

CASPER: a Case-Based Poker-Bot

Ian Watson and Jonathan Rubin

Department of Computer Science
University of Auckland
New Zealand
ian@cs.auckland.ac.nz

Abstract. This paper investigates the use of the case-based reasoning methodology applied to the game of Texas hold'em poker. The development of a CASE-based Poker playER (CASPER) is described. CASPER uses knowledge of previous poker scenarios to inform its betting decisions. CASPER improves upon previous case-based reasoning approaches to poker and is able to play evenly against the University of Alberta's Pokibots and Simbots, from which it acquired its case-bases and updates previously published research by showing that CASPER plays profitably against human online competitors for play money. However, against online players for real money CASPER is not profitable. The reasons for this are discussed.

Key words: Case-Based Reasoning, Game AI, Poker

1. Introduction

The game of poker provides an interesting environment to investigate how to handle uncertain knowledge and issues of chance and deception in hostile environments. Games in general offer a well suited domain for investigation and experimentation due to the fact that a game is usually composed of several well defined rules which players must adhere to. Most games have precise goals and objectives which players must met to succeed. 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. 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 AI research has mainly focused on games such as chess, checkers and backgammon. These are examples of games which contain perfect information. The entire state of the game is accessible by both players at any point in the game, e.g. 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 and effective evaluation functions [Schaeffer *et al.*, 1992]. 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.

The focus of this paper is to investigate the application of Case-Based Reasoning (CBR) to the game of poker. We have developed a poker-bot, called CASPER (CASE-based Poker playER), that uses knowledge of past poker experiences to make betting decisions. CASPER plays the variation of the game known as “limit Texas Hold’em” and has been tested against other poker bots and real players.

2. Related Work

Over the last few years there has been a dramatic increase in the popularity of the game of Texas hold’em. This growing popularity has also sparked an interest in the AI community with increased attempts to construct poker robots (or bots), i.e. computerised poker players who play the game based on various algorithms or heuristics. Recent approaches to poker research can be classified into three broad categories:

1. **Heuristic rule-based systems:** which use various pieces of information, such as the cards a player holds and the amount of money being wagered, to inform a betting strategy.
2. **Simulation/Enumeration-based approaches:** which consist of playing out many scenarios from a certain point in the hand and obtaining the expected value of different decisions.
3. **Game-theoretic solutions:** which attempt to produce optimal strategies by constructing the game tree in which game states are represented as nodes and an agents possible decisions are represented as arcs.

The University of Alberta Poker Research Group are currently leading the way with poker related research, having investigated all of the above approaches. Perhaps the most well known outcome of their efforts are the poker bots nicknamed Loki/Poki [Schaeffer *et al.*, 1999; Billings *et al.*, 2002]. Loki originally used expert defined rules to inform a betting decision. While expert defined rule-based systems can produce poker programs of reasonable quality [Billings *et al.*, 2002], various limitations are also present. As with any knowledge-based system a domain expert is required to provide the rules for the system. In a strategically complex game such as Texas hold’em it becomes impossible to write rules for all the scenarios which can occur. Moreover, given the dynamic, nondeterministic structure of the game any rigid rule-based system is unable to exploit weak opposition and is likely to be exploited by any opposition with a reasonable degree of strength. Finally, any additions to a rule-based system of moderate size become difficult to implement and test [Billings *et al.*, 1999]

Loki was rewritten and renamed Poki [Davidson 2002]. A simulation-based betting strategy was developed which consisted of playing out many scenarios from a certain point in the hand and obtaining the expected value (EV) of different decisions. A simulation-based betting strategy is analogous to selective search in perfect information games.

Both rule-based and simulation-based versions of Poki have been tested by playing real opponents on an IRC poker server. Poki played in both low limit and higher limit games. Poki was a consistent winner in the lower limit games and also performed well in the higher limit games where it faced tougher opposition [Billings

et al., 2002] . More recently the use of game theory has been investigated in the construction of a poker playing bot. The University of Alberta Computer Poker Research Group have attempted to apply game-theoretic analysis to full-scale, two-player poker. The result is a poker bot known as PsOpti that is “able to defeat strong human players and be competitive against world-class opponents” [Billings *et al.*, 2003].

There have also been numerous other contributions to poker research outside the University of Alberta Poker Research Group. Sklansky and Malmuth have detailed various heuristics for different stages of play in the game of Texas hold'em [Sklansky, 1994; Sklansky & Malmuth, 1994]. The purpose of these rules, however, has been to guide human players who are looking to improve their game rather than the construction of a computerised expert system. [Korb *et al.*, 1999] produced a Bayesian Poker Program (BPP) which makes use of Bayesian networks to play five-card stud poker, whilst [Dahl, 2001] investigated the use of reinforcement learning for neural net-based agents playing a simplified version of Texas hold'em.

We have encountered relatively few attempts to apply the principles and techniques of CBR to the game of poker. [Sandven & Tessem, 2006] constructed a case-based learner for Texas hold'em called Casey. Casey began with no poker knowledge and builds up a case-base for all hands that it plays. They report that Casey plays on a par with a simple rule-based system against three opponents, but loses when it faces more opponents. [Salim & Rohwer 2005] have attempted to apply CBR to the area of opponent modeling, i.e., trying to predict the hand strength of an opponent given how that opponent has been observed playing in the past. While CBR seems inherently suited to this particular type of task they report better performance by simply relying on long-term average.

3. Texas Hold'em

Texas hold'em is the variation used to determine the annual World Champion at the World Series of Poker. This version of the game is the most strategically complex and provides a better skill-to-luck ratio than other versions of the game [Sklansky, 1994]. Each hand is played in four stages, these include the preflop, flop, turn and the river. During each round all active players need to make a betting decision. Each betting decision is summarised below:

Fold: A player discards their hand and contributes no money to the pot. Once a player folds they are no longer involved in the current hand, but can still participate in any future hands.

Check/Call: A player contributes the least amount possible to stay in the hand. A check means that the player invests nothing, whereas a call means the player invests the least amount required greater than \$0.

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 able to call, but decides to add more money to the pot this is called a raise.

All betting is controlled by two imposed limits known as the *small bet* and the *big bet*. For example, in a \$10/\$20 game the small bet is \$10 and all betting that occurs during the preflop and the flop are in increments of the small bet. During the turn and the river all betting is in increments of the big bet, \$20. All results detailed in this paper refer to a \$10/\$20 limit game¹. Each of the four game stages are summarised below:

1. **Preflop:** The game of Texas hold'em begins with each player being dealt two hole cards which only they can see. A round of betting occurs. Once a player has made their decision play continues in a clockwise fashion round the table. 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 hand then the player who did not fold their hand wins the pot and the hand is over.
2. **Flop:** Once the preflop betting has completed three community cards are dealt. Community cards are shared by all players at the table. Players use their hole cards along with the community cards to make their best hand. Another round of betting occurs.
3. **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 final stage.
4. **River:** During the river the final community card is dealt preceded by a final round of betting. If at least two players are still active in the hand a showdown occurs in which all players reveal their hole cards and the player with the highest ranking hand wins the entire pot (in the event that more than one player holds the winning hand then the pot is split evenly between these players).

4. CASPER System Overview

CASPER uses CBR to make a betting decision. This means that when it is CASPER's turn to act he evaluates the current state of the game and constructs a target case to represent this information. A target case is composed of a number of features. These features record important game information such as CASPER's hand strength, how many opponents are in the pot, how many opponents still need to act and how much money is in the pot. Once a target case has been constructed CASPER then consults his case-base (i.e. his knowledge of past poker experiences) to try and find similar scenarios which may have been encountered. CASPER's case-base is made up of a collection of cases composed of their own feature values and the action which was taken, i.e. fold, check/call or bet/raise. CASPER uses the k-nearest neighbour algorithm to search the case-base and find the most relevant cases, these are then used to decide what action should be taken.

¹ In "no limit" there is no monetary restriction on the maximum amount a player can bet.

CASPER was implemented using the commercially available product Poker Academy Pro 2.52 and the Meerkat API. The University of Alberta Poker Research Group provides various poker bots with the software including instantiations of Pokibot and the simulation based bot Simbot. These poker bots have been used to generate the training data for CASPER. Approximately 7,000 hands were played between various poker bots and each betting decision witnessed was recorded as a single case (or experience) in CASPER's case-base. Both of Alberta's bots have proven to be profitable against human competition in the past [Davidson 2002], so we hypothesise that the cases obtained are of sufficient quality to enable CASPER to play a competitive game of poker.

5. Case Representation

CASPER searches a separate case-base for each separate stage of a poker hand (i.e. preflop, flop, turn and river). The features that make up a case and describe the state of the game at a particular time are listed and explained in Table 1. These are the indexed features or case vocabulary, which means that they are believed to be predictive of a case's outcome. and by computing local similarity for each feature they are used to retrieve the most similar cases in the case-base. The first eight features are used in all case-bases, whereas the last four features are only used during the postflop stages. Each case has a single outcome that is the betting decision that was made. The *Hand Strength* feature, listed in Table 1, differs somewhat for preflop and postflop stages of the game. During the preflop there exists 169 distinct card groups that a player could be dealt. These card groups were ordered from 1 to 169 based on their hand ranking, where 1 indicates pocket Aces (the best preflop hand) and 169 indicates a 2 and a 7 of different suits (the worst preflop hand). Preflop hand strength was then based on this ordering, whereas postflop hand strength refers to a calculation of immediate hand strength based on the hole cards a player has and the community cards which are present. This value is calculated by enumerating all possible hole cards for a single opponent and recording how many of these hands are better, worse or equal to the current player's hand. For more details on hand strength and potential consult [Billings, Davidson et al. 2002; Davidson 2002].

6. Case Retrieval

Once a target case has been constructed CASPER needs to locate and retrieve the most similar cases it has stored in its case-base. The k-nearest neighbour algorithm is used to compute a similarity value for all cases in the case-base. Each feature has a local similarity metric associated with it that evaluates how similar its value is to a separate case's value, where 1.0 indicates an exact match and 0.0 indicates entirely dissimilar.

Two separate similarity metrics are used depending on the type of feature. The first is the standard Euclidean distance function given by the following equation:

$$si = 1 - \left(\frac{|x_1 - x_2|}{MAX_DIFF} \right) \quad (1)$$

² <http://www.poker-academy.com/>

where: x_1 refers to the target value, x_2 refers to the case value and MAX_DIFF is the greatest difference in values, given by the range in Table 1.

The above Euclidean similarity metric (1) produces smooth, continuous changes in similarity, however, for some features, minor differences in their values produce major changes in actual similarity, e.g. the ‘Bets to call’ feature. For this reason an exponential decay function, given by equation (2), has also been used for some features:

$$s_i = e^{-k(|x_1 - x_2|)} \quad (2)$$

where: x_1 refers to the target value and x_2 refers to the source value and k is a constant that controls the rate of decay.

Global similarity is computed as a weighted linear combination of local similarity, where higher weights are given to features that refer to a player’s hand strength as well as positive and negative potential. The following equation (3) is used to compute the final similarity value for each case in the case-base:

$$\frac{\sum_{i=1}^n w_i x_i}{\sum_{i=1}^n w_i} \quad (3)$$

where: x_i refers to the i th local similarity metric in the range 0.0 to 1.0 and w_i is its associated weight, in the range 0 – 100.

After computing a similarity value for each case in the case-base a descending quick sort of all cases is performed. The actions of all cases which exceed a similarity threshold of 97% are recorded. Each action is summed and then divided by the total number of similar cases which results in the construction of a probability triple (f, c, r) which gives the probability of folding, checking/calling or betting/raising in the current situation. If no cases exceed the similarity threshold then the top 20 similar cases are used. As an example, assume CASPER looks at his hole cards and sees $A\heartsuit-A\spadesuit$. After a search of his preflop case-base the following probability triple is generated: (0.0, 0.1, 0.9). This indicates that given the current situation CASPER should never fold this hand, he should just call the small bet 10% of the time and he should raise 90% of the time. A betting decision is then probabilistically made using the triple which was generated.

Feature:	Type:	Range:	Explanation:
Number of players	int	2 – 10	Number of active players at the beginning of the round (<i>preflop, flop, turn</i> or <i>river</i>).
Relative position	double	0.0 - 1.0	What order the player acts relative to other players at the table. 0.0 means the player is first to act in the round, 1.0 means the player is last to act.
Players in current hand	int	0 – 9	The number of players that have already acted and are still currently in the hand, i.e. players that have checked, bet, called or raised.
Players yet to act	int	0 – 9	The number of players that still need to make a future betting decision.
Bets committed	double	0.0 - 5.0	A multiple of the current bet size a certain player has committed to the pot. Small bets are used during the <i>preflop</i> and <i>flop</i> and big bets are used during the <i>turn</i> and <i>river</i> .
Bets to call	double	0.0 - 5.0	A multiple of the current bet size a certain player has to commit to the pot to stay in the hand. Small bets are used during the <i>preflop</i> and <i>flop</i> and big bets are used during the <i>turn</i> and <i>river</i> .
Pot Odds	double	0.0 - 0.5	The amount to call divided by the amount currently in the pot plus the amount needing to be called, a risk/reward measure.
Hand strength	double	0.0 - 1.0	A numerical measure of the strength of a player's hand. 0.0 is the worst possible hand, whereas 1.0 represents an unbeatable hand ("the nuts").
Positive potential ³	double	0.0 - ~0.40	A numerical measure which represents the chance that a player who does not currently hold the best hand will improve to the best hand after future community cards are dealt.
Negative potential ⁴	double	0.0 - ~0.30	A numerical measure which represents the chance that a player currently holding the best hand no longer holds the best hand after future community cards are dealt.
Small bets in pot	double	0.0 - ~300.0	The total amount in the pot divided by the value of the small bet size.
Previous round bets	int	0 – 5	How many bets or raises occurred during the previous betting round.
Action	char	{f, k, c, b, r}	A character representing the decision that was made. f = fold, k = check, c = call, b = bet, r = raise.

Table 1: Case features. (The first eight features are common for all case-bases while the last four features are only used for *postflop* play. Each case also consists of one outcome or Action.)

7. Results

CASPER was evaluated by playing other poker bots provided through the commercial software product Poker Academy Pro 2.5. CASPER was tested at two separate poker tables. The first table consisted of strong, adaptive poker bots that model their opponents and try to exploit weaknesses. As CASPER has no adaptive qualities of his own he was also tested against non-adaptive, but loose/aggressive opponents. The results from these experiments are reported in detail in Rubin & Watson [2007]. To summarise, the best parameterised version of CASPER made a profit of +0.03 small bets per hand, or \$0.30 for each hand played.

7.1 Real Opponents - Play Money

CASPER02 was tested against real opponents by playing on the 'play money' tables of internet poker websites. Here players can participate in a game of poker using a

³ These features are not used during the *river* as there are no further betting rounds.

bankroll of play money beginning with a starting bankroll of \$1000. All games played at the ‘play money’ table were \$10/\$20 limit games. At each table a minimum of two players and a maximum of nine players could participate in a game of poker. CASPER was tested by playing anywhere between one opponent all the way up to eight opponents. Figure 1 displays the results recorded at ‘play money’ tables for CASPER (using hand picked weights) and CASPERGeneral (using feature weights derived by an evolutionary algorithm described in a future paper) as well as a random opponent which makes random decisions (used as a baseline comparison).

Both CASPER and CASPERGeneral earn consistent profit at the ‘play money’ tables. The results suggest that the use of CASPER with hand picked weights outperforms CASPERGeneral. CASPER earns a profit of \$2.90 for every hand played, followed by CASPERGeneral with a profit of \$2.20 for each hand. Random decisions resulted in exhausting the initial \$1000 bankroll in only 30 hands, losing approximately -\$30.80 for each hand played.

Both Pokibot and Simbot have also been tested against real opponents by playing on Internet Relay Chat (IRC). The IRC server allows bots and humans to challenge each other online using “play money”. Results reported by [Davidson 2002] indicate that Pokibot achieves a profit of +0.22 sb/h, i.e. a profit of \$2.20 per hand, and Simbot achieves a profit of +0.19 sb/h or \$1.90 profit per hand. These results are very similar to those obtained by CASPER, when challenging real opponents for play money. As CASPER used Pokibot and Simbots playing style to build its case-base this result would be expected.




Fig. 1 CASPER vs. Real Opponents at the Online Play Money Tables

Because CASPER was playing against real opponents the time taken to record the results is longer than when challenging computerised opponents. For this reason, fewer hands were able to be played against real opponents. CASPER is also mainly suited to playing poker at a full table, i.e. with nine or ten players present, however the results recorded above consist of anywhere between two and nine players at a table. While we need to take caution in analysing the above results, it is safe to say that CASPER is consistently profitable at the ‘play money’ tables.

7.2 Real Opponents - Real Money

Because there is normally a substantial difference in the type of play at the ‘play money’ tables compared to the ‘real money’ tables it was decided to attempt to get an idea of how CASPER would perform using real money against real opponents.



Fig. 2 CASPER vs. Real Online Opponents at the Real Money Tables

CASPER02 (the large case-base) with hand-picked feature weights that had achieved the best performance at the ‘play money’ tables was used to play at the ‘real money’ tables. The betting limit used was a small bet of \$0.25 and a big bet of \$0.50. CASPER started out with a bank roll of \$100. The results are given in Figure 2.

CASPER achieves a small bet per hand value of -0.07. Therefore, CASPER now loses on average \$0.02 per hand. The results indicate that while CASPER loses money very slowly it is now, nonetheless, unprofitable against these opponents. Due to the fact that real money was being used, fewer hands were able to be played and the experiment was stopped after CASPER had lost approx. \$50. No results are available for Pokibot or Simbot challenging real opponents using real money. Therefore, it is

not possible to evaluate how CASPER performs using real money compared to Pokibot or Simbot.

8. Conclusion and Future Work

In conclusion, CASPER, a case-based poker player has been developed that plays evenly against strong, adaptive poker-bots, plays profitably against non-adaptive poker-bots and against real opponents for play money. Two separate versions of CASPER were tested and the addition of extra cases to the case-base was shown to result in improvements in overall performance. It is interesting to note that CASPER was initially unprofitable against the non-adaptive, aggressive poker-bots. One possible reason for this is that as CASPER was trained using data from players at the adaptive table it perhaps makes sense that they would play evenly, whereas players at the non-adaptive table tend to play much more loosely and aggressively. This means that while CASPER has extensive knowledge about the type of scenarios that often occur at the advanced table, this knowledge is weaker at the non-adaptive table as CASPER runs into situations which it is not familiar with.

Why then was CASPER not profitable on the real money tables. Two hypotheses could explain this. First, that people play poker very differently for real money than for play money. Since CASPER's case-base is derived from poker-bots that are playing for play money these cases are not representative of real money games. The average similarity of cases retrieved when playing against the bots is approx 90%, The average similarity of cases retrieved against real money players is approx. 50%. This is confirmation that CASPER is struggling to find good cases (*i.e.*, similar ones) to retrieve, therefore it's poor performance is not a surprise.

A second possible hypothesis is that the real money poker games are in some way dishonest. Many of these websites are run by organizations operating from countries with poor rule of law and little or no legislation concerning online gambling. Ask yourself a question would you play poker online with your money? This hypothesis could be tested if someone would bank roll CASPER to play 15,000 to 20,000 games online with real money thus enabling it to acquire a large case-base of real money games that it could then use to guide its play.

References

- Billings, D., Burch, N., Davidson A, Holte R, Schaeffer J, Schauenberg T. & Szafron, D. (2003). "Approximating game-theoretic optimal strategies for full-scale poker." *IJCAI*.
- Billings, D., Davidson, A., Schaeffer, J., & Szafron, D.. (2002). "The challenge of poker." *Artificial Intelligence* **134**(1-2): 201-240.
- Billings, D., Peña, L., Schaeffer, J. and Szafron, D. (1999). "Using probabilistic knowledge and simulation to play poker." *Proceedings of the sixteenth national conference on Artificial intelligence and the eleventh Innovative applications of artificial intelligence conference innovative applications of artificial intelligence*: 697-703.
- Dahl, F. A. (2001). *A Reinforcement Learning Algorithm Applied to Simplified Two-Player Texas Hold'em Poker*. *Proceedings of the 12th European Conference on Machine Learning* Springer-Verlag.
- Davidson, A. (2002). *Opponent modeling in poker: Learning and acting in a hostile and uncertain environment*. Master's thesis, University of Alberta.
- Korb, K. B., Nicholson A.E. & Jitnah, N. (1999). "Bayesian poker." *UAI'99 - Proceedings of the 15th International Conference on Uncertainty in Artificial Intelligence*, Sweden: 343-350.

- Rubin, J. & Watson, I. (2007). Investigating the Effectiveness of Applying Case-Based Reasoning to the game of Texas Hold'em. In, Proc. of the 20th. Florida Artificial Intelligence Research Society Conference (FLAIRS), Key West, Florida, May 2007. AAAI Press.
- Salim, M. & Rohwer, P. (2005). Poker Opponent Modeling. Indiana University: Personal communication.
- Sandven, A. & Tessem, B. (2006). A Case-Based Learner for Poker. The Ninth Scandinavian Conference on Artificial Intelligence (SCAI 2006), Helsinki, Finland.
- Schaeffer, J., Billings, D., Peña, L. & Szafron, D. (1999). "Learning to play strong poker." Proceedings of the ICML-99 Workshop on Machine Learning in Game Playing.
- Schaeffer, J., Culberson, J. Treloar, N., Knight, B., Lu, P. & Szafron, D. (1992). "A world championship caliber checkers program." *Artificial Intelligence* 53(2-3): 273-289.
- Sklansky, D. & Malmuth, M. (1994). "Hold'em Poker for Advanced Players." Two Plus Two Publishing, Las Vegas, NV, 2nd ed.
- Sklansky, D. (1994). "The Theory of Poker." Two Plus Two Publishing Las Vegas, NV, 1994 3rd ed.