

On Effort in AI Research A Description along Two Dimensions

Franz-Günter Winkler
Austrian Society for Cognitive Science
Barichgasse 10/2/10, A-1010 Wien, Austria
E-mail:winkler@coams.atc.co.at

Johannes Fürnkranz
Austrian Research Institute for Artificial Intelligence
Schottengasse 3, A-1010 Wien, Austria
E-mail:juffi@ai.univie.ac.at

Abstract

In this paper we describe Artificial Intelligence as research that moves along two different axes: human-compatible knowledge and machine-compatible processing. An analysis of computer chess research along these dimensions shows that AI more and more diverges into an engineering branch and a cognitive branch. As an explanation, we offer a hypothesis about the dependency of research effort on these dimensions. It becomes obvious that the most rewarding projects are the hardest.

"If one could devise a successful chess machine, one would seem to have penetrated to the core of human intellectual endeavour."

Newell, Shaw & Simon

Introduction

We view Artificial Intelligence as a science with the goals of, on the one hand, providing computers with the ability of solving tasks that are commonly perceived as requiring intelligence, and, on the other hand, furthering our understanding of human thinking. The basic problem that AI has to face is to integrate the abstract concepts that human beings usually use for problem solving with the strict algorithmic processing definitions that are required for implementing programs on a computer.


To illustrate these dimensions consider the two problems depicted in figure 1. The first one is a slightly unusual endgame position, but, after some deliberation, a trained chess player will detect many concepts that are familiar from other endgames. For example, it will soon become obvious that black has a dangerous threat with 1...d3, which will either queen the d-pawn or distract white's c-pawn, thus allowing black to queen his own c-pawn. Because of the above drawing chance for black, white's only hope lies in queening one of its pawns on the b-file. To achieve this, white has to conquer the square a6. However, he has no moves that put black into zugzwang, because black can answer all white king moves with king moves b7-a8-b7. A typical maneuver in such position is the so-called *triangle maneuver*, where one king is able to use a 3-cycle to

return to its original square, while the other king is only able to make a 2-cycle. White can therefore try to move its king to e1, playing e1-f2-f1-e1, which would gain one move. When the white king returns to a5, we will have exactly the same position, but with black to move. As he cannot move his king because of white's threat Ka6, he has to move one of his pawns. Then the entire sequence is repeated 11 times until black has no more pawn moves and has to answer 254. Ka5 with Kc8 thus allowing Ka6, followed by a mate in 15.

Trained chess players could solve this problem in their heads (maybe with the exception of the final mate in 15). However, because of their lack of high-level chess-specific concepts and general problem-solving knowledge, this problem is very hard for current computer chess programs, as the solution is too deep to be found with exhaustive search. A brief experiment that we have conducted with FRITZ4 has demonstrated that FRITZ will play Bb1 to prevent black from playing Pd3, bring its king to a1, Ba2, king to e1, Bb1, and finally move the king to f2 in order to capture the pawn on f3. It does not realize that black then has an easy draw by trading its knight for the bishop and playing Pf6-f5 thereafter.


On the other hand, the second position is quite hard for human chess players, because it does not offer many familiar patterns which could be used for narrowing down the search (at least not for chess players that are untrained in the solving of chess problems). The pieces seem to be randomly scattered around the board, thus offering no orientation at all. The only obvious feature is the lack of king safety, so that white should be much better off and probably has a mate in a few moves. However, it is not easy for a human player to find the optimal move (1.Rg3!). In particular, he would more or less have to perform an exhaustive search through the numerous possible threats. On the other hand, positions like these are no different to other positions for a computer chess playing program, which would find the mate in 2 quite easily. In fact, computer chess programs can be expected to play from randomly generated position equally well as from positions that typically arise in a chess game. The reason for this is that

Nenad Petrović, 1969



Mate in 270 moves

Sam Loyd, 1892



Mate in 2 moves

Figure 1: A knowledge-rich and a knowledge poor problem.

these algorithms are not modeled after human memory and thus do not have to rely on the recognition of familiar concepts on the board. Instead their search algorithms are tailored to be efficiently executable on computer hardware.

In this paper, we try to argue that research in AI has diverged into two branches that in an intuitive sense correspond to the two approaches sketched above: On the one hand, the what one might call *cognitive* branch of AI puts a strong emphasis on the psychological validity of computational models, in particular with respect to knowledge representation and memory organization. On the other hand, the *engineering* branch of AI is motivated by solving particular tasks, and is mostly concerned with finding formalizations and software architectures for particular problems that can be efficiently executed on computer hardware. We illustrate this split of AI into two different branches by classifying research in the domain of chess along two different axes: *human-compatible knowledge (HCK)* and *machine-compatible processing (MCP)*. We further argue on the example domain of chess that the successes of AI research can be found along the two axes, but have not yet penetrated into the white area which we consider to contain the core problems of AI. We then ask the question why this is the case, and present a hypothesis to resolve this issue.

Human-compatible knowledge

Chess is probably the game that has been most deeply investigated from a theoretical point of view. Chess books are full of comprehensible knowledge about different aspects of the game. We would like to call such knowledge *human-compatible*, because it enables a chess student to increase his understanding and com-

petence of the game. Nevertheless, it is mostly unclear how the student uses this knowledge for problem-solving.

Human subjects are often able to specify the abstract concepts they use for problem-solving, but are unable to specify the problem-solving process in an exact algorithmic way. For example, a chess player has no problems to explain the reasoning that made him prefer a certain move over other possible continuations. Analyses like "The move b4 gives me a backward pawn on c3, but it prevents black liberation with a5, so that I can attack his weak a6-pawn on the half-open a-file." are full of abstract concepts like *backward pawn*, *half-open file*, etc. that are well-understood by human players. However, it is comparably difficult for human players to specify the thought processes that made them prefer, e.g., the opponent's weak a6-pawn over their own backward pawn on c3. Research in chess psychology (deGroot 1965; Chase & Simon 1972; Holding 1985; deGroot & Gobet 1996) has extensively analyzed verbal thinking-aloud protocols of chess players of different strengths. The results are that differences in playing strength between experts and novices are not so much due to differences in the ability to calculate long move sequences, but to the use of a library of chess patterns and accompanying moves and plans that helps them choose the right moves for deeper investigations. Several authors have even tried to measure the magnitude of this pattern library, resulting in estimates in the range of 5,000 to 10,000 patterns (Simon & Chase 1973; Hayes 1987). Some of these so-called *chunks*¹ are easy

¹Recent research has extended the chunking theory with so-called *templates* (Gobet & Simon 1996), i.e., long-term-memory structures that are quite similar to *scripts* and

to articulate and common to most chess players (like, e.g., *passed pawn*, *skewer*, *minority attack*), while others are presumably subconscious and subjective to individual players. However, even simple concepts like a *knight-fork* are non-trivial to formalize.²

Because of this strong focus on models for memory organization, early AI research has concentrated on the simulation of aspects of the problem-solving process that are closely related to memory, like perception (Simon & Barenfeld 1969) or retrieval (Simon & Gilmarin 1973). Recently, these ideas were re-investigated and integrated into the CHREST program (Gobet 1993), which is the most advanced computational model of a human's chess players memory organization. CHUMP is a variant of this program that is actually able to play a game by retrieving moves that it has previously associated to certain chunks in the program's pattern memory (Gobet & Jansen 1994).

Machine-compatible Processing

AI has soon recognized the difficulty of formalizing human thought in a top-down way (using the human concepts as a starting point), and has instead discovered approaches to solving intelligent tasks, which are more closely modeled after the processing capabilities of a computer. Brute-force chess programs are the best-known example of this line of research. The basic idea of brute-force chess programs dates back to (Shannon 1950) and (Turing 1953), where one can already find many of the ideas that are still used in today's chess programs (like, e.g., search extensions). However early chess programs (see (Newell, Shaw, & Simon 1958) for an overview) relied on highly selective search on a few basic concepts like *material balance*, *center control*, and *king safety*. This selective search was motivated by both hardware limitations and the attempt to model machine chess playing on human chess playing.

However, the somewhat unexpected success of the TECH program (Gillooly 1972) for the first time demonstrated the power of brute-force computing. Further improvements (Slate & Atkin 1983) and the advances in parallel processing (Hyatt, Gower, & Nelson 1985) and chess-specific hardware (Condon & Thompson 1982; Ebeling 1987; Hsu 1987) have eventually lead to the Deep Blue vs. Kasparov challenge, where a computer for the first time won a tournament game against the reigning human chess world champion.

frames, but are based on a detailed psychological model (deGroot & Gobet 1996). For our discussion, the differences in the details of the psychological models of chunks and templates are irrelevant.

²The basic pattern for a fork is a protected knight threatening two higher-valued pieces, like, e.g., rook and queen. However, this simple pattern might not work if the forking knight is pinned. But then again, maybe the knight can give a discovered check...

Two Dimensions of AI Models

The success of brute-force programs is attributable to the fact that their basic architecture is adapted to what computers are good at: fast calculation using only a few isolated chess concepts, which can be evaluated efficiently. Thus the success of these programs depends on *machine-compatible processing (MCP)*. On the other hand, we have seen that human chess players calculate relatively few moves, but rely on a huge pattern library that helps them select the right move. Thus, their success depends on the availability of *human-compatible knowledge (HCK)*. In the following, we classify AI models of the domain of chess with respect to their contribution along either axis.

A perfect chess program that has access to the perfect game-theoretic values of each position (e.g., by exhaustive search until check-mate) would be on the right end of the MCP axis with no contribution on the HCK axis, as its internal knowledge representation does not contribute in any way to a better understanding of human reasoning. The best-known approximation of this principle is the Deep Blue chess program. The other extreme would be an oracle that could derive the best move in each position from general principles and explain this choice in a clear and understandable form. Chess theory can be viewed as an attempt to formalize this knowledge in an understandable way and, in some sense, Garry Kasparov can be viewed as a machine that embodies this knowledge.

A project like CHUMP (see above) is strongly motivated by human memory organization, and its processing is not very compatible with typical computer hardware. For example, the program uses artificial simulations of human long-term and short-term memory. Therefore, it has only made a small step along the MCP axis. Its contribution along the HCK axis is higher, but it is clearly a simplification compared to human memory organization. As a sort of dual example, consider PARADISE (Wilkins 1980), which is a program for solving chess combinations at an abstract level. The main goal of this project was to investigate the extent to which tree search can be guided and controlled with the use of background knowledge (Wilkins 1982). The used concepts are quite abstract and clearly motivated by human knowledge, but the processing is still very machine-compatible, using a systematic best-first search in the space of possible plans.

We view Artificial Intelligence as a science that tries to integrate both aspects: human-compatible knowledge and machine-compatible process definitions. An AI project can be evaluated by its contribution along either axis. Figure 2 positions the "chess models" discussed in the previous paragraphs in this two-dimensional field. We view the overall *value* of a model as the sum of the contributions that it makes along either axis. Intuitively, for a chess program, this value measures the program's competence in terms of playing strength and explanatory power.




Figure 2: Models of chess playing classified along the two dimensions of human-compatible knowledge (HCK) and machine-compatible processing (MCP).

A research program appears as a line progressing through this two-dimensional space. The ideal AI research program should follow the diagonal between the two axes, thus taking into account both aspects equally. The further a research program deviates from this diagonal, the less we are inclined to call it research in AI. Clearly, the progress in chess theory (which moves along the HCK axis) can hardly be regarded as part of AI, because of its lack of MCP. Likewise, we think that a certain minimum amount of human-compatible knowledge is required for a computer program to qualify as AI. Along this dimension, the development that has led to Deep Blue, in our opinion, is at the lower border of AI research, if not beyond.

The example of chess shows us that research has split into two streams: One that proceeds close to the HCK axis and is concerned with a deeper understanding of human problem solving and the domain knowledge on which it relies. The other stream is concerned with the development of faster and better algorithms and hardware heading for an exhaustive search. The little work that has been done in the white area inbetween has only been moderately successful in terms of our combined measure. We believe that this divergence can also be found in other areas of AI like Machine Learning or Natural Language Understanding. Apparently, progress along the axes faces less resistance than progress along the diagonal. The question is "Why?"

A Hypothesis on Research Effort

We interpret figure 2 in the way that moving towards a better model (in terms of the overall value defined above) requires more effort if this improvement in value is made along the diagonal than if one proceeds along

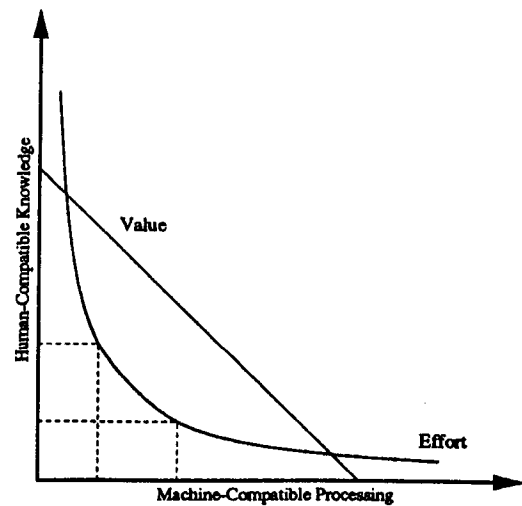


Figure 3: The discrepancy between a class of models that are equal in effort and a class of models that are equal in value.

one of the axes. We think that the effort that is necessary to proceed a step along one axis is proportional to the progress that has already been made along the other axis. Only in the special case of a research program that proceeds along a single dimension (as e.g. the research in computer chess) will the increase in effort be proportional to the increase in value. This leads us to the following hypothesis:

The total effort that has to be spent on an AI problem is proportional to the product of the values along the axes HCK and MCP.

As an illustration consider the curve depicted in figure 3, which shows a line of AI problems which require the same effort, i.e., problems for which the area of the rectangle spanned by the co-ordinates is constant. We believe that the status quo of AI research could be described with such a curve. Relatively few progress has been made along the diagonal, while considerable progress has been made along the axes. On the other hand, consider the line which depicts a class of AI models sharing the same overall value. According to our hypothesis, there is a noticeable discrepancy between these two graphs. There is a difference between the value of an AI model and its required effort.

From these deliberations follows that HCK and MCP cannot be treated independently. A possible way how they are interrelated can be derived within the theory of semantic networks. This theory seems adequate for this task. On the one hand, semantic networks are a psychological model of human memory organization and, on the other hand, have had a considerable influence on AI research.

A semantic network consists of nodes and links that connect such nodes. The nodes are to be understood as

representations of basic concepts like objects or properties, whereas the links represent relations between these concepts. The resulting structure is a model of human semantic memory organization. Processing in such a semantic network is guided by the links (either in the form of explicit rules or by spreading activation). When there are sparse links between nodes that form a linear and mostly hierarchical structure, computation will be quick and easy. When the link structure is dense and cyclic, computation becomes problematic due to circularity and to the exponential growth in the number of possible paths. Thus processing depends on the number and the structure of the links in the network.

Classical semantic memory models (Collins & Quillian 1969; Anderson 1976) rely on the power of hierarchical, sparse network structures, but there is strong critique on this situation. In (Klimesch 1994) Klimesch argues that these conceptions are misleading and he shows that they contradict important experimental data. He instead pleads for highly interconnected and cyclic memory structures in the following sense: The more connections a concept is associated with, the more meaningful is that concept. The expert who really understands what he is doing and who can apply his knowledge in most different situations has a much more interconnected memory structure than the beginner who just learned some strict rules. The interesting thing is that human memory performance becomes better and quicker with increasing connectivity of concepts, which is exactly the reverse for computational processes.³

In the formalism of semantic networks, we can define indicators for HCK and MCP: HCK is high when the link-to-node ratio is high as is given in interconnected and cyclic structures. On the contrary, MCP is high when the node-to-link ratio is high as is given in sparse and linear structures. This construction makes the dependency of the two dimensions explicit by reducing them to more basic concepts.

Steps towards a Re-unification of AI

We take the quotation preceding our paper as an intuitive definition of the goal of AI. The goals expressed in this statement are on the one hand to produce intelligent behavior in the form of a "successful chess machine" and, on the other hand, to "penetrate to the core of human intellectual endeavor". However, the development of AI research, in particular in the domain of chess, has shown that the latter is not an immediate consequence of the former. Achieving functionality does not necessarily increase the understanding of how we achieve this functionality. We think

³One of the authors made extensive experience with this and other facets of problematic computability of interconnected structures when simulating Klimesch's model in his diploma thesis (Winkler 1991).

the original motivation of research in computer chess, namely "merely" to build a successful chess machine, has to be replaced with different goals that require a reconciliation of machine-compatible processing with human-compatible knowledge. However, according to our hypothesis, we cannot expect this to be easy. In the following we would like to give a few examples for rewarding and challenging tasks in the domain of computer chess.

A very rewarding task would be the development of a computable vocabulary of chess concepts in which chess knowledge can be formulated. The characteristics such a representation formalism has to incorporate are that it has to be sufficiently expressive for formulating abstract strategic concepts, that it has to be extensible and can be easily understood by a user (HCK), and that it can be efficiently implemented (MCP). The need for such formalisms has been recognized early in computer chess research. (Zobrist & Carlson 1973) describes an advice-taking chess program which aimed at allowing a chess master to "advise" a playing program in terms of this language. Many formalisms have subsequently been developed in the same spirit (Bratko & Michie 1980; George & Schaeffer 1990), most of them limited to certain endgames (see (Michie & Bratko 1991) for a bibliography). A recent promising step into the right direction can be found in (Donninger 1996), which introduces a very efficient interpreter of an extensible language for expressing certain characteristics of a board position. However, the expressiveness of the language is currently limited to propositional logic, a trade-off that had to be made because of efficiency considerations and the ability to provide a graphical interface that also allows untrained users to formulate rules.

Another promising field for further research could be the discovery of understandable knowledge in chess endgame databases with the goal of enriching chess theory. Consider, for example, Ken Thompson's impressive work on five-men endgame databases, which is now publicly available on three CD-ROMs. The use of these disks allow chess programs to perfectly play the encoded endgames. However, many of these endgame databases are not thoroughly understood by human experts. The most famous example are the attempts of grandmasters to defeat a perfect KQKR database within 50 moves or the attempt of an endgame specialist to defeat a perfect database in the "almost undocumented and very difficult" KBBKN endgame (Roycroft 1988). GM John Nunn's effort to manually extract some of the knowledge that is implicitly contained in these databases resulted in a series of widely acknowledged endgame books (Nunn 1992; 1994b; 1995), but Nunn readily admitted that he does not yet understand all aspects of the databases he analyzed (Nunn 1994a). It would be rewarding to develop algorithms for automatically discovering playing strategies for such endgames (see (Muggleton 1988)

and (Fürnkranz & De Raedt 1997) for some preliminary work). A particularly hard problem is that human-compatible strategies are typically simple, but not necessarily *optimal* in the sense that they require a minimum number of moves. For a machine, it is non-trivial to decide which suboptimal moves contribute to some global progress (and are thus part of a useful strategy) and which suboptimal moves do not improve the position. An attempt to automatically discover a simple playing strategy for the KRK endgame might easily produce the simple strategy "always move your rook away from the enemy king" which will always result in a won position (at least for the next 49 moves), but clearly make no progress towards the goal of mating the opponent's king. Other tasks that could be automatized include the discovery of opening theory mistakes, the automatic detection of particularly promising or unpromising line-ups or middle-game plans in certain types of openings, and many more. One can even imagine facilities that support tournament preparation by analyzing game databases with the aim of unearthing characteristics of the style of individual players and for studying their weaknesses and strengths.

Another obvious point, where chess knowledge would be of considerable importance, and probably the point with the highest commercial potential is the use of high-level chess knowledge in educational chess programs. For example, imagine a program that analyzes a certain position or an entire game on an abstract strategic level, tries to understand your opponent's and your own plans, and provides suggestions on alternative ways to proceed. Some commercial programs already provide such capabilities, but at a very preliminary level that usually is only able to detect tactical, but not strategic mistakes. The ICCA has recognized the potential of such programs, and has created the *The Best Annotation Award* which will be awarded annually for the best computer-generated annotation of a chess game.⁴ However, the competition suffers from a considerable lack of participants. Some preliminary work on using case-based reasoning for a strategic analysis of a given chess position can be found in (Kerner 1994; 1995).

Last but not least, we also believe that additional knowledge can increase the playing strength of current chess programs. However, the motivation to investigate such approaches has significantly declined with the somewhat unexpected success of brute-force programs. In the Introduction, we have already illustrated the weakness of brute-force chess programs in certain endgame positions that require abstract problem-solving and chess-specific knowledge. For some preliminary ideas incorporating strategic long-term knowledge into conventional chess programs see (Kaindl

1982; Opdahl & Tessem 1994; Donniger 1996). However, we are also of the opinion that some of the early approaches to selective search need a re-evaluation in the light of the development of AI in the last 25 years.

Conclusion

In this paper, we described research in AI, in particular in the chess domain, along two axes, human-compatible knowledge and machine-compatible processing. In this framework it became apparent that AI research has diverged into two streams that proceed along these axes, while we believe that the core of AI lies along the diagonal. For us, the reason of this development is that the effort of combining both aspects is considerably harder than the one-dimensional endeavors, and we have offered a hypothesis on research effort that could explain this observation. It follows that these aspects cannot be treated independently.

The development that has led to the Deep Blue vs. Kasparov match and to the first game that a machine has won against the human chess world champion has, in our opinion, demonstrated that proceeding along the "engineering" axis only, hard as it certainly has been, is comparably easy. Consequently, we have proposed a rewarding and challenging set of goals for research in computer chess that concentrates on an integration of knowledge and computability.

References

- Anderson, J. R. 1976. *Language, Memory, and Thought*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Bratko, I., and Michie, D. 1980. A representation of pattern-knowledge in chess endgames. In Clarke, M., ed., *Advances in Computer Chess 2*. Edinburgh University Press. 31-54.
- Chase, W. G., and Simon, H. A. 1972. The mind's eye in chess. In Chase, W., ed., *Visual Information Processing: Proceedings of the 8th Annual Carnegie Psychology Symposium*. New York: Academic Press. Reprinted in Collins (ed.), *Readings in Cognitive Science*, Morgan Kaufmann 1988.
- Collins, A. M., and Quillian, M. R. 1969. Retrieval time from semantic memory. *Journal of Verbal Learning and Verbal Behavior* 8:240-248.
- Condon, J. H., and Thompson, K. 1982. Belle chess hardware. In Clarke, M., ed., *Advances in Computer Chess 3*. Pergamon Press. 45-54.
- deGroot, A. D., and Gobet, F. 1996. *Perception and Memory in Chess*. Assen, The Netherlands: Van Gorcum.
- deGroot, A. D. 1965. *Thought and Choice in Chess*. The Hague: Mouton.
- Donniger, C. 1996. CHE: A graphical language for expressing chess knowledge. *ICCA Journal* 19(4):234-241.

⁴See *ICCA Journal* 15(4):235-236, 1992.

- Ebeling, C. 1987. *All the Right Moves: A VLSI Architecture for Chess*. The ACM Distinguished Dissertation Series. MIT Press.
- Fürnkranz, J., and De Raedt, L. 1997. Learning playing strategies from chess endgame databases: An ILP approach. In preparation.
- George, M., and Schaeffer, J. 1990. Chunking for experience. *ICCA Journal* 13(3):123-132.
- Gillogly, J. J. 1972. The technology chess program. *Artificial Intelligence* 3:145-163.
- Gobet, F., and Jansen, P. 1994. Towards a chess program based on a model of human memory. In Herik, H. J. v. d.; Herschberg, I. S.; and Uiterwijk, J. W. H. M., eds., *Advances in Computer Chess 7*. University of Limburg. 35-60.
- Gobet, F., and Simon, H. A. 1996. Templates in chess memory: A mechanism for recalling several boards. *Cognitive Psychology* 31:1-40.
- Gobet, F. 1993. A computer model of chess memory. In *Proceedings of the 15th Annual Meeting of the Cognitive Science Society*, 463-468.
- Hayes, J. R. 1987. Memory organization and world-class performance. In *Proceedings of the 21st Carnegie-Mellon Symposium in Cognition*.
- Holding, D. H. 1985. *The Psychology of Chess Skill*. Lawrence Erlbaum Associates.
- Hsu, F.-H. 1987. A two-million moves/s CMOS single-chip chess move generator. *IEEE Journal of Solid-State Circuits* 22(5):841-846.
- Hyatt, R. M.; Gower, B. E.; and Nelson, H. L. 1985. Cray Blitz. In Beal, D., ed., *Advances in Computer Chess 4*. Oxford, UK: Pergamon Press. 8-18.
- Kaindl, H. 1982. Positional long-range planning in computer chess. In Clarke, M. R. B., ed., *Advances in Computer Chess 3*. Pergamon Press. 145-167.
- Kerner, Y. 1994. Case-based evaluation in computer chess. In Keane, M.; Haton, J.; and Manago, M., eds., *Topics in Case-Based Reasoning (EWCBB-94)*, Lecture Notes in Artificial Intelligence. Berlin: Springer-Verlag.
- Kerner, Y. 1995. Learning strategies for explanation patterns: Basic game patterns with application to chess. In Aamodt, A., and Veloso, M., eds., *Proceedings of the 1st International Conference on Case-Based Reasoning (ICCB-95)*, Lecture Notes in Artificial Intelligence. Berlin: Springer-Verlag.
- Klimesch, W. 1994. *The Structure of Long-Term Memory: A Connectivity Model of Semantic Processing*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Levy, D. N., ed. 1988. *Computer Chess Compendium*. London: Batsford Ltd.
- Michie, D., and Bratko, I. 1991. Comments to 'chunking for experience'. *ICCA Journal* 18(1):18.
- Muggleton, S. 1988. Inductive acquisition of chess strategies. In Hayes, J. E.; Michie, D.; and Richards, J., eds., *Machine Intelligence 11*. Clarendon Press. chapter 17, 375-387.
- Newell, A.; Shaw, C.; and Simon, H. A. 1958. Chess playing programs and the problem of complexity. *IBM Journal of Research and Development* 2:320-335. Reprinted in (Levy 1988).
- Nunn, J. 1992. *Secrets of Rook Endings*. Batsford.
- Nunn, J. 1994a. Extracting information from endgame databases. In van den Herik, H. J.; Herschberg, I. S.; and Uiterwijk, J. W. H. M., eds., *Advances in Computer Chess 7*. University of Limburg. 19-34.
- Nunn, J. 1994b. *Secrets of Pawnless Endings*. Batsford.
- Nunn, J. 1995. *Secrets of Minor-Piece Endings*. Batsford.
- Opdahl, A. L., and Tessem, B. 1994. Long-term planning in computer chess. In Herik, H. J. v. d.; Herschberg, I. S.; and Uiterwijk, J. W. H. M., eds., *Advances in Computer Chess 7*. University of Limburg.
- Roycroft, A. J. 1988. Expert against oracle. In Hayes, J. E.; Michie, D.; and Richards, J., eds., *Machine Intelligence 11*. Oxford, UK: Oxford University Press. 347-373.
- Shannon, C. E. 1950. Programming a computer for playing chess. *Philosophical Magazine* 41(7):256-275. Reprinted in (Levy 1988).
- Simon, H. A., and Barenfeld, M. 1969. Information-processing analysis of perceptual processes in problem solving. *Psychological Review* 76(5):473-483.
- Simon, H. A., and Chase, W. G. 1973. Skill in chess. *American Scientist* 61(4):394-403.
- Simon, H. A., and Gilmartin, K. 1973. A simulation of memory for chess positions. *Cognitive Psychology* 5:29-46.
- Slate, D. J., and Atkin, L. R. 1983. Chess 4.5 — the northwestern university chess program. In *Chess Skill in Man and Machine*. Springer-Verlag, 2 edition. chapter 4, 82-118.
- Turing, A. M. 1953. Chess. In Bowden, B., ed., *Faster Than Thought*. London: Bitman. 286-295. Reprinted in (Levy 1988).
- Wilkins, D. E. 1980. Using patterns and plans in chess. *Artificial Intelligence* 14(3):165-203.
- Wilkins, D. E. 1982. Using knowledge to control tree search searching. *Artificial Intelligence* 18(1):1-51.
- Winkler, F.-G. 1991. Das Vernetzungsmodell von Klimesch - eine Simulation. Master's thesis, Institute for Medical Cybernetics and Artificial Intelligence, University of Vienna, Vienna, Austria. In German.
- Zobrist, A. L., and Carlson, F. R. 1973. An advice-taking chess computer. *Scientific American* 93-105.