

## **UC Merced**

### **Proceedings of the Annual Meeting of the Cognitive Science Society**

#### **Title**

Skilled like a Person: A Comparison of Human and Computer Game Playing

#### **Permalink**

<https://escholarship.org/uc/item/5xv839dm>

#### **Journal**

Proceedings of the Annual Meeting of the Cognitive Science Society, 17(0)

#### **Authors**

Rattermann, Mary Jo

Epstein, Susan L.

#### **Publication Date**

1995

Peer reviewed

# Skilled like a Person: A Comparison of Human and Computer Game Playing

**Mary Jo Rattermann**  
Cognitive Science and Cultural Studies  
Adele Simmons Hall  
Hampshire College  
Amherst, MA 01002  
mjr@neural.hampshire.edu

**Susan L. Epstein**  
Department of Computer Science  
Hunter College and The Graduate School of the  
City University of New York  
695 Park Avenue  
New York, NY 10021  
epstein@roz.hunter.edu

## Abstract

The subject of this paper is the role of transferable commonsense principles in the acquisition of game-playing expertise. We argue that individuals skilled in a domain develop expertise because they know and apply these principles, and that most game-playing programs do not play like people. The paper describes Hoyle, a model of an expert game player that relies on the use of commonsense principles, limited memory, and useful knowledge to learn to play two-person, perfect information finite-board games expertly. We then describe an experiment in which human subjects played three such games against a computer expert. After playing these games, the subjects evaluated Hoyle's game-playing principles in the context of their own behavior. Verbal protocols and subjects' evaluations revealed considerable overlap between the principles preferred by our subjects and those preferred by Hoyle. Using learning time as a measure of difficulty, the subjects' performance and Hoyle's performance ordered the three games identically. This experiment also revealed differences in the use of game-playing principles between skilled and unskilled players: skilled players judged the game-playing principles to be more effective than did unskilled players, skilled players used several different principles while unskilled players relied on one principle, and skilled players anticipated their opponent's moves while unskilled players merely reacted.

## Introduction

The use of transferable commonsense principles as a hallmark of expertise is the subject of this paper. It is the thesis of this work that an individual skilled in a particular domain is able to develop expertise more thoroughly and more rapidly because she knows and applies a set of commonsense principles. Our research is based in the domain of two-person, perfect information, finite-board games. We argue that most game-playing programs do not play like people, and present a model of expert game playing called Hoyle, which implements commonsense game-playing principles. We describe an experiment that identifies some of the commonsense principles people use, and offers

psychological evidence that the ability to use and evaluate commonsense principles is one of the differences between skilled and unskilled human game-players. Further, we discuss the plausibility of Hoyle as a cognitive model of human expert game playing.

Throughout this paper, *game* denotes a set of rules, playing pieces, and a board, while *contest* distinguishes a single experience playing a game. A *state* is a situation in a game, described completely by the location of the playing pieces on the board and whose turn it is to move. For example, tic-tac-toe is a game at which two people might play a contest, beginning with a state that is an empty board with X to move. The *game tree* for a game is the space of all its possible states, connected by the moves that lead from one state to another. Finally, a *commonsense strategy* is a generally applicable predilection to select an alternative because it satisfies a particular criterion.

## How Human Experts Play Games

Contrary to popular belief, experts do not have distinctive mental abilities (like exceptional powers of concentration, enormous memories, or high IQ's), nor do they perform extensive forward search into the game tree, nor rely on statistical measures of typicality or concrete visual images (Binet, 1894; Charness, 1981; Djakow, Petrowski, & Rudik, 1927; Holding, 1985). What they do have are perceptual focus of attention, carefully organized knowledge, and efficient procedures to manipulate that knowledge.

In particular, game-playing experts summarize some of their knowledge in concepts, both as verbal memories and as unordered spatial patterns called *chunks*. A grandmaster's recall is better than an ordinary person's for chess positions, but only for chess positions that are meaningful, i.e., ones that would arise during the play of a contest (Chase & Simon, 1973). In general, empirical evidence indicates that skilled people in many domains have better memories, but only for meaningful patterns, and that, given the same knowledge, an unskilled person remains unskilled. (See, for example, (Egan & Schwartz, 1979; Engle & Bukstel, 1978; Shneiderman, 1976).) An expert's memory is organized around higher-level concepts, *not* perceptual characteristics,

so that a single board configuration that could arise in either of two different games is chunked and recalled differently, depending upon the game in which it is perceived (Eisenstadt & Kareev, 1975). Finally, chess masters consider both their possible moves and the possible moves of their opponents, and use these plans to direct the search process, making it simpler and more efficient (Holding, 1985).

### How Computer Experts Play Games

The typical game-playing program is very different from the human expert. It plays only one game, e.g., checkers or Go. With a few notable exceptions, such a program does not learn or plan or retain its experience, and it cannot explain its decisions. It relies on large, perfect memory and fast, deep search. The memory describes standard *openings* (early move sequences favored by human experts) and may also include a database indicating the best possible result from many states in the endgame. The program plays the relevant opening as long as it is applicable, and applies the endgame knowledge as soon as possible.

In between, during the middlegame, the program calculates a numerical description of the worth of many possible subsequent states and selects the move that maximizes the mover's potential strength, while minimizing that of her opponent. Each state is rated by the program's heuristic *evaluation function*, a computation based upon game-dependent features selected by the programmer from among those highly regarded by human experts. A *feature* is a description, like "the black king is in check," that is considered a significant component of a static judgment on the strength of any state relative to all other possible states. The typical game-playing program searches as fast and as deep as it can, evaluating encountered states and backing up their values to select the best next move. If a game-playing program were supplied with the ideal evaluation function, or if it could always search to the end of possible contests where the rules determined who had won, then the program would play perfectly. Otherwise, its decision is simply a heuristic approximation, one which selects the current move that maximizes the program's estimated ability to win.

Although this type of massive memory and search program produces an impressive game player, even the best of them are often defeated by the best human players (Anantharaman, Campbell, & Hsu, 1990; Berliner & Ebeling, 1989). Further, the necessary memory requirements and speed of the search algorithm makes this kind of game-playing program psychologically implausible; people simply cannot, and do not, play this way.

### How Hoyle Plays Games

Hoyle is a program that learns to play two-person, perfect information, finite-board games. It is based on a learning and problem-solving architecture called FORR, predicated upon multiple rationales for decision making (Epstein, 1994). The philosophy behind FORR is that a set of possibly overlapping, approximately correct general principles for decision making can collaborate to achieve a pragmatic synergy. The individual reasons identified for a particular skill like game playing are collected, rather than structured,

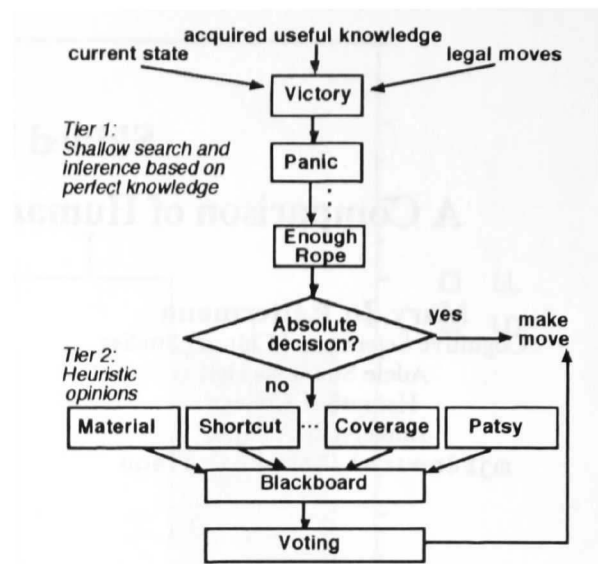


Figure 1: How Hoyle makes decisions.

to permit flexibility in a wide range of tasks. Hoyle learns to play in competition against a hand-crafted, external expert program for each specific new game. Hoyle begins with only the rules for a new game and some general knowledge about game playing, such as taking turns and stopping when someone has won. As seen in Figure 1, whenever it is Hoyle's turn to move, a hierarchy of resource-limited procedures called *Advisors* is provided with the current game state, the legal moves, and any useful knowledge (described below) already acquired about the game. Hoyle has 23 heuristic Advisors, or commonsense principles for game playing, in two tiers. The first tier sequentially attempts to compute a decision based upon correct knowledge, shallow search, and simple inference, such as Victory's "make a move that wins the contest immediately." If no single decision is forthcoming, then the second tier collectively makes many less reliable recommendations based upon narrow viewpoints, like Material's "maximize the number of your markers and minimize the number of your opponent's." Based on the Advisors' responses, a simple arithmetic vote selects a move. During play, Hoyle searches no more than two ply (one move for each contestant) ahead in the contest. Further details on Hoyle are available in (Epstein, 1992).

A FORR-based program learns from its experience to make better decisions based on acquired useful knowledge. *Useful knowledge* is expected to be relevant to future decisions and is probably correct in the full context of the search space. Each item of useful knowledge is associated with at least one learning algorithm, e.g., explanation-based learning or induction. The learning algorithms are highly selective about what they retain, may generalize, and may choose to discard previously acquired knowledge. Examples of Hoyle's useful knowledge include good openings, *significant states* (roots of subtrees whose game-theoretic value is a win for some contestant), and *forks* (overlapping plans that are a function of both the rules and the topology of the game board). Individual Advisors can apply current useful knowledge to construct their recommendations. For example, Open tries to reproduce only early play that has been successful in previous experience, Victory supports

moves to significant win states, and Pitchfork exploits forks both offensively and defensively.

Unlike more traditional game-playing programs, Hoyle does not rely upon massive memory, nor does it perform extensive searches in the game tree during play. Rather, it learns from its previous experience, looks no more than one move ahead for each contestant, uses very little memory, and applies commonsense principles (its Advisors) to the game at hand. Further, and most importantly, Hoyle can become expert in many different games, an ability not possessed by programs designed around vast memory and massive search of game-specific features. Hoyle is also an effective game player; it has learned to play 18 different games expertly in competition with a human or programmed opponent. Hoyle's learning time averages as few as 14 contests and rarely exceeds 100 contests (Epstein, 1993).

Hoyle's primary components (the Advisors, shallow search, and useful knowledge) are more psychologically plausible than those of other expert game-playing programs. Thus, the structure and logic of Hoyle seem ideally suited to guide the study of expertise in humans. Further, several of the board games Hoyle learns are simple enough to be mastered by a subject within an experimental session. Thus, the goal of the following experiment is to investigate whether the process used by Hoyle to become an expert is similar to that used by humans as they become experts.

## The Experiment

### Subjects

The subjects were 8 college students, recruited from the general population and paid for their participation.

### Procedure

The subjects were told they would be playing three different games against a computer expert. They were also told that they were being tape recorded, and that they should verbalize their thought processes while playing the games. To remind the subjects of the importance of verbalization during play, the words "Think out loud!" were displayed next to the game boards. The order of presentation of the three games was counterbalanced across subjects, and for each game the subjects and the computer expert alternated moving first.

The subjects were not told about Hoyle's Advisors until after playing against the computer. They played each game until they reached a criterion of 7 consecutive draws or 25 contests, after which they were switched to the next randomly assigned game. After playing all three games, the subjects were given a list of Hoyle's Advisors and asked to rate their effectiveness, frequency of use, and order of use.

### Materials

Three two-person, perfect information, finite board games: tic-tac-toe, lose tic-tac-toe and achi, were adapted for play on a Macintosh Quadra. Their game boards appear in Figure 2. Each of the programs was an expert, i.e., it was designed to play perfectly, making the best possible move in any state. If there was more than one such move the expert program selected one at random to provide a broad range of high-quality experience. Between two perfect players each of these games is a draw, but the goal in play is always to win while simultaneously preventing the other contestant from doing so. Tic-tac-toe is the classic 3 x 3 grid game where a win is three of your markers (either X's or O's) in a row either vertically, horizontally or diagonally. Lose tic-tac-toe is a variant of tic-tac-toe, in which three of your markers in a row loses the game. Achi is played on the board shown in Figure 2, which is isomorphic to the tic-tac-toe grid. Contestants begin with 4 markers, either white or black, and alternate placing them on the board. Play begins when a player places a marker at the intersection of two or more lines. There are nine such positions (corresponding to the nine positions on a tic-tac-toe grid). Once all eight markers have been placed on the board, a turn consists of moving one's own marker to the single empty position. A win at achi is three of your markers in a row, either vertically, horizontally or diagonally. Play ends in a draw when the same state is cycled through an agreed upon number of times. For Hoyle and our program, that number was 3.

The subjects were also asked to fill out three questionnaires designed to examine their use of 17 of Hoyle's Advisors and their definitions, as in Table 1. (Six Advisors for more complex games were deemed irrelevant in the context of these games and therefore omitted.) Each subject was asked to rate on a scale of 1 to 5 for each game the effectiveness of each Advisor, how frequently they had used each Advisor, and the order in which they had used them.

## Results

We used a contrastive analysis technique to divide the subjects into two groups, skilled and unskilled game players, based on the total number of contests necessary to reach criterion (Chi, Bassok, Lewis, Reimann, & Glaser, 1989). The four skilled game players reached criterion or played 25 contests in all three games after a total of 44 to 51 contests; the four unskilled game players reached criterion or played 25 contests in all three games after a total of 57 to 88 contests. Of the four skilled game players, all reached criterion in tic-tac-toe, three in achi, and two in lose tic-tac-toe. Of the four unskilled game players, three reached criterion in tic-tac-toe, one in achi, and none in lose tic-tac-toe.




Figure 2: The game boards for achi and tic-tac-toe.

Table 1: Hoyle's Advisors and their definitions as renamed and described to the subjects.

Advisor	Definition
Blinders	Select a move to further a simple plan, with no regard for your opponent's plans.
Defense	Select a move to defend against your opponent's simple plans.
Don't Lose	Do not make a move that will result in a loss.
Enough Rope	If your opponent would have a losing move on this board, avoid blocking it, i.e., leave the opponent the opportunity to "hang itself."
Fork	Chose a move giving you more than one opportunity to win, while blocking moves giving your opponent more than one opportunity to win.
Greedy	Make moves that advance more than one winning configuration.
Lookahead	"Look" two moves ahead, and based on this choose your next move.
Panic	If your opponent has a winning move on this board, block it.
Resign	If this is a certain loss, resign.
Victory Again	If you can find a winning move, take it. Repeat winning or drawing moves you have made.
Not Again	Do not repeat previously losing moves you have made.
Copycat	Mimic opponent's winning or drawing moves.
Leery	Avoid moves that your intuition tells you have led to a loss in the past.
Patsy	Reproduce the visual pattern made by the pieces of the winning or drawing player, and avoid the visual patterns made by the pieces of the losing player.
Start	Repeat previously successful observed opponent's opening.
Win	If you remember a certain win, make the winning move.

### Subject Protocols

The goal of this analysis was to determine if the game-playing principles (defined as the procedures used to optimize performance) verbalized by our subjects revealed the same principles embodied by Hoyle's Advisors. To this end, the protocols produced by the subjects during play were first transcribed and any game-playing principles were noted. These principles were then compared to Hoyle's Advisors, and the frequency of Hoyle's Advisors within the subject's game-playing principles was recorded. During the coding procedures the judge was blind to the skill status of the subject. This analysis revealed a remarkable overlap between Hoyle's Advisors and the game-playing principles verbalized by our subjects. Of the 17 relevant Advisors, 10 appeared in the subjects' protocols at least 3 times, suggesting that the

Advisors used by Hoyle are also the game-playing principles used by our subjects. Differences between the manner in which skilled and unskilled players used these principles were also evident. First, skilled game players tended to verbalize several different types of principles when playing against the computer, verbalizing an average of five different Advisors during their play, while unskilled players verbalized an average of three different Advisors. Second, the skilled players used their game-playing principles much more often than the unskilled players -- an average of 16 times, in contrast to an average of 10 times for the unskilled players. Finally, the particular principles verbalized by the two groups also differed; the skilled players verbalized "Lookahead" 25% of the time, and relied on "Defense" 17%, "Again" 11%, and "Patsy" 11%. The unskilled players used "Defense" quite often (30%), and were also very likely to use "Patsy" as a strategy, mentioning it over 17% of the time.

### Subject Ratings – Effectiveness

The effectiveness ratings revealed that the skilled players rated the Advisors as significantly more effective than did their unskilled counterparts. This was confirmed by one-way analysis of variance (Skilled versus Unskilled) performed on the subject's effectiveness ratings, collapsed across Advisor type, ( $F(1,28) = 6.93, p < .05$ ). Given that the skilled players rated the Advisors as more effective overall, we next asked whether there were differences between players of different skill levels within each game, and further, if those differences applied for specific Advisors. A series of item analyses (T-tests,  $\alpha < .10$ ) comparing the effectiveness ratings of skilled and unskilled players for each Advisor revealed that in tic-tac-toe, the skilled and unskilled players differed in their ratings of only two Advisors, in lose tic-tac-toe they differed in their ratings of five Advisors, and in achi they differed in their ratings of seven Advisors.

### Subject Ratings – Frequency

Unlike the effectiveness ratings, there were no significant differences between the skilled and unskilled players in their overall frequency of use ratings for the Advisors. (This was confirmed by one-way analysis of variance, Skilled versus Unskilled,  $F(1, 28) = 2.09, p > .10$ ) Item analyses, however, revealed a difference between skilled and unskilled players for 10 of Hoyle's Advisors, with the skilled players rating the frequency of use significantly higher for six of the Advisors, and the unskilled players rating the frequency of use significantly higher for the remaining four. Interestingly, one of the Advisors rated higher by the unskilled players was Patsy, a result mirrored by the analysis of their protocols in which Patsy was also used more often by the unskilled players.

### Comparing Hoyle to Humans

The analysis of the subjects' protocols revealed that the majority of the principles used by Hoyle are in fact used by humans when acquiring expertise in game playing. Further evidence is provided by the comparison of the Tier 2 Advisors Hoyle finds most *significant* (consistently supportive of expert decisions) and *relevant* (makes

recommendations) to the Tier 2 Advisors used most often by our subjects, as reflected by their protocols. (Because Hoyle treats its Tier 1 Advisors as always significant, only those in Tier 2 were analyzed here.) This analysis revealed that two of the three Advisors most frequently used by our subjects, Defense and Again, were also rated as among the most significant by Hoyle for all three games used in this experiment. Further analysis performed on the subjects' ratings of the Advisors for frequency of use and effectiveness also revealed a significant overlap in the use of Defense and Again.

## Discussion

The results of this experiment support our conjecture that Hoyle is a plausible model of human game playing. The analysis of the subjects' protocols revealed a significant overlap between the Advisors used by Hoyle and the game-playing principles used by our subjects. Further, there was a significant overlap between the Tier 2 Advisors that Hoyle found significant and relevant and the subjects' ratings of these Advisors as highly effective and frequently used. This finding supports the proposal that the commonsense principles Hoyle finds useful when learning to play a game are also the principles humans find useful. A further commonality between Hoyle and our human subjects is found by analyzing the number of contests required to learn to play each of the three games. This analysis revealed that for both Hoyle and our subjects lose tic-tac-toe was the most difficult game to learn (196 contests for our subjects), followed by achi (153), and then tic-tac-toe (99). Hoyle reached criterion in these games considerably faster (X contests for lose tic-tac-toe, X contests for achi and 16 contests for tic-tac-toe) but revealed the same ordering of difficulty.

The similarity between Hoyle's performance and human performance has implications for computer modeling of human expertise. Although the use of massive memory and fast search algorithms is a very effective way to produce expert performance in computers, such programs do not readily transfer to new games, nor is their dependence on extensive search psychologically valid. This research suggests that Hoyle, with its use of commonsense game-playing principles, limited memory, and useful knowledge, is a plausible model of human performance. Thus, Hoyle is psychologically plausible, is as effective in learning new games as the traditional computer models of expertise (Epstein, 1992), and has the ability to transfer its skill to new games.

The use of Hoyle's Advisors as a tool to study the acquisition of expertise in humans has revealed several differences between skilled and unskilled game players, supported by the subject's comments before they were told about Advisors:

- Skilled players use a variety of different game-playing principles and use them frequently when learning to play a new game. Unskilled players do not use these principles as frequently, and when they do, tend to repeat the same principles. This conclusion is supported by an unskilled player who noted, "(My) tendency is to have three

*strategies.... retain about two of them and then screw up on (the other)."*

- Skilled players rated Hoyle's Advisors as significantly more effective overall than did unskilled players, suggesting that they recognize the usefulness of game-playing principles while unskilled players do not. This appreciation of the effectiveness of these principles may be a component of skilled game playing; when learning a new game the skilled player may be more likely to look for, and recognize, beneficial principles when they occur during the course of a contest. In contrast, the unskilled player may not look for, or notice, beneficial principles, and thus will not have a repertoire of principles at his disposal. This conjecture is supported by the finding that skilled players used many different Advisors while unskilled players tended to use one or two Advisors repeatedly.

- The analysis of the subjects' protocols suggests a difference in the playing styles of skilled and unskilled players. Specifically, skilled players were likely to "Lookahead" several moves and choose their plays based on their analysis. Unskilled players, in contrast, were likely to respond with a good "Defense" to the computer's offensive moves. This is particularly evident in an unskilled subject's statement: "*at this point I don't even have a strategy.... now I'm just kind of responding to its moves.*" This suggests that as expertise is acquired the focus shifts from reacting to the behavior of your opponent to anticipating its moves and playing accordingly.

- The protocols revealed that skilled players placed Patsy in the top four Advisors used when learning these games. Hoyle too found Patsy to be a highly relevant Advisor, ranked approximately the same way as skilled players value it. (Further work with FORR suggests that later there is a shift back to pattern-based reactivity, one that was neither anticipated nor measured for in this research (Epstein & Gelfand, 1995).)

- Item analyses of the effectiveness ratings revealed that the differences between the ratings of skilled and unskilled players increased as the games became less familiar, with the skilled players rating the Advisors as more effective. This suggests that as the games became less familiar skilled players relied more on Advisors than did the unskilled players.

- An analysis of the number of contests to criterion revealed that the easiest game to learn was tic-tac-toe, followed by achi (the most unfamiliar game), and lastly, by lose tic-tac-toe. This result suggests a role for analogical transfer between achi and tic-tac-toe -- two different games that share isomorphic rule structures. This conjecture is supported by statements made by two of the skilled game players: "*In the first game I didn't know it was tic-tac-toe because.... I was looking for these squares and I didn't think there was a place like it in the corners,*" and "*it's sort of like a polka dot version of tic-tac-toe.*"

Finally, some of Hoyle's Advisors have also been found to be valid for extremely novice game players: six- to nine-year-old children. Crowley and Siegler (1993) found that when playing tic-tac-toe, grade school children used principles such as "Win," "Fork," and "Block." While the children could flexibly use these principles in response to

task demands, they also tended to adhere to a set order of use. Both of these characteristics are similar to the performance of Hoyle, and suggest that, with the proper adjustments, Hoyle could be modified to play like a child. There is also evidence that people use FORR-like Advisors in chip design (Biswas, Goldman, Fisher, Bhuva, & Glewwe, 1995).

### Conclusion

Previous work (e.g., Epstein, 1993) has shown that Hoyle is an effective computational game-player that efficiently becomes an expert in several different games. This research supports the conjecture that Hoyle, a game learning program that uses commonsense game-playing principles, limited memory, and useful knowledge, is also a plausible model of the acquisition of expertise in humans. Additionally, this work reveals differences between skilled and unskilled players: skilled players use a variety of game-playing principles and use them frequently while unskilled players rely on one or two principles, skilled players anticipate their opponent's moves while unskilled players merely react, and skilled players may notice relational similarity and use that similarity to their advantage.

### Acknowledgments

This work was supported in part by a Hampshire College Faculty Development Grant to the first author and NSF grant #9001936 to the second author. We wish to thank Mike O'Neal for testing the human subjects, and Lee Spector and Jason Juneau for designing and implementing the expert game-playing programs used in this work.

### References

- Anantharaman, T., Campbell, M. S., & Hsu, F.-h. (1990). Singular extensions: Adding selectivity to brute-force searching. *Artificial Intelligence*, 43, 99-10.
- Berliner, H. & Ebeling, C. (1989). Pattern knowledge and search: The SUPREM architecture. *Artificial Intelligence*, 38, 161-198.
- Biswas, G., Goldman, S., Fisher, D., Bhuva, B., & Glewwe, G. (1995). Assessing design activity in complex CMOS circuit design. In P. Nichols, S. Chipman, & R. Brennan, (Eds) *Cognitively diagnostic assessment*. Hillsdale, NJ: Erlbaum.
- Binet, A. (1894). *Psychologie des grands calculateurs et joueurs d'échecs*. Paris: Hachette.
- Charness, N. (1981). Search in chess: Age and skill differences. *Journal of Experimental Psychology: Human Perception and Performance*, 7, 467-476.
- Chase, W. G. & Simon, H. A. (1973). The mind's eye in chess. In W. G. Chase (Ed.), *Visual information processing* (pp. 215-281). New York: Academic Press.
- Chi, M. T. H., Bassok, M., Lewis, M. W., Reimann, P. & Glaser, R. (1989). Self-explanations: How students study and use examples in learning to solve problems. *Cognitive Science*, 13, 145-182.
- Crowley, K. & Siegler, R. S. (1993). Flexible strategy use in young children's tic-tac-toe. *Cognitive Science*, 17, 531-561.

- Djakow, I. N., Petrowski, N. W. & Rudik, P. A. (1927). *Psychologie des schachspiels*. Berlin: de Gruyter.
- Egan, D. E. & Schwartz, B. J. (1979). Chunking in recall of symbolic drawings. *Memory and Cognition*, 7, 149-158.
- Engle, R. W. & Bukstel, L. (1978). Memory processes among bridge players of differing expertise. *American Journal of Psychology*, 91, 673-689.
- Eisenstadt, M. & Kareev, Y. (1975). Aspects of human problem solving: The use of internal representations. In D. A. Norman, & D. E. Rumelhart (Ed.), *Explorations in cognition* (pp. 308-346). San Francisco: Freeman.
- Epstein, S. L. (1992). Prior knowledge strengthens learning to control search in weak theory domains. *International Journal of Intelligent Systems*, 7, 547-586.
- Epstein, S. L. (1993). The Hoyle Learning Experiments (CS-TR-93-02). Hunter College, Department of Computer Science.
- Epstein, S. L. (1994). For the right reasons: The FORR architecture for learning in a skill domain. *Cognitive Science*, 18, 479-511.
- Epstein, S.L. & Gelfand, J. (1995). Learning spatially-oriented game-playing agents through experience. In *Proceedings of the 17th Annual Meeting of the Cognitive Science Society*, Hillsdale, NJ: Erlbaum.
- Holding, D. (1985). *The psychology of chess skill*. Hillsdale, NJ: Lawrence Erlbaum.
- Shneiderman, B. (1976). Exploratory experiments in programmer behavior. *International Journal of Computer and Information Sciences*, 5, 123-143.