


Game theorists crack poker

An 'essentially unbeatable' algorithm for the popular card game points to strategies for solving real-life problems without having complete information.

Philip Ball

08 January 2015



A new computer algorithm can play one of the most popular variants of poker essentially perfectly. Its creators say that it is virtually “incapable of losing against any opponent in a fair game”.

This is a step beyond a computer program that can beat top human players, as IBM's chess-playing computer Deep Blue famously did in 1997 against Garry Kasparov, at the time the game's world champion. The poker program devised by computer scientist Michael Bowling and his colleagues at the University of Alberta in Edmonton, Canada, along with Finnish software developer Oskari Tammelin, plays perfectly, to all intents and purposes.

That means that this particular variant of poker, called heads-up limit hold'em (HULHE), can be considered solved. The algorithm is described in a paper in *Science*¹.

The strategy the authors have computed is so close to perfect “as to render pointless further work on this game”, says Eric Jackson, a computer-poker researcher based in Menlo Park, California.

“I think that it will come as a surprise to experts that a game this big has been solved this soon,” Jackson adds.

A few other popular games have been solved before. In particular, in 2007 a team from the same computer-science department at Alberta — including Neil Burch, a co-author of the latest study — cracked draughts, also known as checkers².

But poker is harder to solve than draughts. Chess and draughts are examples of perfect-information games, in which players have

complete knowledge of all past events and of the present situation in a game. In poker, in contrast, there are some things a player does not know: most crucially, which cards the other player has been dealt. The class of games with imperfect information is especially interesting to economists and game theorists, because it includes practical problems such as finding optimal strategies for auctions and negotiations.

With regret

In poker, the main challenge is dealing with the immense number of possible ways that a game can be played. Bowling and colleagues have looked at one of the most popular forms, called Texas hold'em. With just two players, the game becomes heads-up, and it is a 'limit' game when it has fixed bet sizes and a fixed number of raises. There are 3.16×10^{17} states that HULHE can reach, and 3.19×10^{14} possible points at which a player must make a decision.

Bowling and colleagues designed their algorithm so that it would learn from experience, getting to its champion-level skills required playing more than 1,500 games. At the beginning, it made its decisions randomly, but then it updated itself by attaching a 'regret' value to each decision, depending on how poorly it fared.

This procedure, known as counterfactual regret minimization, has been widely adopted in the Annual Computer Poker Competition, which has run since 2006. But Bowling and colleagues have improved it by allowing the algorithm to re-evaluate decisions considered to be poor in earlier training rounds.

The other crucial innovation was the handling of the vast amounts of information that need to be stored to develop and use the strategy, which is of the order of 262 terabytes. This volume of data demands disk storage, which is slow to access. The researchers figured out a data-compression method that reduces the volume to a more manageable 11 terabytes and which adds only 5% to the computation time from the use of disk storage.

"I think the counterfactual regret algorithm is the major advance," says computer scientist Jonathan Shapiro at the University of Manchester, UK. "But they have done several other very clever things to make this problem computationally feasible."

Bluffing game

As part of its developing strategy, the computer learned to inject a certain dose of bluffing into its plays. Although bluffing seems like a very human, psychological element of the game, it is in fact part of game theory — and, typically, of computer poker. "Bluffing falls out of the mathematics of the game," says Bowling, and you can calculate how often you should bluff to obtain best results.

Of course, no poker algorithm can be mathematically guaranteed to win every game, because the game involves a large element of chance based on the hand you're dealt. But Bowling and his colleagues have demonstrated that their algorithm always wins in the long run.

The problem is only what the researchers call 'essentially solved', meaning that there is an extremely small margin by which, in theory, the computer might be beaten by skill rather than chance. But this margin is negligible in practice.

Bowling says that the approach might be useful in real-life situations when one has to make decisions with incomplete information — for example, for managing a portfolio of investments. The team is now focusing on applying their approach to medical decision-making, in collaboration with diabetes specialists.

Nature | doi:10.1038/nature.2015.16683

References

1. Bowling, M., Burch, N., Johanson, M. & Tammelin, O. *Science* **347**, 145–149 (2015).
2. Schaeffer, J. *et al. Science* **317**, 1518–1522 (2007).