THE WALL STREET JOURNAL.

FRIDAY, JULY 12, 2019

Computers Can Now Bluff Like a Poker Champ. Better, Actually.

A new artificial intelligence program is so advanced at a key human skill — deception — that it wiped out five human poker players with one lousy hand.

By Daniela Hernandez

A super-powered computer performed trillions of calculations to hone an elusive skill: bluffing. Then it proceeded to wipe out five human poker champions with one lousy hand.

"It [felt] very hopeless. You don't feel like there's anything you can do to win," said Jason Les, who's played poker professionally for 15 years and fell repeatedly to the poker bot's bluffs.

Researchers at Facebook Inc. and Carnegie Mellon University have built the most effective non-human bluffer to date. The bot, called Pluribus, is a leap forward in the capabilities of artificial intelligence — and also a cunning player of Texas Hold 'em.

Artificial intelligence systems developed at academic and corporate labs have a disturbingly good track record of beating humans at their own games. Pluribus, described in a paper published in the journal Science, follows in the digital footsteps of superhuman AIs that have vanquished humans at games like checkers, chess, "Jeopardy!", Dota-2 and Go. Two years ago, another AI system developed at Carnegie Mellon, called Libratus, even bested a poker star.

But Libratus beat just one human at a time in a two-player game. Pluribus beat five opponents at once — without breaking a sweat.

According to its creators, the new bot uses less than 128 gigabytes of memory while playing, and ran on two chips. By comparison, Libratus used 100 chips in its one-on-one poker matches. AlphaGo, developed by Alphabet Inc.'s Google DeepMind, used 1,920 chips against a human Go player. International Business Machines Corp.'s Deep Blue used 480 custom-designed chips against top chess champ Garry Kasparov, according to the paper.

Carnegie Mellon University Professor of Computer Science Tuomas Sandholm, in partnership with researchers at Facebook Inc., developed Pluribus. The artificial intelligence system reasoned and bluffed in real time to best five human poker champions at their own game. PHOTO: CARNEGIE MELLON UNIVERSITY

DeepMind declined to comment. IBM did not immediately respond to a request for comment.

Pluribus couldn't try to predict the endgame; playing poker against multiple opponents meant it had to be able to reason in real time, said Noam Brown, a Facebook AI research scientist and one of Pluribus's creators.

Multiplayer poker is considered less a game than an artform that requires a multitude of skills, especially the ability to read human interactions and leverage that knowledge to exploit mistakes and weaknesses.

Pluribus developed its winning poker strategy and superb bluffing skills by playing trillions of hands against five other clones of itself, said Dr. Brown.

After each round, it analyzed its decisions. If these resulted in wins, the bot would be more likely to opt for such moves in the future.

(over please)

Pluribus's digital brain realized it could win by making a bet when it had a weak hand by forcing its opponent to fold, which also taught it that it should bluff in future plays, said Dr. Brown. It then used those lessons to make real-time decisions when battling top human players, all of whom had earned more than $1 million playing professionally, according to the paper.

"People have this notion that [bluffing] is a very human ability — that it's about looking into the other person's eyes," Dr. Brown said. "It's really about math, and this is what's going on here. We can create an AI algorithm that can bluff better than any human."

But can they learn to love?

In one game against five human players, Pluribus was dealt an ace of clubs and a two of clubs, not a great hand. It started by raising $250, a standard move. Two humans called, the other two folded, for a total pot size of $800. Three community cards, a jack of spades, a five of diamonds, and a king of clubs, were dealt. (In Texas Hold 'em, players share a set of cards.) One of the humans checked. Pluribus bluffed and bet $800. Another player folded, while the other called.

The next card was a three of hearts, which killed the bot's chances of getting a flush. It still had a slim chance at a straight if it got lucky with the final card. It chose to bluff with a $2,400 bet, and the last standing human, Linus Loeliger, called. The final community card was an eight of spades. The AI's ace was most likely a losing hand, but it went all in, betting all its chips, worth $6,550. Mr. Loeliger folded. He had a ten of diamonds and a king of diamonds, a strong hand. He would have won.

Mr. Loeliger could not be reached for comment.

Unpredictability is what made playing Pluribus difficult, professional poker players said. It's also at the core of the advance, experts said. It exploits the very essence of poker, uncertainty, by using math.

Mr. Les, the other poker pro who lost to Pluribus and was also beaten by its predecessor, Libratus, said the new bot's moves were aggressive. AI development "is moving more rapidly than people realize."

Scientists are interested in building AIs that can play games like poker and StarCraft where uncertainty abounds because they are microcosms of the real world, which is unpredictable. Traditionally, AI has struggled in uncertain situations, limiting the range of applications to which it can be applied, according to AI experts.

In poker, "there's hidden information, and to make matters worse, your adversary know things you don't know," said Tuomas Sandholm, a Carnegie Mellon University professor and a Pluribus developer. "You have to think about whether your adversary is trying to deceive you."

The more adversaries, the more hidden information an AI needs to contend with. Previous poker-playing bots also had the ability to bluff, Dr. Sandholm said, but it's much harder to bluff successfully when juggling multiple opponents.

Poker allows researchers to test algorithmic strategies for dealing with the unknown and to build the foundations for software that can sleuth out fraud and deception in real-world settings. Dr. Sandholm is involved in two startups, Strategy Robot Inc. and Optimized Markets Inc., that are leveraging technology similar to that baked into Pluribus for applications in defense, financial services, gaming and health care. Facebook says it doesn't have immediate plans to commercialize the technology.

"A good AI has a ridiculously unfair advantage against humans: They don't get tired. They don't get hungry. They don't deal with emotions," said Michael "Gags" Gagliano, a professional poker player with 11 years of experience who also lost to Pluribus. Playing successfully hinges on skill, but also on how players deal with fatigue and stress, Mr. Gagliano said. Exploiting their lack of mental and physical stamina to force them to make a mistake is a big part of the game.

"The algorithm is not doing that. It's just sitting and waiting," he said. "Any time you slip up ... it's going to collect money in that situation. It's really difficult."

Mr. Gagliano said the experience made him more cognizant that poker, like the rest of our lives, is becoming increasingly more about the data.

"It made me think about the theory behind poker, and the math of each situation," said Mr. Gagliano. "Instead of playing the player, you're sticking to the stats."