Applying Artificial Intelligence Techniques to the Game of Texas Hold‘em.

Jonathan Rubin[1] and Ian Watson[2]
Department of Computer Science, University of Auckland, New Zealand
[1] jrub001@ec.auckland.ac.nz
[2] ian@cs.auckland.ac.nz

Overview

This extended abstract refers to research which is currently being undertaken as part of a Masters thesis into the investigation of artificial intelligence techniques applied to the game of poker. This research will focus specifically on a variation of poker known as Texas hold‘em (described below). Texas hold‘em provides a non-deterministic, hostile environment in which players must deal with incomplete information and uncertainty. Decision making in this type of environment has not been well addressed by A.I games research in the past and it is believed that advances made in games of this sort will also reap rewards in real-world problems as well.

Games and Artificial Intelligence

Games provide a well suited domain for artificial intelligence research. This is due to the fact that a game is usually composed of several well defined rules which players must adhere to. For a large majority of games the rules imposed are quite simple, yet the game play itself involves a large number of very complex strategies. This is especially true of games such as chess and checkers which offer opportunities to make very sophisticated and intricate plays. This statement is also true of the game of Texas Hold‘em and is nicely summed up by a popular quote coined by Mike Sexton which states “The name of the game is No Limit Texas Hold ‘em, the game that takes a minute to learn but a lifetime to master”. Another reason why games offer a beneficial environment for artificial intelligence research is the fact that goals and objectives of the game are clearly defined. This is advantageous to research as a performance metric is implicitly embedded in the game. Success can easily be measured by factors such as the amount of games won, the ability to beat certain opponents or, as in the game of poker, the amount of money won.

Up until recently artificial intelligence research has mainly focused on games such as chess and checkers. Successes like *Deep Thought*, *Deep Blue* and *Chinook* are usually the first to come to mind when contemplating A.I and games. Games such as chess, checkers and backgammon are classified *as two-person, zero-sum* games with *perfect information*. This means that there is one winner and one loser (zero-sum) and the entire state of the game is accessible by both players at any point in the game (perfect information), i.e. both players can look down upon the board and see all the information they need to make their playing decisions. These types of games have achieved their success through the use of fast hardware processing speeds, selective search, effective evaluation functions and better opening books and endgame databases. While these achievements are remarkable, their scope is rather limited. They offer little insight into other areas where A.I. techniques may be useful.

Games such as poker on the other hand are classified as stochastic, imperfect information games. The game involves elements of chance, the actual cards which are dealt, and hidden information in the form of other player’s *hole cards* (cards which only they can see). This ensures that players now need to make decisions with uncertain information present. This is still an open research question in the A.I community and research efforts are likely to be beneficial outside the realm of poker itself. For A.I. to be useful for most real world problems, challenges that imperfect information and a stochastic environment offers need to be addressed.

The Game of Poker

There are numerous variations of the game of poker available. The games differ by various aspects such as the number of *hole cards* dealt (cards which only the player can see and use to make their best hand), the number of community cards dealt (cards which all players can see and use to make their best hand), the order in which players bet and the limits imposed on a players bet. There are two variations which control the amount that a player may bet: *limit* and *no limit.* In a *limit* game player's bets are restricted to a certain amount; this amount usually doubles in later rounds of betting. Conversely, in *no limit* there is no restriction on the amount that a player can bet. A player's betting decision can be to *fold*, *check*, *call*, *bet* or *raise*. These are described below:

Fold: A player can *fold* their cards if they are facing a bet by another player, but they don't wish to match the bet. Once a player *folds* they are no longer involved in the current hand, but can still participate in any future hands.

Check/Call: When it comes time for a player to make his/her decision they can *check* if there have been no bets made by other players. *Checking* means the player does not need to invest any of their money into the pot to stay in the current hand. If, however, an opponent has made a bet then a player can *call* the bet by adding to the pot the exact value of the current bet. By contributing their own money to the pot they are able to stay in the current hand.

Bet/Raise: A player can invest their own money to the pot over and above what is needed to stay in the current round. If the player is able to *check*, but they decide to add money to the pot this is called a *bet*. If a player is facing a bet from an opponent, but instead of deciding to just *call* the bet they decide to add more money to the pot then this is called a *raise*.

The Game of Texas Hold ‘em

In the game of Texas hold‘em players are dealt two *hole cards* and five community cards are used in total. This strikes the right balance in terms of information availability (Harrington and Robertie, 2005) and offers opportunities for better strategic play than other poker variations allow for. Texas hold‘em also offers a better skill-to-luck ratio than is offered by other forms of poker. An expert hold‘em player has more of an advantage because the best hand holds up more often than in any other poker variation (Sklansky and Malmuth, 1994). Play in hold‘em proceeds in the following stages: *preflop*, *flop*, *turn* and the *river*. These are described below:

Preflop: The game of Texas hold‘em begins with each player being dealt two *hole cards* which only they can see. Betting order is determined by assigning one player at the table the status of *dealer*. Betting proceeds round the table in a clockwise manner. The minimum size of a bet is determined by the *big blind*. If a player, wishes to play then they must pay at least the *big blind* to enter into the pot. As long as there are at least two players left then play continues to the next stage. During any stage of the game if all players, except one, fold their hands then the player who did not fold his/her hand wins the pot (without having to reveal their *hole cards*) and the hand is over.

Flop: Once the *preflop* betting has completed three community cards are dealt. Players use their *hole cards* along with the community cards to make their best hand. Another round of betting occurs. The player classified as *dealer* is always the last to act (if the *dealer* is no longer in the hand the first active player to the right of the *dealer* becomes the last player to act). As long as there are at least two players left then play continues to the next stage.

Turn: The *turn* involves the drawing of one more *community card*. Once again players use any combination of their *hole cards* and the community cards to make their best hand. Another round of betting occurs and as long as there are at least two players left then play continues to the next stage.

River: During the *river* the final community card is dealt proceeded by a final round of betting. If at least two players are still active in the hand a *showdown* occurs in which both players reveal their hole cards and the player with the highest ranking hand wins the entire pot (if both players hold hands of the same value then the pot is split between both players).

Proposed Solution

Any attempt to develop a strong poker player needs to address many areas of the game. Several key components required for strong poker play have been identified (Billings et al, 2001). These include **hand strength**, **hand potential**, **betting strategy**, **bluffing**, **unpredictability** and **opponent modeling**. The strength of a particular hand and the potential strength of a hand needs to be determined given the *hole cards* that a player possesses, the current community cards, the type and number of opponents the player is up against and the likely cards these opponents might be holding. The ability to vary play is also an essential requirement for strong play. Any static strategy that is unable to adapt to the game conditions will be at risk of being exploited by strong opponents. Conversely, a strong player needs to be able to spot weaknesses in their opponents play and successfully exploit those weaknesses. This ensures the need for an *opponent modeling* component. Other issues such as *bluffing* (trying to deceive opponents about the strength of a hand by playing a weak hand strongly) and *slow-playing/trapping* (trying to deceive opponents about the strength of a hand by playing a strong hand weakly) also need to be considered. A strong player needs to know when *bluffing* or *slow-playing* may be successful and when they won't be, as well as knowing if an opponent is *bluffing* them or trying to *trap* them.

At the present the use of case-based reasoning is being considered to handle *opponent modeling*. Case-based reasoning attempts to solve new problems by reusing or adapting solutions to old problems (Watson and Marir, 1994). We believe this type of approach is well suited to *opponent modeling* as most players tend to not vary their play too much. Keeping track of how a particular opponent has been playing will provide useful information when making decisions about how to act against that opponent. We will also investigate a case-based reasoning approach to the *preflop* and *postflop* stage of the game. With this approach we hope to overcome deficiencies encountered in previous research (Billings, 2001) which employed the use of a static expert system. Case-based reasoning should be able to improve such an approach due to its ability to learn. We also hope to investigate other machine learning approaches such as using neural networks to predict which *hole cards* an opponent may be holding and using this information to inform a betting decision. It is hoped that with the combination of the above and other techniques we can construct a program that plays strong poker.

References

[1] Darse Billings, Computer Poker, M.Sc. research essay, University of Alberta, 1995.

[2] Darse Billings , Lourdes Peña , Jonathan Schaeffer , Duane Szafron, Learning to play strong poker, Machines that learn to play games, Nova Science Publishers, Inc., Commack, NY, 2001.

[3] D. Harrington and B. Robertie. Harrington on Hold ‘em. Expert Strategy For No-Limit Tournaments. Volume 1: Strategic play. Two Plus Two Publishing, 2005.

[4] Schauenberg, Terence, Opponent Modelling and Search in Poker, M.Sc. thesis, University of Alberta, 2006.

[5] D. Sklansky. The Theory of Poker. Two Plus Two Publishing, 1992.

[6] D. Sklansky and M. Malmuth. Hold'em Poker for Advanced Players. Two Plus Two Publishing, 2nd edition, 1994.

[7] Watson, I., & Marir, F. (1994). Case-Based Reasoning: A Review. The Knowledge Engineering Review, Vol. 9 No. 4: pp. 355-381