Computer poker program sets its own Texas Hold'em strategy

July 7 2006

A Carnegie Mellon University computer scientist has demonstrated that you don't necessarily need to know much about poker to create a computer program that can play a winning hand of Texas Hold'Em. A knowledge of game theory, not the specialized expertise of a human poker player, is at the heart of the poker robot called GS1 developed by Tuomas Sandholm, director of Carnegie Mellon's Agent-Mediated Electronic Marketplaces Lab, and graduate student Andrew Gilpin.

Though not yet the equal of the best human players, GS1 outperformed the two leading "pokerbots" in playing heads-up, limit Texas Hold'Em in tests at Carnegie Mellon earlier this year. Both of GS1's opponents were commercially available programs that, like other pokerbots, incorporate the expertise of human poker players. GS1, by contrast, develops its strategy after performing an automated analysis of poker rules. Sandholm and Gilpin have since developed an improved version of their game-theory-based program, called GS2, which will compete in the American

Association for Artificial Intelligence's first Computer Poker Competition during the 21st National Conference on Artificial Intelligence July 16-20 in Boston.

Much as computer chess was an early test of artificial intelligence (AI), computer poker has emerged as an even greater AI challenge. "Poker is a very complex game," said Sandholm, a professor of computer science in Carnegie Mellon's School of Computer Science. "Computer poker

programs really require sophisticated technology."

Unlike chess, where the status of all of the chess pieces is known to both players, poker forces players to make decisions based on incomplete information. "You don't know what the other guy is holding," Sandholm explained. And the sheer number of possible combinations of cards dealt, cards on the table and bets in two-player Texas Hold'Em games -- 1018, or a billion times a billion -- makes it impossible for even the fastest computers to fully analyze every hand.

This element of uncertainty and the vagaries of luck inherent in randomly dealt cards actually make poker a better test of AI's prowess than chess. "A lot of real-world situations have uncertainty in them and you have to deal with the uncertainty," Sandholm said. An algorithm (sequence of steps) that can capably play poker might also be useful in electronic commerce applications, such as sequential negotiation and auctions, he said.

Electronic commerce is a major research focus for Sandholm. He has developed the fastest algorithms for matching supply and demand, which can now be expressed in significantly more detail than before. He is the founder, chairman and chief scientist of CombineNet, a company that helps Fortune 1000 organizations save money and time on procurement. More than $20 billion has been sourced through CombineNet's system, generating in excess of $2.5 billion in savings for customers.

Using AI techniques to automatically set rules for electronic commerce is another direction Sandholm has pioneered. These programs generate mechanisms that can govern electronic auctions, elections or negotiations.

In his computer poker research, Sandholm has developed pokerbots that precompute the strategies for playing the first two rounds of Texas

Hold'Em, the so-called "pre-flop" and "flop" rounds, when players are first dealt two cards and then three additional cards are positioned face-up. For the third and fourth betting rounds, the "turn" and the "river," his pokerbots update the probability of each possible hand by taking into account betting as well as the revealed cards. The strategy for those rounds is then computed in real-time for the setting at hand.

To reduce the computational complexity, GS1 and GS2 automatically recognize strategically equivalent hands. For instance, 25,989,600 distinct hands are possible in the second round, but only about a million are strategically different. That's still too many to compute, so the pokerbots group strategically similar hands together. The end result is 2,465 groups, a small enough number to allow computational analysis.

Source: Carnegie Mellon University