

Chinook: The World Man-Machine Checkers Champion

Jonathan Schaeffer, Robert Lake, Paul Lu and Martin Bryant

ABSTRACT

In 1992, the seemingly unbeatable World Checker Champion, Dr. Marion Tinsley, defended his title against the computer program Chinook. After an intense, tightly-contested match, Tinsley fought back from behind to win the match by scoring four wins to Chinook's two, with 33 draws. This was the first time in history that a human World Champion defended his title against a computer. This article reports on the progress of the checkers (8×8 draughts) program Chinook since 1992. Two years of research and development on the program culminated in a rematch with Tinsley in August 1994. In that match, after six games (all draws), Tinsley resigned the match and the World Championship title to Chinook citing health concerns. Chinook has since defended its title in two subsequent matches. This is the first time in history that a computer has won a human World Championship.

Keywords: game-playing, alpha-beta, heuristic search, checkers.

1. Background

The most well-known checkers-playing program was Samuel's effort in the period around 1955-1965 (Samuel 1967; Samuel 1959). This work remains a milestone in artificial intelligence research. Samuel's program reportedly beat a master and "solved" the game of checkers. Both journalistic claims are false, but they served to create the impression that there was nothing of scientific interest left in the game (Samuel himself made no such claims). Consequently, most subsequent games-related research turned to chess. Other than a program from Duke University in the 1970's (Truscott 1979), little attention was paid to achieving a World Championship caliber checkers program.

The Chinook project started in 1989 as an attempt to better understand heuristic search by using an experimental domain that is simpler than chess. Early success in exhibition games against master and grandmaster opposition convinced us that we could build a world-class checker player. After petitioning the American Checker Federation (A.C.F.), Chinook was allowed to play in the 1990 U.S. Championship. This biennial event attracts the best players in the world, with the winner earning the right to play a match for the World Championship.

Chinook came an undefeated second in the tournament behind the World Champion, Dr. Marion Tinsley. The four individual games between Chinook and Tinsley were drawn. This placing earned Chinook the right to challenge Tinsley for the World Championship, the first time a computer had earned such a right in any game.

Tinsley emerged as a dominant player in the late 1940's. He was touted as the next

challenger for the World Championship when, in 1950, an elementary blunder cost him the chance to play for the title. After quietly retrenching, Tinsley re-emerged as an unbeatable force, winning the World Championship in 1955. Over the period 1950-1992, Tinsley lost an incredible total of only five games! He played eight matches for the World Championship, winning each usually by a large margin. He was as close to perfection as was possible in a human. No one has dominated any game of skill for such a prolonged period of time as Tinsley has.

In December 1990, Tinsley visited Edmonton to play an exhibition match against Chinook. The final score was one win for Tinsley, none for Chinook and 13 draws. This was the closest anyone had come to defeating Tinsley in a match in over 40 years.

The A.C.F. and E.D.A. (English Draughts Association) refused to sanction the Chinook-Tinsley World Championship match, stating they did not want a computer playing for a human title. Tinsley confided in us that he was bored playing humans because they were no opposition to him. He relished the opportunity to play the computer because of its aggressive style and willingness to take chances. When the A.C.F. and E.D.A. refused to budge on this issue, Tinsley resigned his World Championship title and then immediately signed an agreement to play the Chinook match. The A.C.F. and E.D.A. hastily re-thought their position and created a new "Man-Machine" World Championship: the best human against the best computer. Since Tinsley was the undisputed human champion and Chinook had properly earned the right to challenge, the new Man-Machine title was a convenient way of keeping the checker federations happy.

In August 1992, Chinook and Tinsley finally did battle (Schaeffer et al. 1993a). The match was held in London, England, and was sponsored by Silicon Graphics, Inc. After initially falling behind, Tinsley rebounded and emerged victorious, scoring four wins to Chinook's two with 33 draws. The two Chinook wins represented only the sixth and seventh times that Tinsley had lost in the previous 42 years.

Although Tinsley had a close call in London, he was eager to play the computer again. August 1994 was set as the date for a rematch.

2. Chinook 1994

Within a few weeks of the end of the 1992 Chinook-Tinsley match, preparations began for the rematch. Chinook's performance consists of four aspects (Schaeffer et al. 1993b; Schaeffer et al. 1992):

- (a) Search - algorithms for traversing through the $O(10^{20})$ search space.
- (b) Evaluation function - deciding how good a position is.
- (c) Endgame databases - perfect information on which endgame positions are won, lost and drawn.
- (d) Opening book - a database of opening lines to start a game.

Compared to the program that played in 1992, important advances have been made to the program in each of these areas.

2.1. Search

Chinook uses a parallel alpha-beta search algorithm with all the latest enhancements (Lu, 1993). Chinook 1992 had an average minimum search depth of 19 ply excluding search extensions (one ply equals one move by one player). Chinook 1994 ran on a 16 processor Silicon Graphics Challenge computer (150 MHz processors) with one gigabyte of RAM. The machine had twice as many processors, each five times faster than in 1992, and four times as much RAM. The improved hardware allowed the program to search a minimum of an additional two ply deeper.

As has been seen with chess programs, deeper search in checkers translates into stronger play. Some studies in chess have indicated that there might be a linear relationship between search depth and performance (Thompson, 1982). Experiments in Chinook show that there comes a point where increased search depth provides diminishing returns (Schaeffer et al. 1993b). The data suggests that Chinook has already reached the point where little is to be gained by deeper searching; improved performance must come through knowledge enhancements instead of machine speed.

Extensive work was done on improving the quality of the search tree built by Chinook. We have seen positions that Tinsley can correctly analyze, even though the solution requires over 60 ply of search! Fixed-depth alpha-beta search is woefully inadequate here. As in chess, the alpha-beta framework has been modified to selectively extend interesting lines of play (search extensions) and to curtail effort on clearly inferior lines of play (forward pruning). A variety of static (heuristic) and dynamic (based on properties of the search tree) techniques were employed to create a more "intelligent" searcher.

One of the major problems in chess and checker programs is deciding when it is appropriate to stop analyzing a line of play and apply the evaluation function. Ideally, the evaluated position should be quiescent; in other words it should be relatively stable with no surprises waiting around the corner. Instrumenting the search algorithm allowed us to identify positions where a heuristic evaluation was likely to make a gross misassessment due to hidden tactics. A program was created to generalize the features of these positions, to search for any commonalities. The results were quite surprising: one type of tactical combination accounted for roughly 50% of the errors. A further two types accounted for an additional 25% of the cases. These latter two features were easily accounted for by adding simple patterns to the evaluation function.

The dominant source of error involved a so-called "two-for-one" combination. Here the key move would result in a checker being captured, but a double jump ensued for a net gain of one checker. Hard-coding the knowledge required to capture one instance of this pattern was easy to do. However, generalizing that knowledge to cover all the possible scenarios and exceptions proved to be extremely difficult and error prone. To properly identify all the conditions necessary for the combination involved consideration of 28 of the 32 squares on the board. The alternative to using a static set of static rules was to perform dynamic search. Given that the majority of evaluated positions do not have the two-for-one pattern present, the cost of doing the (expensive) search was prohibitive.

All the information about these patterns was captured in what we call *tactical tables*. The essential information required in the position to determine whether the two-for-one is present or not were abstracted into 8,000,000 possibilities. For each scenario, an abstract board position was created containing constraints on each of the squares (such

as a white piece must be on this square, or another square must be occupied). A search was performed on this abstract board position to see if the two-for-one succeeded. For each possible abstract board, the result of the search, the two-for-one works or not, was added to the tactical table. Since each possibility required only one bit, the entire table required only one megabyte of additional storage.

When a position is evaluated, the tactical table is consulted. The essential features of the position are mapped into an index that is used to access the appropriate bit in the tactical table. If the bit isn't set, then the evaluation proceeds normally. If the bit is set, then we know that the side to move is winning a checker. The program then extends the search to resolve the tactics. It can't simply conclude that the checker is being safely won. It's possible that having played the sequence of moves to win the checker leads to a position where the other side has sufficient positional compensation for the lost material, or has, in turn, its own two-for-one regaining the material.

The addition of the tactical table and the additional patterns reduced the occurrences of major errors in the evaluation function by roughly 75%. This significantly improved the stability of the program.

The cumulative effect of deep selective search is overwhelming for most human players: Chinook does not make tactical mistakes.

2.2. Evaluation Function

The knowledge in Chinook's evaluation function has remained largely unchanged since 1992. In the two years subsequent to the London 1992 match, all the individual knowledge components were thoroughly tested for utility and completeness. To our surprise, this uncovered a number of (glaring) errors in the code and some important gaps in the program's knowledge. Some of these errors could have been fatal in a game. One of the best/worst properties of the alpha-beta algorithm is that its max/min behavior hides evaluation function errors. Evaluation function bugs can remain hidden for years. Often errors are only detected when a bad score is propagated to the root of the search tree.

The evaluation function remains a linear combination of roughly two dozen major components, each of which contains several heuristic weights. Attempts have been made to tune the weights of the evaluation function through automated processes, such as using linear equations, neural nets and genetic algorithms. Although these traditional artificial intelligence approaches appear promising in the literature, in our experience they cannot compete with the results of hand-tuning (although in other domains, such as Othello (Lee and Mahajan 1990) and backgammon (Tesauro 1995), some of these techniques have been effective). We wish never to have to hand-tune an evaluation function again!

2.3. Endgame Databases

Chinook 1992 had access to a database containing the value (win, loss, draw) for all positions with seven or fewer pieces on the board, and a small percentage of the 8-piece databases. The databases contained roughly 40 billion positions, compressed into two gigabytes of data. The entire database was available to the program during a game, and was compressed in such a way as to allow rapid decompression at runtime (Lake et al. 1994).

After the 1992 match, errors were discovered in some of the endgame databases. In

retrospect, Chinook was fortunate not to lose game one of the 1992 match, as that result hinged on the erroneous values.

After the match, the databases were repaired and work continued with the goal of building the complete 8-piece databases. The 4×4 subset of the 8-piece database (all positions with four pieces against four pieces, any combination of kings and checkers) was completed in 1994, adding an additional 111 billion positions. For the 1994 matches the 1-7-piece databases and the 4×4 subset of the 8-piece databases, 148 billion positions, were used. Subsequently, the rest of the 8-piece databases have been computed (440 billion positions compressed into 6 gigabytes of disk) and the 9-piece databases are under way.

One of the goals of this project is to eventually construct a perfect checkers player. Although building a database of 440 billion positions may seem to be a primitive approach, currently there is not really much of an alternative. Attempts to use standard artificial intelligence approaches (genetic algorithms (van Belle 1995), neural nets and function optimization) have produced solutions with unacceptably high error rates. On the other hand, brute-force enumeration can have no error. As well, enumeration of the data (with no gaps) allows for implicit addressing which significantly benefits the data compression. For our application, artificial intelligence solutions have more error, poorer data compression and are more expensive to access at runtime.

The knowledge implicit in the 4×4 database greatly exceeds even Tinsley's capabilities. For example, Grandmaster Don Lafferty played into an ending he thought was drawn, only to discover to his surprise that Chinook knew it was a win. Our databases have overturned a number of long-held misconceptions about some standard positions in the checkers literature. In the 1994 U.S. Championship, Chinook's endgame database was used to adjudicate the game between Tinsley and Grandmaster Elbert Lowder. Tinsley was trying for the win; the databases said it was a draw. This was the first time in any game that the word of a computer has been used to decide the outcome of a World Champion's game.

Chinook's deep searches frequently hit positions in the endgame databases. Usually, within a few moves of the start of a game, the program becomes I/O bound on accessing the databases, rather than CPU bound as most game-playing programs are. In most games, Chinook can determine the final result of the game within the first 10 moves; usually a draw. When the program realizes the game is drawn, the problem now becomes one of how to best proceed. All draws are not equal. The program has to choose between drawing moves to maximize the chances of human error.

Chinook was enhanced to differentiate between "strong" draws, ones where there is a chance of inducing the opponent into an error, and "weak" draws, ones where the program has little chances of winning. Instead of returning a score of 0 (equality) for a database draw position, a restricted depth search (with the databases disabled) is performed to get an idea of the non-database minimax score. The resulting score is scaled to allow draws to return scores in the range of +1/3 to -1/10 of a checker. Thus if the program has a choice of playing a "strong" draw versus playing a non-drawing move that yielded a small advantage, the program prefers the former. Most game-playing programs would choose the latter.

2.4. Opening Book

The last 125 years has seen the publication of a large body of literature on checkers. Much of it has to do with the analysis of the opening. Human players study this literature so that they know the best sequences of opening moves. One of the most important parts of grandmaster preparation is to find surprises, or "cooks", in the openings. These are new moves that force the opponent on to their own resources, usually in a complicated position, thus increasing their chances of error.

The biggest weakness in Chinook 1992's play was in the opening. Opening mistakes were responsible for three of the losses to Tinsley (the fourth loss was a forfeit in a drawn position). To solve this problem, we entered into an agreement with Martin Bryant, author of the commercially available checkers program Colossus. Over the years, he built up a large database (or book) of opening play that was acquired largely from the published literature. To reduce the probability of an error (all too common in the literature), some of the positions were checked for correctness by his program (using a deep search).

Once we obtained the Colossus database, Chinook spent several months verifying each position. The searches were to a minimum of 19 ply deep and used all the endgame databases. The result was that several hundred positions in the book were corrected (usually the result of incorrect human assessments of endgame positions), most of them being positions where one side missed a saving draw.

The verification data was filtered to look for "new" moves in the openings. Given Tinsley's phenomenal memory, one is unlikely to beat him by repeating the "best" moves in the literature. The only way to win against Tinsley is to surprise him. We needed to find unusual moves that were on main lines of play (so that we had a reasonable chance of getting him to play that line) that were either new or offbeat. This meant searching for new moves, unusual second or third best moves, and old moves that have faded from popularity.

Chinook's analysis of the openings was only a small part of the effort. The program flagged over two thousand potential cooks. Each one of these cooks had to be assessed by hand based on their likelihood of appearing in a game and the potential benefits. Less than 100 were flagged for additional investigation.

There are 144 openings in checkers (the first 3 moves are randomly selected at the start of a game). In 8 openings, we found so-called "gold" moves - moves that if we had a chance to play them against Tinsley there was a good chance Chinook would win the game. We had roughly 30 "silvers" - new moves that we knew eventually led to a draw, but may surprise Tinsley and might induce a mistake.

In summary, compared to Chinook 1992, Chinook 1994 searched better and deeper, evaluated positions better, had access to more and better quality endgame databases, and had access to 12 times as much (and better quality) opening knowledge. Given the close result from 1992 and all the advances to Chinook in the interim, how could we miss this time around?

3. Pre-match Preparation

Preparing for a World Championship match against a human is not only a programming exercise. Psychology is an integral part of pre-match preparation whether the match be man versus man or man versus machine. Several things were done to keep Chinook's true strength hidden from Tinsley.

A test match was played in April 1994 against the official World Champion Derek Oldbury (Oldbury took over the title when Tinsley resigned it in 1991). The final score was three wins for Chinook, no losses and nine draws. Interestingly, in all three wins Chinook announced the game as drawn. However, Chinook successfully made the draws as difficult to achieve as possible and eventually Oldbury stumbled. It was widely known that the match had been played. However, by agreement with Oldbury, the final score and the games were kept secret until after the Tinsley match. Several people, including Tinsley, tried to uncover this information, but to no avail. All Tinsley knew was that we won.

In July, an 18-game match with Don Lafferty ended in 18 draws. Tinsley came to Edmonton to observe the Chinook-Lafferty match, hoping to get some inside information on the program. In our opinion, Lafferty is the second strongest human player. As Tinsley's *protégé*, he has access to most of Tinsley's opening secrets. Lafferty is a cautious player who seldom takes risks. In tournaments, he rarely loses a game, but usually wins enough to capture a top placing. In contrast, a player like Grandmaster Ron King loses several games a tournament, but his aggressive style ensures many wins and a frequent first-place finish.

At the end of July, Chinook played in the Southern States Championship in Dallas Texas. The program did not lose a game and finished in clear first ahead of several grandmasters including Ron King (Lafferty and Tinsley did not play). The following week, the U.S. Championship was held at the same site. As usual, almost all the best players in the world were participating.

Round one produced a surprise, as Tinsley was paired against Chinook. All four games played were drawn, although Tinsley put Chinook under a lot of pressure in game three. Postmortem analysis revealed a serious bug in Chinook's evaluation function that was quickly repaired. In the fourth round, Chinook drew its match with Lafferty. Going into the last round, Chinook was tied for first with Lafferty, with Tinsley a half point behind. In the final round, Chinook again drew with Lafferty allowing Tinsley to climb into a three way tie for first place with a last round victory. This was the first time in over 40 years that Tinsley hadn't finished alone in first place.

In all the games played leading up to the World Championship match, Chinook's new opening book was concealed. The program was set up to "think" on every move and not play instantly, as most programs do when still in the opening. In effect, no one knew in advance how extensive our opening preparation was. As well, all our opening cooks were disabled, ensuring that we did not accidentally reveal some of our prepared surprises.

In 1992, we switched evaluation functions on Tinsley. All the published Chinook games had been played with an old evaluation function. For the match we used one that Tinsley had never seen before, with different strengths and weaknesses than Tinsley had prepared for. For 1994, we did something similar. Before the start of the match, several

parameters were altered to give the program a slightly different style of play. Whether the differences were discernible to Tinsley is not known.

Prior to the World Championship match, Tinsley expressed confidence that he had figured out Chinook's weaknesses and that he could exploit them. He did not know that he would be playing an adversary with a different evaluation function, large opening book and faster computer. These changes effectively meant that most of Tinsley's pre-match study of Chinook was probably in vain.

4. The 1994 Man-Machine World Championship

The match was held at the Computer Museum in Boston beginning August 15, 1994. Silicon Graphics again sponsored the match. David Levy and Raymond Keene organized and directed the event.

The first six games of the match were drawn. Only in game two did anyone have any winning chances. Tinsley made a mistake early and had to struggle hard to save the game. Afterwards, he confided that everywhere he looked he found only losing moves. One move, however, had a ray of hope and Tinsley "placed his faith in God" and hoped for a miracle; he had found the only path to safety.

One minute before the start of game seven, Tinsley revealed that he had not slept well the night before and wanted to see a doctor. Play was canceled for that day, and Tinsley was taken to the hospital for tests. Nothing untoward was revealed and Tinsley indicated he would play the next day. The following morning, Thursday August 18, Tinsley unexpectedly resigned the match and the World Championship title citing health concerns. Attempts to convince him otherwise, or to postpone the match, failed. On the Friday he was admitted to hospital for tests and the following Tuesday he was diagnosed with cancer.

In accordance with the rules, Chinook was declared the Man-Machine World Champion. Since we were in Boston, and everything was set up to play a match, the organizers hastily flew in Don Lafferty to challenge for Chinook's title. The match ended with one win apiece and 18 draws. He scored in game eight when a bug caused a large misassessment of a position[†]. Chinook wore Lafferty's defenses down in game 10. With the exception of game two where Chinook came close to winning, the rest of the games were uneventful draws.

Although Chinook was World Champion, the way we achieved it did not rest well with us. We won the title on forfeit over Tinsley, and retained the title by drawing the match with Lafferty. Fortunately, Lafferty was eager for a rematch, which was held in Petal, Mississippi, in January, 1995. The final score was one win to none in favor of Chinook, with 31 games drawn. We had several opportunities for modifying the program to increase its chances of winning a game, however we felt it was more important not to lose than it was to win.

In January 1995 Tinsley said that he felt good and was interested in playing a

[†] The bug was uncovered two months later. Warren Smith (NEC) was using the Chinook code for some of his research and found an error in the program's checkers knowledge. An "or" condition was inadvertently expressed as an "and". The bug had been introduced when fixing the program after the Chinook-Tinsley U.S. Championship match.

rematch with Chinook that summer "to regain my title." Within a week of that conversation, he suffered a relapse from which he never recovered. Tinsley died on April 3, 1995. With his passing, we lost not only a feared adversary but also a friend. Every member of our team had the deepest respect and admiration for Tinsley. It was a privilege to know him.

5. Conclusions

The Chinook team's success is due to Marion Tinsley. When faced with the challenge of a computer, Tinsley could have said "no" and had the backing of the checker federations. Instead, Tinsley was only interested in playing the best, whether it be man or machine. Our project might have died in 1990 because of a lack of human competition. With Tinsley's support, Chinook was able to compete in the major checker tournaments and play matches against the top players in the world. We can only hope that his example sets a precedent for other games, such as chess where, so far, computers are banned from officially sanctioned events.

The unfortunate events of Boston leave unresolved the question as to whether Chinook 1994 was better than Tinsley. The lead-up to the 1992 and 1994 matches gives an indication of how Chinook had grown in strength. In 1994/1995, Chinook played 152 games against master/grandmaster opposition and 44 tournament games against commercially-available programs (Colossus, Checkers and Sage Draughts, all strong programs). In this run, Chinook lost only one game and was in serious trouble in one other. Both problems were directly attributable to programming errors. In contrast, in 1992 Chinook lost seven of 99 tournament/match games. Of these games, Chinook was outplayed three times, lost two due to programming errors, lost one on forfeit and lost one because of our underestimation of the importance of a piece of checkers knowledge.

Is Chinook the best player in the world? With the victory over Lafferty, both the A.C.F. and E.D.A. recognize Chinook's title. The next human-only World Championship will be played between Don Lafferty and Ron King. Chinook's lifetime record against King is seven wins, no losses and nine draws. Lafferty too has a winning record against King. Whether King or Lafferty wins, the record suggests that Chinook is the better player. As Chinook continues to improve on a daily basis, and with no new Marion Tinsleys on the horizon, Chinook's dominance can only grow.

Checkers has a rating system comparable to that used in chess. The latest ratings of the American Checker Federation lists the top four players in the world:

Chinook	2712
Ron King	2632
Asa Long	2631
Don Lafferty	2625

The rating system predicts that an 80-point rating differential translates into a 60% winning probability.

The only goal left to achieve is solving the game of checkers: we want to build a perfect checkers player. Towards this goal, endgame database construction continues. At the time of this writing all the 8-piece databases are complete (440 billion positions)

and the 9-piece computation has begun. Every additional database we compute simplifies the task of solving the game. Because of a lack of time, we have not yet begun to research the algorithms needed to efficiently search a space of size $O(10^{20})$.

Over the past 6 years, we have been consumed with the quest of defeating the unbeatable Tinsley. It was Tinsley's fervent desire to never lose a match to Chinook. It is ironic that in the end, both parties got their wish, but not in the way they wanted it. Chinook became World Champion by defeating Tinsley by forfeit. Tinsley's record remains intact; by resigning in the middle of an even match, his record shows that he never lost a match to Chinook over the board.

Chinook represents a large artificial intelligence software engineering project. Given the enormous amount of time and resources we have expended for achieving dominance over humans in a "simple" problem domain such as checkers, one wonders how much knowledge engineering is required for harder, more mainstream artificial intelligence problems.

Come visit our Web site and play a game against Chinook:
<http://web.cs.ualberta.ca/~chinook>.

Postscript

In May 1995, Ron King formally challenged Chinook for the World Man-Machine title. Since then, there have been several postponements on King's end. This match will not occur until after the King-Lafferty match, possibly in March 1996.

Acknowledgments

Many people have contributed to the Chinook program including Norman Treloar, Brent Knight, Duane Szafron, Joe Culberson and Steve Sutphen.

We are grateful to many people in Silicon Graphics, including Bob Bishop, Benoit Marchand, Ashley Mooser and Joe Gilberto. David Levy and Raymond Keene organized the Boston event. Charles Walker sponsored the 1995 Lafferty match. Ken Thompson kindly provided us access to computing equipment for the 1995 match.

This work was sponsored in part by a grant from the Natural Sciences and Engineering Research Council of Canada. The Netherlands Organization for Scientific Research provided financial assistance for Jonathan Schaeffer to work for 6 months in 1994 at the Department of Computer Science, University of Limburg, Maastricht. The help (and friendship) of Jaap van den Herik is greatly appreciated.

References

Lake, R.; Schaeffer, J.; and Lu, P. 1994. Solving Large Retrograde Analysis Problems Using a Network of Workstations. In *Advances in Computer Chess VII*, eds. H.J. van den Herik, I.S. Herschberg and J.W.H.M. Uiterwijk, 135-162. Maastricht, Netherlands: University of Limburg.

Lee, K.F.; and Mahajan, S. 1990. The Development of a World Class Othello Program. *Artificial Intelligence* 43(1):21-36.

Lu, P. 1993; Parallel Search of Narrow Game Trees. M.Sc. thesis, Dept. of Computer

Science, Univ. of Alberta.

Samuel, A.L. 1959. Some Studies in Machine Learning Using the Game of Checkers. *IBM Journal of Research and Development* 3(3):210-229. See also 1963, *Computers and Thought*, eds. E.A. Feigenbaum and J. Feldman, McGraw-Hill, 71-105.

Samuel, A.L. 1967. Some Studies in Machine Learning Using the Game of Checkers II - Recent Progress. *IBM Journal of Research and Development* 11(6):601-617.

Schaeffer, J.; Culberson, J.; Treloar, N.; Knight, B.; Lu, P.; and Szafron, D. 1992. A World Championship Caliber Checkers Program. *Artificial Intelligence* 53(2-3):273-290.

Schaeffer, J.; Treloar, N.; Lu, P.; and Lake, R. 1993a. Man Versus Machine for the World Checkers Championship. *AI Magazine* 14(2):28-35.

Schaeffer, J.; Lu, P.; Szafron, D.; and Lake, R. 1993b. A Re-examination of Brute-force Search. In proceedings of the AAAI Fall Symposium, Games: Planning and Learning (Report FS9302), 51-58. Menlo Park, Calif: AAAI.

Tesauro, G. 1995. Temporal Difference Learning and TD-Gammon. *Communications of the ACM* 38(3):58-68.

Thompson, K. 1982. Computer Chess Strength. In *Advances in Computer Chess 3*, ed. M.R.B. Clarke, 55-56. Oxford: Pergamon Press.

Truscott, T. 1979. The Duke Checker Program. *Journal of Recreational Mathematics* 12(4):241-247.

van Belle, T. 1995. A New Approach to Genetic-Based Automatic Feature Discovery. M.Sc. thesis, Dept. of Computer Science, Univ. of Alberta.

Appendix

The following game was the only decisive result in the 1995 Man-Machine World Championship match between Don Lafferty and Chinook. The game was played on January 17 at the International Checker Hall of Fame in Petal, Mississippi. The game is given using chess algebraic notation, where columns are labeled "a" to "h" (left to right) and rows are numbered "1" to "8" (from bottom to top). Black is at the top of the board and moves first.

Black: Lafferty - White: Chinook

1. d6-c5 c3-d4 2. e7-d6 b2-c3 3. f6-g5 g3-f4 This move was added to the opening book two days before the match started. Although not new (it hasn't been used in Grandmaster play for over 50 years), it catches Lafferty by surprise. **4. f8-e7** Lafferty thought for over 10 minutes on this move. After the game, he was asked why he didn't play the "obvious" g7-f6. Lafferty replied that f8-e7 and g7-f6 looked equally attractive and that he had a

hard time deciding. In pre-match preparation, Chinook was able to prove that g7-f6 leads to a draw. **f2-g3** Chinook is now out of its book. With a 21-ply search, Chinook switches from h2-g3 to f2-g3 and realizes it has a sizable advantage. **5. g5-h4 e1-f2 6. g7-f6** This appears to be a mistake. Chinook expected b6-a5. Chinook's evaluation of the position grows from move to move now. **f4-g5 7. h6-f4 g3-g7 8. h8-f6 h2-g3 9. b6-a5** Apparently forced. The Black position is critical. White is threatening to push a checker through from e3-f4-g5-h6-g7-f8 and get a king. Unless Black can prevent this, the White king will be able to leisurely attack Black's checkers from behind. **d4-b6 10. a7-c5 c3-d4 11. c5-b4** Some annotators of this game have suggested that this is the losing move. Chinook disagrees. If Chinook is right, then the game may have been lost on move six. It is incredible to think that a player as strong as Lafferty could be losing in only six moves. This is a strong endorsement for preparing opening books! **a3-c5 12. d6-b4 g3-f4 13. c7-d6** b8-a7 was expected, but a deep search done while Lafferty was thinking revealed it as losing. **f4-e5 14. d6-f4 e3-g5** Now it is obvious that White is getting a king. **15. b4-c3 d2-b4 16. a5-c3 d4-c5** Chinook announce that it thinks the game is won. **17. d8-c7 g5-h6 18. f6-g5 h6-g7** And now the program announces it has seen to the end of the game. **19. g5-f4 g7-f8 20. e7-f6 c1-d2 21. c3-e1 f8-e7 22. e1-g3 e7-e3** After the game, Lafferty said he played for this ending. He knew this position was bad, but he thought it was drawable. This illustrates how powerful the combination of deep search and endgame databases can be; Chinook had envisioned this position many moves previously. In the following, a "*"ed move is the only move to preserve the win. These moves are played instantly. **23. b8-a7 a1-b2* 24. g3-f2 e3-d4* 25. f2-e1 b2-a3 d4-c3** also wins. **26. h4-g3 d4-c3* 27. g3-h2 c3-d4* 28. e1-d2 a3-b4* 29. d2-c1 d4-e5* 30. c1-b2 b4-a5*** Lafferty completed this move with seconds remaining on his clock. You must play 30 moves in an hour; failure to do so results in an immediate forfeit. **31. b2-c3 c5-d6* 32. c3-d4 e5-c3* 33. c7-e5 c3-b4* 34. e5-d4 b4-c5* 35. d4-e3** Black resigns.

Jonathan Schaeffer

Department of Computing Science
University of Alberta
Edmonton, Alberta
Canada T6G 2H1
jonathan@cs.ualberta.ca
403-492-3851

Jonathan Schaeffer is a professor of computing science at the University of Alberta. He received his B.Sc. degree from the University of Toronto (1979) and his M.Math (1980) and Ph.D. (1986) degrees from the University of Waterloo. His research interests include heuristic search and parallel computing.

Robert Lake

Department of Computing Science
University of Alberta
Edmonton, Alberta
Canada T6G 2H1

Robert Lake is a lead research analyst in the Department of Computing Science at the University of Alberta. He received his B.Sc. (1979) and M.Sc. (1990) degrees from the University of Alberta. His research interests include computer animation of human figures, distributed systems and virtual reality.

Paul Lu

Department of Computer Science University of Toronto Toronto, Ontario Canada M5S
1A4

Paul Lu holds a B.Sc. (1991) and M.Sc. (1993) from the University of Alberta. Currently he is a Ph.D. student in the Department of Computer Science at the University of Toronto. His research interests include parallel and distributed computing, parallel algorithms and operating systems.

Martin Bryant

71 Hunstanton Drive,
Brandlesholme, Bury,
Lancashire BL8 1XH
England

Martin Bryant is an independent software consultant and has developed several commercially available games programs.