

---

## Intelligent Behavior in Humans and Machines

---

**Pat Langley**

PATRICK.W.LANGLEY@GMAIL.COM

Silicon Valley Campus, Carnegie Mellon University, Moffett Field, CA 94305 USA

Computer Science Department, University of Auckland, Private Bag 92019, Auckland, New Zealand

### Abstract

In this paper, I review the role of cognitive psychology in the origins of artificial intelligence and in the continuing pursuit of its initial objectives. I consider some key ideas about representation, performance, and learning that had their inception in computational models of human behavior, and I argue that this approach to developing intelligent artifacts, although no longer common, has an important place in cognitive systems. Not only will research in this paradigm help us understand the nature of human cognition, but findings from psychology can serve as useful heuristics to guide our search for accounts of intelligence. I present some constraints of this sort that future research should incorporate, and I claim that another psychological notion – cognitive architecture – is especially relevant to developing unified theories of the mind. Finally, I suggest ways to encourage renewed interaction between AI and cognitive psychology to the advantage of both disciplines.

### 1. Introduction

In its early days, artificial intelligence was closely allied with the study of human cognition, to the benefit of both fields. Many early AI researchers were concerned with using computers to model the nature of people's thinking, while others freely borrowed ideas from psychology in their construction of intelligent artifacts. Over the past 25 years, this link has largely been broken, with very few of the field's researchers showing concern with results from psychology or even taking inspiration from human behavior. I maintain that this trend is an unfortunate one which has hurt our ability to pursue two of AI's original goals: to understand the nature of the human mind and to achieve artifacts that exhibit human-level intelligence.

In this essay, I review some early accomplishments of AI in representation, performance, and learning that benefited from an interest in human behavior. After this, I give examples of the field's current disconnection from cognitive psychology and suggest some reasons for this development. I claim that AI still has many insights to gain from the study of human cognition, and that results in this area can serve as useful constraints on intelligent artifacts. In this context, I argue that the paradigm of *cognitive systems*, a movement that incorporates many ideas from psychology, and especially the subfield of cognitive architectures, offer promising paths toward developing more complete theories of intelligence. In closing, I propose some steps that we can take to remedy the current undesirable situation.

## 2. Early Links Between AI and Psychology

As it emerged in the 1950s, artificial intelligence incorporated ideas from a variety of sources and pursued multiple goals, but a central insight was that we might use computers to reproduce the complex forms of cognition observed in humans. Some researchers took human intelligence as an inspiration and source of ideas without attempting to model its details. Other scientists, including Herbert Simon and Allen Newell, generally seen as two of the field's co-founders, viewed themselves as cognitive psychologists who used AI systems to model the mechanisms that underlie human thought. This stance did not dominate the field, but it was acknowledged and respected even by those who did not adopt it, and its influence was clear throughout AI's initial phase. The paradigm was pursued vigorously at the Carnegie Institute of Technology, where Newell and Simon based their work, and it was well represented in collections like *Computers and Thought* (Feigenbaum & Feldman, 1963) and *Semantic Information Processing* (Minsky, 1969).

For example, much of the early work on knowledge representation was carried out by scientists who were interested in the structure and organization of human knowledge. Thus, Feigenbaum (1963) developed discrimination networks as a model of human long-term memory and evaluated his EPAM system in terms of its ability to match established psychological phenomena. Hovland and Hunt (1960) introduced the closely related formalism of decision trees to model human knowledge about concepts. Quillian (1968) proposed semantic networks as a framework for encoding knowledge used in language, whereas Schank and Abelson's (1977) introduction of scripts was motivated by similar goals. Newell's (1973) proposal for production systems as a formalism for knowledge that controls sequential cognitive behavior was tied closely to results from cognitive psychology. Not all work in this area was motivated by psychological concerns (e.g., there was considerable work on logical approaches), but research of this sort was acknowledged as interesting and had a strong impact on the field.

Similarly, studies of human problem solving had a major influence on early AI research. Newell, Shaw, and Simon's (1958) *Logic Theorist*, arguably the first implemented AI system, which introduced the key ideas of heuristic search and backward chaining, emerged from efforts to model human reasoning on logic tasks.<sup>1</sup> Later work by the same team led to means-ends analysis (Newell, Shaw, & Simon, 1960), an approach to problem solving and planning that was implicated in verbal protocols of humans puzzle solving. Even some more advanced methods of heuristic search that are cast as resulting from pure algorithmic analysis have their roots in psychology. For example, the widely used technique of iterative deepening bears a close relationship to progressive deepening, a search method that De Groot (1965) observed in human chess players.

On a related note, research on knowledge-based decision making and reasoning, which emerged during the 1980s, incorporated many ideas from cognitive psychology. The key method for developing expert systems (Waterman, 1986) involved interviewing human experts to determine the knowledge they used when making decisions. Two related movements – qualitative physics (Kuipers, 1994) and model-based reasoning (Gentner & Stevens, 1983) – incorporated insights into the way humans reasoned about complex physical situations and devices. Yet another theoretical framework that was active at this time, reasoning by analogy (e.g., Gentner & Forbus, 1991), had even

---

1. Simon (1981) explicitly acknowledges an intellectual debt to earlier work by the psychologist Otto Selz.

closer ties to results from experiments in psychology. Finally, during the same period, research on natural language understanding borrowed many ideas from structural linguistics, which studied the character of human grammatical knowledge.

Early AI research on learning also had close ties to computational models of human learning. Two of the first systems, Feigenbaum's (1963) EPAM and Hovland and Hunt's (1960) CLS, directly attempted to model results from psychological experiments on memorization and concept acquisition. Later work by Anzai and Simon (1979), which modeled human learning on the Tower of Hanoi, launched the 1980s movement on learning in problem solving, and much of the work in this paradigm was influenced by ideas from psychology (e.g., Minton et al., 1989) or made direct contact with known phenomena (e.g., Jones & VanLehn, 1994). Computational models of syntax acquisition typically had close ties with linguistic theories (e.g., Berwick, 1979) or results from developmental linguistics (e.g., Langley, 1983). In general, early efforts in machine learning reflected the diverse nature of human learning, dealing with a broad range of capabilities observed in people even when not attempting to model the details.

The alliance between psychological and non-psychological AI had benefits for both sides. Careful studies of human cognition often suggested avenues for building intelligent systems, whereas programs developed from other perspectives often suggested ways to model people's behavior. There was broad agreement that humans were our only examples of general intelligent systems, and that the primary goal of AI was to reproduce this unique capability with digital computers. Members of each paradigm knew about the others' results, and they exchanged ideas and results to their mutual advantage.

### **3. The Unbalanced State of Modern Artificial Intelligence**

Despite these obvious benefits, the past 25 years have seen an increasing shift in AI research away from concerns with modeling human cognition and a decreasing familiarity with results from psychology. What began as a healthy balance between two perspectives on AI research, sometimes both held by an individual scientist, has gradually become a one-sided community that believes AI and psychology have little to offer each other. Here I examine this trend in some detail, first describing the changes that have occurred in a number of research areas and then considering some reasons for these transformations.

Initial research on knowledge representation, much of which was linked to work in natural language understanding, led to notations such as semantic networks and scripts, which borrowed ideas from cognitive psychology and which influenced that field in return. However, over time, many researchers in this area became more formally oriented and settled on logic as the proper notation to express content. They became more concerned with guarantees about efficiently processing than with matching the flexibility and power observed in human knowledge structures, leading to representational frameworks that made little contact with psychological theories. Similar concerns with mathematical underpinnings have led to the more recent emphasis on probabilistic frameworks.

The earliest AI systems for problem solving and planning drew on methods like means-ends analysis, which were implicated directly in human cognition. For mainly formal reasons, these were gradually replaced with algorithms that constructed partial-order plans, which involved more complex reasoning but retained the means-ends notion of chaining backward from goals. More

recently, these have been replaced with techniques that reformulate the planning task in terms of constraint satisfaction or that carry out extensive forward search for possible paths, neither of which bear much resemblance to problem solving in humans.

Early work on natural language processing aimed to produce the same deep understanding of sentences and discourse as humans exhibit. However, constructing systems with this ability was time consuming and their behavior was fragile. Over time, they have been largely replaced with statistical methods that have limited connection to results from linguistics or psycholinguistics. More important, they often focus on tasks like information retrieval and information extraction that achieve near-term practical goals but that have little overlap with the broad capabilities found in human processing of natural language.

Research in machine learning initially addressed a wide range of performance tasks, including problem solving, reasoning, diagnosis, natural language, and visual interpretation. Many approaches reproduced the incremental nature of human learning, and they combined background knowledge with experience to produce rates of improvement similar to those found in people. However, during the 1990s, work in this area gradually narrowed to focus almost exclusively on supervised induction for classification and reinforcement learning for reactive control. This shift was accompanied by an increased emphasis on knowledge-lean statistical methods that require large amounts of data and that learn far more slowly than humans.

#### **4. Reasons for the Shift in Research Style**

There are many reasons for these developments, some of them involving technological advances. Faster computer processors and larger memories have made possible new methods that operate in somewhat different ways than people. Current chess-playing systems invariably retain many more alternatives in memory, and look much deeper ahead, than human players. Similarly, many supervised induction methods process much larger data sets, and consider a much larger space of hypotheses, than human learners. This does not mean they might fare even better by incorporating ideas from psychology, but these approaches have led to genuine scientific advances.

Unfortunately, other factors have revolved around historical accidents and sociological trends. Most academic AI researchers have come to reside in computer science departments, many of which grew out of mathematics units and which have a strong formalist bent. Many faculty in such departments view connections to psychology with suspicion, making them reluctant to hire scientists who attempt to link the two fields. For such researchers, mathematical tractability becomes a key concern, leading them to restrict their work to problems they can handle analytically, rather than ones that humans tackle heuristically. Graduate students are inculcated in this view, so that new PhDs pass on the bias when they take positions.

The mathematical orientation of many AI researchers is closely associated with an emphasis on approaches that guarantee finding the best solutions to problems. This concern with *optimality* stands in stark contrast with early AI research, which incorporated the idea of *satisficing* (Simon, 1955) from studies of human decision making. The notion of satisficing is linked to a reliance on heuristics, which are not guaranteed to find the best (or even any) solutions, but which usually produce acceptable results with reasonable effort. Formalists who insist on optimality are often

willing to restrict their attention to classes of problems for which they can guarantee such solutions, whereas those willing to use heuristic methods are willing to tackle more complex tasks.

Another factor that has encouraged a narrowing of research scope has been the commercial success of AI technology. Methods for diagnosis, supervised learning, scheduling, and planning have all found widespread uses, many of them supported by content made available on the Web. The financial success and commercial relevance of these techniques have led many academics to focus on narrowly defined problems like text classification, causing a bias toward near-term applications and an explosion of work on “niche AI” rather than on complete intelligent systems. Component algorithms are also much easier to evaluate experimentally, a lesson that has been reinforced by the problem repositories and competitions that have become common in recent years.

Taken together, these influences have discouraged most AI scientists from making contact with ideas from psychology. Research on computational models of human cognition still exists, but it is almost invariably carried out in psychology departments. This work is seldom referred to as artificial intelligence or published in that field’s typical outlets, even when the systems developed meet traditional AI criteria. One reason is that this area has developed its own narrowly defined standards for success that focus on quantitative fits to reaction times or error rates observed in humans, rather than the ability to carry out complex cognitive tasks. This leads modelers to introduce parameters and features that hold little interest to non-psychological researchers, strengthening the impression that models of human thinking have little to offer the broader AI community.

However, some scientists have instead concentrated on achieving qualitative matches to human behavior. For example, Cassimatis (2004) describes such a theory of language processing, whereas Langley and Rogers (2005) report a theory of human problem solving that extends Newell et al.’s (1958, 1960) early work. Their goal has been not to develop narrow models that fit the detailed results of psychological experiments, which to them seems premature, but rather to construct models that have the same broad functionality as we find in people. As Cassimatis, Bello, and Langley (2008) have argued, research in this paradigm has more to offer mainstream artificial intelligence, and it comes much closer in spirit to the field’s early efforts at modeling complex human cognition.

In summary, over the past few decades, AI has become increasingly focused on narrowly defined problems that have immediate practical applications or that are amenable to formal analysis. Pre-occupations with processing speed, predictive accuracy, and optimality have drawn attention away from the flexible forms of intelligent behavior at which humans excel. Research on computational modeling does occur in psychology, but it emphasizes quantitative fits to experimental results and also exhibits a narrowness in scope. The concerns that dominate both fields differ markedly from those pursued at the dawn of artificial intelligence.

## **5. Benefits of Renewed Interchange**

Clearly, the approach to developing intelligent systems adopted by early AI researchers is no longer a very common one, but I maintain that it still has an important place in the field. The reasons for pursuing research in this paradigm are the same as they were 50 years ago, at the outset of the discipline. However, since many AI scientists appear to have forgotten them, I will repeat them at the risk of stating the obvious to those with longer memories. In addition, I will discuss one path that holds special promise for bridging the rift.

First, research along these lines will help us understand the nature of human cognition. This is a worthwhile goal in its own right, not only because it has implications for education and other applications, but because intelligent human behavior comprises an important set of phenomena that demand scientific explanation. Understanding the human mind in all its complexity remains an open and challenging problem that deserves far more attention, and computational models of cognition offer a natural way to tackle it.

Second, findings from psychology can serve as useful heuristics to guide our development of intelligent artifacts. They reveal the ways people represent their knowledge about the world, the processes they employ to retrieve and use that knowledge, and the mechanisms by which they acquire it from experience and instruction. AI researchers must make decisions about these issues when designing intelligent systems, and psychological results about representation, performance, and learning provide reasonable candidates to consider. Humans remain our only example of general intelligent systems and, at the very least, insights about how they operate should receive serious consideration in the design of intelligent artifacts. Future AI research would benefit substantially from increased reliance on such design heuristics.

Third, observations of human capabilities can serve as an important source of challenging tasks for AI research. For example, we know that humans understand language at a much deeper level than current systems, and attempting to reproduce this ability, even by means quite different than those people use, is a worthwhile endeavor. Humans can also generate and execute sophisticated plans that trade off many competing factors and that achieve only the most important goals. They have the ability to learn complex knowledge structures in an incremental manner from few experiences, and they exhibit creative acts like composing musical pieces, designing complex devices, and devising scientific theories. Most current AI research sets its sights too low by focusing on simpler tasks like classification and reactive control, many of which can hardly be said to involve intelligence. Psychological studies reveal the impressive abilities of human cognition, and thus serve to challenge the field and pose new problems that require extensions to existing technology.

I have made similar arguments elsewhere in my characterization of the *cognitive systems* paradigm (Langley, 2012), which adopts connections to psychology as one of its central tenets. This movement incorporates a number of other ideas, including a focus on high-level processing, a reliance on structured representations, the use of heuristic methods, and an emphasis on system-level research, but respect for findings about human cognition is another major theme that guides work in this area. Not all cognitive systems research aims to model human thinking, but it often takes into account results on this topic.

An important subfield of cognitive systems, which explores the notion of *cognitive architectures*, is especially relevant to developing unified theories of the mind. A cognitive architecture (Newell, 1990; Langley, Laird, & Rogers, 2009) specifies aspects of an intelligent system that are stable over time, much as in a building's architecture. These include the memories that store perceptions, beliefs, goals, and knowledge, the representation of elements that are contained in these memories, the performance mechanisms that use them, and the learning processes that build on them. Such a framework typically comes with a programming language and software environment that supports the efficient construction of knowledge-based systems.

Most research on cognitive architectures shares a number of important theoretical assumptions. These include claims that:

- short-term memories are distinct from long-term memories, in that the former contain dynamic information and the latter store more stable content;
- both short-term and long-term memories contain symbolic list structures that can be composed dynamically during performance and learning;
- the architecture accesses elements in long-term memory by matching their patterns against elements in short-term memory;
- cognition operates in cycles that retrieve relevant long-term structures, then use selected elements to carry out mental or physical actions; and
- learning is incremental, being tightly interleaved with performance, and involves the monotonic addition of symbolic structures to long-term memory.

Most of these ideas have their origins in theories of human memory, problem solving, reasoning, and skill acquisition. They are widespread in research on cognitive architectures, but they remain relatively rare in other branches of artificial intelligence to their great detriment. The benefits of this approach include access to a variety of constraints from psychology about the nature of intelligent systems, an emphasis on demonstrating the generality of architectures across a variety of domains, and a focus on systems-level research that moves beyond component algorithms toward unified theories of intelligence.

Research on cognitive architectures varies widely in the degree to which it attempts to match psychological data. ACT-R (Anderson & Lebiere, 1998) and EPIC (Kieras & Meyer, 1997) aim for quantitative fits to reaction time and error data, whereas PRODIGY (Minton et al., 1989) incorporates selected mechanisms like means-ends analysis but otherwise makes little contact with human behavior. Architectures like Soar (Laird, Newell, & Rosenbloom, 1987; Newell, 1990; Laird, 2012) and ICARUS (Langley & Choi, 2006; Langley, Choi, & Rogers, 2009) take a middle position, drawing on many psychological ideas but also emphasizing their strength as flexible AI systems. What they hold in common is an acknowledgement of their debt to theoretical concepts from cognitive psychology and a concern with the same intellectual abilities as humans. Developing cognitive architectures is not the only approach within the cognitive systems paradigm that incorporates ideas from psychology, but it constitutes an important role model.

## 6. Healing the Intellectual Rift

If we assume, for the sake of argument, that artificial intelligence would indeed benefit from renewed links to cognitive psychology, then there remains the question of how we can achieve this objective. As is usually the case for complex problems, the solution will involve many steps, each of which will take us some distance toward our goal. I will discuss these independently, although naturally they will interact in some ways.

Clearly, one response must involve education. Few AI courses discuss connections to cognitive psychology, and few students read papers they cannot access on the Web, so they are often unaware of the older literature. We should provide a broader education in AI that incorporates concepts from

cognitive psychology, many of which are far more relevant to the field's original agenda than ones from mainstream computer science. My course at Arizona State University on artificial intelligence and cognitive systems (<http://www.cogsys.org/courses/langley/aicogsys11/>) adopted this perspective, but we need many others in the same spirit. Tutorials and summer schools on similar topics would complement university classes by giving more condensed presentations in shorter periods.

We should also encourage more research that links AI and cognitive psychology. This requires both conferences at which to present results in this tradition and journals in which to publish them. *Advances in Cognitive Systems* aims to serve both functions by holding an annual refereed conference and publishing papers in an electronic journal, but other venues also exist, such as the International Conference on Cognitive Modeling and the journal *Cognitive Systems Research*. Funding agencies can also foster AI research that incorporates insights from psychology. In the US, both the Office of Naval Research and the National Science Foundation have long supported such work, and DARPA's initiative on cognitive systems encouraged research along these lines further. Nevertheless, the field would benefit from expanded funding, by these and other agencies, of scientific work that blurs the boundaries between AI and psychological modeling.

In summary, the original vision of AI was to understand the principles that support intelligent behavior and to use them to construct computational systems with the same breadth of abilities as humans. Many early systems doubled as models of human cognition, while others made effective use of ideas from psychology. The last few decades have seen far less research in this tradition, and work by psychologists along these lines is seldom acknowledged by the AI community. Nevertheless, observations of intelligent behavior in humans can provide an important source of mechanisms and tasks to drive research. Without them, AI seems likely to become a set of narrow, specialized subfields that have little to tell us about the nature of the mind. Over the next few decades, AI must reestablish connections with cognitive psychology if it hopes to achieve its initial goal of creating intelligent systems that exhibit the same capabilities as humans.

## Acknowledgements

This essay expands on my presentation at the 2006 Dartmouth Artificial Intelligence Conference. The analyses reported here were supported in part by DARPA under agreement number FA8750-05-2-0283 and in part by Grant N00014-10-1-0487 from ONR, which are not responsible for its contents. Discussions with Herbert Simon, Allen Newell, John Anderson, David Nicholas, John Laird, Randolph Jones, and many others over the years have influenced the ideas I have presented.

## References

- Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. Mahwah, NJ: Lawrence Erlbaum.
- Anzai, Y., & Simon, H. A. (1979). The theory of learning by doing. *Psychological Review*, 86, 124–140.
- Berwick, R. (1979). Learning structural descriptions of grammar rules from examples. *Proceedings of the Sixth International Conference on Artificial Intelligence* (pp. 56–58). Tokyo: Morgan Kaufmann.

- Brachman, R., & Lemnios, Z. (2002). DARPA's new cognitive systems vision. *Computing Research News*, 14, 1.
- Cassimatis, N. L. (2004). Grammatical processing using the mechanisms of physical inferences. *Proceedings of the Twentieth-Sixth Annual Conference of the Cognitive Science Society*. Chicago.
- Cassimatis, N. L., Bello, P., & Langley, P. (2008). Ability, breadth and parsimony in computational models of higher-order cognition. *Cognitive Science*, 32, 1304–1322.
- de Groot, A. D. (1965). *Thought and choice in chess*. The Hague: Mouton.
- Feigenbaum, E. A. (1963). The simulation of verbal learning behavior. In E. A. Feigenbaum & J. Feldman (Eds.), *Computers and thought*. New York: McGraw-Hill.
- Feigenbaum, E. A., & Feldman, J. (Eds.) (1963). *Computers and thought*. New York: McGraw-Hill.
- Gentner, D., & Forbus, K. (1991). MAC/FAC: A model of similarity-based retrieval. *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society* (pp. 504–509). Chicago: Lawrence Erlbaum.
- Gentner, D., & Stevens, A. L. (1983). *Mental models*. Hillsdale, NJ: Lawrence Erlbaum.
- Hovland, C. I. & Hunt, E. A. (1960). The computer simulation of concept attainment. *Behavioral Science*, 5, 265–267.
- Jones, R. M., & VanLehn, K. (1994). Acquisition of children's addition strategies: A model of impasse-free, knowledge-level learning. *Machine Learning*, 16, 11–36.
- Kieras, D., & Meyer, D. E. (1997). An overview of the EPIC architecture for cognition and performance with application to human-computer interaction. *Human-Computer Interaction*, 12, 391–438.
- Kuipers, B. (1994). *Qualitative reasoning: Modeling and simulation with incomplete knowledge*. Cambridge, MA: MIT Press.
- Laird, J. E. (2012). *The Soar cognitive architecture*. Cambridge, MA: MIT Press.
- Laird, J. E., Newell, A., & Rosenbloom, P. S. (1987). Soar: An architecture for general intelligence. *Artificial Intelligence*, 33, 1–64.
- Langley, P. (1982). Language acquisition through error recovery. *Cognition and Brain Theory*, 5, 211–255.
- Langley, P. (2012). The cognitive systems paradigm. *Advances in Cognitive Systems*, 1, 3–13.
- Langley, P., & Choi, D. (2006). A unified cognitive architecture for physical agents. *Proceedings of the Twenty-First National Conference on Artificial Intelligence*. Boston: AAAI Press.
- Langley, P., Choi, D., & Rogers, S. (2009). Acquisition of hierarchical reactive skills in a unified cognitive architecture. *Cognitive Systems Research*, 10, 316–332.
- Langley, P., Laird, J. E., & Rogers, S. (2009). Cognitive architectures: Research issues and challenges. *Cognitive Systems Research*, 10, 141–160.
- Langley, P., & Rogers, S. (2005). An extended theory of human problem solving. *Proceedings of the Twenty-seventh Annual Meeting of the Cognitive Science Society*. Stresa, Italy.
- Minsky, M. (Ed.) (1969). *Semantic information processing*. Cambridge, MA: MIT Press.

- Minton, S., Carbonell, J. G., Knoblock, C. A., Kuokka, D., Etzioni, O., & Gil, Y. (1989). Explanation-based learning: A problem solving perspective. *Artificial Intelligence*, 40, 63–118.
- Newell, A. (1973). Production systems: Models of control structures. In W. G. Chase (Ed.), *Visual information processing*. New York: Academic Press.
- Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
- Newell, A., Shaw, J. C., & Simon, H. A. (1958). Elements of a theory of human problem solving. *Psychological Review*, 65, 151–166.
- Newell, A., Shaw, J. C., & Simon, H. A. (1960). Report on a general problem-solving program for a computer. *Information Processing: Proceedings of the International Conference on Information Processing* (pp. 256–264). UNESCO House, Paris.
- Quillian, M. R. (1968). Semantic memory. In M. Minsky (Ed.), *Semantic information processing*. Cambridge, MA: MIT Press.
- Schank, R., & Abelson, R. (1977). *Scripts, plans, goals, and understanding*. Hillsdale, NJ: Lawrence Erlbaum.
- Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69, 99–118.
- Simon, H. A. (1981). Otto Selz and information-processing psychology. In N. H. Frijda & A. de Groot (Eds.), *Otto Selz: His contribution to psychology*. The Hague: Mouton.
- Waterman, D. A. (1986). *A guide to expert systems*. Reading, MA: Addison-Wesley.