Computer Poker

Darse Billings
Department of Computing Science
University of Alberta

August 30, 1995

Abstract

Games are an interesting and challenging domain for computer science research, having the nice characteristics of a clearly defined set of rules and a specific goal. Developing a program to play a strategic game well often involves the application of theoretical concepts to practical situations, and the relative success of that endeavour can be measured with quantifiable results.

The game of poker is logistically simple yet strategically complex, and offers many properties not exhibited by chess, checkers, and most other well-studied games. Most importantly, poker is a non-deterministic game with imperfect (hidden) information. Handling unreliable or incomplete information is a fundamental problem in computer science, and poker provides an excellent domain for investigating problems of decision making under conditions of uncertainty.

Somewhat surprisingly, the potential benefits of studying poker have been largely overlooked by computer scientists and game researchers. In this essay we survey what resources are available for academic researchers, and lay some foundations for the scientific exploration of this fascinating game.

1 Introduction: Why Study Poker?

The area of computer strategic game playing has a history almost as long as computer science itself. Strategic games have proven to be a worthwhile domain for study, and many prominent figures in computer science have devoted time to the study of computer chess and other skill-testing games. Some of these luminaries include John Von Neumann, Alan Turing, Claude Shannon, John McCarthy, Donald Knuth, Alan Newell, Herbert Simon, Arthur Samuel, Donald Michie, and Ken Thompson.

Games generally have a clearly defined set of rules and a specific goal, and these properties are highly desirable for research in computer science. Most games of skill feature simple logistics yet complex strategy, which allows many interesting experiments to be undertaken.

The most basic question is how well a computer can be taught to play a given game. The fact that it is possible to accurately measure the degree of success of a particular approach, by comparing the results of many games played against programs based on rival ideas, is an important characteristic of games. Indeed, many areas of Artificial Intelligence have greatly benefited from having such a well-defined metric for progress.

1.1 Games with Perfect Information

Chess is by far the most well studied strategic game in the field of computer game playing. Many advances in computer science have resulted from this research, including developments in problem specification, exhaustive search techniques, parallel algorithms, and others. The body of literature on computer chess is quite extensive, including dedicated journals and magazines [1, 10] and collections of important papers [5, 11].

Until recently, the concentration of research on computer chess has overshadowed all other games. But with the dominating success of fast brute-force search techniques, other approaches to the problem have been unable to compete in terms of practical results. Consequently, interest in many other strategic games has grown dramatically over the past few years, with each new game offering new challenges and requiring novel approaches. The research on this broader scope of games has also been facilitated by collections of reprinted articles [6] and original material [7, 8, 9].

The best studied class of games are somewhat similar to chess in nature. These games are categorized as *two-player deterministic zero-sum games with perfect information*, and include Go, chess, checkers, Othello, Go-Moku, and tic-tac-toe. Search techniques have proven to be quite effective for many of these games, provided the search space is relatively small, or can be restricted with good heuristics. Programs based primarily on efficient search engines are approaching or have surpassed the level of the best human players for many games in this category.

Complementing the search progress is the technique of retrograde analysis, typically performing an exhaustive enumeration of all positions, working backwards from a final position. This has had a significant impact on the playing strength of chess [95, 92] and checkers programs [54], and was instrumental in solving several non-trivial games, including $4 \times 4 \times 4$ Tic-Tac-Toe [63, 16, 18], Connect-Four [13, 98, 14], Go-Moku [99, 17, 18, 19], and Nine Men's Morris [45, 46]. Other games in this category are either close to being solved or have programs which play stronger than the best human players, including Othello [66, 47], Awari [15, 18], and Renju [18, 19].

In recent years, the game of checkers has been the focus of intense investigation by researchers at the University of Alberta, headed by Jonathan Schaeffer [78]. The resulting program, *Chinook*, quickly achieved world-class strength and earned the right to challenge the human world champion, Marion Tinsley, but lost that 1992 match by a narrow margin [79, 80, 81, 82]. In the rematch of 1994, Chinook became the first computer program to win the official world championship for a game of skill, and decisively defended that title in 1995.

Chess appears to be more challenging than these other games. Despite being the flagship of computer game playing research, the best programs have achieved only a modest degree of success against the best human players. The strongest program, *Deep Thought*, plays at the level of a human grandmaster and is currently ranked among the top 200 players in the world [49]. While most researchers believe that a computer program will eventually overtake the human world chess champion, there is a disparity of opinion on how long it will take to achieve that milestone [18].

In contrast to the high skill level attained in the above games, the same techniques have failed, and appear to be hopeless, for the game of Go. The vast number of possible positions on the 19×19 Go board and the high density of reasonable alternatives at each move make the effective search space prohibitively large. Achieving a respectable level of play by the standard techniques is computationally intractable. Considerable work has been done on Go by computer scientists who are also highly accomplished Go players. Despite this effort and expertise, the strongest programs have only reached the level of weak amateur [108, 67, 96, 21, 103, 85, 29, 30, 53, 64, 31, 65].

Fundamentally different problems arise with games having a *non-deterministic* element, such as the roll of dice or the drawing of cards. Backgammon is an example of a game which preserves all of the above characteristics but is non-deterministic. Conventional search techniques are not immediately applicable, since the breadth of all possible rolls makes it difficult to restrict the search space, and does not account for the stochastic nature of the game [25]. Recently, the game of backgammon has been the basis for a highly successful application of neural nets, resulting in a program of world-class strength [93, 94].

1.2 Games with Imperfect Information

Beyond this lies the realm of games having *imperfect information*, where certain relevant details are withheld from the players, or where knowledge is not reliable. Computer science, like mathematics, is based on the processing of perfect information, which tends to separate it from the less well-defined problems that are commonplace in the real world. The study of games like poker and bridge could be highly valuable, since problems involving imperfect information are of fundamental importance to all areas of computer science. Indeed, the difficulties associated with handling incomplete or uncertain information are now receiving an increasing amount of attention in many other computer science research domains.

Surprisingly difficult problems can also arise in games with more than two players. In some cases, an allowable degree of cooperation between players can result in complicated issues of diplomacy, and other complex implications. These interesting questions are only beginning to be addressed [48].

Within the field of computer strategic game playing, it is becoming increasingly popular to investigate games which involve reasoning under conditions of uncertainty [43, 26, 12]. Most of this work has been done in relation to the game of Bridge, where there is imperfect information both in the bidding and in the play of the hand. While some interesting results have been obtained, the overall playing level of current Bridge programs is still very weak compared to human players [22, 56, 102, 57, 55, 44, 97, 26, 12, 83]. This may be due, at least in part, to the logistic complications involved with bidding, interpreting bids, and applying information from the auction to the play of the cards.

Technically, Scrabble is also a game of imperfect information, since each player does not know what tiles are held by another player. However, it appears this property does not have a significant degree of strategic importance, in practice. Other compensating strengths, such as having perfect memory of the complete official dictionary and the ability to search all legal moves, have resulted in programs which are on par with the best human players [84, 100].

The game of poker involves all of the above extensions to traditional studies in computer game playing, being a non-deterministic multi-player zero-sum game with imperfect information. Unlike Scrabble, the imperfect information cannot simply be ignored, since handling this aspect of the game is fundamental to competent play. As noted, poker has well-defined parameters (with much simpler logistics than Bridge), and a wealth of strategic complexity. This favourable combination of good properties suggests that research into poker could be quite fruitful. The potential benefits of studying small, precisely defined questions involving limited degrees of imperfect information are manifest.

Specific features of poker that are not prominent in other games include risk management, the necessity of bluffing, the implications of multiple opponents, discerning deception, and deducing the styles of other players in order to exploit their weaknesses.

Somewhat surprisingly, very little academic research has been done on the game of poker. This seems to be a serious oversight, especially for researchers in the field of computer game playing. This essay will survey the diverse resources that are available for such research, and make some well-informed recommendations on how these investigations might be performed.

2 Literature Survey

Since research into poker is not yet well developed, this literature review will encompass a broad scope of loosely related work, including game theoretic analysis of simplified poker models, other mathematical studies on restricted problems, the classic instructive books on the game and, of course, the limited work that has been done by computer scientists. We will also note resources other than the written word, which may be useful to future researchers.

2.1 Game Theoretic Analysis

Preliminary academic research into poker actually started very early in the computer age. In the book "Theory of Games and Economic Behavior" (the founding work of game theory [101]), John von Neumann and Oskar Morgenstern used mathematical models to analyze certain greatly simplified games of "poker". Among other things, they demonstrated the fact that bluffing is an absolutely essential component of poker, and that any sound strategy must include bluffing with a certain frequency. While this was interesting, and useful as an example of the application of game theoretic principles, the games studied were too far removed from real poker to be of much practical value.

Other fundamental works into the study of simplified poker were developed by John Nash and Lloyd Shapley [61] and by Samuel Karlin [51, 52]. Collections of related papers on the theory of games are also available [2, 3, 4], as well as an excellent treatise on the analysis of all games [23, 24].

An attempt to adapt these mathematical models to more realistic versions of poker was made by Newman [62], but with only a limited degree of success. More recently, this approach has been revisited and more fully explained by Sakaguchi [71]. Beyond this, there has been little development of the original ideas, probably because they were originally intended as a lesson in the use of game theory, rather than as a serious investigation of poker dynamics.

Consequently, the models which have been developed to date are severely limited with respect to the real game of poker, and are of little use to the practical problem of writing a computer algorithm to play a strong game of poker. Nevertheless, general game theoretic notions can be applied to the practical problem, and the original references may be helpful in directing that method of thought.

There are at least two potentially useful ideas stemming from game theory. The first is the techniques used for determining certain optimal betting strategies. The second is the utilization of optimal bluffing and calling strategies. In both cases, the 'pure' solutions to the simplified problems must be adapted to be applicable to the real game, but the underlying principles constitute a solid starting point from which to develop a sound approach.

An optimal betting strategy for pot-limit poker was developed in a paper by William Cutler [32]. Like previous studies, this analysis was based on a simplified poker game with only one betting round and no draw. However, the analysis method is generalized to include games where any number of re-raises are permitted, which is more realistic than the usual no-raise or one-raise scenario. Furthermore, the manner in which the optimal frequencies were computed should still be applicable to a more realistic poker setting, once the effects of multiple betting rounds and the drawing of cards is taken into account.

We now look at two books which undertake a complete game theoretic analysis of real poker games, albeit with limited degrees of success.

2.1.1 "Poker Strategy: Winning with Game Theory"

An application of game theoretic betting principles is found in the book "Poker Strategy: Winning with Game Theory" by Nesmith Ankeny [20]. In this 1981 book, the author attempts to give a complete near-optimal strategy for the game of Five-Card-Draw. This is an ambitious undertaking, and should be read by any researcher in the area. Nevertheless, understanding this analysis also reveals some of the many limitations inherent with such an approach.

The book provides many good examples of game theory at work. Each phase of the game is broken down into separate problems, and analyzed mathematically. The solutions are presented in the form of betting and calling rules, and supported by tables of hand distributions for that particular phase.

Unfortunately, a lot of valuable playing information is lost by treating each phase of the hand in isolation, and this is not easily corrected with a strongly mathematical approach. Additionally, several fundamental poker principles, such as betting position (when one must act in a given betting round), are not adequately addressed. Despite this, the very conservative approach recommended in this text is likely to break even or show a small profit, even against fairly strong players. On the other hand, it is not especially practical, and is rather self-limiting since predictable play on the part of the weak opponents is not exploited.

This is a moot point in any case, since the analysis is outdated, and of limited usefulness for other reasons. The game of Five-Card-Draw is no longer popular, and the same analysis techniques are unlikely to produce good results for the more strategically complex forms of poker played today, such as Texas Hold'em and Seven-Card-Stud.

2.1.2 "Winning Poker Systems"

A generalized analytic method for comparing the value of different betting decisions was presented by Norman Zadeh in his PhD thesis [104] and expanded upon in his 1974 volume "Winning Poker Systems" [105].

In this important book, Zadeh performs a complete mathematical analysis of several games, including Five-Card-Draw, Lowball-Draw, Stud poker with five, six or seven cards, High-Low Seven-Card-Stud, and High-Low Five-Card-Draw. In each case, he goes beyond the purely mathematical computations, offering "best" strategies to apply in practice. These strategies are often presented in tables, and based on a Bayesian probability analysis given the minimal strength of the opponent's hand.

He also computes the quantitative value of betting position, and incorporates this into the proposed strategy for each game. This is indicative of the completeness of this analysis, as most studies omit this complicating factor, even though it is clearly a critical factor in the actual game.

With the lessons learned from each variation, Zadeh formulates a number of generalized rules, which should serve as reliable guidelines to all poker variations. Since these conclusions are not limited to small models of poker, but to the actual play of the game, they should be of some value to researchers in computer strategic game playing. An appendix outlining some of the computational methods used for the book is also highly valuable, although somewhat brief.

2.1.3 Other Mathematical Studies

Although game theory would seem to be the natural mathematical discipline for the study of poker, a number of other specific mathematical problems arising from the game have also been studied.

Many of these are only tangentially related to the core problems being addressed by strategic game playing, but are still worth looking at, if only for the sake of completeness. Two Japanese mathematicians, Minoru Sakaguchi and Setsuko Sakai, are responsible for most of the work on these loosely related topics. Some of the problems they have looked at include the effects of partial information [68, 72, 73, 74], multi-stage poker [69, 75], the disadvantage of being the first player to act in a given betting round [70], and a few of the subtleties encountered with more realistic poker models [71, 77]. Notwithstanding the highly specialized nature of these problems, a few of their mathematical ideas might be incorporated into algorithmic analysis techniques. More optimistically, the purely mathematical approach may eventually produce some tangible dividends for poker practitioners. For example, in one of their most recent articles, Sakaguchi and Sakai solve (from a purely mathematical standpoint) some of the fundamentally difficult problems in three-person playing scenarios [76].

While these papers may be of limited practical value, it is important to maintain a mathematically precise view of the game. Toward this end, some background in probability theory is essential for academic poker researchers. While this knowledge can be acquired in many ways, one strongly recommended reference is "The Theory of Gambling and Statistical Logic", by Richard Epstein [33].

2.2 Classic Books on Poker

Thousands of books have been written on how to play poker, including some that are over a hundred years old. Unfortunately, the large majority of these books are deeply flawed in terms of strategic advice. The principles of correct poker strategy are both highly complex and subtle, and most authors simply do not fully understand these concepts, or are not capable of expressing them with the written word.

Moreover, the level of human play is generally so low that it is quite possible for a player to be highly successful despite having some serious misconceptions, simply because those weaknesses are never exploited, or because the opponents are making more frequent or more serious mistakes.

Typically, a poker book will offer a collection of playing tips, some of which are valid, and others which are dubious, or highly dependent on the specific conditions of the actual game. The authors of such books should not be judged too harshly, however, since the information may be quite valuable to the intended audience – the average human poker player. But for the purposes of productive academic research, a much more systematic and mathematically rigorous analysis is required. Once that strong theoretical foundation has been established, a completely sound approach to playing strategy can be constructed.

Fortunately, there are a few books which are excellent resources. These have been written by world-class professional poker players who have not hidden any "trade secrets", and who are also excellent writers and teachers. The most prominent authors on poker theory and praxis are David Sklansky (a former actuary), and Mason Malmuth (a professional mathematician) [87, 90, 86, 58, 88, 89, 59, 60, 107, 91]. The treatise by twice world champion Doyle Brunson is also highly valuable [27].

The books reviewed in this section will prove to be indispensable as a source of knowledge about poker in general, and about Texas Hold'em in particular. Nevertheless, it must be noted that these works are intended for the development of human players, and must be interpreted and refined considerably to be applicable to the task of computer programming.

2.2.1 "The Theory of Poker"

Probably the single most important book ever written on poker is "The Theory of Poker" by David Sklansky [87]. Written in 1987, it was the first book to correctly identify many of the underlying strategic principles of poker. These concepts are illustrated with examples from Texas Hold'em, Seven-Card-Stud, Five-Card-Draw, Seven-Card-Lowball, and Lowball-Draw, but they are equally applicable to all variations of poker.

While it is beyond the scope of this essay to present a complete overview of poker theory, a few examples of essential concepts will be given for context. First, after explaining the nature of mathematical expectation, Sklansky states the overriding principle of the game, which he calls "The Fundamental Theorem of Poker".

> **Sklansky's Fundamental Theorem of Poker:**
>
> Every time you play a hand differently from the way you would have played it if you could see all your opponents' cards, they gain; and every time you play your hand the same way you would have played it if you could see all their cards, they lose. Conversely, every time opponents play their hands differently from the way they would have if they

could see all your cards, you gain; and every time they play their hands the same way they would have played if they could see all your cards, you lose.

The Fundamental Theorem is stated in common language, but has a precise mathematical interpretation. The expected value of each decision made during an actual game can be compared to the expectation of the correct decision, based on perfect information. Each player's long term expectation is determined precisely by the relative frequency and severity of these "misplays". On average, a player who makes fewer misplays than her opponents will be a winning player. The theorem may appear to state the obvious, but has many subtle implications to poker strategy, some of which are illustrated in the text.

Other fundamental concepts introduced in this book include "odds" (pot odds, effective odds, implied odds and reverse implied odds), the value of deception, the danger of the free card, the semi-bluff, and the importance of position. Each of these notions can be encorporated into a theoretical framework for understanding the game, and could prove to be substantial strengths for a computer algorithm.

Issues of practical importance are also addressed in the book, such as reading hands, understanding the psychology of poker, and evaluating the profitability of a game. While these topics may be of a less theoretical nature, they are among the many abilities required for play at the highest levels. It is unclear to what degree a computer algorithm can excel at these "human" aspects of the game, or whether it is even necessary to attain world class strength.

Note that this classic book does not attempt to give a step-by-step procedure for playing each game, but instead teaches the player how to *think* correctly about each situation that may arise. This requires considerable effort on the part of the student, but once the principles are fully understood, they are much more reliable, and can be applied to any form of poker, regardless of the particular characteristics or game conditions.

Sklansky also includes a chapter on game theory, as it applies to bluffing and calling. This is done largely for the sake of completeness, and to show that he is aware of such views. He then goes on to explain some of the limitations of such a system, and justifies the more pragmatic approach to bluffing, described in a separate chapter.

2.2.2 "Hold'em Poker for Advanced Players"

Probably the most popular book among poker professionals is "Hold'em Poker for Advanced Players", by David Sklansky and Mason Malmuth [90]. Written in 1988, it was a follow-up to an earlier introductory book by Sklansky [86], but is a much more thorough exposition of Hold'em. While geared toward middle-limit (typically \$20-\$40) Hold'em games, it can also be applied to other limit Hold'em games, since it is based on the same sound general principles established in "The Theory of Poker".

Naturally, each poker variation will have its own unique characteristics, and many of the common problems encountered in Hold'em are addressed in this book.

Since there are only 169 distinct types of two card starting hands in Hold'em, it is possible to prescribe a strong system for the play of the cards in the first betting round. The book classifies each starting hand into one of nine different groups, according to strength and playability. The text then indicates which groups should be played in various betting positions (early, middle, late, or blind positions), and what adjustments should be made for various game conditions, earlier raises, and so on.

This rule based system is strong enough to ensure that even a learning player can handle the opening round at a strategically high level. The method could easily be incorporated into a Hold'em program, although there are certain disadvantages to using a predetermined set of guidelines based on human experience, regardless of how successful they may be. Moreover, an "expert system" cannot be used for later stages of the game, since there are too many distinct situations where any static method will incur serious inadequacies.

Each of the general principles of poker strategy are addressed in the book, as they apply to the game of Hold'em. Many specific examples are given to illustrate the most important concepts, such as semi-bluffing and free card danger. Beyond this, several Hold'em specific situations are considered in depth, from commonly occurring patterns to rarer but critical decisions which can win or lose several bets. A chapter is devoted to the topic of modifying the given strategy to games different in nature from the middle-limit game being used as the standard. Unfortunately, this chapter does not go into much depth, and the onus is put on the reader to deduce the appropriate adjustments.

Despite the popularity of this book, few readers extract the maximum benefit from it. Each passage is written in a clear, concise style, which may actually mask the depth, accuracy, and importance of each idea. To gain a full appreciation for the wisdom contained in the book, it must be studied very closely, applied in practice, and reviewed frequently. It is reasonable to expect that poker researchers will have to attain a complete mastery of the game in order to gain the necessary insights required to write a very strong algorithm. This book will provide the surest route to a deeper understanding of Texas Hold'em, which should then serve as a strong model for all forms of poker.

2.2.3 "Super/System: A Course in Power Poker"

In the opinion of many serious players, the first great book on poker was "Super/System: A Course in Power Poker", produced in 1978 by Doyle Brunson, in conjunction with many of the top professional players of the day [27]. Each chapter addresses a different game, and is written by one of the foremost experts in the field. Despite the long history of poker literature, many of the most important concepts of poker strategy appeared in print for the first time in this book. While somewhat dated, it is still a valuable reference to players and researchers alike, because of the breadth of topics and poker formats discussed.

The opening chapter on general poker strategy was written by Brunson and discusses his philosophy and approach to the game. He introduces many of the topics which are to be discussed throughout the book, as well as some more nebulous views that would not fit elsewhere in the tome. While his opinions may or may not have any particular relevance to how a computer should play the game of poker, it is nevertheless important to determine, in theoretical terms, why a particular approach is highly successful against human opponents, while another is less suitable.

The key difference between Brunson's style and that of many theoreticians is his aggressiveness. While other authors suggest that aggressiveness is an important trait of top level players, the true value of this characteristic is belied by their conservative recommendations. Brunson makes it clear that his success is directly attributable to his willingness to take risks – betting or raising where other players would prefer a safer option.

It is now known that there are solid theoretical explanations for why a "fast" aggressive style is superior to a more conservative approach, but these observations have never been properly accounted for in previous studies.

Game theoretic studies generally have not considered situations where there is to be a draw of cards followed by another betting round. One of the consequences of this characteristic of real poker is that there is considerable risk involved in *not* betting when one has the best hand. Moreover, from a practical point of view, an opponent is far more likely to commit a misplay (such as folding a stronger hand) when confronted with a bet or raise, rather than a check or call. This suggests that an aggressive style is bound to be a more successful strategy than a more conservative approach, in practice.

Super/System provides a vast amount of poker knowledge, but understanding the value and necessity of aggressiveness is perhaps the book's most important contribution to poker praxis.

The first chapter dedicated to a specific game addresses Five-Card-Draw, and is written by Mike Caro, who is recognized as one of the foremost authors on the game of poker. In addition to sound basic strategy, Caro discusses other ways in which a strong player can gain an advantage over the opponents. One of these is through the use of "tells" – mannerisms exhibited by other players at the table which reveal information about their hands. Caro is considered to be one of the world's premier players at the art of poker psychology and behaviour, and later authored "The Book of Tells", in which he uses numerous photographs to illustrate the traits he considers to be valuable clues. Unfortunately, a computer algorithm is unlikely to be able to obtain or utilize such information, but neither is it vulnerable to such an attack.

The next game to be looked at is Seven-Card-Stud – one of the most enduring and challenging of all poker variations. The author is David "Chip" Reese, who is widely acknowledged to be among the best players in the world. Among the advanced concepts he addresses are the effect of the ante size, obtaining free cards, betting or check-raising a mediocre hand to narrow the field, the strength of concealed pairs, and adapting to unusual game conditions. This excellent treatment of Stud should be a definite asset to future researchers of this important poker variation.

The next chapter, authored by Joey Hawthorne, deals with three popular forms of Lowball poker, where the objective is to make the weakest possible hand. Some practical advice is given which may be useful to the casual player, but this section is not considered to be as strong as the other chapters of the book.

The fourth specialized chapter investigates High-Low Seven-Card-Stud, and is written by David Sklansky. Unfortunately, the particular game he addresses is no longer very popular, having given way to a variant where an Eight low or better is required to win the low half of the pot. Despite this fact, most of the information in this chapter remains indispensable, and was the first complete exposition of the special theoretical properties of high-low poker. Written in his usual clear and concise style, Sklansky demonstrates the unexpectedly high value of “scooping” (winning both halves of the pot), which thereby dictates all aspects of high-low strategy. Another general concept is the computation of the “escape probability”, which determines the odds necessary to call with a weak hand in the hope of winning one half of the pot. Each element of the game (with no qualifier for low) is taken in turn, from playable starting hands to common playing errors. The chapter ends with a side discussion of declaration strategy, although this is not used in standard games, which are played “cards speak” (no declaration).

The first of two chapters on Texas Hold'em is written by Bobby Baldwin, who was one of the youngest players ever to win the world championship. This chapter investigates the limit poker form of Hold'em. While it has since been superseded by “Texas Hold'em for the Advanced Player”, this chapter still provides a useful introduction to the game, without having to invest as much time and effort. It also complements the work of Sklansky and Malmuth, nicely illustrating some of the finer points of the game, including the varied play of drawing hands.

The final chapter is on No-Limit Hold'em, written by Doyle Brunson. No-Limit is a form of poker where a player may bet any amount at any time, up to and including the entire stack of chips he has on the table (called “table stakes”). Pot-Limit is another form of “big bet” poker, which is quite similar in nature to no-limit, since it is usually appropriate to choose a bet of about the size pot.

No-Limit poker is strategically very different from the more common limit format. Since a player's whole stack is always in jeopardy, the risks and the rewards are much greater in no-limit poker. While limit poker is a patient exploitation of small advantages, no-limit emphasizes “going for the kill” (trapping an opponent for their whole stack). To be played at the level of the best humans, no-limit also requires a greater knowledge of the opponent, better judgement, and more courage. Because of the emphasis on these more “human” characteristics, limit poker may be more suitable to early studies in computer game playing. However, this does not suggest that computers will never be able to compete at the highest level in no-limit poker. On the contrary, a risk-neutral algorithm that is incapable of fear may eventually have a decided advantage over human opponents.

Because of the decreased importance of the technical aspects of the game, much of the chapter on no-limit Hold'em is spent discussing the appropriate mental attitude a player should adopt. While this may not be of much use to a mathematically oriented researcher, it does provide a valuable counter-point to be considered, and put into perspective.

After an extensive glossary, the book concludes with an appendix of poker probabilities and statistics. These tables, computed by Mike Caro, go far beyond the usual fare, exploring many interesting questions of practical value. Included are many game scenarios for Draw poker (with and without the Joker used as a limited wild card), Hold'em (starting hands, connecting with the board, long shots), Seven-Card-Stud (starting hands, chance of improving), and all forms of Lowball (with and without the Joker wild card). These tables form the basis of correct strategy, and are commendable both for their relevance and correctness.

2.2.4 "Gambling Theory and Other Topics"

A useful introduction into deeper theoretical principles of poker is given in the book "Gambling Theory and Other Topics" by Mason Malmuth [58]. In this collection of essays, the author explores several interesting and subtle questions within the domain of poker, using mathematical techniques to augment his theoretical insights.

Among the many topics addressed are non-self-weighting strategies, variance in poker, the fallacies of money management, skills needed to become a world-class player, debunking pseudo-theory, strategy in poker tournaments, and a historical perspective of gambling strategy in wars.

Also of significant value to poker researchers is an appendix of Malmuth's opinions of over one hundred popular book's on poker and gambling. Each book is rated on a scale from one to ten, and a paragraph outlines the books strengths, weaknesses, and important contributions. He does not rate his own material, although it it clear he believes it to be among the best information available on these subjects. This section could be used to determine whether a particular poker book is likely to contain correct and useful information for research purposes.

2.2.5 Other Books

While the aforementioned texts contain most of the high quality information available on the game of poker, several other books may be of some tangential value to researchers.

For the study of other popular poker variations, the primary source of knowledge is again the work of Sklansky and Malmuth, along with their colleague, Ray Zee. Other titles by Sklansky include "Sklansky on Poker" [88] and "Getting the Best of It" [89]. Other books by Malmuth include "Winning Concepts in Draw and Lowball" [59] and "Poker Essays" [60]. Zee authored one book in the series, "High-Low-Split Poker, Seven-Card Stud and Omaha Eight-or-Better for Advanced Players" [107] and all three men co-authored the definitive work "Seven-Card Stud for Advanced Players" [91].

A recent book with some practical value is "Winning Low Limit Hold'em", by Lee Jones [50]. This book uses the fundamental principles set down by Sklansky and Malmuth, and shows how they apply to typical low limit games. This is an important contribution because many of the recommendations made in the context of middle limit games are not appropriate against weaker opposition (who generally call much more often than they should). The ability to adapt to the specific game conditions is an essential aspect of mastering the game of poker, and this book provides a good example of what adjustments are necessary for one commonly occurring type of game. The author is also accessible via the Internet, and is responsive to questions, being a frequent contributor to the rec.gambling.poker newsgroup.

One book which may be of particular interest to computer scientists is "Percentage Hold'em: The Book of Numbers", by Justin Case (*aka* Will Hyde) and Steve Jacobs [28]. This book gives the complete results of hundreds of computer simulations, pertaining to the winning rates of each starting hand in Hold'em. Each type of hand was played off against a number of randomly generated hands, ranging from one to nine opponents, and the winning probability (and return on investment) was measured in each case. Every experiment was performed 500,000 times to ensure a high degree of statistical significance. Although it is now known that there was a small systemic error in the computation process (tied hands were counted as wins, rather than partial wins), the results are nevertheless interesting, and are available online. Unfortunately, there are severe limitations to the applicability of such numbers, resulting from the underlying, highly unrealistic, assumption that all players will call every round, to the showdown.

2.3 Previous Computer Science Studies

Virtually all of the early investigations into computer poker, dating from 1961 to 1980, were done by Professor Nicolas Findler and his associates. In a series of papers and reports, he used the game of poker as a model for a variety of studies which would now fall into the broad category of Artificial Intelligence [35, 36, 37, 38, 39, 40, 42].

Findler must be given credit for recognizing the potential rewards of research into poker, as a model for decision making under conditions of uncertainty. Indeed, many of the high-level ideas he expressed as motivations for his research are still poignant and insightful.

Historically, Findler was one of the founding researchers in machine learning and computer game playing [34]. During these formative stages of computer science, he was a leader in the area of algorithmic representation of human cognitive processes [36]. Since the origins of his computer poker research are tied to the social sciences, it is not surprising that his approaches were largely based on psychological precepts of human thought, rather than mathematically oriented analysis.

Unfortunately, by modern standards, the techniques he used to explore these issues can be criticized for being rather arbitrary, and the results obtained have little to offer present day researchers. After 30 additional years of progress in computer science, we can now use the advantage of hindsight to develop more rigorous scientific methodologies. These new methods, some of which are discussed in the last section of this essay, will hopefully ensure a better line of computer poker research, and produce results with lasting value.

Findler's first paper on the topic of computer poker was published in 1961, and was a brief preliminary report on the feasibility of designing an algorithm to enable a digital computer to play poker [35]. The introductory sentence is indicative of his background, and is also suggestive of the broad scope of interests addressed by such research:

> The machine simulation of human behavior in the mental states of uncertainty, such as estimation, prediction, choice, risk taking, decision making, makes more comprehensive these difficult conceptual and logical problems for the social scientist, psychologist, military strategist, etc.

Just over three pages in length, this article gives a succinct logical framework for a computer program to play the game of draw poker. The details are in the form of a pseudo-code algorithm, along with a flowchart depiction.

The poker variation addressed is not entirely realistic, being a form of two-person Draw poker without any opening requirements or ante, but it is still closer to real poker than any previously existing mathematical model. The omission of ante considerations suggests that Findler may not have had a deep understanding of poker theory, since the game is strategically trivial under this condition.[1]

Strictly speaking, the program was not complete, since it only accounted for the case where the machine always plays last. The algorithm also incorporated certain weak strategies, such as betting in exact proportion to the strength of its hand (which would convey far too much information to an intelligent opponent). Despite these obvious flaws, the paper made the essential point that in principle it is possible to program a computer to play a legal game of poker. Thus, with such a simple set of rules and logistics, many potentially interesting and unique questions could be studied within the domain of poker.

By 1972, Findler's poker research had matured considerably [37]. His view of the potential value from studying poker is clearly expressed in the abstract of that paper:

> A realistic approach to the study of the problems of decision making seems to be to select a reasonably complex and rich environment, in which the relevant variables are easy to identify and relatively small in number. This idea will hopefully lead to the construction of a computer program that may serve as a rigorous model and theory, with high descriptive, explicatory, and predictive power of the information processes involved. One would expect to draw, on this basis, some conclusions of fairly universal validity concerning decision making in general.

In the first part of the paper, the use of poker as a decision making environment is justified. He describes many of the unique opportunities provided by the poker domain, and these ideas are still valuable today. And while we may fault some of the methods used to explore these questions, his clear vision of the higher-level objectives of such research is well worth reading.

As part of the review of previous poker research, Findler provides an excellent summary of the mathematical and game theoretic models studied up to 1970. Since most of the interest in simplified poker models occurred in the 1950s, the majority of the important papers are summarized in tabular form, comparing the following properties:

[1] with nothing in the pot, the optimal strategy for each player is to bet only if she has a royal flush, and to fold otherwise.

– Number of players (usually two),

– Probability distribution of the hands (continuous or discrete, etc),

– Number of possible ranks for hands,

– Number of raises allowed before the draw (if one exists),

– Number of possible bids (betting actions),

– Timing of bids (alternate or simultaneous),

– Whether drawing cards is included (usually not), and

– Possible number of discards (usually not applicable),

– Number of raises allowed after the draw (usually not applicable).

This summary table is recommended for anyone wishing to review the extent and limitations of the game theoretic studies referenced earlier in this section.

Findler recognized the fundamentally important role of mathematical models, but also noted their severe shortcomings with respect to the practical problem of programming a computer poker player. His conclusions in 1972 are much the same as our conclusions today:

> As can be seen from Table 1, the mathematical approach has not been successful so far because the presently available tools of the theory of games are not sufficiently rich for the complexity of poker. Simplifying assumptions change the character of the game to varying extents. The effect of the assumptions becomes obvious when one compares the sometimes sharply conflicting recommendations obtained from different models. [...] We can, however, say that the mathematical models are valuable inasmuch as they provide a conceptual framework of investigation. They pinpoint certain problems, shed light on the complexities of poker, and make a clear account of their simplifying assumptions.

The second part of the 1972 paper discusses the empirical explorations conducted by Findler and his associates. It is here that we encounter many of the weaknesses (by modern standards) in his research methods. It is certainly not our intention to denigrate the work Dr. Findler performed, but merely to point out that his specific objectives were not the same as our current priorities, and to achieve his goals he did not require the more rigorous scientific approach we now recommend.

For example, Findler performed some simple Monte Carlo simulations to estimate certain commonly arising probabilities, and used those results to extrapolate a number of game-specific heuristics for Five-Card Draw. Judging from the *non-sequitur* conclusions and the *ad hoc* rules that were drawn, it is clear that this was not a "double-blind" procedure of inference, but was heavily influenced by the preconceived notions of the observer. The result was a rather arbitrary set of guesses, each of which may or may not have had significant value. It is no coincidence that some of the heuristics obtained were in agreement with human perception, but a telling observation is that they also contained many of the same limitations and misconceptions held by many of the so-called poker experts of the day.

To be fair, we note that Findler's research was not geared toward producing the strongest possible poker program, but was focused on simulating the thought processes of good, mediocre and poor human players. While these may be interesting problems in their own right, they impose severe limitations to the ultimate performance of a program based on these ideas. Not only is the representation cumbersome and imprecise, but the overdependence on the human model may also result in many of the flaws of the human approach to poker being carried over to the proposed algorithms.

A follow-up to this research was published in 1977 [40]. This paper again begins with a discussion of the higher-level objectives and potential rewards of poker research. These ideas are basically the same as appear in the 1972 article, with some passages essentially being transcribed, although a few concepts have been given greater attention as a consequence of the empirical developments. In particular, the failure of static "formula" methods is noted and the necessity of dynamic, adaptive procedures is expounded.

This view is supported by the relative success of learning algorithms in direct competition against simple-minded fixed rule programs. While this conclusion is not in question, it is not clear that the experimental procedures are completely valid and the results lack quantitative detail. Findler lists and describes six types of static programs and seven kinds of learning players or learning components that had been developed. Each of these simple strategies is given a subjective evaluation, both in terms of approximate asymptotic quality of play and, in the case of learning players, rate of improvement. These ratings are relative and are on a numerical scale of 0 to 100, with 100 representing the strongest program among the restricted sample. Unfortunately, no absolute measure of strength is indicated, beyond the anecdotal observation of some reasonable behavior being exhibited.

In the discussion that follows the presentation of results, the various ideas behind each static and learning player are contrasted. The importance of adaptive algorithms is reasserted, and five new learning strategies are proposed. All of these suggestions are listed in the table of results and given ratings, despite the fact that they are indicted as "under development". The fact that these fairly straight-forward ideas had not yet been implemented suggests that the overall system may not have been very conducive to the representation and execution of novel strategies. Moreover, each of these more advanced concepts represents a fairly radical shift in the basic conception of poker strategy, and therefore they would not necessarily benefit from the implementation of earlier ideas.

In the final analysis, it is the opinion of this writer that the body of research performed by Findler and his associates is of limited usefulness to future investigations into poker. While it is worthwhile to review the work he did, the goal of producing high performance poker programs will require not only different techniques, but a completely new theoretical framework from which to build upon.

The more formal structure and computer-oriented protocols advocated later in this essay are designed to allow certain incremental embellishments to be accomplished easily, thereby avoiding many developmental obstacles. By abandoning the human model in favour of a strongly mathematical view of poker, we can formulate a series of fine-grained probabilistic refinements to be considered at each stage of program development, and each of these is entirely compatible with the (more simplistic) existing systems.

2.4 Other Resources

Several resources other than the written word are available to researchers interested in computer poker. Many of these are accessible through the Internet, using the World Wide Web (WWW), anonymous ftp, or various other protocols.

The Usenet newsgroup "rec.gambling.poker" is dedicated to the analysis and discussion of a wide variety of poker topics. Many authorities on the game are among the regular contributors to this newsgroup, including well-known authors, theoreticians, casino directors, and numerous professional and expert players. The majority of articles posted to this forum are interesting and well thought out, resulting in a much higher 'signal-to-noise' ratio than is typical for most newsgroups.

Archive sites for this newsgroup are available for the perusal of past articles [126, 127]. A "Frequently Asked Questions" (FAQ) file covers a wide variety of common topics, including the rules of Hold'em and other variations, good books and magazines, poker probability questions, and an extensive glossary of poker terminology [128]. A WWW home page for rec.gambling.poker also serves as a starting point for locating many related Internet sites

[129].

Of particular interest is a computer server which enables anyone to play poker against other people from around the world. This program exists on a isolated branch of the Internet Relay Chat (IRC) system, and features a variety of poker games and players at all skill levels [124].

Residing on this server are a number of "bots" – automated programs which facilitate the play of games, analyze the probabilistic outcomes of each situation, and record a summary of each hand played [124, 121, 122] . The information acquired from this observation is compiled into databases, which provide a valuable and unique source of raw data on poker playing [116]. Michael Maurer, the author and maintainer of the observer programs and IRC poker databases, has processed some of this information in a number of different ways [118, 119, 120]. For example, this knowledge was used to conduct a study of mathematical variance in poker games, varying in skill level, number of players, and player styles [123]. Such a complete study has never before been possible, and many other inquiries could conceivably be addressed with this body of information.

The IRC server is also amenable to hosting poker playing programs – allowing them to compete against human opponents or other programs with a minimal amount of supervision. This could serve as a useful testing ground during the development of poker algorithms, and as a battleground for competition between rival programs.

In the spirit of "Percentage Hold'em" [28], Michael Maurer has also conducted simulation experiments for the evaluation of starting hands for several other games, including Hi-Low Hold'em and Hi-Low Seven-Card-Stud [111]. His WWW poker page contains many other interesting and useful features [117].

Some high quality software support is also freely available via the Internet. An extremely fast winning hand evaluator written in C is available for Unix systems, and could serve at the core of higher level algorithms [112, 113]. Superseding some of the "Percentage Hold'em" information, programs are now available which can compute *exact* values for head-to-head match-ups of Hold'em starting hands. These programs use highly optimized hand comparators, and determine all possible outcomes by exhaustive enumeration [114]. Simulator programs for comparing hands in Hold'em and Seven-Card Stud are also available [125].

There are several commercial poker programs available on the market [115]. The level of play demonstrated by these programs is generally quite poor, although some modest improvements have been attained over the past few years. Unfortunately, the proprietary nature of these products makes them of limited use to academic researchers. Moreover, it appears that the business oriented developers of these programs devote much more effort to enhancing certain selling features, such as attractive graphics and user interfaces, rather than the greater challenge of improving the playing strength.

Some free demonstration versions of commercial programs can be downloaded from Internet sites, including a fully functional and relatively strong program for Seven-Card-Stud [109].

Finally, one of the most invaluable resources to the poker researcher is human expert consultation. Authoring a program that exhibits superior skill will likely involve solving a myriad of practical problems. The most effective means of attaining the the insight necessary to overcome these obstacles is through personal mastery of the game. For this reason, a complete understanding of the theoretical principles and strategic dynamics of the game is an essential quality for pioneers in poker research. Fortunately, much of this ability can be achieved fairly quickly, even by a newcomer to poker, through the study of the classic books reviewed in this chapter, and with regular practice.

3 Recommendations for Academic Research

3.1 Scientific Research Philosophy

The opportunity exists to structure poker research in such a way that each person's contributions will be relevant and usable by others in the field. In view of the limited usefulness of previous poker research, we believe a more scientific approach is both desirable and necessary.

In keeping with the well established paradigm of the natural sciences, steady progress can be made within a research discipline through the careful design of objective experiments and quantifiable results. It is the belief of this author that such rigour has often been lacking in previous investigations, although this is somewhat understandable given that computer science, and Artificial Intelligence in particular, are still in their formative years. We now look at some general methodologies and specific techniques applicable to the study of poker that are more akin to the classic scientific method.

As with the physical sciences, an important step is to eliminate the subjective and often unreliable perceptions of humans, in favour of a more objective methodology and interpretation of results.

To some, this might seem to be an absurd idea, since poker is usually regarded to be heavily dependent on human elements, such as psychology and bravado. The key point is that the human approach to poker is not necessarily how poker should be played. Indeed, perfect poker is based on certain probabilistic facts and fundamental strategic principles which may only be clouded by focusing on the particular biases and vagaries of human praxis. By seeking out the underlying truths of poker, it is possible to achieve higher standards than those of current human expertise. By concentrating on the "physics" first, we stand to learn more about this fascinating game, and about the abstract problems of decision making in conditions of uncertainty.

Some of the concrete steps we can take in this direction include the standardization of the topic, and the search for *general* solutions rather than game-specific methods. For example, betting strategies should be based on probabilistic evaluations of hand strength and potential, rather than case-by-case situations which may arise in a particular poker variation. The same betting principles can then be applied directly to any other poker variant, provided the "front-end" of the new program produces the same form of numerical evaluations.

Conversely, researchers should avoid rule-based expert systems whenever possible, because the number of subtle and distinct cases which must be considered in a game like Hold'em is prohibitively large. More to the point, even if a set of inference rules can be developed to attain an adequate level of play, it will not teach us much about the general nature of poker or decision processes.

For this reason, we believe future poker researchers should strive to use computer-oriented techniques whenever feasible and reasonable. There are many examples from Artificial Intelligence and computer strategic game playing which demonstrate that steady improvements in performance can be attained by exploiting the natural strengths of computers, while de-emphasizing the human model. Although it may sometimes be a significant challenge to cast problems of imperfect information in terms of high speed computation and perfect memory, these methods promise to eventually produce the most natural solutions to the task, from a digital logic point of view.

We now look at some specific examples of this philosophy in more detail.

3.2 Choice of Poker Variation and Betting Structure

A necessary first step is the selection of one or two primary poker games and betting structures to study. Poker is reknown for having an unbounded number of variations and formats, but for research purposes it is necessary to focus our attention on a particular standard. This allows certain foundations to be established and used by everyone in the field. Independent results can then be readily compared, and the body of research can become more directed.

The *"Drosophila melanogaster"* we recommend for poker research is the game of Limit Texas Hold'em with a 2-2-4-4 unit betting structure and blinds of 1 and 2 units. A more detailed description of this game can be found in any of the classic books referenced in the previous section.

There are many reasons to choose this particular variation above other candidates. First, it is the most popular poker game played in casinos and public card rooms, so it is well known to all serious players. This ensures that many human expert players are available for consultation and evaluation of programs and ideas. More importantly, there is also a good source of literature on the game, as described in the previous section.

Moreover, Hold'em has particularly simple logistics and is (arguably) the most strategically complex form of poker among popular games. It is generally accepted that among common variations, Hold'em has the smallest ratio of luck to skill (although all forms of poker have a naturally high element of chance and degree of variance). Partly for these reasons, Hold'em is the game played to determine the World Champion at the annual World Series of Poker Championship.

The recommended betting structure follows the most common pattern seen in North American casinos and card rooms, and the same basic structure is used for low, medium, and high stakes tables. For example, a "\$2-\$4" game has two blinds of \$1 and \$2, bets and raises of \$2 for the first two betting rounds, and \$4 for the last two rounds. A high stakes "\$50-\$100" game uses the same format, but with all values scaled up by a factor of 25. Since most betting strategies are dependent only on the size of the bet in relation to the size of the pot, the two games are strategically equivalent (although normally very different in character and in strength of the players).

Another advantage of studying Hold'em is that it is a member of a family of closely related games, called "flop games". Each variation in this family uses the same method of community board cards, and differ only in the number of private hole cards and rules for forming hands. This allows generalized strategies for Hold'em to be applied to other variations with limited modifications. This will enable us to see if those principles are indeed transferable, rather than game-specific.

The best alternative choice for poker research is the game of Seven-Card Stud. Top level players generally consider it to be the only other popular game having a strategic complexity in the same class as Hold'em. Again, there is a good supply of expert players and correct literature on the game. Seven-Card Stud is slightly more involved from a technical standpoint, since part of the basic strategy involves tracking all exposed cards and evaluating the partial information exhibited by each opponent's hand. These characteristics can lead to interesting new problems for study, but since they are not fundamental to poker in general, they seem better suited to future research.

Many of the past studies on poker have suffered from a poor choice of variation. Five-Card Draw poker was a natural choice for study in the 1960's and 1970's, since it was one of the most popular forms of poker at that time, but it has since lost much of it's appeal. Moreover, Five-Card Draw is a rather simple game from a strategic perspective, when compared to Hold'em or Seven-Card Stud. This is unfortunate for researchers because certain *ad hoc* rule-based approaches to the game can be fairly successful without properly addressing certain poker fundamentals. Such methods should prove to be less adequate for the more challenging games of Hold'em and Seven-Card Stud, allowing more robust techniques to emerge as the best methods for high-performance results.

3.3 Executive Dealer Program for Algorithm Development

Because of the simple rules and structure of poker, it is a fairly straight-forward task to write a program to facilitate competition between rival playing algorithms. This executive program can handle all of the mechanics of the game such as shuffling and dealing the cards, providing private and public information, prompting for betting actions, determining the winner at the showdown and awarding the pot. Through such an interface thousands of hands can be played at computer speeds, and the relative strengths of different programs can be measured with statistically significant results.

This can also facilitate the systematic development of algorithms. For example, after a new feature is added to a particular program, it can be tested against previous versions to see how much the refinements have improved the playing ability. Similarly, it is possible to quantitatively measure the relative importance of certain poker game characteristics. For instance, it is known that the last player to act in each betting round has an advantage over those in earlier positions. It is fairly straightforward to design experiments which measure the precise value of good position in terms of overall expectation, and an algorithm can be fine-tuned to better utilize this strategic element as a result.

The protocol of a dealer program can be standardized and made publically available, allowing independent researchers to create compatible programs. This might be used, for example, in an open programming competition via the Internet. It is quite conceivable that programs written for fun by hobbyists may contain valuable ideas that can be used to further the field by academic researchers.

Recognizing the need for a statistically significant number of trials to accurately measure playing strength, it is worth noting that testing methods can be developed which result in a much more rapid convergence. Modifying the ideas of duplicate bridge, we can design tournament structures such that each program plays the same set of cards and situations, thereby greatly reducing the element of luck for each experiment.

Note that the a given sequence of deals can be repeated by reselecting the same seed to the random number generator governing the randomization of the cards. For example, suppose ten playing algorithms are to compete against each other for ten thousand hands to achieve a certain degree of confidence in the results. By reordering the seating positions of the participants into ten independent duplicate tournaments of one thousand hands, a higher degree of confidence can be attained with the same amount of work, because each player has the opportunity to play the same set of starting hands and resulting board cards. This does not completely eliminate the element of luck, however, since different players will be controlling the opposing hands in each situation. Nevertheless, this duplicate tournament structure can reduce the natural variance of the game by several orders of magnitude, and thereby reduce the number of trials needed for meaningful results.

To further minimize the effect of relative playing position, the seating assignments for each duplicate tournament should be shuffled rather than just rotated. For example, it would be a practical advantage to always play behind an opponent who is particularly aggressive and unpredictable, so we prefer to mix the playing order as well. This more complicated scheduling task can be accomplished elegantly for ten players by using a cyclic sequence of combinatorial derangements as illustrated in Table 1. The reader may recognize this as being isomorphic to the multiplication of integers modulo 11.

3.4 Enumeration Techniques for Hand Evaluation

In keeping with the theme of computer-oriented techniques, we now show how accurate estimates of certain game-related probabilities can be achieved by enumeration of all possible cases and the appropriate weighting of each case. In the following example, we wish to evaluate the relative strength of a given Hold'em hand, in terms of the probability of currently having the best hand among active players.

Suppose our starting hand is **K♣-Q♣** and the flop is **K♡-7♣-2♠**. Since there are 47 unknown cards in the remaining stock, there are only 1081 possible combinations of hands that our opponents might hold. We can count how many of these are better, equal, or worse than our hand, and determine where our hand is ranked against a random hand. It turns out that 42 combinations are better, 6 are equal and 1033 are worse, which corresponds to a percentile ranking of 1036/1081 = 0.958, or about a 96% chance that our hand will be better than a random opposing hand.

	Seat Number for Each Player									
Round	1	2	3	4	5	6	7	8	9	T
1	1	2	3	4	5	6	7	8	9	T
2	2	4	6	8	T	1	3	5	7	9
3	3	6	9	1	4	7	T	2	5	8
4	4	8	1	5	9	2	6	T	3	7
5	5	T	4	9	3	8	2	7	1	6
6	6	1	7	2	8	3	9	4	T	5
7	7	3	T	6	2	9	5	1	8	4
8	8	5	2	T	7	4	1	9	6	3
9	9	7	5	3	1	T	8	6	4	2
T	T	9	8	7	6	5	4	3	2	1

Table 1: Seating Assignments by Player Number and Round Number

However, this is only a crude measure of hand strength because in actuality, not all opposing hands are equally likely. For example, a hand like **A♣-A♡** would certainly be played to this stage of the game, but it is unlikely that a reasonable opponent would have called a first round bet with **7♠-2♢**. Since there are only 169 types of starting hands in Hold'em, each of the 1081 possible cases can be multiplied by an empirically determined estimate of the probability that a hand would have been played to the current point of the deal.

One implementation of this idea produces a probabilistic distribution of about 23-6-346 better-tied-worse hands, yielding an actual strength of 0.931. The fact that the true strength is lower than the simple ranking indicates that those hands which are currently superior to ours are, in general, slightly more likely to be played by a rational opponent than a random hand. In the absence of other information, we can approximate the probability of having the best hand against two opponents to be $0.931^2 = 0.867$, which is certainly strong enough to bet. This refined technique for hand evaluation turns out to be surprisingly accurate and robust, even when using rather simplistic estimates for each subcase.

3.5 The Use of Game Theory for Betting Strategies

An Example: Optimal Bluffing and Calling Frequencies

Much has been written about game theoretic optimal frequencies for bluffing, and calling a possible bluff. This serves as a nice example of how the underlying principles of game theory can be used as a starting point for a poker algorithm, but then must eventually be transcended to achieve the highest playing levels.

It can be shown that the theoretical maximum guaranteed profit from a given poker situation can be attained by bluffing, or calling a possible bluff, with a predetermined probability. The relative frequency of these actions is based only on the size of the bet in relation to the size of the pot. To ensure the best result against perfect play, the action must be unpredicatable, and one way to accomplish this is by selecting a particular range of hands to act upon, which will occur uniformly at random.

In the following example, we imagine two players involved in a hand of pot-limit Draw poker (where either may bet an amount up to the current size of the pot). Player B has called with a one-card draw to a flush, against Player A who currently has the best hand. To simplify the math, we will assume that Player B will win the showdown if she makes the flush, but will lose otherwise. We will further assume that the probability of completing the flush is exactly 0.20, or one in five. The question is how the hand should be played after the draw.

The first principle is that Player A should not bet, because Player B will simply fold if she missed on the draw, but will call (or raise) if she made the flush. Since there is no profit in Player A betting, we can assume without loss of generality that Player B is first to act after the draw. The correct strategy for Player B is to bet (the size of the pot) whenever she makes the flush, and also to bet occasionally when the draw failed. The optimal frequency of bluffs by Player B and calls by Player A are computed with a game theoretic analysis. For each pair of frequencies the overall expectation (expressed as a fraction of the total pot before the draw) can be calculated. Table 2 gives a sampling of these values over the full range of bluffing and calling frequencies.

	Player B Bluffing Ratio (BR) and Frequency (ABF)									
	0:1	1:4	**1:2**	3:4	1:1	5:4	3:2	2:1	3:1	4:1
CFr	0.00	0.05	**0.10**	0.15	0.20	0.25	0.30	0.40	0.60	0.80
0.00	.20	.25	**.30**	.35	.40	.45	.50	.60	.80	1.00
0.10	.22	.26	**.30**	.34	.38	.42	.46	.54	.70	.86
0.20	.24	.27	**.30**	.33	.36	.39	.42	.48	.60	.72
0.30	.26	.28	**.30**	.32	.34	.36	.38	.42	.50	.58
0.40	.28	.29	**.30**	.31	.32	.33	.34	.36	.40	.44
0.50	**.30**	**.30**	**.30**	**.30**	**.30**	**.30**	**.30**	**.30**	**.30**	**.30**
0.60	.32	.31	**.30**	.29	.28	.27	.26	.24	.20	.16
0.70	.34	.32	**.30**	.28	.26	.24	.22	.18	.10	.02
0.80	.36	.33	**.30**	.27	.24	.21	.18	.12	.00	-.12
0.90	.38	.34	**.30**	.26	.22	.18	.14	.06	-.10	-.26
1.00	.40	.35	**.30**	.25	.20	.15	.10	.00	-.20	-.40

Table 2: Expected Values for a Four-Flush Draw: Bluffing vs Calling Frequencies

Legend:

BR = ratio of bluffs to legitimate bets

ABF = absolute bluff frequency (fraction)

CFr = absolute calling frequency (fraction)

(expected values are expressed as a fraction of the total pot,

given a 0.2 legitimate betting frequency, and pot-sized bet)

If Player B never bluffs and Player A never calls, it has the same effect as having no betting round after the draw, and the expected value is 0.20 of the pot for B and 0.80 for A. We can see from the table that to obtain the guaranteed maximum, Player B should bet 30% of the time – 20% with the flush and an additional 10% as bluffs, selected at random. Player A can always ensure his optimal expectation of 0.70 by calling exactly 50% of the time Player B bets. The bluffing ratio of one bluff for every two legitimate bets and the calling frequency of 50% are a general results for *all* situations in which Player B will bet or bluff the size of the pot. The optimal ratios will change depending on the size of the bet in relation to the pot, but are independent of other factors.

While these are *optimal* strategies, they are not *maximal* strategies. A maximal strategy is directed toward exploiting weaknesses in the opponent, whereas an optimal strategy implicitly assumes perfect play on the part of the opponent.

The game theoretic approach is valid if the opponent is a very strong player, or perhaps an unknown player, but is certainly not the way to maximize net profit in the long run. In a typical game of poker, game theory is not an appropriate strategy, because it also guarantees that a player makes no *more* than the expected value from the particular game situation. This effectively ensures that the opponent also plays optimally, regardless of her approach to the game.

As an example of maximizing strategies, we observe how a strong poker player handles this type of situation. If faced with a bet from a player who never bluffs, a strong player will usually fold a marginal hand, knowing she cannot win. Conversely, she will often call a chronic bluffer, even with only a mediocre holding. In the role of Player B, a strong player will frequently bluff against an overly conservative player, but will seldom try to bluff a player who almost always calls. The net result is an expectation higher than the optimal 0.3, and the table demonstrates just how profitable these strategy adjustments can be in practice.

An algorithm based on game theoretic principles will provide a solid basis for betting strategy. Nevertheless, to advance to the highest levels, a program must be able to understand each opponent's playing style, and be able to adapt to the specific game conditions.

3.6 Problems for Future Research

Beyond these promising foundations there lies many difficult and interesting problems in poker research, from both a theoretical and practical point of view.

In developing a near-optimal game theoretic strategy, many pragmatic issues must be addressed that have never been properly considered in existing game theory studies. For example, the vast majority of mathematical poker models do not account for the drawing of cards and subsequent betting rounds. This has a radical effect both on theoretically correct strategy and practical considerations. Also at issue is the appropriate use of betting history in earlier rounds and previous actions in the current round.

It is clear that a maximal algorithm must observe opponent behavior and make appropriate strategy adjustments. How to best exploit any perceived weakness or predictability is a non-trivial problem. One of the keys to poker mastery is the ability to handle many different game conditions, and the strongest algorithms must have the ability to smoothly adapt to the prevailing characteristics, which may change during the course of a game.

While we believe limit Hold'em is the most natural choice for early study, new problems arise in other poker variations. For example, Sakaguchi and Sakai have proven that handling *partial information*, such as the opponent's face-up cards in the game of Seven-Card Stud, can make a game harder than having no information at all. Hi-Low forms of poker also have their own unique theoretical properties. No-limit poker may be the ultimate challenge within the domain, since it seems to emphasize the more nebulous poker skills, such as in-depth knowledge of the opponent and the ability to make fine judgements.

Finally, many other approaches could be viable for producing strong algorithms, such as genetic algorithms or neural nets. Each of these methods can be developed to be compatible with existing techniques, and the relative success of each paradigm can be settled at the virtual poker table.

References

[1] D F Beal (editor), **Advances in Computer Chess**, Edinburgh University Press, Edinburgh, Scotland, (1977-).

[2] H Kuhn and A W Tucker (editors), **Contributions to the Theory of Games**, vol.1, Princeton University Press, (1950).

[3] H Kuhn and A W Tucker (editors), **Contributions to the Theory of Games**, vol.2, Princeton University Press, (1953).

[4] H Kuhn and A W Tucker (editors), **Contributions to the Theory of Games**, vol.3, Princeton University Press, (1957).

[5] D N L Levy (editor), **Computer Games**, vol.1, Springer-Verlag, New York, (1988).

[6] D N L Levy (editor), **Computer Games**, vol.2, Springer-Verlag, New York, (1988).

[7] D N L Levy and D F Beal (editors), **Heuristic Programming in Artificial Intelligence – The First Computer Olympiad**, Ellis Horwood Limited, Chichester, England, (1989).

[8] D N L Levy and D F Beal (editors), **Heuristic Programming in Artificial Intelligence 2 – The Second Computer Olympiad**, Ellis Horwood Limited, Chichester, England, (1990).

[9] H J van den Herik and L V Allis (editors), **Heuristic Programming in Artificial Intelligence 3 – The Third Computer Olympiad**, Ellis Horwood Limited, Chichester, England, (1992).

[10] H J van den Herik (editor-in-chief), **International Computer Chess Association (ICCA) Journal**, University of Limburg, Delft, Netherlands, (1983-).

[11] T A Marsland and J Schaeffer (editors), **Computers, Chess, and Cognition**, Springer-Verlag, New York, (1990).

[12] R Levinson and S Epstein (editors) **Proceedings of the AAAI Fall Symposium on Games: Planning and Learning**, AAAI Press Technical Report FS93-02, Menlo Park CA, (1993), 59-67.

[13] J D Allen, *A Note on the Computer Solution of Connect-Four*, **Heuristic Programming in Artificial Intelligence 1**, [7], (1989), 134-135.

[14] L V Allis, **A Knowledge-based Approach of Connect-Four. The game is solved: White wins**, M.Sc. Thesis, Free University, Amsterdam, The Netherlands, (1988).

[15] L V Allis, M van der Meulen, and H J van den Herik, *Databases in Awari*, **Heuristic Programming in Artificial Intelligence 2**, [8] (1990), 73-86.

[16] L V Allis and P N A Schoo, *Qubic Solved Again*, **Heuristic Programming in Artificial Intelligence 3**, [9] (1992), 192-204.

[17] L V Allis, H J van den Herik, and M P H Huntjens *Go-Moku Solved by New Search Techniques*, **Proceedings of the 1993 AAAI Fall Symposium on Games: Planning and Learning**, AAAI Press Technical Report FS93-02, Menlo Park CA, (1993).

[18] L V Allis, **Searching for Solutions in Games and Artificial Intelligence**, PhD thesis, Vrije Universteit, Amsterdam, Netherlands, (1994).

[19] L V Allis, M van der Menlen, and H J van den Herik *Proof-Number Search*, **Artificial Intelligence 66**, no.1, (1994), 91-124.

[20] N C Ankeny, **Poker Strategy: Winning with Game Theory**, Basic Books, Inc., New York, (1981).

[21] D B Benson, *Life in the Game of Go*, **Information Sciences 10**, no.1, (1976), 17-29. (Reprinted [6], pp 203-213).

[22] E R Berlekamp, *Programs for Double-Dummy Bridge Problems – A New Strategy for Mechanical Game Playing*, **Journal of the ACM 10**, no.4, (1963), 357-364.

[23] E R Berlekamp, J H Conway, and R K Guy, **Winning Ways for your Mathematical Plays I**, Academic Press, London, England, (1982).

[24] E R Berlekamp, J H Conway, and R K Guy, **Winning Ways for your Mathematical Plays II**, Academic Press, London, England, (1982).

[25] H J Berliner, *Backgammon Computer Program Beats World Champion*, **Computer Games 1**, [5] (1980), 29-43.

[26] J R S Blair, D Mutchler, and C Lin, *Games with Imperfect Information*, **Proceedings of the AAAI Fall Symposium on Games: Planning and Learning**, AAAI Press Technical Report FS93-02, Menlo Park CA, (1993), 59-67.

[27] D Brunson, in collaboration with B Baldwin, M Caro, J Hawthorne, D Reese, and D Sklansky, **Super/System: A Course in Power Poker**, B & G Publishing Co, Inc., Las Vegas NV, (1978).

[28] J Case (*aka* W Hyde) and S Jacobs, **Percentage Hold'em: The Book of Numbers**, The Monitor, Salt Lake City UT (with Whitestone Books, San Jose CA), (1993).

[29] K Chen, *Group Identification in Computer Go*, **Heuristic Programming in Artificial Intelligence 1**, [7], (1989), 195-210.

[30] K Chen, A Kierulf, M Muller and J Nievergelt, *The Design and Evolution of Go Explorer*, **Computers, Chess, and Cognition**, [11] (1990), 271-285.

[31] K Chen, *Attack and Defense*, **Heuristic Programming in Artificial Intelligence 3**, [9] (1992), 146-156.

[32] W H Cutler, *An Optimal Strategy for Pot-Limit Poker*, **American Math Monthly 82**, (1975), 368-376.

[33] R A Epstein, **The Theory of Gambling and Statistical Logic**, Academic Press, New York (1977).

[34] N V Findler, *Programming Games*, **Summarized Proceedings of the First Conference on Automatic Computing and Data Processing**, Australia, (1961), paper B1, 3.3.

[35] N V Findler, *Computer Model of Gambling and Bluffing*, **IRE Transactions on Electronic Computers**, vol.EC-10, no.1, (1961), 5-6. (Reprinted in [6], pp 409-412).

[36] N V Findler, *Some New Approaches to Machine Learning*, **IEEE Transactions on Systems Sciences and Cybernetics**, vol.SSC-5, no.3, (1969), 173-182. (Reprinted in [6], pp 304-324).

[37] N V Findler, H Klein, W Gould, A Kowal, and J Menig *Studies on Decision Making Using the Game of Poker*, **Proceedings of IFIP Congress 1971**, (1972), 1448-1459. (Reprinted in [6], pp 413-429).

[38] N V Findler, H Klein, and Z Levine *Experiments with inductive discovery processes leading to heuristics in a poker program*, **Proceedings of Conference on Cognitive Systems**, Springer-Verlag, Berlin, (1973), 257-266.

[39] N V Findler, H Klein, R C Johnson, A Kowal, Z Levine, and J Menig, *Heuristic programmers and their gambling machines*, **Proceedings of the ACM National Conference**, San Diego CA, (1974), 28-37.

[40] N V Findler, *Studies in Machine Cognition Using the Game of Poker*, **Communications of the ACM 20**, no.4, (1977), 230-245. (Reprinted in [6], pp 430-460).

[41] N V Findler, *Computer Poker*, **Scientific American**, July (1978), 112-119.

[42] N V Findler, *Aspects of Computer Learning*, **Cybernetic Systems 11**, no.1-2 (1980), 67-86.

[43] A Frank, *Brute Force Search in Games of Imperfect Information*, **Heuristic Programming in Artificial Intelligence 2**, [8] (1990), 204-209.

[44] B Gamback, M Rayner, and B Pell, *An Architecture for a Sophisticated Mechanical Bridge Player*, **Heuristic Programming in Artificial Intelligence 2**, [8] (1990), 87-107.

[45] R Gasser, *Applying Retrograde Analysis to Nine Men's Morris*, **Heuristic Programming in Artificial Intelligence 2**, [8] (1990), 161-173.

[46] R Gasser, **Heuristic Search and Retrograde Analysis: their application to Nine Men's Morris**, Diploma thesis, Swiss Federal Institute of Technology, Zurich, Switzerland, (1990).

[47] J Gnodde, **Aida, New Search Techniques Applied to Othello**, M.Sc. Thesis, University of Leiden, Netherlands, (1993).

[48] M R Hall and D E Loeb, *Thoughts on Programming a Diplomat*, **Heuristic Programming in Artificial Intelligence 3**, [9] (1992), 123-145.

[49] F-h Hsu, T S Anantharaman, M S Cambell and A Nowatzyk, *Deep Thought*, **Computers, Chess, and Cognition**, [11], (1990), 55-78.

[50] L Jones, **Winning Low-Limit Hold'em**, ConJelCo, Pittsburgh PA, (1994).

[51] S Karlin and R Restrepo, *Multistage Poker Models*, **Contributions to the Theory of Games**, vol.3 [4], Princeton University Press, (1957), 337-364.

[52] S Karlin, *Mathematical Methods and Theory in Games*, **Programming and Economics**, vol.2, Pergamon Press, London, (1959).

[53] A Kierulf, K H Chen and J Nievergelt, *Smart Game Board and Go Explorer: A Case Study in Software and Knowledge Engineering*, **Communications of the ACM, 33**, no.2, (1990), 152-166.

[54] R Lake, J Schaeffer, and P Lu, *Solving Large Retrograde Analysis Problems Using a Network of Workstations*, **TR 93-13**, Department of Computing Science, University of Alberta, (1993).

[55] D N L Levy, *The Million Pound Bridge Problem*, **Heuristic Programming in Artificial Intelligence 1**, [7] (1989), pp 95-103.

[56] E T Lindelof, **COBRA – The Computer Designed Bidding System**, Gollancz, London, England, (1983).

[57] J M MacLeod, *MICROBRIDGE – A Computer Developed Approach to Bidding*, **Heuristic Programming in Artificial Intelligence 1**, [7] (1989), 81-87.

[58] M Malmuth, **Gambling Theory and Other Topics**, Two Plus Two Publishing, Las Vegas NV, (1987) (1994 4th ed).

[59] M Malmuth, **Winning Concepts in Draw and Lowball**, Two Plus Two Publishing, Las Vegas NV.

[60] M Malmuth, **Poker Essays**, Two Plus Two Publishing, Las Vegas NV.

[61] J F Nash and L S Shapley, *A Simple Three-Person Poker Game*, **Contributions to the Theory of Games**, vol.1 [2], Princeton University Press, (1950), 105-116.

[62] D J Newman, *A model for "real" poker*, **Oper. Res. 7**, (1959), 557-560.

[63] O Patashnik, *Qubic: 4x4x4 Tic-Tac-Toe*, **Mathematics Magazine 53**, (1980), 202-216.

[64] B Pell, *Exploratory Learning in the Game of GO: Initial Results*, **Heuristic Programming in Artificial Intelligence 2**, [8], (1990), 137-152.

[65] R Popma and L V Allis, *Life and Death Refined*, **Heuristic Programming in Artificial Intelligence 3**, [9] (1992), 157-164.

[66] P S Rosenbloom, *A World-Championship-Level Othello Program*, **Computer Games 2**, [6] (1979), 365-408.

[67] J L Ryder, **Heuristic analysis of large trees as generated in the game of Go**, Ph.D. Thesis, Stanford University, (1971).

[68] M Sakaguchi and S Sakai, *Partial Information in a Simplified Two-Person Poker*, **Math. Japonica 26**, no.6, (1981), 695-705.

[69] M Sakaguchi, *A Simplified Two-Person Multistage Poker with Optional Stopping*, **Math. Japonica 28**, no.3 (1983), 287-303.

[70] M Sakaguchi, *A Note on the Disadvantage for the Sente Player in Poker*, **Math. Japonica 29**, no.3 (1984), 483-489.

[71] M Sakaguchi, *On Two-Person "Real" Poker by Newman*, **Math. Japonica 30**, no.3 (1985), 471-483.

[72] M Sakaguchi, *Information Contained in Face-up Cards in Poker*, **Math. Japonica 30**, no.4 (1985), 535-549.

[73] M Sakaguchi, *The Value of Sample Information in High-Hand-Wins Poker*, **Math. Japonica 33**, (1988), 587-607.

[74] M Sakaguchi, *The Value of Sample Information in La Relance Poker*, **Math. Japonica 33**, (1988), 777-800.

[75] M Sakaguchi and S Sakai, *Multistage Poker with Random Amount of Bets*, **Math. Japonica 37**, no.5, (1992), 827-838.

[76] M Sakaguchi and S Sakai, *Solutions of Some Three-Person Stud and Draw Poker*, **Math. Japonica 37**, no.6, (1992), 1147-1160.

[77] S Sakai, *A model for real poker with an upper bound of assets*, **Journal Optim. Theory Appl. 50**, no.1, (1986), 149-163.

[78] J Schaeffer, J Culbertson, N Treloar, B Knight, P Lu, and D Szafron, *Reviving the Game of Checkers*, **Heuristic Programming in Artificial Intelligence 2**, [8] (1990), 119-136. (Also available as a University of Alberta Computer Science department technical report).

[79] J Schaeffer, J Culbertson, N Treloar, B Knight, P Lu, and D Szafron, *A World Championship Caliber Checkers Program*, **Artificial Intelligence 53**, no.2-3, (1992), 273-290.

[80] J Schaeffer, *Man Versus Machine: The Silicon Graphics World Checkers Championship*, **TR 92-19**, Department of Computing Science, University of Alberta, (1992).

[81] J Schaeffer, N Treloar, P Lu, and R Lake, *Man Versus Machine for the World Checkers Championship*, **AI Magazine 14**, no.2, (1993), 28-35. (An abbreviated version of this article also appeared in the ICCA Journal [10], vol.16, no.2, 105-110).

[82] J Schaeffer, *A Re-examination of Brute-Force Search*, **Proceedings of the AAAI Fall Symposium on Games: Planning and Learning**, AAAI Press Technical Report FS93-02, Menlo Park CA, (1993), 51-58.

[83] P N A Schoo, *Optimal Play in a Single Bridge Suit*, in preparation.

[84] S C Shapiro and H R Smith, *A Scrabble Crossword Game-Playing Program*, **Computer Games 1**, (1977), 403-422.

[85] K Shirayanagi, *Knowledge Representation and its Refinement in Go Programs*, **Computers, Chess, and Cognition**, [11] (1989), 287-300.

[86] D Sklansky, **Hold'em Poker**, Gambler's Book Club, Las Vegas NV, (1976).

[87] D Sklansky, **The Theory of Poker**, (formerly titled "Winning Poker"), Two Plus Two Publishing, Las Vegas NV, (1987) (1992 3rd ed).

[88] D Sklansky, **Sklansky on Poker**, Two Plus Two Publishing, Las Vegas NV.

[89] D Sklansky, **Getting the Best of It**, Two Plus Two Publishing, Las Vegas NV.

[90] D Sklansky and M Malmuth, **Hold'em Poker for Advanced Players**, Two Plus Two Publishing, Las Vegas NV, (1988) (1994 2nd ed).

[91] D Sklansky, M Malmuth and R Zee, **Seven-Card Stud for Advanced Players**, Two Plus Two Publishing, Las Vegas NV.

[92] L Stiller, *Parallel Analysis of Certain Endgames*, **ICCA Journal 12**, [10] no.2, (1989), 55-64.

[93] G Tesauro, *NEUROGAMMON: A Neural-Network Backgammon Learning Program*, **Heuristic Programming in Artificial Intelligence 1**, [7] (1989), 78-80.

[94] G Tesauro, *TD-Gammon, A Self-Teaching Backgammon Program, Achieves Master-Level Play*, **Proceedings of the AAAI Fall Symposium on Games: Planning and Learning**, AAAI Press Technical Report FS93-02, Menlo Park CA, (1993), 19-23.

[95] K Thompson, *Retrograde Analysis of Certain Endgames*, **ICCA Journal 9**, [10] no.3, (1986), 131-139.

[96] E O Thorp, and W E Walden, *A Computer-Assisted Study of Go on M x N Boards*, **Computer Games 2**, [6] (1972), 152-181.

[97] T Throop and T Guilfoyle, *A Thrilling Hand*, **Heuristic Programming in Artificial Intelligence 3**, [9] (1992), 27-28.

[98] J W H M Uiterwyk, H J van den Herik, and V Allis, *Knowledge-Based Approach to Connect Four. The Game is Over: White to Move Wins!*, **Heuristic Programming in Artificial Intelligence 1**, [7] (1989), 113-133.

[99] J W H M Uiterwijk, *Knowledge and Strategies in Go-Moku*, **Heuristic Programming in Artificial Intelligence 3**, [9] (1992), 165-179.

[100] I H Ulgee, *Letters beyond Numbers*, **Heuristic Programming in Artificial Intelligence 3**, [9] (1992), 63-66.

[101] J von Neumann and O Morgenstern, **Theory of Games and Economic Behavior**, Princeton Univ. Press, Princeton, (1944).

[102] R Wheen, *Brute Force Programming for Solving Double Dummy Bridge Problems*, **Heuristic Programming in Artificial Intelligence 1**, [7] (1989), 88-94.

[103] B Wilcox, *Reflections on Building Two Go Programs*, **SIGART News 94**, (1985), 29-43.

[104] N Zadeh, PhD thesis, University of California at Berkeley, (1972).

[105] N Zadeh, **Winning Poker Systems**, Prentice Hall, Inc., Englewood Cliffs, NJ, (1974).

[106] N Zadeh, *Computation of optimal poker strategies*, **Operations Res. 25**, no.4, (1977), 541-562.

[107] R Zee, **High-Low-Split Poker, Seven-Card Stud and Omaha Eight-or-Better for Advanced Players**, Two Plus Two Publishing, Las Vegas NV.

[108] A L Zobrist, **Feature extraction and representation for pattern recognition and the game of Go**, Ph.D. Thesis, U of Wisconsin, (1970).

[109] A Andrews, *Seven-Card-Stud Poker for Windows*, rec.gambling.poker Berkely Archive Site [126], 7crdstd2.zip. (More recent versions also available).

[110] S Brecher, *IRC .ircrc Configuration File*, rec.gambling.poker Berkely Archive Site [126], ircrc.poker.

[111] M Maurer, *Simulation Data for Various Poker Games*, available from the author [117].

[112] R T Hashimoto, *A Poker Hand Distribution Generator: fish.c*, rec.gambling.poker Berkely Archive Site [126], fish.c.

[113] C T Matthews, *Rewrite of Hashimoto's fish.c*, rec.gambling.poker Berkely Archive Site [126], poker.tar.gz, or anonymous FTP site: "ftp.cs.unm.edu/pub/poker/poker.tar.gz".

[114] C T Matthews, K Miyake, M Hostetter, *Very Fast Poker Hand Comparators*, rec.gambling.poker Berkely Archive Site [126], poker.tar.gz, or anonymous FTP site: "ftp.cs.unm.edu/pub/poker/poker.tar.gz".

[115] *Gambler's Book Club catalogue*, Gambler's Book Club, Las Vegas NV. Better program titles include: "Workware Master Series Hold'em", "Wilson Software Turbo Texas Hold'em", "Szobodan Seven-Card Stud", and "The World Series of Poker Adventure".

[116] M Maurer, *IRC Poker Database*, WWW Address: "http://www-star.stanford.edu/ maurer/r.g/".

[117] M Maurer, *WWW Home Page*, WWW Address: "http://www-star.stanford.edu/ maurer/".

[118] M Maurer, *Bankroll Variance as a Function of Game Size*, WWW Address: "http://www-star.stanford.edu/ maurer/r.g/pstats4.html".

[119] M Maurer, *Profit Rate as a Function of Position and Game Size*, WWW Address: "http://www-star.stanford.edu/ maurer/r.g/pstats5.html".

[120] M Maurer, *Player Strategy Statistics: Abridged Set*, WWW Address: "http://www-star.stanford.edu/ maurer/r.g/psort2.html".

[121] M Maurer, *IRC Observer Bot Source Code*, WWW Address: "http://www-star.stanford.edu/ maurer/r.g/pobot" (etc).

[122] M Maurer, *IRC Bot Poker Player*, WWW Address: "http://www-star.stanford.edu/ maurer/r.g/pmaubot".

[123] M Maurer, *A Study of Variance Using the IRC Poker Database*, WWW Address: "http://www.conjelco.com/r_g/las_vegas/variance_3-6.html".

[124] T Mummert, *An IRC Poker Dealing Program*, IRC Server location: "vegas.scandal.cs.cmu.edu 6667", WWW Address: "http://www.cs.cmu.edu/ mummert/ircbot.html".

[125] H Ruegg, *7-Card Stud and Texas Hold'em Simulations*, rec.gambling.poker Berkely Archive Site [126], spoker.tar.Z

[126] Usenet newsgroup rec.gambling.poker, *Berkeley Archive Site*, WWW Address: "ftp://ftp.csua.berkeley.edu/pub/rec.gambling/poker" (note: maximum of 15 anonymous users).

[127] Usenet newsgroup rec.gambling.poker, *UIUC Archive Site*, WWW Address: "gopher://gopher.uiuc.edu/1ftp

[128] Usenet newsgroup rec.gambling.poker, *Frequently Asked Questions (FAQ)*, WWW Address: "http://www.conjelco.com/faq/poker.html".

[129] Usenet newsgroup rec.gambling.poker, *WWW Home Page*, WWW Address: "http://www.conjelco.com/r.g.html".