

Evolution of Counter-Strategies: Application of Co-evolution to Texas Hold'em Poker

Thomas Thompson, John Levine and Russell Wotherspoon

Abstract—Texas Hold'em Poker is similar to other poker variants in that our decision process is controlled by outside factors as much as the cards themselves. Factors such as our seating position, stack size, the stage of the tournament and prior bets can strongly influence a player's decision to bet or fold on a given hand of cards. Previous research has explored the use of these factors as means of betting influence through use of a genetic algorithm applied in an evolutionary learning process. However in this previous work, the evolved player performed against scripted opponents at the table. In this paper we describe a co-evolutionary approach where all players on the table are part of the learning process. Results vary wildly between simulations, with further analysis showing that the ability to create robust strategies is difficult given the adversarial dynamic of the game. Despite this, agents are still capable of adhering to guidelines recommended in expert literature.

I. INTRODUCTION

Poker is one of the more challenging games available to researchers due to the nature of the way the game is played. While the actions required are restricted to a small set; check, bet, all-in or fold, the decision process to make these actions is influenced by local data as well as outside factors. Poker in all forms is a game of imperfect information, hence the traditional game-tree approach to adversarial play can no longer be applied. Instead we seek to assess the wealth of a hand; how good is the current hand the player has when taking into consideration any communal cards (depending on which poker variant you play) as well as the cards you can *not* see, which may or may not be in the hand of your opponent. Furthermore, professional poker players insist we must also consider factors beyond the cards in the deck; do I have sufficient chips to survive the next round? What were the actions of other players in this round? Interestingly, even the position of your seat on the table can have a significant effect on decisions [1].

To assess the validity of these claims, work found in [2] sought to apply machine learning algorithms to see whether agents that were provided with this information would make the same decisions as those suggested by authors of poker literature. Using a genetic algorithm tied to an evolutionary process, the authors found that an AI agent would learn the same principles in all but 1 instance. Furthermore, the resulting agents were capable of performing on poker tables consisting of 3 scripted agents; Sklansky Basic, Sklansky Improved and Kill Phil, each inspired by strategies from non-academic poker literature.

Strathclyde Planning Group, University of Strathclyde, Glasgow, G1 1XH, UK, email: `forename.surname@cis.strath.ac.uk`.

In this paper we present an expansion on this work by applying the same principles to a competitive co-evolution model; where agents are now assessed by how well they perform against tables populated by fellow evolving agents. With the intention of discovering whether applying co-evolution to our agents will generate generic agent behaviours capable of playing against different styles of play.

The layout of this paper is as follows, we begin with a brief recap of the Texas Hold'em variant of poker followed by an exploration of related work in Section II. A thorough breakdown of agent representation, evolution cycle design and evaluation methods are given in Section III followed by a report of the most interesting results from our experiments in Section IV. Final conclusions in Section V.

II. TECHNICAL & RESEARCH BACKGROUND

A. Texas Hold'em Poker

Texas Hold'em is one of the most popular variants of poker in America and Europe and perhaps the most widely played, with play ranging from internet poker rooms to the main event at the annual world series of poker. Texas Hold'em is a community card game (i.e. utilises cards dealt face up in the centre of the table that can be used by all players) where the player must score the best hand possible using the 5 community cards as well as the player's 2 hole (hidden) cards. Texas Hold'em accommodates for up to 22 players at a single table, however in tournament situations 2 to 10 players at a table is far more common. A breakdown of the rules of Texas Hold'em and other important information can be found online [3].

A game of Texas Hold'em will follow one of the 3 forms of betting structure; limit, pot limit or no limit. In the opening 2 rounds of limit Hold'em, the sizes of the bets are fixed while in the final rounds the fixed bet size will double. Pot limit restricts all bets to a fixed size, while No-Limit as the name suggests removes all restrictions. Texas Hold'em like all other poker variants, can be played in a ring/cash game or in a tournament. In the former players contest with real money with no predetermined end time and players are given the opportunity to buy themselves back into the game. Meanwhile tournaments play with chips and end once only one player is at the table as no buy-ins are possible. A summary of the differences between the two types of play can be found in Table I.

Texas Hold'em presents an opportune game for strategic and mathematical analysis as a result of certain factors within the game. Namely the nature of the betting system, the imperfect information available to any given player and the

TABLE I

A BREAKDOWN OF THE STRUCTURAL DIFFERENCES BETWEEN RING GAME AND TOURNAMENT POKER [2].

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

decision processes required from pre-flop (the betting round prior to the first 3 community cards being revealed) to river (the point at which the 5th card is placed on the table). All the while maintaining a simplicity that allows for novice players to participate. It is these unique properties that bring many researchers to the (simulated) poker tables.

B. Academic Poker Research

The first traces of research in poker dates back to the 1950's, with work conducted by von Neumann and Morgenstern found in [4] seeking to formulate a framework for strategy selection in a simplified poker variant. The toy domain incorporated 2 to 3 player games with only pared decks of cards being drawn (so as to reduce the space of possible strategies).

However as time has progressed, the interest in mathematical models and strategical analysis has made way for artificial intelligence research. AI researchers are intent on developing robust software agents capable of competing beyond novice level in more realistic versions of poker variants.

The GAMES Group at the University of Alberta have made significant contributions in ten-player limit Texas Hold'em ring games. Initially work found in [5] placed emphasis on rule-based decision making; with a system that combined statistical analysis with a rule inference system to dictate actions. This work was followed by research published in [6] with the authors applying abstraction techniques to model 2-player Texas Hold'em, which is later challenged using linear programming to generate pseudo-optimal strategies.

There are efforts that have attempted a more autonomous approach to agent decision making, with popular work by Barone and Barone [7] and Noble [8], [9] providing a variety of canonical evolution and co-evolution based approaches.

C. Research by Carter

The focus of our work is an expansion of work found in [2], in which a series of experiments are conducted to assess whether the use of game-specific factors in tournament poker play helps generate better players than those relying solely on hand strength. Authors of non-academic poker literature repeatedly state that factors outwith the cards in play are conducive to good strategies should they be incorporated into the decision making process prior to betting. The factors taken into consideration are reduced to binary representations

TABLE II

THE SERIES OF FACTORS, TAKING HAND STRENGTH AS A GIVEN, USED IN [2] AS PART OF THE DECISION MAKING PROCESS. NOTE THE VARIABLE M IS TAKEN FROM [10] AND REPRESENTS THE RATIO OF A PLAYERS STACK SIZE TO THE TOTAL OF THE BLINDS AT THE CURRENT LEVEL OF PLAY.

Factor	Binary Variable
Prior Opponents' Action	No prior bet in the current hand
Tournament Stage	Tournament level ≤ 6
Chip Stack Amount	$M \leq 5$
Seating Position	Early Position

based on information from the world. These are expressed as shown in Table II.

Having applied Monte-Carlo Simulations to train tuples of hand strength with one of the four factors in 10-player winner-takes-all all-in or fold pre-flop Texas Hold'em¹. Results proved promising, with agents learning to tailor behaviours with successful results, then a second series of experiments applied an evolutionary algorithm to combine these factors together into one decision making process. The resulting agents generated interesting results against a variety of scripted agents; the Sklansky Basic and Sklansky Improved strategies (coded from the descriptions given in [1]) and the Kill Phil Rookie strategy (based on readings from [11]). Interestingly, the trained agents followed many of the recommendations suggested in the non-academic literature, many agents played *less* hands under the following circumstances:

- After an opponent has bet.
- In the earlier stages of the tournament.
- When owning large stack sizes.
- Should the player be seated in a late position.

In comparison with suggestions made in the non-academic literature, the agents learned the same rules in three of the four instances; as authors suggest betting should be more active in later seating positions.

Carter reflects on these behaviours, considering them as possible "counter-strategies" to the 3 scripted agents provided. And it is this reflection that is expanded upon in this research. Given that Carter's work focussed on the 3 scripted agents, we have applied a competitive co-evolution model to the problem. It is our intent to assess whether it is possible to generate an all-round counter strategy; whether it is feasible for a player to develop these factors into a behaviour that can counter any player as efficiently as possible.

¹Please refer to Section III of [2] for a breakdown of the changes this version makes to the traditional poker variant

III. EXPERIMENTAL DESIGN

Once again all experiments focus on the ten-player winner-takes-all in or fold pre-flop Texas Hold'em. Each player is given \$1,000 in tournament chips using an 11 tier tournament structure with blinds increasing as the tournament progresses. Each level consists of ten hands with the exception of the final level which holds as many hands as necessary to conclude the tournament.

Agents are encoded as the 16 possible scenarios that can occur within the strategy space. With each gene representing one of the 13 possible hand strength orderings by Sklansky and Chubukov in [12]. These orderings dictate the strength of a hand from the strongest in Group 1, to the poorest in Group 13 and providing the threshold values dictating whether an action is taken in a particular scenario.

In order to assess the players we use a population assessment system referred to as *population sampling*, a concept originally applied in the *EvoTanks* domain as means of assessing fitness relative to the population without having to assess a given solution against the complete genepool [13]. Using a population of candidate solutions of specified size, we specify a sampling rate; the number of different agents that the agent must face against in order to give an adequate fitness assessment. In order to properly express the intricacies of this method we shall walk through an example.

Using a population of 200 candidates with a sampling rate of 20%, our sample size is 40 players. However it is important that this number be a multiple of 9 (as shown later), hence our sampling rate is altered to generate the closest multiple of 9, 18% generating 36 players. From the population we then pull 2 sets of 36 players, *s1* and *s2*. Each player of *s2* is seated at a 10-seat table leaving one space empty (hence the necessity for a multiple of 9 set). At this point we iterate through *s1*, placing each agent in turn at all of the tables populated by players from *s2* for a specified number of tournaments (typically 100), with all players moving to new seats at the beginning of each tournament. Once a player from *s1* has finished playing at a particular table, the agent moves to the next table and the next agent from *s1* is seated. In time this results in all 72 agents being assessed against a variety of players in different seating positions for several tournaments. Once this has been completed, the best players on average from *s1* and *s2* are added to the parent set, with fitness being attributed as the player who wins the most tournaments on average. This process is then repeated until we have filled the parent set (20% of the population size). Afterwards a new population is created using both 1-point crossover at 70% probability and random point mutation with 20% probability.

The drawback of this method is that we are no longer aware of how capable agents are in solving the initial task, only their capabilities in reference to the fellow members of the population. To counter this issue we introduce *global awareness* tests every 20 generations throughout the learning process; where the parent set of the generation is then tested in a series of tournaments of mixed tables of Sklansky Basic,

TABLE III

FITNESS RESULTS FOR A PLAYER EVOLVED USING THE ORIGINAL EVOLUTION STRUCTURE IN [2] DEVELOPED AGAINST ALL 3 SCRIPTS. ONCE TRAINED ON A MIXED TABLE THEY ARE ASSESSED AGAINST TABLES OF ONLY THE SPECIFIED SCRIPT. A RANDOM PLAYERS PERFORMANCE IS ALSO PROVIDED AS A COMPARISON MEASURE.

Opponent	Fitness of chromosome(win%)	Fitness of random (win%)	% Improvement
Sklansky Basic	76.4	21.4	257.01
Sklansky Improved	24.0	7.3	228.77
Kill Phil Rookie	17.4	5.3	228.30

Improved and Kill Phil scripted agents. This allows us to ascertain how effective the agents are in comparison to those trained in [2].

IV. EXPERIMENT SERIES

A. Preliminary Experiment

To provide comparison measures prior to our experimentation, we apply Carter's original evolution practice on a table of mixed scripted agents. The resulting agent was then assessed against tables populated solely by each type of script, with a breakdown of the results shown in Table III. It is clear from this table that an evolved agent across all players will perform better on average than the random players. We hope to use this as a comparison measure against further results as shown in this paper.

B. Initial Co-evolution Experiments

Our first co-evolution experiments applied the design shown in Section III using the specified parameters. As shown in Figure 1, the population fails to generate an arms race dynamic, with the population constantly in flux. Furthermore the global awareness tests also indicate instability and a lack of conducive learning. Further experiments using varied parameters also failed to generate successful results. Due to the scope of material relating to co-evolution and the lack of trends in the data, it was not immediately discernable what the underlying problem was for our players. Hence it was necessary to commit diagnostics on the 'best' agents accrued.

Our first efforts were to evaluate whether our sampling structure was robust enough for the learning process. Using the best chromosomes from the previous experiment, we ran each agent against a table of random players for 10,000 tournaments. This is conducted under the hypothesis that our agents are being beaten in future populations by mere random players, otherwise it is the evaluation system at fault by not clearly assessing a player's abilities. The results from one of our runs can be found in Table IV. These results proved that our agents were not generating even remotely competent strategies, as they failed to beat completely random players on average. While this assisted in verifying our evaluation structure was sound, we still sought to discover why our agents were failing against random opponents.




Fig. 1. A graph depicting the average and best fitness values for our population in both local and global contexts in the initial coevolution experiment. As we can see no conducive learning occurs as the agents fitness values fluctuate rapidly.

TABLE IV

THE NUMBER OF TOURNAMENTS WON (OUT OF 10,000) FOR ONE OF OUR RESULTING CHROMOSOMES FROM OUR FIRST CO-EVOLUTION EXPERIMENT, AGAINST A TABLE OF 9 RANDOM PLAYERS.

Chromosome	Tournaments Won
'Good'	1061
Rand1	1624
Rand2	1136
Rand3	474
Rand4	564
Rand5	1068
Rand6	922
Rand7	1493
Rand8	566
Rand9	1092

C. Discovering why the 'Good' player lost

Having shown that supposedly 'Good' players were losing, it was then a question of discovering why these agents are losing. Upon analysing the hand groups played in the 'Good' chromosome in Table IV, we note that the agent plays loose in early positions with no prior bet later in the tournament with a small stack. Hence the agent is in fact playing a good strategy but far too loosely as many group values were in the high/poor hand regions. This could suggest why the agent fails to perform well against the 3 scripted players in the global awareness tests, as Sklanksy Basic, Sklanksy Improved and Kill Phil play tight, defensive strategies. In order to discover why these values were so high we need to look at the players the agent evolved against.

One means to achieve this was to use our 'Good' player as the seed of a new population trained using canonical evolution, and each candidate solution would be assessed against a table of random players with only the best of set s/l transferring to the parent set. When running on similar parameters as our original co-evolution experiment, our progress plots as shown in Figure 2. We can easily note the difference between this plot and that shown in Figure 1, with the best fitness proving less erratic throughout the learning process. However it still suffers from sharp drops periodically. Furthermore the average fitness of the




Fig. 2. A graph depicting the average and best fitness values for our population when running canonical evolution seeded with our 'good' player. Each agent is assessed on their ability to eliminate random players. As we can see, while there is still heavy fluctuation at times with best fitness, the average appears more stabilised.

population appears a lot more stabilised. It is interesting that the agents learn to play against these agents better than those in the co-evolution population. Since the agents are trying to learn a general solution against random opponents; a solution that is constantly shifting as each generation passes, having exploited specific random players while seeking to overcome others.

Gene analysis of these players show that agents gradually tighten their bidding strategies as learning progressed. The latter agent in Table V accrues greater fitness against random players since it played tighter than the evolved agent in Table III. This table shows the agents tightening again the random players, and it was considered beneficial to continue on this path and assess whether our population evolved under the 'Good' seed performed better than a completely random seed. A breakdown of this is found in Table VI, when evaluated the agents each accrued 32% and despite significant differences in gene values, the average difference was only by +0.1688. We feel that these highlight examples of good *general* poker strategy and not necessarily a good counter-strategy for random players.

Using this information, we sought to return to the co-evolution experiments and attempt one final experiment to see whether agents could be successfully trained.

D. Return to Co-evolution

From the results in our previous section, we come to the conclusion that our mixed table strategies are not robust enough for the type of tournament set-up we have been applying in the co-evolution model. We therefore seek to apply the strategies evolved in the previous section in our global awareness tests. Hence in each instance we now place one of the parent set on a table with 9 of the strategies accrued in our seeded random experiment. We ran the co-evolution experiment from earlier using the same parameters, with only this minor alteration made to the global awareness tests. The plot is shown in Figure 3, and once again we see no improvement in the local/population relative fitness as the average and best values continue to fluctuate erratically. The

TABLE V

THE DIFFERENCE IN GENES OF TWO EVOLVED CHROMOSOMES. THE FIRST IS FROM OUR SCRIPT-TABLE EXPERIMENT AND THE SECOND IS AFTER EVOLVING OUR CO-EVOLVED AGENT AGAINST RANDOM PLAYERS

Gene	Script Trained Agent	'Good' Random Evolved Agent	Diff
1	13	12	-1
2	13	5	-8
3	13	12	-1
4	13	13	0
5	6	4	-2
6	8	3	-5
7	13	8	-5
8	8	11	+3
9	13	12	-1
10	6	3	-3
11	13	11	-2
12	12	9	-3
13	4	0	-4
14	1	2	-1
15	7	5	-2
16	2	2	0

TABLE VI

THE DIFFERENCE IN GENES OF TWO EVOLVED CHROMOSOMES. THE FIRST IS OUR BEST AGENT EVOLVED AGAINST RANDOM PLAYERS USING THE 'GOOD' SEED, WHILE THE SECOND IS THE BEST RANDOM SEED CHROMOSOME EVOLVED AGAINST THE RANDOM PLAYERS.

Gene	From 'Good' Start	Random Start	Difference
1	12	12	0
2	5	4	-1
3	12	13	+1
4	13	7	-6
5	4	2	-2
6	3	6	+3
7	8	7	-1
8	11	11	0
9	12	7	-5
10	3	4	+1
11	11	13	+2
12	9	7	-2
13	0	6	+6
14	2	4	+2
15	5	8	+3
16	2	4	+2

global tests show us the agents were gradually improving, however this is not sustained and they drop again. This was considered worthy of further investigation, since the best players are being exploited by poorer players that have learned to counter their strategies. Hence poor players are being selected for future generations of training. In order to fully assess this, we observed changes in strategy at the 20th (Table VII), 120th (Table VIII), 160th, 200th and 230th generation² compared to recommended play shown in Table IX.

- 20th - 120th (Table VII): While little changes are made in local and global fitness at this point, we note that the genes themselves begin to take shape. With the agent playing looser in genes 9 and 12, this is promising

²Space limitations prevent us from providing complete statistics and gene snapshots, hence a written explanation is given




Fig. 3. A graph depicting the average and best fitness values for our population in both local and global contexts in our second coevolution experiment. Once again no real learning is observed at local level, but at global level we see gradual improvement followed by sharp decline.

as this is what a player would be expected to do. In particular gene 12 is being exploited to steal blinds in order to build up chips. We note that the value of gene 8 creates a subtle tactic that can accrue good takings. This works best in situations where small-stacked opponents may be forced to go all in on mediocre hands due to the blind bets. Hence a player with a large stack could possibly capitalise and knock out enemy players.

- 120th - 160th (Table VIII): During this period we see dramatic increases in global fitness, hence this period is of great interest to us. Only 4 major changes are noted in the genes as shown in Table VIII, it is interesting to note that gene 2's fluctuation contradicts the professional recommendations, becoming more loose in early situations. However we do consider it a by-product of a previous generation as means of counter-strategy and would hypothesise that this will later fluctuate as training continues. Genes 4 and 11 show tightening of the players strategy, while gene 7 presents more loose play in latter periods of the tournament; an action that contradicts professional advice however could be forgiven if using such random play brings in chips for the player.
- 160th - 200th: Once again fluctuations are found in genes 2, 7 and 12 with all changes pointing to a tighter strategy. Gene 12 tightens drastically, perhaps as a result of being caught out too often stealing chips. Having seen 12 take dramatic changes since the first gene snapshot, we expect this to continue throughout. Meanwhile overall behaviours begin to move towards professional recommendations, however the global fitness tests show a drop in fitness.
- 200th - 230th: Very little changes are made during this phase, with all genes fluctuating marginally from the gene values set in the 200th generation.

TABLE VII

A BREAKDOWN OF THE BIGGEST CHANGES IN GENES BETWEEN THE 20TH AND 120TH GENERATION OF OUR REVISED COEVOLUTION EXPERIMENT.

Gene	8	9	12
Seating Position	Early	Late	Late
Prior Bets	Yes	No	No
Tournament Stage	Late	Early	Late
Stack Size	Large	Small	Large
Old Average Value	4.88	5.08	2.355
New Average Value	11.95	10.9	12.74

TABLE VIII

A BREAKDOWN OF THE BIGGEST CHANGES IN GENES BETWEEN THE 120TH AND 160TH GENERATION OF OUR REVISED COEVOLUTION EXPERIMENT.

Gene	2	4	7	11
Seating Position	Early	Early	Early	Late
Prior Bets	No	No	Yes	No
Tournament Stage	Early	Late	Late	Late
Stack Size	Large	Large	Small	Small
Old Average Value	7.47	8.44	7.39	11.72
New Average Value	12.29	2.35	11.5	1.58

In conclusion of our gene analysis, we find that in time, the majority of genes converge on the recommended values as shown in Table IX. In short, the players are learning the best strategies suggested by Carter as well as those found in the professional literature. Interestingly, one gene that Carter highlighted as conflicting with the literature was gene 1, which has shown the same traits here as well. However the problems emerge in areas where strategy is not agreed upon in the general consensus (note the areas in Table IX marked

TABLE IX

THIS TABLE REPRESENTS THE SUGGESTED AND FINAL DECISIONS FOR EACH SCENARIO (REPRESENTED BY EACH GENE) WITH THE NOTABLE FLUCTUATIONS THAT OCCUR. IT IS IMPORTANT TO NOTE THAT WE GROUP GENES BASED ON SIMILARITIES IN GAME FACTORS WITH ONLY MINOR DIFFERENCES IN EACH.

Gene	Gene Group	Recommended Play	Learned Play	Fluctuation in Training
1	1	Tight	Loose	Low
2	2	Varied	Varied	High
3	3	Loose	Loose	Limited
4	6	Varied	Varied	Unstable
5	6	Varied	Varied	Unstable
6	4	Tight	Tight	Low
7	5	Varied	Varied	Cyclic
8	n/a	Tight	Tight	Low
9	n/a	Tight	Tight	Low
10	4	Tight	Tight	Low
11	3	Loose	Loose	Limited
12	6	Varied	Varied	Unstable
13	6	Varied	Varied	Unstable
14	4	Tight	Tight	Low
15	5	Varied	Varied	Cyclic
16	5	Varied	Varied	Cyclic

with strategy ‘varied’³). In certain instances we see either patterns emerging, with cyclic changes in the strategies as the training process continues, to outright erratic behaviour as strategy rapidly fluctuates as time progresses. This is where our agents suffer from their biggest downfall, however it does provide us with interesting insights. While our agents will adhere to principles recommended by the professional literature, the competitive nature of the coevolution model forces components of the strategy to fluctuate rapidly. As a result we are seeing the agents continue to develop counter-strategies in order to push themselves to the top of the population. Since this will only result in a temporary period of dominance, it creates a cycle of counter-strategic learning. Since the population has adhered to specific strategies in a given generation, agents attempt to surpass one another by exploiting these areas of uncertainty. The largest problem that is emerging from these exploitations is that the agents do not necessarily know what are the best actions to take in particular situations as a result of our unsupervised learning approach. Hence as we can see in Figure 3 that while the agents continue to surpass one another the strategies that showed improvement in the global assessments are lost to us.

V. CONCLUSION

In this paper we have shown efforts to generate general strategies to play 10-player winner-takes-all all in or fold pre-flop Texas Hold’em using the non-hand related factors highlighted in professional poker literature. Our sampling approach to generate strategies has shown little promise in providing us what we would consider general strategies. Surprisingly however it seems that the process may well have generated general strategies, it is simply our understanding of the game that is inaccurate.

We conducted this research to be able to play against a variety of players competently. The problem in this is due to the nature of poker as a game, with many outcomes the result of luck more than skill. Our unsupervised training approach sought to tackle this problem and in doing so finds itself trapped in cycles of counter-strategies, as one behaviour becomes dominant in the population a new one will find means to counter it. While this helps the agents achieve their task within local scope, the overall task is failing as the constant shifts in strategy mean that a lot of information is lost and seldom recovered.

We would be interested in pursuing this work further, possibly by introducing more supervised learning methods. The key in improving this process is for the agents to recognise strategies in these grey areas that are generating improvement in the global assessments. How this could be applied is at present left unanswered.

In closing, one of the most important results taken from this work is that regardless of the difficulties in strategy formulation, numerous factors highlighted in Carter’s work as

³It is recommended in some instances that fluctuating between aggressive and defensive play in certain situations can benefit the player.

well as that found in the professional literature is reproduced in our agents. While we can not claim that this gives more credibility to the professional recommendations, it possibly adds more weight to their argument.

ACKNOWLEDGMENTS

The authors wish to thank Richard Carter for his assistance that proved integral to the progress of this research.

REFERENCES

- [1] D. Sklansky, *Tournament Poker for Advanced Players*. Two Plus Two Pub, 2002.
- [2] R. Carter and J. Levine, "An Investigation into Tournament Poker Strategy using Evolutionary Algorithms," *Computational Intelligence and Games, 2007. CIG 2007. IEEE Symposium on*, pp. 117–124, 2007.
- [3] R. Ciaffone, "Robert's rules of poker," web Address: "<http://www.worldseriesofpoker.com>".
- [4] J. von Neumann and O. Morgenstern, "Theory of games and economic behavior." 1953.
- [5] D. Billings, A. Davidson, J. Schaeffer, and D. Szafron, "The challenge of poker," *Artificial Intelligence*, vol. 134, no. 1-2, pp. 201–240, 2002.
- [6] D. Billings, N. Burch, A. Davidson, R. Holte, J. Schaeffer, T. Schauenberg, and D. Szafron, "Approximating game-theoretic optimal strategies for full-scale poker," *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI)*, 2003.
- [7] L. Barone and L. While, "Evolving Computer Opponents to Play a Game of Simplified Poker," *proceedings of the 1998 International Conference on Evolutionary Computation (ICEC98)*, pp. 108–113, 1998.
- [8] J. Noble and R. Watson, "Pareto coevolution: Using performance against coevolved opponents in a game as dimensions for Pareto selection," *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO*, pp. 493–500, 2001.
- [9] J. Noble, "Finding robust Texas Holdem poker strategies using pareto coevolution and deterministic crowding," *Proceedings of the 2002 International Conference on Machine Learning and Applications (ICMLA-02)*.
- [10] D. Harrington and B. Robertie, "Harrington on Hold em," *Expert Strategy For No-Limit Tournaments*, vol. 1.
- [11] B. Rodman and L. Nelson, *Kill Phil - The Fast Track to Success in No-Limit Hold'em Poker Tournaments*. Huntintdon Press, 2005.
- [12] V. Chubukov, "Sklansky-chubukov hand rankings," web Address: "<http://www.decf.berkeley.edu/chubukov/rankings.html>".
- [13] T. Thompson, J. Levine, and G. Hayes, "EvoTanks: Co-Evolutionary Development of Game-Playing Agents," *Computational Intelligence and Games, 2007. CIG 2007. IEEE Symposium on*, pp. 328–333, 2007.