

FROM MINIMAX TO MANHATTAN

*T.A. Marsland and Y. Björnsson*¹

University of Alberta
Edmonton, Alberta, Canada

Abstract

The thinking process for playing chess by computer is significantly different from that used by humans. Although computer hardware and software have evolved considerably, computers still have difficulties to understand the more elaborate chess concepts. In this paper we look at the technology behind today's chess programs, its current status and how it has developed. Finally, we explore some directions for the future.

1. INTRODUCTION

One of the most spectacular shows ever hosted on Broadway took place in May 1997. Human supremacy in chess was challenged by a computer. Garry Kasparov, the World Champion and one of the greatest chess players in history, defended the honor of mankind. The challenger, DEEP BLUE, was a computer specially designed for playing chess, a masterpiece of science and engineering. Independent of the outcome of the match, the significance of this event would be the same. A combined effort of research in artificial intelligence (AI) and computer hardware development, has made it possible for computers to match humans in a game that has for centuries been considered a measure of human intelligence.

Problem solving in AI often reduces to a process of tackling three main issues: representing knowledge, searching for solutions, and using knowledge to direct the search. A hallmark of intelligence is the use of knowledge to make search problems more tractable. The most studied search problems are those that arise in game playing. For decades chess has been viewed as an important testbed for

¹ Department of Computing Science, University of Alberta, Edmonton, Alberta, Canada T6G 2H1. Email: {tony, yngvi}@cs.ualberta.ca.

machine intelligence, with some researchers regarding computer-chess as its *Drosophila* (cf. Michie, 1980). Early chess programs were inspired by concepts from human thinking. However, with steadily increasing computer power – making brute-force methods more attractive – and with the difficulties of implementing human-like reasoning, computer-chess research took a different direction. Most of the effort went into search techniques, while less attention was paid to the acquisition and use of knowledge.

This paper opens with a brief history of computer chess, followed by an overview of the current methodology. Finally, we consider the future of computer chess and its technology, and offer some direction for future work on software aspects.

2. THE PAST: A BRIEF HISTORY

Development of computer-chess techniques has been strongly influenced by advances in search algorithms and computer hardware. The evolution of these two has often gone together, more powerful hardware allowing more complex search strategies.

2.1 Search Algorithms

Half a century ago von Neumann and Morgenstern (1944) proposed *minimax* as a method to decide which move to make in chess. Using the negamax framework, the minimax algorithm is described recursively as

$$MinMax(n) = MAX_i(-MinMax(n_i))$$

where the set $\{n_i\}$ are the successors of node n . If n does not have successors, then it is a leaf node and $MinMax(n)$ is defined as the merit of node n from the perspective of the side that has the move. In principle, minimax can be used to determine an optimal action in every chess position, simply by expanding all continuations to the end and backing up the precise value of the outcome; in this case a win, a loss, or a draw. In practice, this is of course infeasible because of the astronomical number of possible continuations.

In the 1950s early pioneers of computer chess, like Claude Shannon (1950) and Alan Turing (1953), proposed various search strategies based on the minimax principle, and laid the foundation of an active research area. Faced with the large search space, early chess programs used highly selective search methods, expanding just a few plausible moves in each position. This human-like approach proves difficult; the programs often neglect good moves – throwing the

baby out with the bathwater. The quality of play by the early chess programs was therefore poor.

Over the years, numerous algorithms were proposed to back up minimax values in game trees, most notably *Alpha-Beta*. Knuth and Moore (1975) worked the algorithm into the negamax framework and described its full pruning capabilities. However, the underlying idea was known in 1958, and weaker forms of it were implemented in such early chess programs as NSS, by Newell, Shaw, and Simon (1958). The method uses an effective cut-off technique based on the observation that parts of the tree expanded by the minimax algorithm are irrelevant for proving the value of the tree, and thus need not be visited. The basic idea behind the pruning is that having already found one way of refuting the opponent's play, there is no need to look for others. The Alpha-Beta algorithm is a major improvement over exhaustive minimax. In the best case, the number of nodes visited is about the square root of the number seen by minimax, and it never searches more nodes.

As more and more powerful computers became available, and the full pruning capabilities of Alpha-Beta became better known, brute-force search programs based on the method replaced other approaches. During the past two decades, much effort was invested in enhancing the efficiency of the search. By expanding the expected best line of play first (the principal variation) one need only show that the alternatives are inferior. This observation led to more efficient variants, such as *Principal Variation Search* (Marsland and Campbell, 1982) and *NegaScout* (Reinefeld, 1983). These variants rely on having a good move-ordering scheme that expands promising moves early. Various techniques to improve the move ordering and to make the search more efficient were developed, for example, iterative-deepening, use of transposition-tables, and forward pruning. From now on we will refer to these techniques as search enhancements, and call Alpha-Beta – particularly its variants that employ such enhancements – by the generic name *Null Window Search* (NWS), thus reflecting the essence of the more effective implementations. For a detailed description we refer to Plaat (1996).

Minimaxing is not the only way to back up values in game trees, nor do the enhanced methods necessarily traverse the smallest trees. Various other algorithms have been proposed and implemented, but none have found wide use in practice, because they are not time-efficient. For an extended overview of such search alternatives we refer to Junghanns (1998).

2.2 Hardware

As computers evolved it became clear that increasing search speed dramatically improves the playing strength of even the simplest chess programs. Using only

elementary chess knowledge, the programs could play reasonably well by applying extensive search. The quest for more speed took two directions: on the one hand developing special-purpose hardware for playing chess, and on the other adapting the search to run on many general-purpose processors in parallel.

The first special-purpose high-speed circuitry for playing chess was developed by Joe Condon and Ken Thompson (1982) in the late 1970s, and was used in his chess program BELLE. The chess-specific hardware provided fast move generation. In the 1980s more high-speed chess circuitry saw the light of day, like the developments at CMU for HITECH and DEEP THOUGHT. Not only were the new custom chips capable of move generation, but they also provide position evaluation and search. In the early 1980s chess programs running on multi-processor systems started to compete in computer-chess tournaments. Although the first multi-processor based chess programs ran on only few processors, today's systems use hundreds. Searching in parallel poses many new programming problems, and many techniques have been developed to allow NWS algorithms to traverse game trees in parallel.

The DEEP BLUE program combines both of the above approaches; it runs on many general-purpose processors, where each processor has access to several high-speed special-purpose chess circuits.

2.3 Other Improvements

Although improvements in search algorithms and computer hardware have drastically increased chess programs' playing strength, other sources have contributed as well. For instance, use of openings books and endgame databases have proved very useful. Last, but not least, decades of experience have shown programmers what kind of chess knowledge most benefits the program, and how it can efficiently be represented. For chess programs, just as for humans, the knowledge acquisition period needed to achieve mastery is long and tedious.

3. THE PRESENT: CURRENT TECHNOLOGY

It is safe to state that today's best chess programs are playing at the Grand Master (GM) level. This success comes from gradual improvements over the last half century; a direct result of active research in computer chess. Advances in search algorithms, parallel computing, special-purpose chess hardware, knowledge representation, and opening/endgame databases all contribute to the increased playing strength.

3.1 Search Algorithms

Most contemporary programs use some form of NWS (see Figure 1). It keeps track of a lower bound α and an upper bound β of the value a player can hope to achieve. During the search, whenever the value returned from a subtree exceeds the upper bound β , a cut-off occurs. Intuitively this means that we have found a refutation of the current line of play, and therefore our opponent will not select this variation, but will instead play the line that achieves the value indicated by the upper bound β . The algorithm first expands the so-called principal variation with a full $\alpha\beta$ window, and then the remaining moves with a minimal window around the score the principal variation returned. This results in efficient searches and only occasionally, when an alternative move really turns out to be better, is a re-search with a wider window necessary.

When equipped with enhancements such as a transposition table (accessed by functions *LookTT()* and *AddTT()* in Figure 1), move-ordering (*AddCredit()*, *NextMove()*), search extension/pruning schemes (*Extend()*, *NullMove()*), and quiescence search (*QS()*), the algorithm is extremely efficient. Thus the fastest chess programs look ahead about 5 to 7 moves for each side (i.e., 10 to 14 plies) in complicated middle-game positions, and search lines of special interest even more deeply. In the endgame and in other less complicated positions a much deeper search is possible.

3.2 Search Enhancements

Various search enhancements are essential for the NWS algorithm to achieve exceptional performance. Many of them focus on allowing the algorithm to be more selective in how the search effort is spent. The selectivity is introduced by varying the search horizon; some lines are searched to a depth that is greater than the nominal depth, while others are terminated at a shallower depth. The real task is to identify which alternatives are worth considering further and which can be pruned off. In Figure 1 three common methods to vary the search horizon are shown: quiescence search, extending forcing lines, and a forward-pruning heuristic called null move.

Having searched some initial position to a designated maximum depth, some of the positions that arise are volatile and hard to evaluate. For example, if we evaluate a position where a capture has just occurred, without giving the opponent opportunity to re-capture, a huge error may be introduced. Therefore, all captures for both sides are played out before a position is evaluated. Sometimes other moves that can drastically change the evaluation are also made, for example promotions. This resolution of dynamic factors is called the *quiescence search* phase.

```

int NWS( int height, int alpha, int beta ) {
  MOVE move, tt_move=NullMove, best_move=NullMove;
  int score, best_score, bound, ttbound;
  TT tt;

  //Check for draw by repetition or 50 move rule
  if ( IsDraw() ) return DRAW;
  if ( Extend() ) height++; //Extend forced moves

  //Call quiescence search if horizon is reached
  if ( height == 0 ) return QS( alpha, beta );

  //Look the position up in the transposition table
  if ( LookTT( &tt ) ) {
    if ( tt.height >= height ) {
      switch ( tt.bound ) {
        case TValue: return tt.score;          break;
        case UBound: beta = MIN(tt.score,beta);break;
        case LBound: alpha = MAX(tt.score,alpha);
      }
      if ( alpha >= beta ) return tt.score;
    }
    best_move = tt_move = tt.move;
  }

  //Do a null-move search with search reduction R>1
  //if not in check and last move not a null-move
  if ( (height > R) && NullmoveIsOK() ) {
    Make( NullMove );
    score = -NWS( height-1-R, -beta-1, -beta );
    Retract( NullMove );
    if ( score >= beta ) return beta;
  }

  //Get the first move if it wasn't found in the TT.
  //If no legal moves, either a mate or a stalemate.
  if ( best_move==NullMove && !NextMove(&best_move) )
    return IsCheckmate() ? LOSS : DRAW;

  //Search the principal move before the others
  Make( best_move );
  best_score = -NWS( height-1, -beta, -alpha );
  Retract( best_move );
  while ( (best_score < beta) && NextMove(&move) ) {
    if ( move == tt_move ) continue;
    bound = MAX( alpha, best_score );
    Make( move );
    score = -NWS( height-1, -bound-1, -bound );
    if ( (score > bound) && (score < beta) )
      score = -NWS( height-1, -beta, -score );
    Retract( move );
  }
}

```

```

if ( score > best_score ) {
    best_move = move; //A new best move is found
    best_score = score;
}
}
//Update the TT and move ordering information
ttbound = (best_score <= alpha) ? UBound :
           (best_score >= beta) ? LBound : TValue;
AddTT( height, best_move, best_score, ttbound );
AddCredit( ttbound, height, best_move );
return best_score;
}

```

Figure 1: NWS enhanced with TT and Null-move.

A common practice is to extend the search along “forced” lines, such as certain re-captures, or when one side is moving out of check. In Figure 1 the function *Extend()* checks for forced moves. If a decision is made to extend, then the maximum depth of search for this line is increased by one.

As opposed to search extensions, pruning schemes select branches that are to be searched to less than nominal depth. This is referred to as forward pruning, and can involve erroneous decisions by prematurely truncating good lines. The null-move heuristic is the commonest forward-pruning method used in contemporary chess programs. The underlying idea is that in chess it is almost always beneficial to make a move rather than to pass. Therefore, if the score received by giving up a move is still good enough to cause a cut-off, it is very likely that at least one legal move will too. When the position is assessed using a shallower search than we would search otherwise, considerable search effort is saved. The constant R shown in Figure 1 decides the search reduction for the null-move searches, and is typically set to 2. Some precautions are necessary when using a null move, for instance, we can not allow a null move to be made when the side to move is in check, nor can two consecutive null moves be allowed. In chess it is almost always safe to make the assumption that making a move will improve the position, but for a special case, called *zugzwang* positions, this is not true. Most chess programs turn off the null-move heuristic when entering the endgame, because *zugzwang* positions are more likely to arise there.

The *transposition table* is used to store information about positions that have been visited before, either during an earlier iteration or when different move sequences transpose into the same position. Typical items to store in a transposition table are: the merit value of a position, type of the value (i.e., true value, upper bound, or lower bound), the height of the subtree the position was searched to establish the value, the best move in the position, and a hash key for the position. This key is necessary because hashing is used to index the table.

Whenever a node is visited, the transposition table is checked to see if it was seen before, and if so, whether it was searched sufficiently deeply so that we can reuse the merit value from the table. The merit can be used to determine a true value for the position, thus eliminating the need to search the position further, or to adjust the current $\alpha\beta$ bounds, which can also lead to a cut-off. The use of a transposition table can reduce the search space significantly, especially in endgames where transpositions occur frequently. The table can only hold a small portion of actual positions searched, therefore various replacement schemes are used to decide which positions to keep; one scheme uses two-level transposition tables. A two-level table stores for each entry, not only the most recent position hashed into that entry but also the position that was searched the deepest. For an adequate comparison between the various schemes we refer to Breuker (1998). In Figure 1 the functions *LookTT()* and *AddTT()* retrieve and insert entries into the transposition table, respectively.

NWS-like search algorithms perform best when good moves are expanded first. Therefore, various *move-ordering heuristics* have been developed. One such is to store the best move in the transposition table. When a node is revisited this move is always tried first; the rationale being that a move previously found to be good in one position is also likely to be good when the position is searched to a greater depth. Brute-force methods explore all possible moves in each position, many of which can easily be refuted by an obvious capture. Chess programs therefore often order capture moves early in the move list. The killer move (Slate and Atkin, 1977) and the history heuristic (Schaeffer, 1989) are more sophisticated move ordering schemes that are widely used. Both are based on the idea that a move that is good in one position is often good elsewhere. The former keeps track of a few best moves found at each height level in the tree, while the second keeps a global table for all possible moves. The best move found in each position gets a credit, shown in Figure 1 as a function call *AddCredit()*, and when moves are generated by *NextMove()*, captures are done first and then the remainder in order by the credit they have.

When using a depth-first search it is necessary to decide beforehand how deep to search. This makes it difficult to estimate how long the search will take. However, by gradually increasing the search depth one can better decide how long the search will take and when to stop searching. This is called *iterative deepening*. Because the time for each iteration grows exponentially with increased search depth, the effort spent in earlier iterations is relatively small compared to the last iteration. The search overhead introduced by iterating on the search depth is therefore small. Furthermore, when used in combination with a transposition table, the principal variation is kept between iterations. This leads to a better move ordering, often resulting in iterative deepening searching fewer nodes in total than the non-iterative approach. The technique of iterative-

deepening search later found its way into other AI domains, such as theorem proving and single-agent search.

3.3 Hardware

Today's state-of-the-art PC is extremely powerful. The commercial chess programs running on them are also quite sophisticated, and the playing strength approaches the GM level. The state-of-the-art computer-chess hardware is undoubtedly DEEP BLUE. It uses 256 to 512 chess-specific processors, and under tournament time controls is capable of searching 50 to 100 billion positions for each move. In the world of computer-chess research IBM's DEEP BLUE project is comparable to the well-known Manhattan project.

3.4 Knowledge Representation

The chess knowledge used in contemporary chess programs is primitive by human standards. The programs have a notion of basic chess concepts such as material, mobility, pawn structures, king safety, weak/strong squares, space, center control, and development. A good description of a typical knowledge encoded in chess programs can be found in CRAY BLITZ (Hyatt, Gower, and Nelson, 1990), whose evaluation function is described in some detail.

More elaborate chess concepts, such as long-term planning, piece co-ordination, theme consistency, imbalances, long-term weakness, are totally missing from their vocabulary. It would be impossible for human players to achieve master strength without this knowledge. The chess programs partially compensate for this lack by using deep exhaustive searches.

3.5 Opening Books and Endgame Databases

Just as for humans, good opening play is an important part of a strong chess program. The best programs invariably have opening books prepared by chess masters. Not only do the masters manually prepare new openings lines, but computers are also used to verify existing analysis in the hope of finding flaws or further improvements. The opening book must guide the program to positions that are not only favorable, but also well suited to the computer's play. This often implies avoiding closed positions in favor of more open and tactical ones. The importance of having a good opening book is evident from the fact that the top commercial programs specifically prepare opening lines against other top programs. Therefore, to prevent both chess programs and human players from repeatedly exploiting the same opening mistakes, some programs automatically update their opening book during a game; deleting lines that lead to inferior positions soon after deviating from the book, while adding others that showed

promise. Similar concepts were tried in earlier chess programs such as BEBE (Scherzer, Scherzer, and Tjaden, 1990). It stored positions from its previous games in a special long-term transposition table, thus preventing the program from making the same mistake twice.

Pre-calculated endgame databases are also available, for endgames with 5 or fewer pieces. When arriving at a position in the database, the programs play the best moves from there without any search. A database construction is described in detail by Van den Herik and Herschberg (1985). The endgame databases by Ken Thompson (1986) are made publicly and are therefore used in many chess programs.

4. THE FUTURE: WHAT IS THE RIGHT WAY

The benefit of each additional ply diminishes as we search more deeply. This implies that the utility of searching more deeply will become less and less. Therefore, we believe that there will be a shift in emphasis; less effort will be spent on search algorithms and more on the knowledge part.

4.1 Search

Search is, and will remain, fundamental to chess-playing programs – after all, tactical combinations are a big part of the game. This is one reason why chess programs have an advantage over human players. By exhaustively looking at all alternatives the programs often discover unexpected tactics or brilliant defending resources that are hidden to the human eye. Further search improvements can come from one of two sources:

- *Searching faster*
The search may benefit from increasing computing speed, but this will not be as dominating a factor as before. It can be more beneficial to invest the extra computing speed in incorporating additional knowledge. Research in parallel search is, and will, stay important, although not of ultimate importance to the AI community.
- *Searching smarter*
Alternatively, more selective-search algorithms can be developed, that search “smarter” by more aggressively extending branches of interest while pruning others of lesser potential. Selective-search methods are likely to attract increased attention from the research community. One of the fundamental questions of AI – how to use knowledge to direct search – has to be revisited.

4.2 Knowledge

Further major improvements must come from additional chess knowledge. Questions like, what additional knowledge is important, how can it be acquired and effectively implemented, need to be addressed. These issues have been largely neglected in the past.

As mentioned before, knowledge acquisition is a long and tedious task. Methods to make this task easier, like tools for automatically extracting relevant chess knowledge from big databases, will become increasingly useful.

The long-range planning component of chess programs should also be worked on. An old adage says that it is better to play with a weak plan in mind than no plan at all. Deep searches certainly give the programs some notion of a plan, but often it is too superficial and short sighted. It is sometimes embarrassing to watch a well-prepared human opponent make even the strongest chess programs look like absolute beginners; simply by applying an uncomputer-like style of play.

Ideas are the tools of the chess player. Players do not reinvent the wheel every time a game is played. The same themes occur over and over again, and during a game players recall similar positions from their own and other people's games, and reuse the ideas from there. Often these ideas are based on long-term plans that are outside the scope of any search, they are simply memorized. In general, humans have an extraordinary talent for adapting existing knowledge and applying it to similar situations. It is called reasoning by analogy and reflects the way they play chess. So, chess might become an active testbed for Case-Based Reasoning methods that experiment with such memory based reasoning. In the long run, chess programs could greatly benefit from such methods, allowing them to reuse plans and adapt them to different over-the-board situations.

4.3 Chess

An interesting question to ask is whether chess programs will influence the way chess is played in the future. Chess, like most other things is receptive to changes, both from fashion and from new discoveries. Computers have already been used to discover new facts about various chess endgames, and to improve opening theory. The question here is more on the lines of whether they will change the style of human play. Humans sometimes deliberately play slightly inferior moves, to surprise the opponent, to complicate matters, or to seek some counter chances.

The psychological effect of a surprise move is of little value against a computer program. One can argue that the programs have certain psychological advantages over the humans. For example, in the first game of the 1996 Kasparov vs. DEEP BLUE match, DEEP BLUE showed a total disregard for Kasparov's attack, calmly grabbing a Pawn on the Queenside while its king fortress was under seemingly disastrous threats. A human player might have shown the World Champion a little more respect by taking his attack seriously. In the last game of the same match, Kasparov had a winning piece sacrifice on h7, but after thinking for a long time he played differently, apparently deferring to the judgment of the computer.

5. CONCLUSIONS

It is inevitable that chess programs will one day out-smart even the strongest humans. The future programs will continue to rely on extensive search, but whether that approach alone is sufficient to reach the top is moot. More sophisticated methods are needed, that allow the programs to understand the more elaborate chess concepts. However, even though a pure brute-force approach will suffice to make a world-class chess program, the program will be of limited interest unless it can share its knowledge. For a program to be an effective tutor, it must be able to explain various chess-related concepts in a manner that is understandable and natural to us humans, instead of simply returning a raw numeric score indicating a merit of a chess position. Therefore, disregarding the outcome of the Garry Kasparov vs. DEEP BLUE match, we believe that computer chess will remain an important testbed for AI applications, not only in search but also in fields such as reasoning, knowledge representation, and knowledge acquisition.

6. REFERENCES

- Breuker, D.M. (1998). *Memory versus Search in Games*. Ph.D. Thesis, Universiteit Maastricht. ISBN 90-9012006-8.
- Condon, J.H. and Thompson, K. (1982). Belle Chess Hardware. *Advances in Computer Chess 3* (ed. M.R.B. Clarke), pp. 45-54. Pergamon Press, Oxford. ISBN 0-08-026898-6.
- Herik, H.J. van den and Herschberg, I.S. (1985). The Construction of an Omniscient Endgame Data Base. *ICCA Journal*, Vol. 8, No. 2, pp. 66-87. ISSN 0920-234X.

- Hyatt, R.M., Gower, A.E., and Nelson, H.L. (1990). CRAY BLITZ. *Computers, Chess, and Cognition* (eds. T.A. Marsland and J. Schaeffer), pp. 111-130. Springer-Verlag, New York. ISBN 0-387-97415-6.
- Junghanns, A. (1998). Are there Practical Alternatives to Alpha-Beta? *ICCA Journal*, Vol. 21, No. 1, pp. 14-32. ISSN 0920-234X.
- Knuth, D.E. and Moore, R.W. (1975). An Analysis of Alpha-Beta Pruning. *Artificial Intelligence* Vol. 6, No. 4, pp. 293-326. ISSN 0004-3702.
- Marsland T.A., and Campbell, M.S. (1982). Parallel Search of Strongly Ordered Game Trees. *ACM Computing Surveys* Vol. 14, No. 4, pp. 533-551.
- Michie, D. (1980). Chess with Computers. *Interdisciplinary Science Reviews*. Vol. 5, No. 3, pp. 215-227. ISBN 0308-0188.
- Neumann, J. von and Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Second Edition, 1947. Princeton University Press, Princeton, N.Y.
- Newell, A., Shaw J.C., and Simon, H.A. (1958). Chess Playing Programs and the Problem of Complexity. *IBM Journal of Research and Development*, Vol. 4, No. 2, pp. 320-335. Reprinted (1963) in *Computers and Thought* (eds. E.A. Feigenbaum and J. Feldman), pp. 39-70. McGraw-Hill, New York, N.Y.
- Plaat, A. (1996). Research, Re: Search & Re-Search. Ph.D. Thesis. Erasmus University Rotterdam. Thesis Publishers Amsterdam.
- Reinefeld, A. (1983). An Improvement to the Scout Tree Search Algorithm. *ICCA Journal* Vol. 6, No. 4, pp. 4-14. ISSN 0920-234X.
- Schaeffer, J. (1989). The History Heuristic and Alpha-Beta Search Enhancements in Practice. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 11, No. 11, pp. 1203-1212. ISSN 0162-8828.
- Scherzer, T., Scherzer, L. and Tjaden, D. (1990). Learning in Bebe. *Computers, Chess, and Cognition* (eds. T.A. Marsland and J. Schaeffer), pp. 197-216. Springer-Verlag, New York. ISBN 0-387-97415-6.
- Shannon, C.E. (1950). Programming a Computer for Playing Chess. *Philosophical Magazine*, Vol. 41, No. 7, pp. 256-275. Reprinted (1988) in *Computer Games I* (ed. D.N.L. Levy), pp. 81-88. Springer-Verlag, New York. ISBN 0-387-96496-7.

Slate, D.J. and Atkin, L.R. (1977). CHESS 4,5 - The Northwestern University Chess Program. *Chess Skill in Man and Machine* (ed P.W. Frey), pp. 82-118. Springer-Verlag, New York. 2nd ed. 1983. ISBN 0-387-90815-3. Reprinted (1988) in *Computer Chess Compendium* (ed. D.N.L. Levy), pp. 80-103; B.T. Batsford, London.

Thompson, K. (1986). Retrograde Analysis of Certain Endgames. *ICCA Journal*, Vol. 9, No. 3, pp. 131-139. ISSN 0920-234X.

Turing, A.M. (1953). Digital Computers Applied to Games. *Faster Than Thought*, (ed. B.V. Bowden), pp. 286-297. Pitman, London.