IN THE NEWS

Features Editor: **Lee Garber**, lgarber@computer.org

Anteing up for Poker Bots

George Lawton

For every system that involves making money, someone will usually figure out a way to use technology to defeat it. This is exactly what has happened with online poker, an increasingly popular and potentially lucrative activity.

Some online poker players now use AI-based bots specially designed to play poker. Others use AI-enabled analysis tools that can help identify weak players or players using weak strategies. Such tools are openly available online.

Because of the disadvantage this creates for players not using technical assistance, the largest online poker sites—including PokerStars, Full Tilt Poker, and PartyPoker—monitor their sites for poker bots and shut them down when they can. In some cases, poker sites have also confiscated bot users' winnings and given them to the players they beat.

"Poker is a contest between humans of various abilities. It is not a test of programming skills, and therefore, software that plays on behalf of a player is prohibited," said PokerStars security manager Michael Josem.

According to Michael Bowling, associate professor of computer science at the University of Alberta, serious poker research could improve AI techniques for strategic analysis and improved decision making in negotiations and auctions.

He added, "Poker is the quintessential game for understanding how to deal with a situation in which you don't have all of the information. It is something that humans do brilliantly. The insights we gain and the algorithms we are developing in poker research can be applied [in other areas]."

Inside Poker Bots

Poker bots are not new, but until recently, they required too much computational power to play well. Humans could beat early poker bots using techniques such as bluffing.

Arizona State University professor Nicolas Findler started the early poker-science research in 1961 but was constrained by existing computers' performance limitations.

In 1984 Mike Caro, a professional poker player and computer programmer, demonstrated Orac, a poker-playing program at the World Series of Poker.

In 1995, Darse Billings, as part of his doctoral research at the University of Alberta, conducted research into implementing AI algorithms capable of playing poker in real time.

Since then, companies such as WinHoldem, Shank Bot, OpenHoldem, and ICM-Bot have released poker bots for sale online. In some cases, these tools automate play by automatically folding bad hands or using better card-drawing strategies. The services provide preconfigured templates to make it easier to connect to different poker-playing services.

A second group of tools, which includes Poker Edge, Poker Crusher, IdleMiner, and HandHQ, use AI to analyze the way various participants play. This is designed to let users identify weaker opponents or choose winning strategies against specific players.

In general, online poker sites permit these tools if players use them to

1541-1672/11/$26.00 © 2011 IEEE
Published by the IEEE Computer Society

analyze their own games, noted Josem. But they are prohibited when users exchange electronic databases of other participants' playing histories.

Services such as SharkScope, Top Shark, and Smart Buddy use external sources of information, analyze and aggregate players' histories, and provide it to customers as a service. These services don't use bots, which are fully or partially automated, but they provide tools that use AI decision-making algorithms to provide recommendations on who to play with.

Other tools, including Poker Calculator Pro, SNG Endgame Tools, and Poker Academy, use AI techniques to help a player understand the odds of different hands and different plays. According to Josem, these tools are acceptable for offline training but are not permitted on major online poker sites in real-time game play.

Major poker sites declined to comment on how they are able to detect when players are employing poker bots. Poker-bot vendors also declined to comment for this article and would not say how successful their products have been at making money for users.

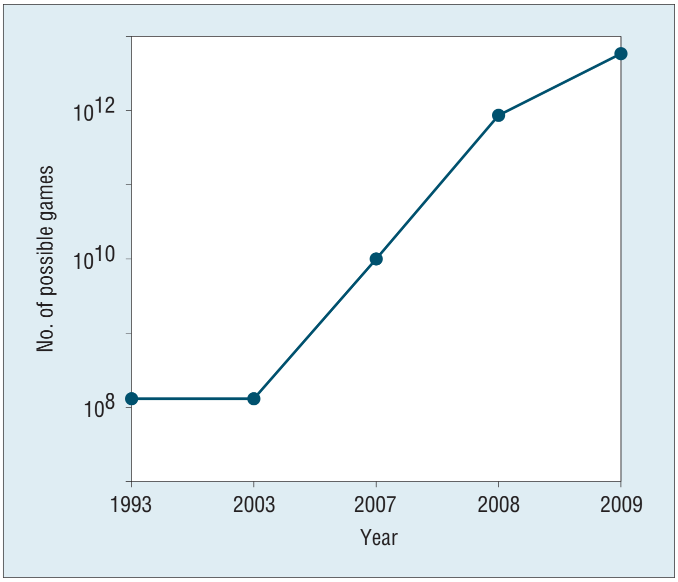

Figure 1. During the last few years, state-of-the-art techniques have been able to determine the best moves in games that have many possible states, such as poker. The number of possible games on the chart indicates games with the greatest number of possible states that state-of-the-art techniques can solve.

Poker-Science Advances

Poker consists of a family of games. Bots' applications and challenges vary depending on the type of game and betting structure used. Most of the science has been based on understanding two-player, winner-take-all scenarios.

One problem is that the underlying mathematics is more complicated for poker than for other games, Bowling noted. For example, a full game of checkers has about 10^{18} possible game states. But it is relatively easy to determine the best strategy because only a set number of moves can be made and there is no information hidden from either player.

A two-player game of Texas Hold'em poker with betting limits also has 10^{18} possible states, representing all combinations of cards and bets. (Games with no betting limits have more variability, so the number of possible states jumps to 10^{60}.) However, it is not perfectly solvable because of factors such as bluffing and cards not visible to opponents, which makes computing the merits of different moves and strategies more challenging.

In poker research, and decision theory in general, algorithms are judged based on the Nash equilibrium, a statistical approach designed to determine the likelihood that a strategy is guaranteed to win against a "perfect" opponent, one who always makes the best move possible.

A bot's success in poker is quantified by how much money a chosen strategy could keep a user from losing against a perfect player. Bowling said his latest poker bot, running on a 24-core server with 256 Gbytes of RAM, computes strategies that lose an average of $1 per bet against a perfect opponent. However, he noted, the bot would do better against average poker players.

Reducing the Possibilities

One research avenue involves reducing the number of possible moves that poker bots must consider before computing the perfect strategy. In a two-player Texas Hold'em game with betting limits, current poker bots' algorithms can whittle down the number of possible game states that must be considered by a factor of 1 million, from the maximum 10^{18} to 10^{12}, which a large computer can readily compute.

These algorithms use a combination of statistical and symbolic techniques to find patterns—in this case, poker hands—with roughly the same significance. This reduces the number of patterns the evaluation algorithms must consider.

Over time, poker bots have gotten better at figuring out the best moves for games with many possible states, as Figure 1 shows.

There are ways to further reduce the number of states that bots must consider. For example, there is less of a difference in significance between bets of $30 and $31 than between wagers of $30 and $40. Therefore, considering plays based on larger categories of bets, rather than individual amounts, could reduce the number of possible game states that must be computed.

Reducing Memory Requirements

The algorithms for evaluating different poker strategies are even more limited by RAM requirements than by processing capabilities, according to Bowling. He said the biggest poker-research advancements have been in memory management. Researchers have reduced the amount of memory required to compute the merits of all possible moves in a game.

Older algorithms used linear programming approaches—a mathematical technique, based on linear equations, for determining how to achieve the optimal outcome for some list of requirements—with memory requirements that didn't scale well as the number of game states grew.

Researchers have since used counterfactual regret minimization AI techniques that let poker bots focus on a smaller set of moves with a higher probability of long-term success. These techniques don't work as well for no-limit or multiplayer games.

Computing the perfect two-player game with betting limits would require 2 to 4 Pbytes of memory using existing algorithms, which is beyond most systems' capabilities. The latest generation of bots can compute a game that loses $1 per hand using only 30 to 40 Gbytes of RAM, compared to a terabyte for older algorithms.

> The latest generation of poker bots can compute a game that loses $1 per hand using only 30 to 40 Gbytes of RAM, compared to a terabyte for older algorithms.

Faster Strategic Analysis of Other Players

One problem for poker systems is the number of hands required to evaluate the capabilities and weaknesses of other players. In the early days of poker research, algorithms required 40,000 hands to accomplish this.

Using variance-reduction techniques, poker researchers have built tools that can assess players or strategies in only 500 hands. These approaches accomplish this by identifying information that can be ignored and by measuring the frequency and magnitude of differences between a player's actions and those of a perfect player.

Bowling expects some of the advances in his poker research to help create better AI agents for e-commerce. He said many AI algorithms from poker research also apply to other forms of exchange, such as auctions, that involve incomplete information about participants' activities. Companies or individuals could use the agents to bid on products, such as on eBay, or create investment strategies in financial markets.

Meanwhile, some people might use poker to learn more about decision making in general, said Josem. "Poker is an outstanding way for people to learn about risk assessment," he explained, "especially in making good decisions with incomplete information."

Using AI to Catch Criminals

George Lawton

One of the oldest law-enforcement techniques for catching criminals involves an artist's sketch of a suspect based on witness testimony.

Now, a team of academic researchers has developed AI algorithms for better matching forensically derived police sketches to criminals' mug shots (see Figure 2). The team includes Michigan State University Professor Anil Jain, director of the Department of Computer Science and Engineering's Biometrics Research Group, as well as doctoral students Brendan Klare of MSU and Zhifeng Li of Northwestern University.

Forensic sketches represent a different abstraction of a face than a photo. The drawings portray the general appearance of specific features more than an exact photographic representation of a face, explained Officer Rick Horwood, an East Lansing Police Department sketch artist.

Trained officers can use these sketches to get an idea of a suspect's appearance, which can help with recognition and subsequent apprehension, he said.

Seeking to develop a way to correlate a sketch and a photo, MSU researchers developed the Local Feature Discriminant Analysis pattern-matching algorithm. LFDA is a customized adaption of local discriminant analysis—a statistical technique used to determine which known classification or group an item or object belongs to on the basis of its characteristics or essential features—employed in other AI research areas.

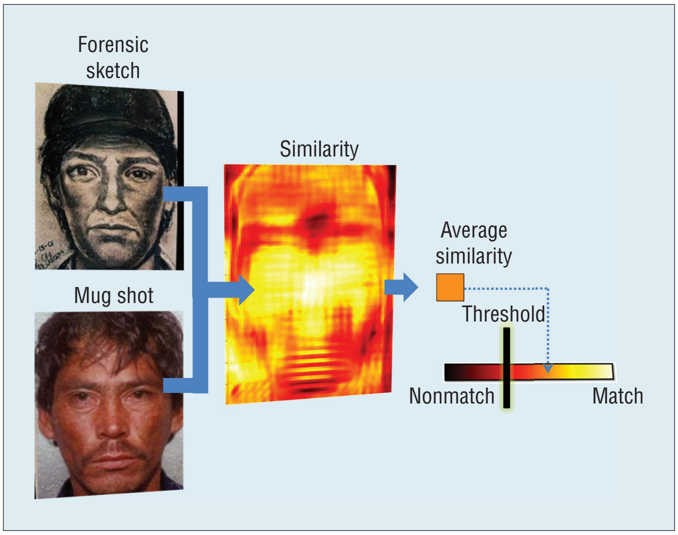

Figure 2. The Local Feature Discriminant Analysis pattern-matching algorithm. Michigan State University developers trained the LFDA algorithm to generate a list of possible suspects based on forensically derived police sketches.

The research team trained LFDA with hundreds of examples of matching faces from mug shots and forensic sketches.

LFDA uses rules to transform the number of variations possible in an image—photograph or sketch—into a smaller subset of data optimized for identifying different classes of objects such as faces. This makes the pattern-matching process more efficient and lets the system generate a list of possible suspects with less computational horsepower, noted Klare.

Police often use facial-recognition software to generate a list of the 50 top suspects in a particular crime. Klare said LFDA matched the correct mug shot to one of the top 50 picks 32.65 percent of the time, compared to only 8.16 percent for the leading facial-recognition application, Cognitec System's FaceVACS. When race and gender information is included in the matching process, the MSU system's accuracy increased to 44.9 percent, compared to 26.53 percent for FaceVACS.

The researchers are now trying to enable their tools to correlate infrared images taken by night-vision cameras with those taken by traditional cameras and to match 3D images of people taken by multicamera arrays with mug shots.

According to Klare, the same basic techniques could create tools to more accurately age photos of suspects or missing persons. Older tools used models of aging processes. The MSU team is working with thousands of sets of images of older and younger versions of the same people to develop more accurate algorithms.

The researchers have no immediate plan to commercialize their system.

AI Improves Robotic Sensors

George Lawton

A Massachusetts Institute of Technology and University of Southern California project has developed a set of AI learning algorithms that improve the capabilities of robotic sensors that could help in responses to disasters such as the massive Tohoku earthquake and subsequent tsunami that struck Japan on 11 March.

A key advance has been improvements in algorithms for maximizing the effectiveness of autonomous robot sensors' geographic distribution. Another has been reducing uncertainty about their measurements' accuracy in areas undergoing rapidly changing conditions, such as those where a natural disaster has occurred, noted MIT professor of computer science and electrical engineering Daniela Rus.

The researchers have conducted several experiments with their algorithms via ground robots, flying drones, and robot submarines.

Better Sampling

The first of the researchers' adaptive-sampling algorithms improve mobile-sensor platforms (such as those used for environmental testing in the ocean) that monitor a large area for long periods. Adaptive-sampling systems select regions to monitor based on values of interest observed during a sampling survey.

The new algorithms direct sensors to the parts of the environment with the greatest presence of data of interest—such as water temperature or the concentration of specific chemicals or organisms—to make sure they go where most needed or useful. The algorithms project an array of polygons onto a model of the environment and calculate a level of importance for each polygon.

"Our work is guaranteed to distribute the sensing devices to approximate the event distribution," said Rus.

In addition to enabling robot sensors to respond to existing events, the algorithms would help them redistribute themselves in response to new events. According to Rus, during a testing mission, the machines automatically went to places where oceanographers independently detected activity targeted for monitoring.

The MIT researchers are exploring ways to use the same types of algorithms to optimize the distribution of other types of mobile sensors, such as cameras mounted in drone aircraft. There might be advantages to being able to accurately move a camera sensor to different places or altitudes to better detect events of interest. In case of a disaster such as the Japanese earthquake and subsequent nuclear reactor leak, the new algorithms could help direct radiation sensors to the areas experiencing the biggest problems.

Calculating Uncertainty

The MIT team is also identifying AI algorithms for reducing the uncertainty of sensors' measurements during changes in the environment. In this case, AI-based sensor controllers create an environmental model and the rate of change of variables of interest, such as temperature, chemical density, salinity, or radioactivity.

Based on this, the MIT algorithm computes the way robots should change speed along their routes. The robots could slow and gather more information in areas of frequent activity or could traverse a region more often but collect less data when more samples are appropriate.

Researchers anticipate that these kinds of robots will make environmental sampling easier and open the door for new kinds of oceanographic research.

The algorithms also automatically and independently direct robots to places where the variables of interest change the most frequently. They then calculate the level of uncertainty of measurements in different areas.

The robot models the environment in which it operates and updates the model based on changes in environmental measurements. Changes in the variables it measures can follow the model. When this happens, the algorithm can take fewer measurements. When this doesn't happen, which can occur during disasters, the algorithm assigns a higher level of uncertainty to measurements from that location.

"This class of algorithms directs the robots to spend more time sampling the regions of the highest uncertainty," explained Rus. This is desirable because these are the areas where disasters and other anomalous events have the greatest impact.

The new algorithms have been deployed for plankton studies in a set of ocean-monitoring robots by Teledyne Webb Research, a company that develops ocean instrumentation. These robots autonomously monitor ocean environments for long periods. Each airplane-shaped robot has an inflatable bladder on its nose. The device rises or falls and thereby catches ocean currents as the bladder fills or empties, respectively. Algorithms can control the bladder and thereby direct the robot to move as desired.

The research team is also working on a controller for autonomous robotic taxis that dynamically send the vehicles to locations with highest demand, as part of the Singapore-MIT Alliance for Research and Technology (Smart).

In the long run, researchers anticipate that these kinds of robots will make environmental sampling easier and open the door for new kinds of oceanographic research.

Gaurav S. Sukhatme, professor of computer science and codirector of USC's Robotics Research Lab, said the research will lead to tools that even oceanographers without robotics expertise can use effectively.

The 10th International Semantic Web Conference

ISWC 2011

October 23-27, 2011
Bonn, Germany

ISWC 2011

ISWC is the major international forum where the latest research results and technical innovations on all aspects of the Semantic Web are presented.

In Bonn, scientists from all over the world will meet to present novel and innovative research that demonstrates new trends and addresses scientific challenges for the Semantic Web.

Bonn is easy to get to being within a little more than one hour's reach from two international airports (Frankfurt am Main and Cologne). Venue: Maritim Convention Center.

Key Topics

- Management of Semantic Web Data
- Natural Language Processing
- Ontologies and Semantics
- Semantic Web Engineering
- Social Semantic Web
- User Interfaces to the Semantic Web
- Applications of the Semantic Web

Keynote Speakers

- Sandy Pentland
  Massachusetts Institute of Technology (MIT), US
- Gerhard Weikum
  Max-Planck-Institute for Informatics (MPI), DE
- Frank van Harmelen
  Vrije Universiteit Amsterdam, NL
- Dinner Talk: Nigel Shadbolt
  University of Southampton, UK

Main Events

- Research Track presents original, principled research papers addressing both the theory and practice of Semantic Web technologies.
- Semantic Web In Use Track presents papers describing innovative use of Semantic Web technologies in real-life applications.
- Doctoral Consortium provides an opportunity for doctoral students to present their current progress and future plans and to get advice on their work.
- Posters and Demos present the latest developments in the field and showcase Semantic Web systems. Submissions due August 15, 2011.
- 6 Tutorials offer hands-on guidance on using key Semantic Web technologies.
- 16 Interactive Workshops discuss hot topics and new research areas.
- Semantic Web Challenge provides a forum for the best Semantic Web applications to compete.
- 2nd Linked Data-a-thon The main goal of the competition is to develop innovative applications that showcase the benefits of Linked Data ... in two weeks!
- Industrial Track Submissions due August 1, 2011.
- Outrageous Ideas Track
  Submissions due September 5, 2011.

Join the ISWC 2011 Community

twitter www.twitter.com/iswc2011

facebook www.facebook.com/iswc2011

WeST People and Knowledge Networks

gesis Leibniz Institute for the Social Sciences

Platinum Sponsors

Artificial Intelligence

fluid Operations

Gold Sponsors

YAHOO! RESEARCH

Microsoft Research

ontoprise know how to use Know-how

poolparty

http://iswc2011.semanticweb.org
E-mail: confsec@uni-koblenz.de