

Playing Your Cards Right

Poker comes out of the back room and into the computer science lab

By IVARS PETERSON

Darse Billings plays poker for a living.

Just a few years ago, he was a computer science student at the University of Alberta in Edmonton. Interested in games, he had chosen the development of a poker-playing computer program as his thesis topic.

"I discovered that very little work had been done on computer poker despite the many potential benefits of that research," Billings says. Although he did not create a working program, he learned enough about the game along the way to begin playing professionally—which he has done since completing his thesis in 1995.

"Understanding the theory and mathematics of poker gives you a solid foundation," he contends. "This alone can put you ahead of the vast majority of players. Beyond that, you learn methods of analysis. It teaches you how to think about a given poker situation and that enables you to make more effective decisions in the heat of battle."

Billings now also serves as a consultant to a team of researchers at Alberta intent on developing a program that can play poker at the level of the best human players. "It's wonderful working with him because he understands computers and he understands poker," says Jonathan Schaeffer, who heads the group.

In an earlier effort, Schaeffer and his coworkers had created a checker-playing computer program named Chinook, which could beat the world's top players (SN: 7/20/91, p. 40).

The current version of the Alberta poker program—named Loki for the Norse god of mischief and chaos—already plays a strong game, Schaeffer says. "But it isn't

ready yet to win a world championship."

The first public demonstration of Loki is scheduled to take place later this month at the Fifteenth National Conference on Artificial Intelligence in Madison, Wisc.

Games have long been an important focus of research in computer science and artificial intelligence. With clearly defined rules and specific goals, "they are wonderful domains for testing ideas," says Matthew L. Ginsberg of the University of Oregon in Eugene.

Researchers have programmed computers to play chess, backgammon, bridge, Go, Othello, Scrabble, and numerous other games—in several instances eventually achieving a world championship level of play (SN: 8/2/97, p. 76; 8/16/97, p. 100).

Poker is an example of a game of incomplete information in which chance plays a role. Whereas a chess player sees the disposition of all the pieces all the time, a poker player sees only some of the cards—drawn or dealt from a shuffled deck—that are in play.

"You don't know what cards your opponent has," Schaeffer says. "All you can do is make educated guesses."

"One of the fundamental problems in computer science is how to deal with information that may be erroneous, unknown, or incomplete," Billings adds. Poker provides an excellent domain for investigating problems of decision making under uncertain conditions. Researchers can study such issues as risk assessment and management (betting strategy), opponent modeling (exploiting weaknesses in an opponent's play), and deception (bluffing).



Early efforts to model poker, some dating back to the 1950s, were generally unrealistic and rather limited, Billings says. Moreover, although there are many commercial poker-playing computer programs now on the market, they range widely in quality and lack the strategic flexibility and learning capability that a world-class player must have.

"Loki is probably better than commercial programs at evaluating the strength and potential of its cards," Billings notes. "Loki is also better at observing each opponent's actions over time and then adjusting its play accordingly."

Hundreds of books have been written about how to play poker. Billings contends that the vast majority give flawed advice on strategy. The typical level of human play is so low, however, that a contender can be highly successful despite serious misconceptions, he argues. Players often fail to exploit their opponents' weaknesses, or they make more serious errors themselves.

In recent years, books by experts such as David Sklansky and Mason Malmuth have pioneered a more systematic, mathematical approach to the game. Their analyses serve as a useful foundation for computer poker.

"In theory, poker is a game of probabilities, where the expected value of each betting decision can be determined by applying certain fundamental principles," Billings says. "In practice, there is much more to the game than the math can convey. Poker is a 'black art' because it is played against humans who are far from perfect."

Schaeffer and his coworkers have con-

centrated on a popular poker variant called Texas Hold'em, which is commonly played in casinos and is the version featured at the annual world series of poker. Each player holds two cards that are hidden from the others, and those cards can be combined with any three out of another five cards—dealt face-up in the middle of the table and shared by all players—to make the best possible five-card hand.

To the Alberta group, the key elements of the game are card probabilities, betting strategy, bluffing approach, and opponent modeling.

"In the chess and checkers world, you can safely assume that your opponents don't make mistakes, and you just play the best possible move," Schaeffer says. In poker, "you've got to watch what your opponent does, make inferences, and exploit that knowledge."

If you're playing against a cautious player who bets heavily on a certain hand, for example, you should fold, because he or she probably holds very good cards. Such a response isn't necessarily a good strategy against an opponent who consistently bets aggressively. That player could be bluffing a lot of the time.

"We've been working on categorizing and understanding styles of play," Schaeffer says. "Unfortunately, it gets more difficult in games against top players because they mix up their play deliberately to confuse you."

Being unpredictable is a good strategy, he adds. "If our computer is going to beat the world champion, it's got to do something similar."

Opponent modeling may not be essential to building a winning poker program, says computer scientist Daphne Koller of Stanford University.

In chess, human players often take into account an opponent's idiosyncrasies, past history, and state of mind, Koller argues. "You could make the case that opponent modeling is also very important in chess," she adds. "Nevertheless, chess-playing computers don't do that, and they do very well despite that limitation."

Koller is developing a poker-playing program based on game theory, which offers mathematical techniques for determining optimal strategies. Her approach exploits a remarkably efficient algorithm that she developed for searching through long sequences of possible moves to find the best one in games involving incomplete information. That algorithm applies to a wide variety of games.

"It's the kind of search algorithm we have long had for chess," Koller says.

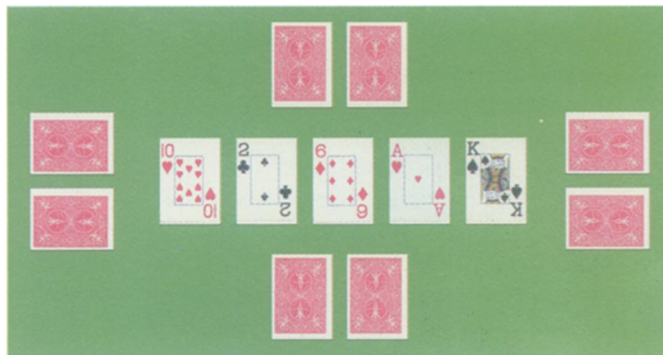
Poker serves as a test case to see if

Koller's algorithm can be scaled up to larger games, including those that economists and social scientists study to model human behavior in a variety of contexts.

"For me, this is more an exercise in pushing the boundaries of game theory than in specifically solving poker," Koller says.

Koller and Schaeffer "are taking totally different approaches," Ginsberg notes. Koller's program uses an extensive search for the best possible strategy, whereas Schaeffer's program depends heavily on exploiting weaknesses in play.

Interestingly, Koller has found that bluffing, which is generally considered an innately human capability, emerges naturally from her search program. "You would think that bluffing is something that you do to take advantage of an



The computer program Loki plays a strong, but not championship-quality, game of the popular poker variant called Texas Hold'em, in which a player's two cards can be combined with any three of the five dealt face-up on the table.

opponent," she says. Game theory also calls for bluffing to obtain optimal results.

It would be interesting to see how Koller's poker player matches up against Loki. Such a confrontation is probably at least a year away, however. Unable to find the time necessary to complete the project, Koller is now looking for student help to finish the programming.

After several rounds of successive improvements in probability calculations, betting and bluffing strategies, and opponent modeling, Loki can play reasonably well.

At the moment, "the limiting factor is Loki's unsophisticated betting strategy," Billings says. For example, when the program holds a strong hand, it always makes a bet instead of occasionally holding back. Good opponents can take advantage of such a pattern.

"It's absolutely amazing how strong players can pick up on any kind of pattern in your play and exploit it," Schaeffer says. "Whatever your decision-making process, you've got to make random choices sometimes."

"The next major thrust in the research is to have Loki weigh the relative merits of different betting plans and to randomize those actions in a manner that con-

ceals information from the opposition," Billings remarks.

Loki has already faced a variety of opponents on the Internet, and it wins consistently. The problem is that these human players have just pride, rather than money, at stake.

"It's very difficult to evaluate the program," Schaeffer says. "In chess and checkers, you can go to a tournament, pay the entry fee, and play to see how good the program is. Poker is different because it's meaningless without money." If a computer plays for money, however, who's going to pay for the program's losses?

Further improvements may make it worthwhile for Schaeffer to bring Loki to Las Vegas. "I would love to say that a year from now, we'll be sitting down with some of the best players in the world, but that remains to be seen," Schaeffer says. Many technical problems have yet to be solved.

Moreover, the sorts of expenses associated with a trip to Las Vegas aren't normally covered by a research grant. "We may need to find someone to bankroll our venture," Schaeffer says.

Meanwhile, at the age of 36, Billings makes a fairly comfortable living on the Edmonton poker circuit—though he has to cope with large fluctuations in income from week to week.

"The modern game of poker is a game of small percentages," Billings says. "It's not like you're hauling off wheelbarrows of money all the time. It's an accumulation over time. You just earn more than you lose—maybe winning in seven sessions, breaking even in one, and losing in two for every 10 that you play."

The lifestyle associated with the game isn't particularly appealing to Billings. "I'm considering abandoning it," he says. He knows that if he were to play in Las Vegas casinos, for example, where the financial rewards can be much greater, he would have to improve his game considerably without a guarantee of an increase in income.

Billings is thinking about writing a book, particularly on how theory can fall short in practice. "I've always been very cognizant of the limitations of the theory and where the theory breaks down," he says. "That has come into much clearer light through the process of mastering the game in real life."

"The real essence of the game is knowing your opponent and adjusting accordingly," he adds.

That's a lesson that Schaeffer and other researchers believe can be transferred from poker to many real-world situations involving conflict and negotiation. □