



TRENDING: CES 2015 RISE OF DRONES NEW SPACE RACE

MORE ▾

VIDEOS **SUBSCRIBE**

THE COMPUTER THAT TAUGHT ITSELF TO BLUFF

TODAY AI CAN PLAY MISTAKE-FREE POKER; WILL IT SOLVE NATIONAL SECURITY NEXT?

By Andrew Rosenblum Posted 4 hours ago

 0 Shares



Poker Pot

Nintendere

Practice makes perfect, even if you

happen to be a piece of artificial intelligence.

That was the premise of an experiment led by Michael Bowling of the University of Alberta, which set up a program called Cepheus to play a billion billion (yes, a billion billion) hands of a poker variant called heads-up limit Texas Hold'Em against itself. Cepheus ran on 4,600 CPUs, considering 6 billion hands per second, learning from each victory, split pot, and defeat. After the equivalent of 1,000 years of CPU time during 70 actual days, Cepheus had played more poker than that played by the entire human race. In a paper published in *Science*, the Bowling team announced that with Cepheus, they had effectively "solved" heads-up limit Texas Hold'Em -- meaning that the program's decisions were so close to perfect that there was no way to see if a theoretically perfect human playing 200 hands an hour 12 hours a day over 70 years could do better.

The somewhat arcane statistician's definition of "solving" the game is necessary because Cepheus is not actually unbeatable in every hand – precisely because there is an irreducible element of luck in the game. A cunning pro can lose to a rank amateur if dealt a cruddy hand.

The program's decisions are so close to perfect that there is no way to see if a theoretically perfect human could do better.

"The worst case scenario is probably when you have a good hand, and your opponent has a better hand," says Mike Johanson, a co-author on the study. "You think you're going to win so you bet a lot and then lose a lot of money."

But once the effects of statistical noise vanish over thousands of hands, Cepheus' skill guarantees that it will not lose money in the long run.

"The first step was to build a program capable of beating human experts, which we did in 2008," says Johanson. "What we're announcing in the paper is that Cepheus is able to play essentially perfectly, without making mistakes."

In the version of Texas Hold'Em used in the study, 2 players ("Heads-up") compete using fixed ("limit") bet sizes, with 2 cards hidden.



What is really new here is that Cepheus had to learn to make decisions in spite of such “imperfect information” about what cards the opponent held.

**Mike Johanson,
Michael Bowling,
and Neil Burch**

*University of
Alberta*

Part of a University of Alberta team that built a computer that solved heads-up limit Texas hold'em poker.

Though computer scientists had previously solved “perfect information” games like Connect Four or Checkers, where the computer had full knowledge of previous moves and possible future outcomes, the Alberta study is the first solution of a nontrivial imperfect information game played by humans.

“And Cepheus had to learn how to play without human expert help,” says Johanson. “We taught it the rules, and it was trained against itself, figuring out this tricky psychological stuff like how to slow-play and to bluff.”

Bluffing occurs when a player has a weak hand but bets aggressively in order to fool the opponent into folding. Slow-play is the opposite; with a strong hand, the player wants to bet conservatively in order to lure

the opponent into staying in the game for additional bets. And the “imperfect information” about what the other player holds is what has traditionally made these psychological tactics so hard for computers to handle. Until now. And those billion billion hands of Texas Hold’Em. It’s safe to say that Cepheus has seen every trick a poker player might conceivably try to pull.

You can play against Cepheus [online](#), or ask it strategy questions.

Though poker is a big business, the game interests computer scientists more as a benchmark. And Johanson anticipates game theorists in other fields using the methodology for other fields in which imperfect information predominates, like negotiation or anti-terrorism.

For example, University of Southern California professor Milind Tambe has developed a game theory tool called [ARMOR](#) used by both Los Angeles International Airport and the Federal Air Marshals to schedule patrols and checkpoints in a way that incorporates randomization but also weights potential threats.

Airport security can be understood as a strategy game

like chess.

For Johanson, airport security can be understood as a strategy game like chess.

“Think of it like chess, except with different-sized armies. The airport has several pieces, representing security guards, maybe bomb-sniffing dogs, and checkpoints. Maybe the terrorist only has one piece, but a good one, like the queen, that can move freely and attack wherever the airport’s weakest.”

As in heads-up limit Texas Hold’Em, both sides in airport security have imperfect information about what moves the opponent will make.

“So you hide the board,” says Johanson. “You don’t know when the terrorist will attack, but you know an attack is probably coming. The terrorist knows there’s security, but doesn’t know exactly where it will be.”

Reconceiving a complex, massive airport like LAX as a chessboard is easier said than done. But the hope is that just as Cepheus taught itself to bluff and slow-play through trillions of rounds of practice, tomorrow’s repetitive AI security simulations will uncover vulnerabilities that we mere

mortals would never have thought to consider.

TAGS: GAMES, GAME THEORY, POKER, CARDS, TEXAS HOLD 'EM, MACHINE LEARNING, ML, AI, ENTERTAINMENT



0 Shares

ENTERTAINMENT

CES 2015: TV MAKERS AND HOLLYWOOD STUDIOS JOIN FORCES TO STANDARDIZE 4K TECHNOLOGY

THE UHD ALLIANCE WILL SET STANDARDS FOR ULTRA-HIGH DEFINITION AND 4K CONTENT

By **Michael Nuñez** Posted January 7, 2015



0 Shares