist at the University of California, Davis, says that data should be made available on publication and also to reviewers, even if on a confidential basis. Without such transparency, assessing a study "is like buying a used car without being able to look under the hood", she says. But by 2012, only 38% of researchers publishing in four APA journals shared their data when asked, Storms and others reported in a 2015 paper (W. Vanpaemel *et al. Collabra* **1**, 3; 2015).

Despite prolonged pressure, the APA hasn't changed its data policies for years, says Jelte Wicherts, a psychologist at Tilburg University in the Netherlands who co-authored the 2006 review and has since criticized the APA's policies and psychology's data-sharing standards. "My hope is they will change their view relating to the openness of data," he says. ■