Proceedings of the 2007 IEEE Symposium on
Computational Intelligence and Games (CIG 2007)

An Investigation into Tournament Poker Strategy using Evolutionary Algorithms

Richard G. Carter
Centre for Intelligent Systems and their Applications
School of Informatics
University of Edinburgh, Scotland
r.g.carter@sms.ed.ac.uk

John Levine
Strathclyde Planning Group
Dept. of Computer and Information Sciences
University of Strathclyde, Scotland
john.levine@cis.strath.ac.uk

Abstract— **In this paper we assess the hypothesis that a strategy including information related to game-specific factors in a poker tournament performs better than one founded on hand strength knowledge alone. Specifically, we demonstrate that the use of information pertaining to opponents' prior actions, the stage of the tournament, one's chip stack size and seating position all contribute towards a statistically significant improvement in the number of tournaments won.**

Additionally, we test the hypothesis that a strategy which combines information from all the aforementioned factors performs better than one which employs only a single factor. We show that an evolutionary algorithm is successfully able to resolve conflicting signals from the specified factors, and that the resulting strategies are statistically stronger.

Keywords: evolutionary algorithms, game playing, tournament poker

I. INTRODUCTION

Researchers in artificial intelligence have long been interested in developing programs for games which are able to compete with and ultimately beat human opposition. Recent successes in the games of checkers, chess and backgammon have encouraged efforts on more complex games such as go, bridge and poker. Card games such as those mentioned present a particular challenge for researchers due the lack of complete information of the game state at any point in time. Opponents' hands must be modelled probabilistically, with inferential conclusions reached by computing the likelihood of each players' holding given their actions in the current and previous games. The development of poker players is especially exacting since the nature of the game demands that each competitor tries to deceive their opponents as to which cards they hold. Research into games of imperfect information such as poker has the potential to be extremely valuable, since reasoning under conditions of uncertainty is typical of many real-world problems.

Early game theoretic investigations of poker have been superseded by significant contemporary contributions. The GAMES Group at the University of Alberta has led the way in combining a strong analytical understanding of the game together with innovative approaches in opponent modelling. Their work over the last ten years has been so successful that a particular poker variant - two-player limit ring game Texas Hold'em - is now practically solved in the game theoretic sense, with continual improvements in the opponent modelling leading to increasingly strong players.

However, many hurdles have yet to be cleared before we reach the point of having a poker World Champion. Three difficulties in particular have to be overcome to take the current state of the art to such a level. Firstly, poker is almost always played between several competing players, rather than two. This means that multiple opponent models must be maintained. Secondly, the main event of the World Series of Poker held annually in Las Vegas uses a form of betting known as no limit. Where limit betting sets pre-specified increments to each player's bet and raise, no limit allows the competitors to bet any amount up to their current chip stack size. The third challenge is that the World Championships play Texas Hold'em in a tournament, rather than ring game format. A tournament setting adds extra complexity to the decision-making process compared to the same situation in a ring game. The nature of tournament play dictates that each player must balance accumulating and protecting their chips to ensure survival. This paper marks the first attempt to understand some of the issues involved in developing a player for Texas Hold'em played within a tournament, so that strong ring game programs may be more easily adapted to the tournament format.

Noted poker authors consistently state in the non-academic poker literature that factors contained within the tournament are important and should be included in the decision-making process before selecting a betting action. Taking hand strength knowledge as a given, the first hypothesis that we test in this research is that the inclusion of information relating to:

- opponents' prior betting actions
- the stage of the tournament
- one's own chip stack amount
- seating position

improves a player's tournament performance against three different static opponents. Through the use of Monte Carlo simulation on a slightly simplified form of no limit Texas Hold'em we find that all four of these factors are statistically significant in their contribution.

The strategies we find determine which hands should be played dependent upon the additional game information. To create a strategy using all the available information, however, we cannot simply combine the suggested betting actions from the single factor results. This is because the single factor strategies are often in conflict in any given situation. Therefore the second hypothesis that we test is that an evolutionary algorithm is able to resolve such disagreements, and that the resulting strategy is stronger than any achieved by employing only one of the factors in isolation. Results contained within this work show this to be true.

II. BACKGROUND

A. Texas Hold'em

Texas Hold'em (also known simply as Hold'em) is the most widely played poker variant. It is this form which is used as the main event in the annual World Series of Poker. A game of Texas Hold'em can be played with up to 22 players, although it is more usual to see between two and ten players at a single table. A comprehensive set of rules for the game of Texas Hold'em is available online [1].

Texas Hold'em typically employs one of three different forms of betting structure: limit, pot limit, or no limit. In limit Hold'em, the size of the bets are fixed amounts. In the first two betting rounds, each bet or raise is a set amount. In the final two rounds the fixed bet size doubles. In pot limit Hold'em a player may wager any amount up to the size of the pot. No-limit Hold'em removes this restriction by allowing each player to bet any amount up to their stack size.

As with all poker games, Hold'em can be played either as a ring game (also known as cash game) or in a tournament. In a ring game the players contest pots with real money and no predetermined end time. A poker tournament, on the other hand, is played with tournament chips and ends once the game has been reduced to a single player. In a ring game players may continually enter and exit the table, and players who lose all their chips are able to purchase more to continue in the game. By contrast, the players in a tournament usually buy a set number of tournament chips before the game and are eliminated from the competition if their stack size reaches zero. The major differences between ring game and tournament play are summarised in Table I.

TABLE I
STRUCTURAL DIFFERENCES BETWEEN RING GAME AND TOURNAMENT POKER

Difference	**Ring Game**	**Tournament**
Entry fee	Variable	Tournament cost
Chips	Money replacement	Game tokens
Blinds	Fixed	Rising schedule
Number of players	Limited to a table	Unlimited
Game exit	Player discretion	Zero chips
Profit and loss	On each hand	Based on finish

B. Previous Poker Research

The first academic investigations into the game of poker were undertaken in the mid-twentieth century. Early pioneers in the field of game theory, such as von Neumann and Morgenstern [2], employed greatly simplified poker variants to formulate a framework for strategy selection in non-cooperative environments. The toy pokers that were examined were typically only two- or three-player games, and used pared decks of cards to reduce the space of possible strategies.

More recent investigations into poker have focused on developing computer programs which are able to play more realistic variants to a high standard. The most advanced and successful work on computer poker play to date has been produced by the Game-playing, Analytical methods, Minimax search, and Empirical Studies (GAMES) Group at the University of Alberta, led by Jonathan Schaeffer.

The GAMES Groups' poker research started with the development of programs designed to play ten-player limit Texas Hold'em in a ring game, and employed a combination of statistical measures and expert rules to effect decision-making [3]. Billings *et al.* [4] then turned to two-player limit Hold'em, for which they were able to derive "pseudo-optimal" strategies. Recent efforts by the GAMES Group target the addition of opponent modelling to the game theoretic foundation, with greatly improved results [5].

Outside of academia, many books have been written on strategy for play in Texas Hold'em. Whilst poker literature for ring game play abounds, relatively little has been written on Texas Hold'em tournament strategy. Those that exist discuss the differences between ring game and tournament play, and expound on how certain factors affect strategic considerations.

The most common strategic messages in the non-academic writings relate to how one's range of playable starting hands should increase throughout a tournament, the importance of stack size in betting decisions, and how the payoff structure of the tournament determines correct strategy. The purpose of this research is to empirically validate some of these assertions.

C. Evolutionary Algorithms Applied to Poker

Evolutionary algorithms have been commonly applied to search for strong strategies in a wide variety of different games. One notable example of this is the development of the checkers player *Blondie24* by Kumar Chellapilla and David Fogel [6].

The first, and to date most extensive attempts to apply evolutionary computation to poker have been performed by Luigi Barone and Lyndon While [7], [8]. The authors develop poker players that are able to adapt strategically given inputs from their environment, such as their own hand strength, seating position, bet size, and a measure of their opponents' playing styles. The experiments performed in their research use Texas Hold'em with limit betting, and are employed within a ring game format.

Graham Kendall and Mark Willdig [9] employ evolutionary methods to learn to play a simplified draw poker game. Candidates from the population are played at tables containing opponents of different styles, and the adaptive players are seen to adjust their strategy appropriately to each situation.

Texas Hold'em has also been used in research by Jason Noble [10]. Rather than play against static opposition, his work seeks to develop strong poker players through self-play. However, the primary focus of these studies is on the comparison of different co-evolutionary methods, and poker is simply used as the test bed for the experiments.

The programs developed by the GAMES Group have been shown to perform well against professional poker players, but no such achievements were reported in the work on evolutionary poker programs. Whilst it presently seems that the approach of marrying game theoretic understanding with opponent modelling has the greatest promise in achieving the strongest poker programs, results from evolutionary methods have highlighted some of the strategic problems faced in playing poker games.

III. Specification and Implementation

All the experiments performed focus on ten-player winner-takes-all all in or fold pre-flop Texas Hold'em. This format makes one simplification to an authentic poker tournament. The restriction of the players' betting actions to all in or fold reduces the strategy space to a tractable size. Substituting the more flexible betting choice with a binary decision shifts the focus to the more general and important strategic question of when to bet, rather than how much. Since players will have either bet all of their chips or folded before the flop, the resulting betting rounds are redundant. All in or fold pre-flop betting strategies have been proposed for use in real tournaments, and we use three of these taken from the poker literature as the opponents within our experiments.

Whilst it is more usual for ten-player tournaments to employ a percentage payout structure, the "shootout" format with a single winner can also be employed. The winner-takes-all design only credits a player for finishing first, and hence second place is equivalent to finishing last. Establishing the competitions in this way ensures that we assess each strategy's ability to win tournaments, and not just their capacity for tournament survival.

All players start with $1,000 in tournament chips. The tournaments use eleven levels, with the blinds increasing as the tournament progresses. Each level consists of ten hands, except for the final level which is used until the tournament concludes.

There were a total of three different opponents used across all the simulations. These encoded the Sklansky Basic strategy, Sklansky's Improved strategy, and the Kill Phil Rookie strategy, and are explained below.

The first two strategies are taken from the book "Tournament Poker for Advanced Players" [11]. The Sklansky Basic strategy is highly restrictive, and will only bet the very best starting hands once another player has already bet into the pot. If no other player has yet bet, this strategy will move all in with a slightly larger subset of hands. The Sklansky Improved strategy is similarly restrained in the hands it will play if an opponent has entered the pot. Where this strategy differs from the first is in its use of the ratio of the player's stack to the total amount of the blinds to determine playable hands when no other player has yet made a bet.

The third opponent employed is based on one from the book "Kill Phil" [12]. This book contains several strategies of increasing complexity, and these experiments use the simplest, so-called "Rookie", strategy. Similar to the two Sklansky strategies, the Kill Phil Rookie strategy contains instructions on which hands are playable depending upon whether or not an opponent has yet bet into the pot. The major difference in this strategy is that the classification of playable hands is determined by the number of players remaining and the tournament level.

For all the experiments, we seat our test player at a table against nine similar opponents from either of the three mentioned above.

Note that we do not incorporate any opponent modelling within these experiments. Whilst we acknowledge that this is an essential element of strong poker strategy in any setting, here we seek more general results regarding the dynamics of tournament play.

By discriminating between "suited" and "offsuit" hands there are 169 different starting hands in Texas Hold'em. Using the hand strength ordering of Sklansky and Chubukov [13], we create thirteen groups of thirteen hands: the strongest in Group 1 down to the weakest in Group 13. Note that the classes are not all of precisely the same size due to the varying frequencies of pairs, suited, and unsuited hands. Whilst the classification employed here is extremely coarse, it is sufficient to show the dynamics of tournament strategy that we seek.

The first suite of experiments use Monte Carlo simulation to assess strategies which employ either hand strength knowledge alone, or this in conjunction with another factor believed to influence decision-making. All strategies take the form of a threshold value between zero and thirteen. In measuring the performance of players using hand strength knowledge alone, a strategy (x) represents a player who will move all in with cards in the groups higher than or equal to *x*, and fold otherwise. When we incorporate one extra game factor the representation becomes two-dimensional, (x, y). The *x*-value is used if the binary variable is true, and the *y*-value is used if the binary variable is false. The binary variables used are summarised in Table II.

TABLE II

Binary variables used to assess decision factors

Factor	Binary Variable
Opponents' actions	No prior bet in the current hand
Tournament stage	Tournament level ≤ 6
Chip stack amount	M ≤ 5
Seating position	Early position

The variable *M* used to assess chip stack size is taken from Harrington [14] and is simply the ratio of one's stack to the total of the blinds at the current level. Early and late position are determined relative to the dealer, such that a player in the first half of those required to act is deemed to be in early position.

The simulations sequentially test every possible strategy over a series of 200 tournaments. A player scores one point for winning a tournament, and zero otherwise.

To extend the aforementioned representation into multiple dimensions for each possible binary variable presents problems due to the computational time required to cover all of the strategies. Instead, we employ an evolutionary algorithm which is able to encode an action for any possible combination of the four factors.

A strategy within the evolutionary algorithm is encoded as a chromosome of sixteen real numbers, each value taken from the interval $[0, 14)$. A single gene determines the hands played in one of the 2^4 scenario combinations of the four binary variables. When required to act the evolutionary player refers to the relevant gene, and moves all in if the floor of the value is less than or equal to the group containing their hand.

The evolutions each commence with a population of twenty randomly generated strategies. This comparatively small population was found to be sufficient in producing significant results. Individuals within the population reproduce according to tournament selection, with two elites passing through to the next generation unaltered. Reproduction occurs at a rate of 70%, and employs uniform crossover with equal weighting between the parents. Mutation applies a Gaussian shock to a randomly selected allele with a standard deviation of 2, and reflection occurs at the upper and lower bounds. Many of these shocks do not result in a change in an individual's strategy due to the representation used. In these experiments we use a relatively high mutation rate of 20%, although smaller values were found to produce similar results.

IV. Experimental Results

A. *Monte Carlo Simulation Results*

Results from the first set of experiments, in which the test players act according to hand strength alone, are plotted in Figure 1.

The graph clearly shows that betting with either too many hand groups or too few results in reduced performance. Against all three opposing strategies the highest number of tournament wins arise when the player moves all in with around the top four groups of starting hands.

There is, unfortunately, a large amount of noise seen in this plot. This can be reduced by simulating each strategy over a larger number of tournaments than the 200 used here. Further simulations were performed with the players each playing 1,000 tournaments, and more clearly defined peaks were observed. We present the results over 200 tournaments for comparison with those from the inclusion of the binary variables which follow. In these latter cases, simulating all 196 (x, y) strategy pairs over 1,000 tournaments would be extremely time-consuming.

The next suite of experiments expands the strategy representation from one to two threshold values. Each value corresponds to playable hand groups depending on the state of the additional binary variable. Here again each possible (x, y) strategy, $(x, y \in \{0, 1, ..., 13\})$ is played in 200 tournaments against a table of each of the three adversarial strategies. Figures 2, 3, 4, and 5 show the results achieved against the Sklansky Basic strategy. Plots showing scores against the Sklansky Improved and Kill Phil Rookie strategies can be found at the authors' website [15], and all show similar surfaces.

There are two interesting aspects to these graphs. Firstly, all are asymmetric in the line $y = x$. This shows that the information contained within the binary variable is of consequence to tournament performance. For comparison purposes a control experiment was performed in which the binary variable was chosen to be the result of a coin toss. The plot of this experiment in Figure 6 shows much greater symmetry, subject to noise.

The second important aspect of these results concerns the number of tournaments won by incorporating the additional information. The peaks of those graphs representing the inclusion of prior bet knowledge, tournament level, chip stack size, and seating position are far larger than that achieved through hand strength knowledge alone. The peak from the control experiment utilising a coin toss is comparable to hand strength only, as expected.

The strategies leading to the best results against the Sklansky Basic (SB), Sklansky Improved (SI), and Kill Phil Rookie (KPR) strategies are summarised in Table III.

TABLE III

Best strategies found by Monte Carlo simulations against various opponents

Game Knowledge	SB	SI	KPR
Hand strength only	(4)	(3) & (5)	(4)
Hand and prior bet	$(9, 0)$	$(13, 3)$	$(12, 3)$
Hand and level	$(1, 12)$	$(1, 11)$	$(3, 12)$
Hand and stack	$(5, 0)$	$(12, 2)$	$(8, 2)$
Hand and position	$(8, 2)$	$(13, 2)$	$(11, 1)$

It is evident from these results that the state of a binary variable polarises the test player's best strategy in all cases against all opponents. The nature of these strategies also bear favourable comparison to suggested tournament play in the non-academic poker literature.

If no opponent has yet bet the best strategy is to play many hands. If the pot has already been opened, however, a lesser number of hands should be played. Poker authors state that a player should be much tighter if an opponent has already bet into the pot.

Knowledge of tournament level affects the best play by being tighter in the early levels and looser in the late ones. These strategies also follow poker author's recommendations. They state that at the start of a tournament one should be more concerned about being eliminated and so should play to protect chips. Then in the later levels one should play more hands to try to steal the increased blinds.

A player with a small number of chips is seen to move all in with more hand groups than one with a relatively large stack. The escalating blinds used in tournaments mean that

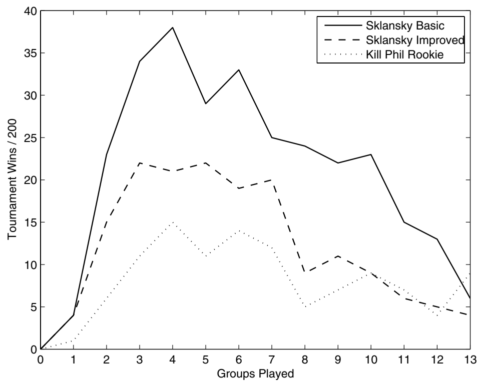

Fig. 1. Tournament wins against various opponents using hand strength knowledge only

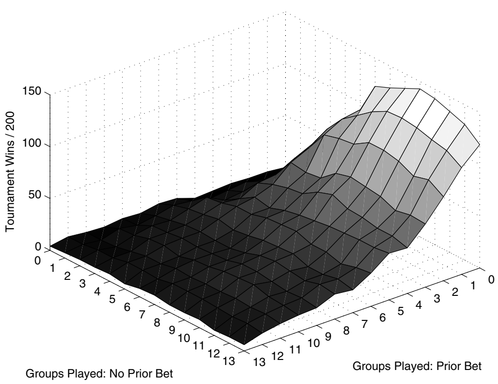

Fig. 2. Tournament wins against Sklansky Basic opponents for strategies incorporating opponents' prior bet knowledge

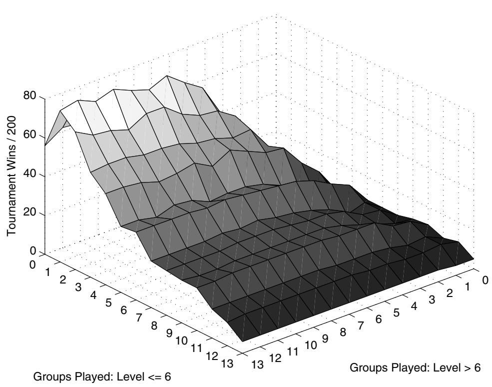

Fig. 3. Tournament wins against Sklansky Basic opponents for strategies incorporating tournament level knowledge

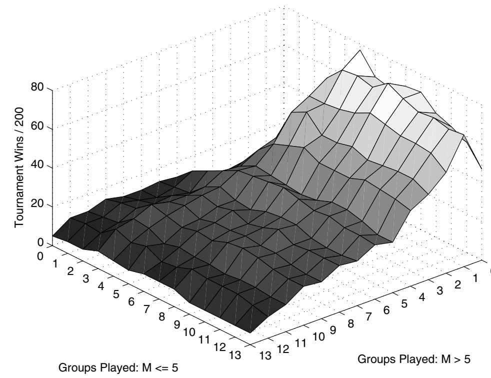

Fig. 4. Tournament wins against Sklansky Basic opponents for strategies incorporating chip stack knowledge

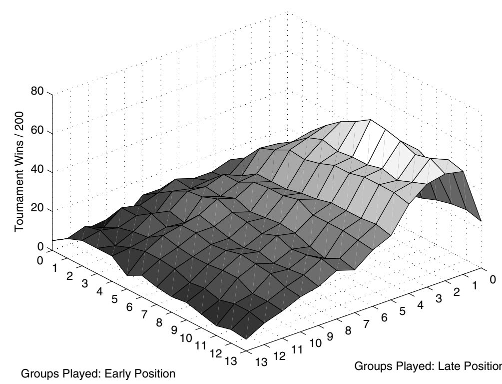

Fig. 5. Tournament wins against Sklansky Basic opponents for strategies incorporating seating position knowledge

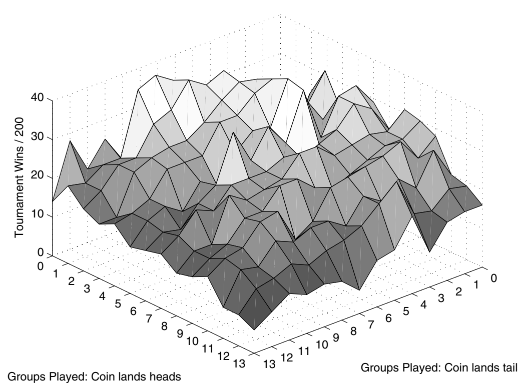

Fig. 6. Tournament wins against Sklansky Basic opponents for strategies incorporating coin toss knowledge (control experiment)

persistently waiting for strong hands leads a player to being anted away. Again the best strategies found here conform to the suggested style.

The only results here that are in conflict with poker authors concern those from the inclusion of seating position knowledge. Here we see that a player in early position bets more hands than one in a later seat. It is more commonly suggested that correct strategy is the opposite of this. A late seat means that there are less players to act after oneself, and so specifically in an unopened pot a bet with a slightly worse hand than normal can prove to be profitable. The reason for the disagreement in our experiments is likely to be two-fold. Firstly, all the opposing strategies play a more restricted set of hands once a player has entered the pot. When our test player is in early position it seems to be able to take advantage of this by betting weaker hands and essentially "scaring" the opposition. The second likely cause is that we are not able to distinguish between those occasions where the test player is first into a pot. It could still be the case that in these scenarios our player is playing more aggressively.

It is this desire to see all factors in combination that leads to the next set of experiments. The results previously seen in Table III contain potential conflicts for a strategy seeking to incorporate all game-related factors. For example, suppose we are playing the Sklansky Basic opponents and are in the latter half of a tournament with a large stack size. The best strategies found would suggest playing any hand in the top twelve groups due to the tournament level, but contradictorily that we should fold all hands since we have a large stack.

The simulation-based approach is expedient in the case of only one decision factor, but as we include more the time taken to cover the strategy space grows exponentially. For this reason we turn towards an evolutionary approach.

B. Evolutionary Algorithm Results

Candidates within a population encode strategies that dictate which hand groups to play in any of the 16 possible scenarios within the strategy hypercube. The guided stochastic search then moves the population in the direction of stronger solutions.

A randomly seeded population of twenty candidate solutions were each played in 200 tournaments per generation against tables consisting of each of the three opponents. The average population fitness per generation, and the incremental global best solutions are shown for the evolutions against Sklansky Basic, Sklansky Improved, and Kill Phil Rookie opposition in Figures 7, 8, and 9 respectively.

All of these plots show that the average population fitnesses are quick to increase in the first ten generations of the evolution. Strong solutions are more likely to be maintained, whilst weaker strategies fail to reproduce. In the early generations the rate of finding new global best strategies is at its fastest. These results are in keeping with generally observed trends in the use of evolutionary algorithms.

Improvement in the global best solution against Sklansky Basic opponents noticeably levels off by generation 25. The evolutions against the other two opponents were still finding

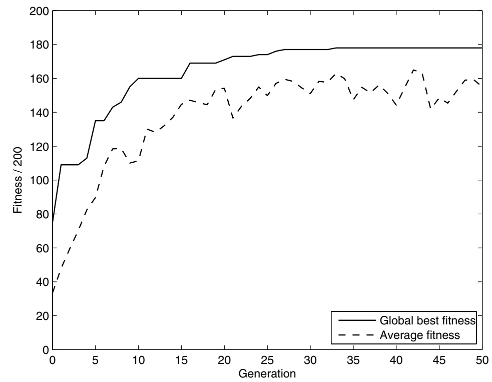

Fig. 7. Evolution against Sklansky Basic opponents

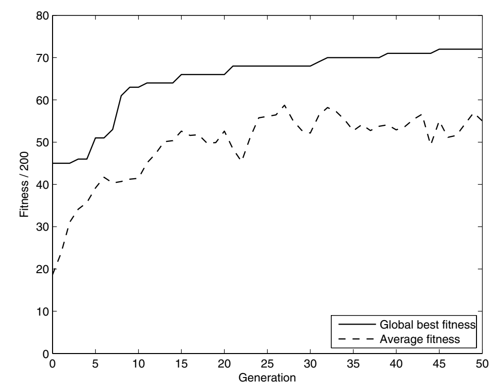

Fig. 8. Evolution against Sklansky Improved opponents

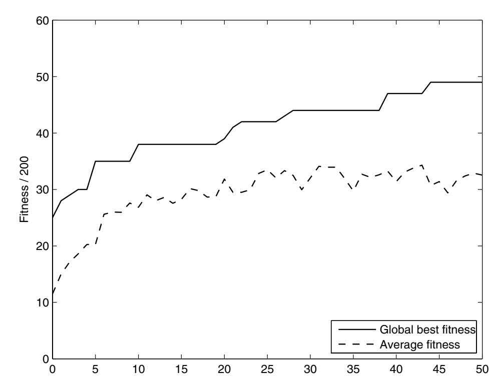

Fig. 9. Evolution against Kill Phil Rookie opponents

new global bests approximately every five generations when the runs were terminated. It is expected that even stronger strategies would have been uncovered had the number of generations been increased.

The most remarkable aspect of these graphs is in observing the number of tournament wins achieved by the best found strategies. At the termination of the run, the best solution against Sklansky Basic had won 178 out of 200 tournaments, a rate of almost 90%. This figure is a marked improvement on the comparable 38 tournament wins gained using hand strength alone shown previously in Figure 1. The evolutions against Sklansky Improved and Kill Phil Rookie opposition also show marked improvements, with both more than tripling the top score found based only on hand strength.

Next, we took all the best strategies found using the different representations and played them off against their respective opponents over 5,000 further tournaments. Performing these experiments over a larger number of tournaments allows us to gain a greater level of statistical surety in the scores attained. The number of tournament wins in these enlarged experiments are shown in Table IV.

TABLE IV
TOURNAMENT WINS (OUT OF 5,000) OF BEST FOUND STRATEGIES AGAINST VARIOUS OPPONENTS

Game Knowledge	SB	SI	KPR
Hand strength only	1,004	565	319
Hand and prior bet	3,562	1,095	595
Hand and level	1,997	746	574
Hand and stack	1,637	819	534
Hand and position	1,417	715	460
Hand and all factors	4,340	1,165	786

For all three opponents the knowledge of whether or not another player has bet into the pot shows the largest gains. Intuitively, we would expect knowledge of one's opponents actions to be the most important factor since all players are in competition with one another for the money in the pot. Least improvement comes from seating position knowledge. The lesser importance of this is most probably due to the all in or fold nature of the tournaments. Poker authors have stated the much of the positional advantage is reduced by such strategies.

Statistical analysis was performed on these results using the Z-test for the equality of two proportions. Firstly, we tested the null hypothesis that the proportion of tournaments won by the strategies incorporating a single binary variable was the same as those using hand strength alone against each respective opponent. All were found to be rejected as is shown in Table V. From this we conclude that the inclusion of any of these extra pieces of game information significantly improves tournament performance.

Following this we tested the null hypothesis that the strategies found by the evolutionary algorithms won the same proportion of tournaments as those achieved in the best case of hand strength and one other factor (in each case this factor is knowledge of an opponent's prior bet). Table VI reveals that all null hypotheses were rejected at the 95% confidence level. Hence we conclude that the incorporation of knowledge from all four factors has produced strategies with a markedly higher win rate compared to those utilising only a single factor.

TABLE V
P-VALUES FOR THE PROPORTION OF TOURNAMENTS WON WITH THE INCLUSION OF A GAME FACTOR COMPARED TO HAND STRENGTH ALONE

Factor	SB	SI	KPR
Prior bet	< 0.00001	< 0.00001	< 0.00001
Level	< 0.00001	< 0.00001	< 0.00001
Stack	< 0.00001	< 0.00001	< 0.00001
Position	< 0.00001	< 0.00001	< 0.00001

TABLE VI
P-VALUES FOR THE PROPORTION OF TOURNAMENTS WON WITH THE INCLUSION OF ALL GAME FACTORS COMPARED TO HAND STRENGTH AND PRIOR BET KNOWLEDGE

	SB	SI	KPR
Evolved player	< 0.00001	< 0.05	< 0.00001

It should be remembered that the opposition strategies are able to act on exact knowledge of their hand, stack size, and tournament level. Our test players are severely limited in the granularity of the information they can base decisions on due to the coarse classifications employed. The results given are perhaps even more impressive in the light of this.

To analyse the nature of the best evolutionary strategies found we calculated the average hand group played for each possible state of the four binary variables. These values are presented in Table VII.

TABLE VII
AVERAGE ALLELE VALUES IN GLOBAL BEST STRATEGIES AGAINST EACH OPPONENT

Scenario	SB	SI	KPR
No prior bet	8.5	7.6	8.4
Prior bet	2.1	5.9	5.0
Level ≤ 6	4.5	5.0	5.0
Level > 6	6.1	8.5	8.4
M ≤ 5	6.9	6.8	8.0
M > 5	3.8	6.8	5.4
Early position	6.6	8.6	8.0
Late position	4.0	4.9	5.4

An inspection of these results shows the same trends as those previously found in Table III. The strategies typically play less hands in the following situations: after an opponent has bet, in the early levels of a tournament, with a large stack size, and in late position.

It is interesting to note that these best strategies have managed to resolve conflicts in the signals given by each separate game factor. The strategies found by the evolutionary algorithm manage to implicitly weigh the importance of each piece of information in formulating a betting action depending upon the situation.

V. Further Work

The framework used in the evolutionary experiments is extremely amenable to expansion. For example, any other factor thought to influence decision-making could be readily incorporated into the representation.

One possible enhancement would be to increase the resolution on the factors already employed. An enlarged number of classifications of stack size, tournament level, and seating position would further refine the strategy space.

The hand groupings used in these experiments are too coarse to be of direct benefit in a real poker game. Splitting out the upper groups into better defined classes, and condensing the bad hands into a smaller number of larger groups would better reflect the differences in starting hand potential.

An interesting topic for future investigation is to better understand the relative importance of each factor dependent upon tournament level. For example, some poker authors state that stack size becomes more important than hand strength towards the end of a tournament.

Clearly the most important continuation of this research is the removal of the all in or fold betting restriction. As well as being able to bet fractional amounts of one's stack, the strategies called for would also have to encode for post-flop play. It remains an open question whether an unrestricted Texas Hold'em tournament strategy could be found within such a framework.

VI. Conclusions

In this research we set out to show that information available to players related to their opponents' prior actions, the stage of the tournament, chip stack size, and seating position are all important elements in the strategy of a Texas Hold'em tournament player. By comparing the performance of players who use knowledge of their hand strength alone to those who incorporate each of the above factors we have shown that there is a statistically significant improvement in the number of tournaments won by the latter.

We then demonstrated that an evolutionary algorithm is able to resolve conflicting signals from the decision-making factors. The strategies which combine all the available information are seen to perform to a statistically higher standard than those which use only one piece of knowledge in conjunction with their hand strength.

The "strategies" we have derived in this research should perhaps technically be termed "counter-strategies", since each is specific to their own particular opponent. However, the interpretations of the counter-strategies found (i.e. play more hands if no opponent has bet, if one has a small stack, late in the tournament etc.) are the same against all three opponents of increasing complexity. It is compelling to suggest that these tactics, which mirror the guidance given in the non-academic poker literature, should underpin the strategy of a competent player against any given opponent(s). Whilst we cannot go as far as to claim this outright, our results do lend weight to that argument.

At present the approach of game theoretic understanding and opponent modelling has been seen to yield stronger poker programs than any that have been found by evolutionary algorithms alone. However, transitioning these limit ring game strategies to a no limit tournament setting requires additional understanding of the complexities of tournament play. The results presented in this research show that an evolutionary approach is well suited to the task of analysing tournament strategy, and that it can be used to complement other forms of computer poker research.

Acknowledgments

The authors would like to thank the University of Alberta's GAMES Group for the publicly available Texas Hold'em Java code [16] which was integral to this research.

References

[1] R. Ciaffone, "Robert's rules of poker," WWW Address: "http://www.worldseriesofpoker.com".

[2] J. von Neumann and O. Morgenstern, *Theory of Games and Economic Behaviour*. Princeton, NJ: Princeton University Press, 1953.

[3] D. Billings, A. Davidson, J. Schaeffer, and D. Szafron, "The challenge of poker," *Artificial Intelligence Journal*, vol. 134, no. 1, pp. 201–240, 2001.

[4] D. Billings, N. Burch, A. Davidson, R. Holte, J. Schaeffer, T. Schauenberg, and D. Szafron, "Approximating game-theoretic optimal strategies for full-scale poker," in *Eighteenth International Joint Conference on Artificial Intelligence*, Acapulco, Mexico, 2003, pp. 899–905.

[5] F. Southey, M. Bowling, B. Larson, C. Piccione, N. Burch, D. Billings, and C. Rayner, "Bayes' bluff: Opponent modelling in poker," in *Proceedings of the Twenty-First Conference on Uncertainty in Artificial Intelligence*, Edinburgh, Scotland, 2005, pp. 550–558.

[6] K. Chellapilla and D. B. Fogel, "Evolving an expert checkers playing program without using human expertise," *IEEE Transactions on Evolutionary Computation*, vol. 5, no. 4, pp. 422–428, 2001.

[7] L. Barone and L. While, "Evolving computer opponents to play a game of simplified poker," in *Proceedings of the 1998 International Conference on Evolutionary Computation*, Anchorage, AK, 1998, pp. 108–113.

[8] ——, "An adaptive learning model for simplified poker using evolutionary algorithms," in *Proceedings of the Congress on Evolutionary Computation*, vol. 1, 1999, pp. 153–160.

[9] G. Kendall and M. Willdig, "An investigation of an adaptive poker player," in *14th Australian Joint Conference on Artificial Intelligence*, Adelaide, Australia, 2001, pp. 189–200.

[10] J. Noble, "Finding robust Texas Hold'em poker strategies using pareto coevolution and deterministic crowding," in *Proceedings of the 2002 International Conference on Machine Learning and Applications (ICMLA 02)*, Las Vegas, NV, 2002, pp. 233–239.

[11] D. Sklansky, *Tournament Poker for Advanced Players*. Las Vegas, NV: Two Plus Two Publishing, 2004.

[12] B. Rodman and L. Nelson, *Kill Phil - The Fast Track to Success in No-Limit Hold'em Poker Tournaments*. Las Vegas, NV: Huntintdon Press, 2005.

[13] V. Chubukov, "Sklansky-Chubukov hand rankings," 2006, WWW Address: "http://www.decf.berkeley.edu/ chubukov/rankings.html".

[14] D. Harrington and B. Robertie, *Harrington on Hold'em: Expert Strategy for No-Limit Tournaments; Volume II: The Endgame*. Las Vegas, NV: Two Plus Two Publishing, 2005.

[15] R. Carter, "Monte Carlo and evolutionary algorithm results for CIG'07," 2006, WWW Address: "http://homepages.inf.ed.ac.uk/s0343921/".

[16] U. of Alberta GAMES Group, "Computer poker (Texas Hold'em) Java source code," 2000, WWW Address: "http://spaz.ca/aaron/poker/src/".