

IJCAI 1991 Panel:
The Current and Future Role of Chess
in Artificial Intelligence and Machine Learning Research

by

Robert Levinson (Chairperson)
Computer and Information Sciences
Applied Sciences Building
University of California at Santa Cruz
Santa Cruz, CA 95064
(408)459-2087
ARPANET:levinson@cis.ucsc.edu

AND

Feng Hsiung-Hsu
PO Box 704
Yorktown Heights, NY 10598

AND

Jonathan Schaeffer
Dept. of Computing Science
University of Alberta
Edmonton, Alberta
Canada T6G 2H1

AND

T. Anthony Marsland
Computing Science Dept.
University of Alberta
Edmonton, Alberta
Canada T6G 2H1

And

David E. Wilkins
SRI International EJ227
333 Ravenswood Ave
Menlo Park, Ca 94025

in the Proceedings of the 1991 International Joint Conference on Artificial Intelligence,
(Sydney, Australia), pp. 547–552, 1991

Abstract

Our great researchers John McCarthy, Allen Newell, Claude Shannon and Herb Simon, Ken Thompson and Alan Turing have each put significant effort into computer chess research. It seems that now that computers have reached the grandmaster level and are beginning to vie for the World Championship the AI community should pause to evaluate what significance chess has to the evolving objectives of AI, what contributions have been made to date, and what can be expected in the future. Despite the general interest in chess amongst computer scientists and the significant progress in the last twenty years, there seems to be a lack of appreciation for the field in the AI community. On one hand this is the fruit of success (brute force works, why work on anything else?), but also the result of a focus on performance above all else in the chess community. Also, chess has proved to be too challenging for many of the AI techniques that have been thrown at it. We wish to promote chess as the unique testbed that our founding researchers recognized it to be and increase awareness of its contribution to date.

1 Panel Summary

The factors that make chess an excellent domain for AI research include the following:

- Richness of problem-solving situations.
- Accurate measures of progress through competition and ratings. Due to the well-structuredness is easy to monitor and record progress.
- Chess has been around for centuries - the basics are well-understood internationally, expertise is readily available and is (generally!) beyond proprietary or nationalistic interests. Has been considered a “game of intelligence.” Many players of the game feel mentally “stretched.”
- Detailed psychological studies of chess playing exist. These studies indicate that human players use reasoning modes quite different from the current chess implementations. Further, the reasoning modes are also used in many other problem-solving domains.
- Excellent test bed for uncertainty management schemes - the basis of most expert problem-solving. The well-definedness and discreteness of the game have led many to ignore this.

The above factors make chess a useful tool regardless of the strength of the current programs. Due to the success of the current methods there is a vast arena of other methods that have not been explored. The most obvious lack is in the application and development of machine learning techniques to chess, but other areas, including knowledge representation and compilation, planning and control, also seem to be applicable. AI researchers should be encouraged to use chess as a testbed for their techniques, with the understanding that chess is

not the end in itself. Chess may provide the avenue by which bridges may be built between cognitive science, AI and connectionist modeling.

With the current and future battle for the World Human-Computer Championship the AI community should be made more sensitive to the issues involved and their bearing on intelligence research: Is search sufficient? How much detailed chess knowledge is required? How is this knowledge implemented and incorporated with search? We are fortunate to have a World Champion who promotes creativity over the chess board and is willing to face the challenge from computers head-on.

The members of the panel and the presentations have been designed to address these topics in a way that we think will support our objectives to help make chess an important and respected AI tool in this new decade.

Jonathan Schaeffer will be speaking on the contributions of the current chess research to AI, especially in the area of search algorithms. He will emphasize those areas of computer chess research that have been ignored due to a competitive/engineering approach rather than a scientific one.

Feng-Hsiung Hsu of the Deep Thought team will discuss the role of knowledge in current chess programming and argue that more responsibility for the knowledge should be put on the machines themselves.

Tony Marsland will discuss specific open research issues in computer chess that will require AI solutions.

Robert Levinson will describe an alternative model of chess computation. Morph is a self-learning pattern-oriented chess program. In Morph's "cognitively-inspired" learning framework knowledge must be learned incrementally from experience, without many examples being stored (and very little guidance as to relevant features).

David Wilkins will provide balance to the discussion by pointing out the limitations of chess and claiming that Go is a better domain. He will also describe a new type of games tournament that prevents the human tailoring of evaluation functions and encourages the use of learning and more robust approaches.

The timing for this panel is particularly good with the current World Championship having completed (with the first half in New York), a more powerful Deep Thought on the scene, a recent article in Scientific American, and new books by M. Newborn and D. Levy [16], and T. Marsland and J. Schaeffer [17].

2 Presentations

2.1 Computer Chess: Science or Engineering?

Jonathan Schaeffer
University of Alberta

Research into artificial intelligence using chess as the application domain has produced several important contributions to AI:

- The power of brute-force search. Chess has clearly shown that simple brute-force approaches can go a long way.
- Iterative search. Some of the ideas developed for alpha-beta search, iterative deepening in particular, are applicable to other search domains.
- The inadequacy of conventional AI techniques for real-time computation. No competitive computer chess program uses AI languages or knowledge representation methods. Why? They are too slow for a real-time, high performance application.

Although these (and other, lesser contributions) have enhanced our knowledge, it is not clear whether the effort expended justifies the results obtained.

It is easy to question the usefulness of computer chess research. It is important to distinguish between computer chess research and research using chess as a testbed. Unfortunately, the latter has evolved into the former. An entirely new field of “computer chess” has evolved, with the emphasis on chess performance and chess research - not generally of much interest to the AI community. There is a much deserved credibility problem here. The unfortunate correlation between program speed and performance encourages short-term projects (speeding up a move generator 10%) at the sacrifice of long-term research projects (such as chess programs that learn).

After over 30 years of work on chess programs, where are the scientific advances in:

- knowledge-based search algorithms? Alpha-beta simplifies the program task, but the exponential search limits what can be achieved.
- error analysis? No one understands the interactions of search and knowledge and the errors they introduce.
- tool development? With the right tool, work that might take days could be done in minutes. No tools are being developed to help build chess programs. For example, why isn't someone working on tools for defining chess knowledge?

If the community were committed to research, many of these problems would have been addressed by now. Sadly, most of the work currently being done on computer chess programs is engineering, not science. For example, the engineering of special-purpose VLSI chips to increase the speed of a chess program only underlines the importance chess programmers attach to speed.

In my opinion, conventional computer-chess methods will yield little of further interest to the AI community. I believe they will be inadequate to defeat the human World Champion in a match for a long time to come. It is still very easy to set up a position that the computer has no

idea what is going on - even if you sped the machine up 1000-fold. The current computer chess work will only underscore the need for better ways of adding and manipulating knowledge reliably.

The defeat of the human World Chess Champion sooner rather than later will help artificial intelligence. This will help to reestablish chess as an ideal problem domain for experimenting with the fundamental problems of artificial intelligence.

Many of the opinions in this paper are elaborated in [7] and [17] pp. 259–268.

2.2 Designing an almost “knowledge-free” chess machine: “Expert inputs” are sometimes harmful

Feng Hsiung-Hsu
IBM

Experience from the chess machine Deep Thought indicates that inputs from chess experts, while generally useful, cannot be trusted completely. A good example of this is Deep Thought’s evaluation function. Several iterations of modifications by capable human chess experts had failed to produce significant improvements and occasionally even had a negative impact on the machine’s performance. Human experts, in this case, along with their expertise, introduced some of their own prejudices into the program. One way of solving this problem is to limit the type and the amount of expert inputs allowed into the program; in other words, having an almost “knowledge-free” machine. The availability of on-line high quality chess game databases makes this an attractive approach. Instead of having the value of, say, an isolated pawn set by human experts either explicitly or in functional form, one can simply tell the program that isolated pawns are important features and statistical procedures, with some additional expert inputs, can then be used to decide the functional form and the proper weighting of the features in question.

2.3 Open Problems and lessons for AI from computer chess

T. Anthony Marsland
University of Alberta

Based on predicted advances in computer technology, particularly the faster speeds and increasing memory of low cost systems, it is reasonable to assume that within the next decade the World Chess Champion will lose an informal game to a computer, and within twenty-five years lose a 12-game match. The early losses will reflect more breaks in concentration at first and later a recognition of the inevitable, as arose when trains started to out-pace runners or forklifts out-hoisted weightlifters. Although the defeat of humans by machines will be significant, it will mean neither the destruction of chess as a pass-time and learning medium, nor the

end of interest in computer chess per se. Rather it will focus attention even more sharply on precisely how and why humans can become so expert at selecting sound (often optimal) variations in seemingly complex situations, without resorting to the exhaustive techniques used by computer programs.

Some fundamental AI questions that will remain are:

- Given a patient and seemingly perfect teacher (that is a superior chess-playing machine), how should one use it to “teach” an AI-based learning program about strategies for playing chess (given that the rules of chess themselves are already perfectly known)?
- A related but perhaps simpler problem comes from the realm of endgame play. Given a perfect N-piece database holding an optimal move for each position (or even perhaps only the length of the optimal sequence from that position), develop a program which can deduce a sound set of rules or strategies for playing the endgame perfectly (or at least better than any other expert).
- Given endgame positions which cannot be solved by search or databases alone, deduce a plan or playing strategy that will transform the position into a known (win/draw) state. One class of positions of this type are where the remaining pieces are held to relatively few squares. Progress can only be made by a freeing move which converts a short-term loss of material (or perhaps position) to the achievement of a later, more significant goal. One example is from the Duchess-Chaos game (1979)[9] which is still thought to be beyond brute force search. Related examples abound, for instance giving up a passed pawn on one side of the board in order to win a pawn race on the other.
- Given a well-defined threat (for example mate) deep in the tree, identify un-examined moves at an earlier level along the current path which have the potential to explicitly deny the threat. This is a form of dynamic re-ordering of moves, but is also (if no potential denials exist) a good forward pruning criteria—providing evidence to abandon this line of play.
- Given a well-defined result involving a few pieces (e.g. an exchange of material), note the sphere of influence of the participants and by analogy determine those other moves that do not affect that sphere, and hence do not alter the significant outcome of the exchange. Again a forward pruning criterion.

The first two projects rely on perfect domain knowledge and the availability of an un-tiring teacher to whom questions can be posed. The need for convergence to a solution within some arbitrary or unreasonably short time-frame will thus be eliminated.

The learning mechanisms used will have the benefit of drawing on results obtained by exhaustive means, that is knowing that a solution exists. However it is clear that over the years humans have developed techniques that allow them to reduce the search space through judicious use of forward pruning (that is, by temporarily abandoning certain variations) and

either deducing by analogy that further consideration would be irrelevant, or (upon questioning the validity of the pruning) force reconsideration of the omitted lines. Pruning by analogy is a powerful general-purpose tool and if developed satisfactorily for a perfect information game like chess would almost certainly be applicable to related decision-tree searches. For example Horacek showed by deduction how the search space for a simple pawn endgame could be reduced [12].

The remaining three problems are linked closely to formal or probabilistic pruning methods. Computer chess is computationally expensive enough that one can afford to expend considerable time eliminating parts of the search space by deduction. On the one hand, the intent is to be more selective about variations that are to be expanded fully. Methods like the null-move heuristic [3], conspiracy number search and singular extensions [2] all do this by expanding non-quiet lines of play. On the other hand, the more formal probabilistic methods [23] attempt to limit the width of search at any node by estimating the probability that a better move exists in the moves that remain to be searched (so called Fischer set in [9]). In effect this problem requires looking again at the method of analogies [1]. It is remarkable that no significant improvement has been made to that method, despite the passage of 15 years. Not even attempts to implement simple forms of the idea in serious chess programs. The fundamental work here is to determine how best to make the method of analogies pay for itself. In this era of faster processors and parallel computation this must be an area that is ripe for exploration.

To these general methods one must explore more fully the power of simple rote-learning techniques. These have been shown to be effective [29] in both avoiding known losses in the opening and in extending the the effective search depth during play of replicated games. To this we should add the expanded use of other memory tables (typically hash-access tables for speed) to record or hold values for known pawn formations and so help in the preservation of long-term positional factors. A useful part of this work could be the playing of all known games and checking each move for weakness or unseen error. By this means an inventory of “innovations” or surprises can be built at modest cost. It is not trivial work because it requires making a plausible re-construction from imperfect data [18]. Finally, once errors are found, a backtracking mechanism will be needed to find the best place earlier in the game-tree path to correct (avoid) the flaw which follows.

One important lesson for the AI community is the importance of competitive testing and performance comparison of algorithms. In a sense 20 years of computer chess championships has provided a long-running series of experiments proving conclusively that progress has been made, identifying clearly those methods that have been effective and making a direct comparison from year to year possible. In principle theorem proving programs could be tested the same way, as indeed could language translation systems. These forms of comparison are standard for pattern recognition systems, why not for natural language understanding. In conclusion AI would benefit if more of its work were done on a direct competitive basis to identify more sharply those methods that are truly generally applicable.

2.4 Morph: An adaptive, pattern-oriented chess system

Robert Levinson

University of California

Although chess computers now are competitive at master and grandmaster levels, that is where their resemblance to human players ends. Psychological evidence indicates that human chess players search very few positions, and base their positional assessments on structural/perceptual patterns learned through experience. Morph is a computer chess program that has been developed to be more consistent with the cognitive models.

The main objectives of the project are to demonstrate capacity of the system to learn, to deepen our understanding of the interaction of knowledge and search, and to build bridges in this area between AI and cognitive science.

The current model of chess programming came into its own in the 70's and early 80's and has been refined ever since [32]. The main characteristic of the model is the use of brute-force alpha-beta minimax search with selective extensions for special situations such as forcing variations. This has been further enhanced by special purpose hardware. This model has been so successful that little else has been tried.

The alternative AI approaches have not fared well due to the expense in applying the "knowledge" that had been supplied to the system. Those times in recent years that chess has been applied as a testbed [8, 27, 19, 21, 22, 30, 33, 26, 20] only a small sub-domain of the game was used, so that fundamental efficiency issues that AI must grapple with have been largely unaddressed. However, we feel that there is a third approach that neither relies on search or the symbolic computation approach of knowledge-oriented AI: what we shall call the "pattern-oriented approach." In this approach configurations of interaction between squares and pieces are stored along with their significance. A uniform (and hence efficient method) is used to combine the significances in a given position to reach a final evaluation for that position. That such an approach is possible is evidenced by psychological models of human chess play [6, 25, 31]:

Morph¹ is a system developed over the past 3 years that implements the pattern oriented approach [14, 13]. It is not conceivable that the detailed knowledge required to evaluate positions in this way could be supplied directly to the system, thus learning is required.

To strengthen the connections with the cognitive literature the system's knowledge is to come from its own playing experience, no sets of pre-classified examples are given and beyond its chess pattern representation scheme little chess knowledge such as the fact that having pieces is valuable (leave alone their values) has been provided to the system. Further, the system is limited to using only 1-ply of search.²

¹The name "Morph" comes from the Greek *morph* meaning form and the chess great, Paul Morphy.

²Though nothing in the method except perhaps efficiency, prevents deeper search.

3 System Design

Morph makes a move by generating all legal successors of the current position, evaluating each position using the current pattern database and choosing the position that is considered least favorable to the opponent. The system is designed so that after each game patterns are created, deleted and generalized and weights are changed to make its evaluations more accurate (in its view) based on the outcome of the game.

For a complete report on the system, please see [15].

There are two types of patterns stored in the Morph system. In addition to graph patterns (which represent attacks and defends relationships between pieces and squares) Morph stores “material” patterns—vectors that give the relative material difference between the players, e.g. “up 2 pawns and down 1 rook,” “even material,” etc.

Along with each pattern is stored a weight that reflects the significance of the pattern. A weight is a real number in $[0,1]$ that is stored with each pattern as an estimate of the expected true minmax evaluation of states that satisfy the pattern.

3.0.1 Results with Morph

There have been many encouraging signs in the three months since Morph was fully implemented. For details see [15], attend the panel and write the author.

3.0.2 Relationship of Morph to other approaches

The chess system combines threads of a variety of machine-learning techniques that have been successful in other settings. It is this combination and exactly what is done to achieve it that is the basis for Morph’s contributions. The learning-method areas and their involvement in Morph include genetic algorithms [10], neural nets (weight updating) [28], temporal-difference learning, explanation-based generalization (EBG), and similarity-based learning. To combine these methods some design constraints usually associated with these methods are relaxed.

In summary, in addition to Morph’s unique combination of methods, what distinguishes it are:

- A uniform representation of search knowledge.
- A syntactic approach to playing and learning.
- An attempt to play a complete game of chess rather than a small subdomain.
- Rejection of a learning-by-examples framework for an experiential framework that is more cognitively-inspired.

- Responsibility for feature discovery given to the system.
- Non-reliance on search (though at some point a small amount of guided search may be incorporated, bringing us even closer to the cognitive model).

3.1 Chess Was Good for AI Research.

David E. Wilkins
SRI International

Over the years, chess has proven to be a fertile ground for ideas and techniques that have spread to other areas of AI. These include database enumeration techniques [4], chunking [5], search techniques (minimax, alpha-beta, iterative deepening), and the utility of information [11]. Considering the lack of funding for chess, it is significant that it has produced so many results.

Chess has been fertile because it provides a complex reasoning problem from a simple domain with a built-in performance criteria. The simple domain permits research to progress with little initial overhead. Having a hostile opponent adds complexity to the reasoning. In many domains (natural language understanding comes to mind), progress can be hindered by lack of performance criteria — it can be hard to tell whether or not the latest thesis is an improvement on the current state of the art. Chess provides precise answers to performance questions.

However, hardware advances have made chess a less fertile ground for addressing the basic issues of AI. The game is small enough that brute-force search techniques have dominated competitive computer chess, and I see little AI interest in squeezing out the last few hundred points on the chess ratings, except for the psychological impact of having a computer beat the human world champion.

Obviously, many basic issues in AI are not naturally addressed in a game-playing environment and should be explored in other domains. These include communication, forming models of one's environment, sensor analysis and integration, and (perhaps) reasoning about uncertainty. In addition, real-world domains force AI researchers to address issues such as economy of scale, noise, real-time response, failed actions, novel phenomena, and multiple agents — issues that can be ignored in chess.

Of the AI areas well-suited to a game-playing domain, there are better domains than chess. In particular, Go has all the advantages of chess but provides more complex reasoning and an even simpler domain. A successful symbolic Go program would have to plan, would have to use goal-directed search, would encourage machine learning, and would promote visual reasoning — all basic AI research issues that are ignored in competitive computer chess.

Even better than Go may be an event where programs compete against each other, but are given a description of the game to be played at the beginning of the match. [24]

Chess is particularly well suited to modification, so one could, for example, have a competition on a chess board with chess pieces where the match begins by giving the programs a declarative statement of how the pieces move, how they capture, the initial position, and what the objective of the game is. The programs would have to play this newly defined game under time constraints. A longish series of games could be required. This would require machine learning and a robust symbolic problem-solving capability that is not tailored to a specific game. For each new game, the programs would have to learn evaluation functions (if needed), learn what goals are advantageous to attempt, and learn heuristics or features for selecting moves. Brute force techniques would be disadvantaged by the lack of opportunity to fine tune an evaluation function and a quiescence search to the game at hand.

The author is indebted to Barney Pell for many of the ideas here.

References

- [1] G. Adelson-Velsky, V. Arlazarov, and M. Donsky. Some methods of controlling the tree search in chess programs. *Artificial Intelligence*, 6(4):361–371, 1975.
- [2] T. Anantharaman, M. S. Campbell, and F. H. Hsu. Singular extensions: Adding selectivity to brute force searching. *Journal of International Computer Chess Association*, 11(4):135–143, 1988.
- [3] D. Beal. Experiments with the null move. In *Advances in Computer Chess 5*, pages 65–79. Elsevier, 1989.
- [4] I. Bratko, D. Kopek, and D. Michie. Pattern-based representation of chess end-game knowledge. *Computer Journal*, 21(2):149–153, May 1978.
- [5] Murray S. Campbell. *Chunking as an Abstraction Mechanism*. PhD thesis, Carnegie Mellon University, 1988.
- [6] A. D. de Groot. *Thought and Choice in Chess*. The Hague, 1965.
- [7] M. Donskoy and J. Schaeffer. Perspectives on falling from grace. *Journal of the International Computer Chess Association*, 12(3):259–268, 1989.
- [8] N. S. Flann and T. G. Dietterich. A study of explanation-based methods for inductive learning. *Machine Learning*, 4:187–226, 1989.
- [9] P. W. Frey, editor. *Chess Skill in Man and Machine*. Springer-Verlag, 1983.
- [10] D. E. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley, Reading, MA, 1989.
- [11] I. J. Good. Dynamic probability, computer chess, and the measurement of knowledge. In E. W. Elcock and D. Michie, editors, *Machine Intelligence 8*, pages 139–150. Ellis Horwood, Chichester, 1977.

- [12] H. Horacek. Knowledge based move selection and evaluation to guide the search in chess pawn endings. *Journal of International Computer Chess Association*, 6(3):20–37, 1983.
- [13] R. Levinson. Pattern formation, associative recall and search: A proposal. Technical Report UCSC-CRL-89-22, University of California at Santa Cruz, 1989.
- [14] R. Levinson. A self-learning, pattern-oriented chess program. *International Computer Chess Association Journal*, 12(4):207–215, December 1989.
- [15] R. Levinson and R. Snyder. Adaptive pattern oriented chess. In *Proceedings of the Ninth National Conference on Artificial Intelligence*, 1991.
- [16] David Levy and Monty Newborn. *How Computers Play Chess*. Computer Science Press, New York, NY, 1991.
- [17] T. Anthony Marsland and Jonathan Schaeffer, editors. *Computers, Chess and Cognition*. Springer-Verlag, 1990.
- [18] T.A. Marsland. Workshop report: Theory and practice in computer chess. *ICCA*, 10(4):205–210, 1987.
- [19] R. S. Michalski and P. Negri. An experiment on inductive learning in chess end games. In E. W. Elock and D. Michie, editors, *Machine Representation of Knowledge, Machine Intelligence*, volume 8, pages 175–192. Ellis Horwood, 1977.
- [20] S. Minton. Constraint based generalization- learning game playing plans from single examples. In *Proceedings of AAAI*. AAAI, 1984.
- [21] T Niblett and A. Shapiro. Automatic induction of classification rules for chess endgames. Technical Report MIP-R-129, Machine Intelligence Research Unit, University of Edinburgh, 1981.
- [22] P. O’Rorke. A comparative study of inductive learning systems AQ11 and ID3. Technical Report Intelligent Systems Group Report No. 81-14, Department of computer Science, University of Illinois at Urbana-Champaign, 1981.
- [23] A.J. Palay. *Searching with Probabilities*. Pitman, 1985.
- [24] Barney Pell. A computer game-learning tournament. (In Preparation), 1990.
- [25] H. Pflieger and G. Treppner. *Chess: The Mechanics of the Mind*. The Crowood Press, North Pomfret, VT, 1987.
- [26] J. Pitrat. A program for learning to play chess. In *Pattern Recognition and Artificial Intelligence*. Academic Press, 1976.
- [27] J. R. Quinlan. Learning efficient classification procedures and their application to chess end games. In R. S. Michalski, J. G. Carbonell, and T. M. Mitchell, editors, *Machine Learning*. Morgan Kaufmann, San Mateo, CA, 1983.

- [28] E. D. Rumelhart and J. L. McClelland. *Parallel Distributed Processing*, volume 1–2. MIT Press, 1986.
- [29] T. Scherzer, L. Scherzer, and D. Tjaden. Learning in Bebe. In T. A. Marsland and J. Schaeffer, editors, *Computer, Chess and Cognition*, chapter 12, pages 197–216. Springer-Verlag, 1990.
- [30] A. D. Shapiro. *Structured Induction in Expert Systems*. Turing Institute Press with Addison-Wesley, 1987.
- [31] H. A. Simon and K. Gilmartin. A simulation of memory for chess positions. *Cognitive Psychology*, 5(1):29–46, 1973.
- [32] D. J. Slate and L. R. Atkin. Chess 4.5—the northwestern university chess program. In P. W. Frey, editor, *Chess Skill in Man and Machine*, pages 82–118. Springer-Verlag, 1977.
- [33] P. Tadepalli. Lazy explanation-based learning: A solution to the intractable theory problem. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, Detroit, MI, 1989. Morgan Kaufmann.