A Memory-Based Approach to Two-Player Texas Hold'em

Jonathan Rubin and Ian Watson

Department of Computer Science
University of Auckland, New Zealand
http://www.cs.auckland.ac.nz/research/gameai
jrub001@aucklanduni.ac.nz
ian@cs.auckland.ac.nz

Abstract. A Case-Based Reasoning system, nicknamed SARTRE, that uses a memory-based approach to play two-player, limit Texas Hold'em is introduced. SARTRE records hand histories from strong players and attempts to re-use this information to handle novel situations. SARTRE'S case features and their representations are described, followed by the results obtained when challenging a world-class computerised opponent. Our experimental methodology attempts to address how well SARTRE'S performance can approximate the performance of the expert player, who SARTRE originally derived the experience-base from.

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

1 Introduction

Poker has been identified as a useful domain for Artificial Intelligence research [1]. As the number of researchers working within the environment of Computer Poker has increased, so too has the development of strong poker robots (or *poker-bots*) which play increasingly more sophisticated strategies [2, 4]. A beneficial result of the increased attention paid to computer poker has been the creation of the Annual Computer Poker Competition (CPC) [11], where researchers can evaluate their systems by challenging other computerised opponents to the game of `Texas Hold'em` poker.

Competitors of past CPC's can typically be characterised into two broad categories. Firstly, those systems that attempt to approximate a *Nash-equilibrium* strategy [2, 4]. A *Nash-equilibrium* strategy guarantees that no matter what playing style an opponent adopts, they will never win more than what the *equilibrium* strategy guarantees [2]. This type of strategy can be said to favour not losing, rather than looking for ways to win. At present, *equilibrium* strategies may only be approximated for the game of `Texas Hold'em` due to the incredibly large size of the game tree. On the other hand, an *exploitative* [3] strategy will attempt to win by maximising profits and exploiting weaker competition. This approach requires a system to model their opponent's play. As this strategy deviates from the *equilibrium*, the system may be prone to exploitation itself.

Our research currently looks into the use of memory in game AI. Rather than relying on game-theoretic principles to construct *near-equilibrium* strategies, the goal of this research is to investigate whether hand histories from strong poker players can be re-used within a Case-Based Reasoning (CBR) framework to achieve a similar performance? CBR is an AI methodology that uses solutions to past problems to solve new problems [7]. A collection of experiences is recorded, which consists of problems and solutions. When a novel situation is encountered a CBR system attempts to retrieve similar experiences from its experience-base and re-use or adapt the solutions to solve the new problem.

SARTRE (Similarity Assessment Reasoning for Texas hold'em via Recall of Experience) is the latest outcome of our research that attempts to address the above question. SARTRE differs from a previous system we developed, CASPER [8], in that it is specifically designed to play 2-player poker, whereas CASPER was more suited to challenge multiple opponents.

SARTRE'S *experience-base* is generated by observing and recording hand histories from the strongest opponents of past CPC's. In 2008 the University of Alberta's Hyperborean-eq took out first place in the limit Hold'em competition [11]. Hyperborean-eq plays a *fixed, near-equilibrium* strategy.

The remainder of this paper proceeds as follows. The game of Texas Hold'em is discussed in Section 2. Section 3 provides an overview of the SARTRE system, followed by Section 4, the experimental results and finally the discussion and conclusion in Section 5.

2 Texas Hold'em

Currently our research focuses around the `Texas Hold'em` variation of poker. At present, `Texas Hold'em` is the most popular form of poker as well as being the most strategically complex [5]. In `Texas Hold'em` play is broken down into four main stages: *preflop, flop, turn* and *river*. For a full description of each stage of the game consult [6].

SARTRE is a `heads-up, limit` poker-bot. This means SARTRE will only ever challenge one opponent at any one time and betting will be capped at certain limits during each round of play. Factors such as challenging multiple opponents or handling a `no-limit` betting structure pose extra challenging research problems for poker playing agents. `Heads-up, limit` poker simplifies these tasks, however, it still preserves the key qualities and structure of other more complicated variants. It also offers its own unique challenges, for example, in `heads-up` play both players need to play a lot more hands in order to be profitable [5]. Players therefore need to play weaker hands than they would play at a full table (i.e. approx. 9 players), it then becomes more important to determine whether an opponent actually has a valuable hand, or not, more often than would be required at a full table. As only one opponent is available during each game more opportunity exists to model and adapt to your opponents play. It makes sense that two-player, `limit` poker should be investigated first [2] be-

fore focusing effort on more complicated concepts such as `no-limit` betting and multiple opponents.

3 SARTRE: System Overview

SARTRE makes decisions by retrieving similar cases from its *experience-base*. The authors have hand picked three key factors to represent case features that SARTRE uses to determine a solution for a particular case.

1. The previous betting for the current hand.
2. The current strength of SARTRE'S hand given by combining personal *hole cards* with the publicly available board cards.
3. Information about the state of the current community cards, called the *texture of the board*.

As SARTRE is a computer program the information required needs to be easily recognised and able to be reasoned about algorithmically. Qualitative feature descriptions have been favoured over quantitative descriptions as they are more likely to be used by an expert, human player. Each case feature is described in more detail below, including the representation we have chosen to implement for the SARTRE system.

3.1 The previous betting for the current hand.

The type of betting that can occur at each decision point in a hand consists of a fold *(f)*, check/call *(c)*, or bet/raise *(r)*. A combination of these symbols corresponds to all the decisions made during a particular hand. We have chosen to represent each betting pattern as a path within a betting tree. A betting tree succinctly enumerates all betting combinations up until a certain point in the hand. A path within this tree represents the actual decisions that were made by each player during this hand. Fig. 1. represents a situation where SARTRE'S opponent has made a bet on the *flop* and it is now SARTRE'S turn to act.

Given this representation, we can calculate the similarity between two separate trees (a target tree and a source tree) by comparing the betting path within each tree. If the betting path in the target tree is exactly the same as the betting path within the source tree a similarity value of 1.0 is assigned. Currently, SARTRE will simply assign a value of 0.0 to any betting paths that are not exactly similar, however, we plan to investigate less stringent approaches for future implementations. For example, if one betting path mostly resembles that of another, with a small number of variations, a similarity value close to (but less than) 1.0 could be assigned.

3.2 The current strength of SARTRE'S hand.

The second case feature used to determine a betting action is a qualitative category describing SARTRE'S personal hand. During the *pre-flop* SARTRE'S

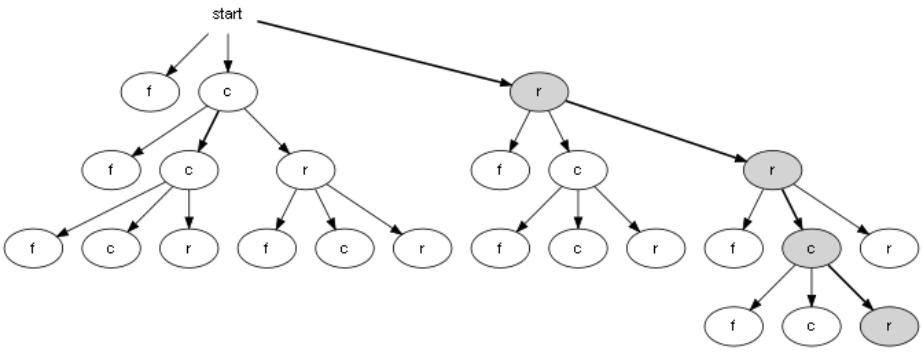

Fig. 1. A tree that describes betting decisions for two players during a hand of Texas Hold'em Poker. The highlighted nodes are the actual decisions that were made by each player.

hand simply consists of its personal *hole cards*, whereas for the *post-flop* stages of play SARTRE'S hand is constructed by combining its *hole cards* with the publicly available *community cards*, the best 5 card combination is used.

SARTRE'S best 5 cards are mapped to a category that describes the hand. The classic hand categories in poker include *no-pair, one-pair, two-pair, three-of-a-kind, straight, flush, full-house, four-of-a-kind* and finally a *straight-flush*. Each category has a greater strength than the previous one, where a *straight-flush*, consisting of the cards `Ten, Jack, Queen, King, Ace`, represents the highest rank possible (i.e. a *Royal Flush*).

During the *flop* and the *turn* all the community cards have yet to be dealt and therefore a player's hand has the ability to improve from one category to another, depending on which card is drawn next. It is therefore too simplistic to only consider the current hand category, so further classification is required for hands with the potential to improve. These types of hands are called *drawing hands* (in poker terminology). SARTRE considers two types of drawing hands: *flush draws* & *straight draws*. An example mapping is illustrated in Fig. 2.

The hand categories SARTRE uses to classify cards were decided upon by the authors. Fig. 2. shows a combination of two categories, one which represents the current hand category: `overcards` (i.e. no pair has been made, but both *hole cards* have a higher rank than the community cards). Appended to this category is a separate drawing category: `ace-high-flush-draw-uses-both`, that indicates the strength of the current hand has the ability to improve to a *flush*. The "ace-high" portion of this category further specialises this category by indicating the strength of the possible *flush*.

Currently a simple rule-based system is used to decide which category a combination of cards belongs to. Similarity for this feature is currently either 1.0 when the category of the target case is exactly that of the source case, otherwise it is 0.0 when the categories are distinct.

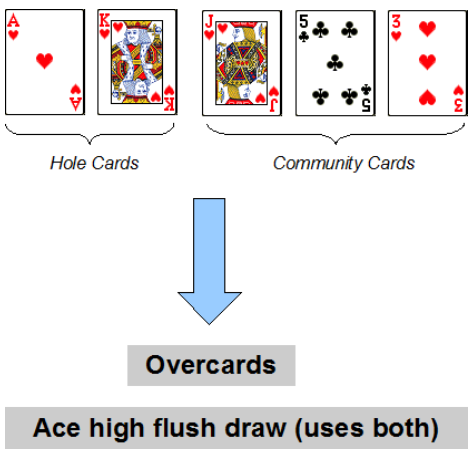

Fig. 2. Mapping a combination of five cards to a category that represents the current hand rank and the drawing strength of this hand.

3.3 The texture of the board.

The final indexed feature attempts to summarise the state of the community cards without considering the *hole cards* of a player. The *texture of the board* refers to salient information a human poker player would usually notice about the public cards, such as whether a *flush* is possible. Once again a set of qualitative categories were hand-picked by the authors to map various boards into. Some categories used by SARTRE'S current implementation that refer to flush and straight possibilities are *Is-Flush-Possible* (where three cards of the same suit are showing), *Is-Flush-Highly-Possible* (where four cards of the same suit are showing) & *Is-Straight-Possible* (where three consecutive card values are showing), *Is-Straight-Highly-Possible* (where four consecutive card values are showing).

If two boards are mapped into the same category, they are given a similarity value of 1.0, whereas boards that map to separate categories have a similarity of 0.0.

3.4 SARTRE'S Experience-Base

SARTRE'S *experience-base* is generated by analysing information from the logs of previous CPC matches involving Hyperborean-eq. For each hand played in the game log at least one new case is added to SARTRE'S *experience-base*. Each feature described above is assigned into an appropriate category to represent the situation. The decision that Hyperboean-eq made is recorded and acts as the solution for that particular case. The final outcome of that decision is also recorded.

The current version of SARTRE uses just over 1 million cases in total, these are sub-divided into different stages of the game as follows: Preflop cases: 201335, Flop cases: 300577, Turn cases: 281529, River cases: 216597.

When it is time for SARTRE to make a decision, the *experience-base* is consulted and the most similar cases are retrieved, along with their solutions. A *probability triple* is then constructed by summing the number of times each decision was made and dividing by the total decisions. SARTRE then probabilistically selects a decision based on the values within the triple.

4 Experimental Results

Experimental results were obtained for SARTRE using a 3.00 GHz Intel Core 2 Duo CPU with 4.00 GB of Memory (RAM). SARTRE challenged two separate computerized opponents: FellOmen2 [4] & BluffBot [9], both were chosen because they are freely available. FellOmen2 is currently a world-class poker-bot, finishing second equal in the 2008 AAAI Computer Poker Competition [11]. FellOmen2 uses a co-evolutionary strategy, to create a *near-equilibrium* solution [4]. The limit version of BluffBot finished second in the 2006 AAAI CPC and, by today's standards, is not a world-class poker-bot [10]. BluffBot attempts to approach a *Nash-equilibrium* strategy using game-theoretic methods, similar to [2].

All matches played were `limit, heads-up`, Texas Hold'em. The betting structure was \$2/\$4, meaning all bets made during the *preflop* and the *flop* were in increments of \$2 and all betting on the *turn* and *river* were in increments of \$4. As FellOmen2 and BluffBot were made available in different platforms, two separate poker environments were used to obtain results, described in detail below:

AAAI Computer Poker Competition poker server Version 2.3.1. Using the poker server software, duplicate matches were able to be played. Duplicate matches proceed by playing N hands in a forward direction, then each competitor's memory is reset and the hands are replayed in the reverse direction, i.e. each player now plays the hands that were dealt to their opponent on the forward run. This has the effect of decreasing the inherent variance involved with poker, as one player will not receive a set of better hands than another player. Once the duplicate match is complete the total profit/loss for each direction is summed and the competitor with a positive bankroll is determined the winner. SARTRE challenged FellOmen2 by playing 6 separate duplicate matches, using $N = 3000$, for a total of 36,000 hands.

Poker Academy Pro 2.5 BluffBot was only available to challenge using the commercial application Poker Academy[1]. Poker Academy doesn't allow a duplicate match structure to be played as described above. Instead, all matches played using Poker Academy proceeded in a forward direction and no reduction

[1] http://www.poker-academy.com/poker-software

of variance took place. SARTRE challenged BluffBot by playing a total of 30,000 hands.

4.1 SARTRE Vs. FellOmen2

Fig. 3. plots SARTRE'S bankroll for each of the 6 duplicate matches played against FellOmen2. Table 1. provides a summary of the overall outcome. The figures refer to SARTRE'S bankroll.

Table 1. Sartre Vs. FellOmen2 Summary

	Total Hands	Forward	Reverse	Final Outcome
Round1	6000	532	-827	-295
Round2	6000	-204	-292	-496
Round3	6000	-869	-261	-1130
Round4	6000	549	-1208	-659
Round5	6000	109	-900	-791
Round6	6000	226	-1063	-837
Total	36000	343	-6859	-6516

From the above results we can calculate that, on average, SARTRE loses -2.92 ± 0.5 big bets per 100 hands (BB/100) to FellOmen2. BB/100 is a value commonly used in poker to measure a players success, without considering the stakes the player is playing at. As the big bet was \$4 this means that, on average for each duplicate run, SARTRE will lose \$11.60 $\pm$ \$2 for every 100 hands played against FellOmen2. Generally, any positive BB/100 value, over a large sample, is considered good, whereas, a player who always folds would lose -37.5 BB/100. During the 2008 CPC, Hyperborean-eq achieved an average value of $+1.205 \pm 0.15$ BB/100 when challenging FellOmen2. An independent samples t-test gives $p < 0.00001$, hence a significant difference is observed between the average profit/loss of Hyperborean and SARTRE when challenging FellOmen2.

4.2 2008 CPC Competitors Vs. FellOmen2

It is also interesting to consider the results of other competitors who challenged FellOmen2, during the 2008 CPC. Table 2., lists the final outcome of matches played against FellOmen2 for each of the 9 competitors in the `limit Hold'em` competition [11, 4].

Our experiments show that SARTRE'S win rate against FellOmen2 was -2.92 BB/100 which would place SARTRE 6th, in Table. 2., between GS4-Beta and PokeMinn2.

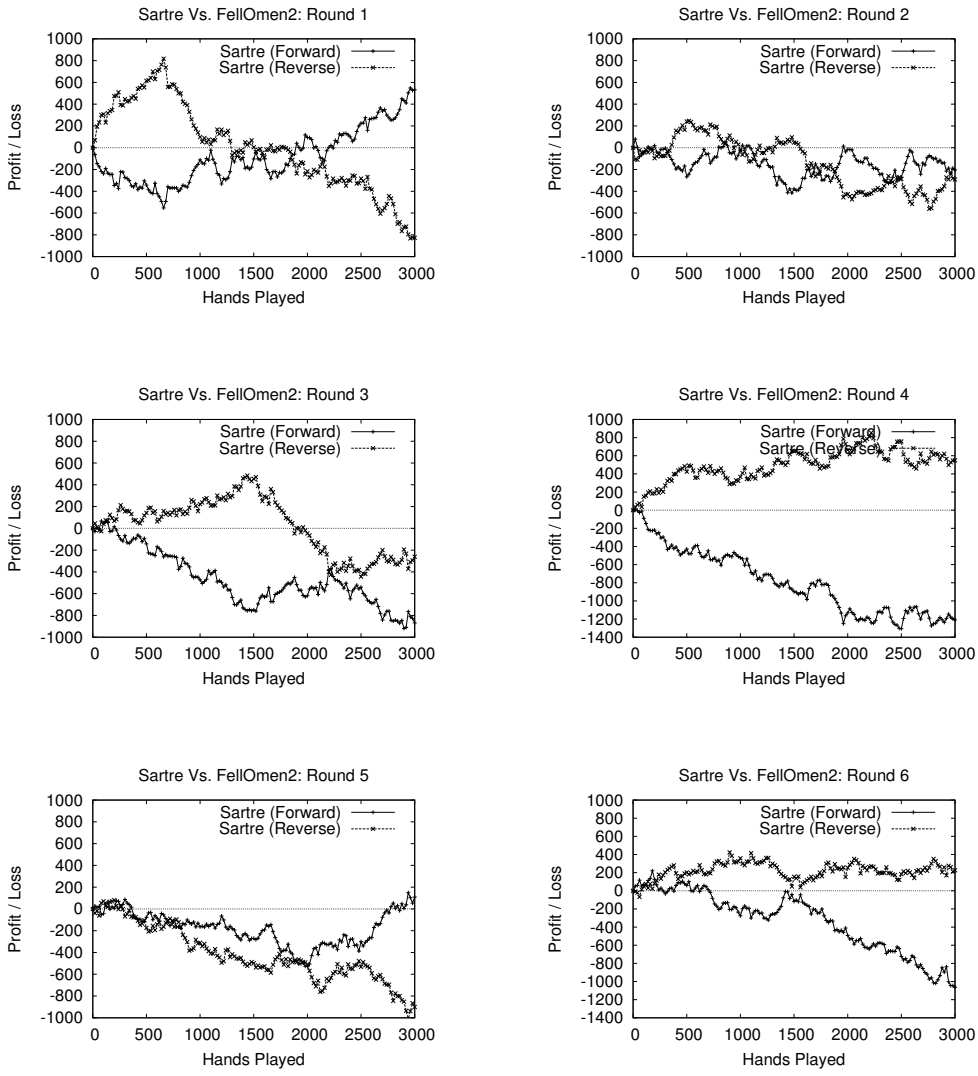

Fig. 3. Sartre Vs. FellOmen2

Table 2. FellOmen2 Vs Opponents from 2008 AAAI CPC

Place	Name	Win rate against FellOmen2 (BB/100)
1	Hyperborean-eq	1.2
2	Fell Omen 2	0
2	Hyperborean-on	0.2
2	GGValuta	-0.15
5	GS4-Beta	-1
6	PokeMinn 2	-7.65
7	PokeMinn 1	-7.7
8	GUS	-23.35
9	Dr. Sahbak	-26.6

4.3 SARTRE Vs. BluffBot

The above results only represent SARTRE'S performance against one specific opponent, FellOmen2. To further evaluate the system, SARTRE challenged a separate, computerised opponent. The next opponent SARTRE faced was BluffBot. SARTRE played 30,000 hands against BluffBot and the outcome is illustrated in Fig. 4.

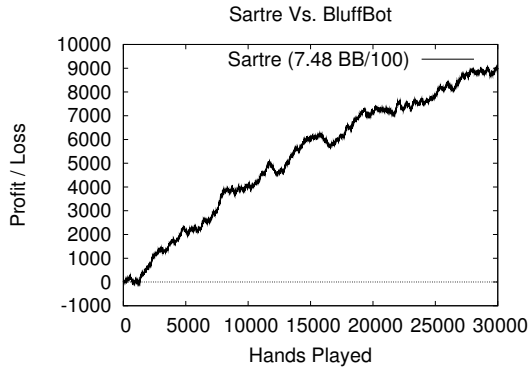

Fig. 4. Sartre Vs. BluffBot

The platform that BluffBot was made available on did not allow for the duplicate match structure that was used when SARTRE challenged FellOmen2, so caution must be used in interpreting the results. However, it is safe to say that Fig. 4. clearly illustrates a profitable trend for SARTRE. SARTRE achieves a win rate of +7.48 BB/100 against BluffBot.

5 Discussion and Conclusion

From the results it is clear that SARTRE has not reached the quality of performance that Hyperborean-eq exhibits, as Hyperborean-eq is profitable against FellOmen2, but SARTRE is unprofitable. Some possible reasons for this include:

- The hand strength feature needs to be improved. Presently, a large combination of dissimilar hands are mapped into one category. This results in detailed information being lost which could degrade the level of play.
- There are still many situations where case retrieval is sparse. For one match (chosen at random) against FellOmen2 the results indicated that out of a total of 3769 river decisions made by SARTRE, for 357 (9.47%) of these, SARTRE was unable to retrieve any similar cases. When SARTRE cannot retrieve a similar case a crude strategy of always *checking/calling* is adopted.

However, while SARTRE does not yet achieve the level of play of Hyperborean-eq, the system still appears to play reasonably strong poker. SARTRE was profitable against BluffBot and appears to perform better than four other competitors of the 2008 CPC when challenging FellOmen2.

Acknowledgements Thanks to Teppo Salonen & Ian Fellows for supplying BluffBot and FellOmen2.

References

1. Billings, D., Davidson, A., Schaeffer, J., Szafron, D.: The challenge of poker. Artificial Intelligence (134). 201-240 (2002)
2. Billings, D., Burch, N., Davidson, A., Holte, R., Schaeffer, J., Schauenberg, T., Szafron, D.: Approximating Game-Theoretic Optimal Strategies for Full-scale Poker. IJCAI-03. Morgan Kaufmann. 661–668 (2003)
3. Johanson, M., Zinkevich, M., Bowling, M.: Computing Robust Counter-Strategies. Advances in Neural Information Processing Systems (NIPS) 20. MIT Press (2007)
4. Fellows, I.: Ian Fellows Poker Research. World Wide Web. http://thefell.googlepages.com/poker. (2008)
5. Harrington, D., Robertie, B.: Harrington on Hold'em: Expert Strategy for No-Limit Tournaments Volume II: The Endgame. Two Plus Two Publishing, Las Vegas, Nevada (2005)
6. PokerListings.com.: Poker Rules. http://www.pokerlistings.com/poker-rules (2009)
7. Riesbeck, C., Schank, R.: Inside Case-Based Reasoning. L. Erlbaum Associates Inc. Hillsdale, NJ, USA. (1989)
8. Watson, I., Rubin, J.: CASPER: A Case-Based Poker-Bot AI 2008: Advances in Artificial Intelligence, 21st Australasian Joint Conference on Artificial Intelligence Springer. 5360, 594–600 (2008)
9. Salonen, T.: The BluffBot webpage. World Wide Web. http://www.bluffbot.com/ (2009)
10. Salonen, T.: personal communication. (2009)
11. University of Alberta.: The Annual Computer Poker Competition. World Wide Web. http://www.cs.ualberta.ca/ pokert/ (2009)