

## Team reveals inner workings of victorious AI: Libratus AI defeated top pros in 20 days of poker play

December 18 2017



Jason Les, a professional poker player specializing in heads-up, no-limit Texas Hold'em, is watched by Tuomas Sandholm, professor of computer science at Carnegie Mellon University, as play gets underway in the Brains Vs. AI competition in January 2017. Libratus, an AI developed at Carnegie Mellon, beat Les and three other pros during the 20-day competition in Pittsburgh. Credit: Carnegie Mellon University



Libratus, an artificial intelligence that defeated four top professional poker players in no-limit Texas Hold'em earlier this year, uses a threepronged approach to master a game with more decision points than atoms in the universe, researchers at Carnegie Mellon University report.

In a paper being published online today by the journal *Science*, Tuomas Sandholm, professor of computer science, and Noam Brown, a Ph.D. student in the Computer Science Department, detail how their AI was able to achieve superhuman performance by breaking the <u>game</u> into computationally manageable parts and, based on its opponents' game play, fix potential weaknesses in its <u>strategy</u> during the competition.

AI programs have defeated top humans in checkers, chess and Go—all challenging games, but ones in which both players know the exact state of the game at all times. Poker players, by contrast, contend with hidden information—what cards their opponents hold and whether an opponent is bluffing.

In a 20-day competition involving 120,000 hands at Rivers Casino in Pittsburgh during January 2017, Libratus became the first AI to defeat top human players at head's up no-limit Texas Hold'em—the primary benchmark and long-standing challenge problem for imperfectinformation game-solving by AIs.

Libratus beat each of the players individually in the two-player game and collectively amassed more than \$1.8 million in chips. Measured in millibig blinds per hand (mbb/hand), a standard used by imperfect-information game AI researchers, Libratus decisively defeated the humans by 147 mmb/hand. In poker lingo, this is 14.7 big blinds per game



"The techniques in Libratus do not use expert domain knowledge or human data and are not specific to poker," Sandholm and Brown said in the paper. "Thus they apply to a host of imperfect-information games." Such hidden information is ubiquitous in real-world strategic interactions, they noted, including business negotiation, cybersecurity, finance, strategic pricing and military applications.

Libratus includes three main modules, the first of which computes an abstraction of the game that is smaller and easier to solve than by considering all 10161 (the number 1 followed by 161 zeroes) possible decision points in the game. It then creates its own detailed strategy for the early rounds of Texas Hold'em and a coarse strategy for the later rounds. This strategy is called the blueprint strategy.

One example of these abstractions in poker is grouping similar hands together and treating them identically.

"Intuitively, there is little difference between a King-high flush and a Queen-high flush," Brown said. "Treating those hands as identical reduces the complexity of the game and thus makes it computationally easier." In the same vein, similar bet sizes also can be grouped together.

But in the final rounds of the game, a second module constructs a new, finer-grained abstraction based on the state of play. It also computes a strategy for this subgame in real-time that balances strategies across different subgames using the blueprint strategy for guidance—something that needs to be done to achieve safe subgame solving. During the January competition, Libratus performed this computation using the Pittsburgh Supercomputing Center's Bridges computer.

Whenever an opponent makes a move that is not in the abstraction, the module computes a solution to this subgame that includes the opponent's move. Sandholm and Brown call this nested subgame solving.



DeepStack, an AI created by the University of Alberta to play heads-up, no-limit Texas Hold'em, also includes a similar algorithm, called continual re-solving; DeepStack has yet to be tested against top professional players, however.

The third module is designed to improve the blueprint strategy as competition proceeds. Typically, Sandholm said, AIs use machine learning to find mistakes in the opponent's strategy and exploit them. But that also opens the AI to exploitation if the opponent shifts strategy.

Instead, Libratus' self-improver module analyzes opponents' bet sizes to detect potential holes in Libratus' blueprint strategy. Libratus then adds these missing decision branches, computes strategies for them, and adds them to the blueprint.

In addition to beating the human pros, Libratus was evaluated against the best prior poker AIs. These included Baby Tartanian8, a bot developed by Sandholm and Brown that won the 2016 Annual Computer Poker Competition held in conjunction with the Association for the Advancement of Artificial Intelligence Annual Conference.

Whereas Baby Tartanian8 beat the next two strongest AIs in the competition by 12 (plus/minus 10) mbb/hand and 24 (plus/minus 20) mbb/hand, Libratus bested Baby Tartanian8 by 63 (plus/minus 28) mbb/hand. DeepStack has not been tested against other AIs, the authors noted.

"The techniques that we developed are largely domain independent and can thus be applied to other strategic imperfect-information interactions, including non-recreational applications," Sandholm and Brown concluded. "Due to the ubiquity of hidden information in real-world strategic interactions, we believe the paradigm introduced in Libratus will be critical to the future growth and widespread application of AI."



The technology has been exclusively licensed to Strategic Machine, Inc., a company founded by Sandholm to apply strategic reasoning technologies to many different applications.

A paper by Brown and Sandholm regarding nested subgame solving recently won a Best Paper award at the Neural Information Processing Systems (NIPS 2017) conference. Libratus received the HPCwire Reader's Choice Award for Best Use of AI at the 2017 International Conference for High Performance Computing, Networking, Storage and Analysis (SC17).

**More information:** Noam Brown et al. Superhuman AI for heads-up no-limit poker: Libratus beats top professionals, *Science* (2017). <u>DOI:</u> <u>10.1126/science.aao1733</u>

Provided by Carnegie Mellon University

Citation: Team reveals inner workings of victorious AI: Libratus AI defeated top pros in 20 days of poker play (2017, December 18) retrieved 28 April 2024 from https://techxplore.com/news/2017-12-team-reveals-victorious-ai-libratus.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.