
The Cognitive Systems Paradigm

Pat Langley

PATRICK.W.LANGLEY@GMAIL.COM

Computing Science and Engineering, Arizona State University, Tempe, AZ 85287 USA

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

Abstract

In this essay, I review the motivations behind the cognitive systems movement and attempt to characterize the paradigm. I propose six features that distinguish research in this framework from other approaches to artificial intelligence, after which I present some positive and negative examples in an effort to clarify the field's boundaries. In closing, I also discuss some avenues for encouraging research in this critical area.

1. Introduction

The early days of artificial intelligence were guided by a common vision: understanding and re-producing, in computational systems, the full range of intelligent behavior that we observe in humans. Many researchers continued to share these aims until the 1980s and 1990s, when AI began to fragment into a variety of specialized subdisciplines, each with far more limited objectives. This produced progress in each area, but, in the process, many abandoned the field's original goal. Rather than creating intelligent systems with the same breadth and flexibility as humans, most recent research has produced impressive but narrow idiot savants.

The field's central goal was to understand the nature of the mind. This is one of the core aims of science, on an equal footing with questions about the nature of the universe, the nature of matter, and the nature of life. As such, it deserves the same respect and attention it received during discipline's initial periods. However, since mainstream AI has largely abandoned this goal, we require a new name for research that remains committed to the original vision. For this purpose, I propose the phrase *cognitive systems*, which Brachman and Lemnios (2002) championed at DARPA in their efforts to encourage research in this early tradition. As we will see later, this label incorporates some key ideas behind the movement.

In this article, I attempt to characterize the cognitive systems paradigm. I review six assumptions that were widely adopted during the AI's first three decades, in each case explaining why the idea made sense, why it has fallen into disfavor among many researchers, and why we should reaffirm it in order to make progress toward understanding intelligence. After this, I attempt to further clarify the paradigm by considering positive and negative examples of cognitive systems research, and then conclude with suggestions for fostering research activity in this important area. Note that cognitive systems is not, at heart, a new movement, but rather a continuation of an old, established, and productive tradition that deserves more attention than it has received in recent years.

2. High-Level Cognition

Early research in AI revolved around the study of *high-level cognition*, a feature that separated it from its sister fields of pattern recognition and robotics. When we say that humans or machines exhibit intelligence, we are not referring to their ability to recognize concepts, perceive objects, or execute complex motor skills, which they share with other animals like dogs and cats. These abilities are clearly important for agents that operate in physical environments, but they are not distinguishing features of intelligent systems. Rather, by *intelligence* we mean the capacity to engage in abstract thought that goes beyond immediate perceptions and actions. In humans, this includes the ability to carry out complex, multi-step reasoning, understand the meaning of natural language, design innovative artifacts, generate plans that achieve goals, and even reason about their own reasoning.

During AI's first three decades, much of the discipline's research dealt with these issues, and the progress during that period arguably increased our understanding of the mind. Researchers addressed tasks like proving theorems in logic (Newell, Shaw, & Simon, 1957) and geometry (Gelernter, 1959), solving problems in symbolic integration (Slagle, 1962) and physics (Novak, 1977), playing games like chess (Greenblatt, Eastlake, & Crocker, 1967), reading sentences and responding (Winograd, 1972), and generating novel plans to achieve goals (Sacerdoti, 1974). These and many other efforts addressed the types of abstract thought that we associate with human intelligence, and thus dealt with high-level cognition.

This idea is still active in some AI subfields, such as planning and automated reasoning, but they have developed their own specialized methods that exhibit good performance at the expense of limited flexibility. Even these subareas have redirected some of their attention to lower-level tasks like reactive control, while others have nearly abandoned their initial concern with high-level cognition. For instance, machine learning focuses almost exclusively on classification tasks (Langley, 2011), whereas natural language processing has largely replaced its original emphasis on understanding with text classification and information retrieval.

Such shifts in emphasis have produced short-term gains, with clear performance improvements on the narrowly defined tasks studied in each subfield. They have also led to many successful applications in both the commercial and public sectors. But advances on these fronts, however impressive, tell us little about the general nature of intelligence, and we need far more research on the broad set of high-level cognitive abilities that make humans distinctive. For this reason, an emphasis on high-level mental processing is both one of the key aims and one of the defining features of the cognitive systems paradigm.

This emphasis on high-level cognition does not mean there is no place for research on other aspects of mental processing. Reasoning, problem solving, and language understanding could not occur without more basic mechanisms, such as pattern matching and memory retrieval, to support them. Advances on these and other facets of 'elementary information processing' are relevant to cognitive systems when they are demonstrated in the context of high-level capabilities. The same holds for research peripheral processes involved in perception and action, which are relevant when studied in terms of their relationship to problem solving, reasoning, language understanding, and other high-level mental activities.

3. Structured Representations

Early AI researchers also assumed that *structured representations* play a central role in intelligence, which in turn depend on the ability to represent, create, and interpret content encoded in such representations. This position is closely related to the fundamental insight – arguably the foundation of the 1956 AI revolution – that computers are not simply numeric calculators but rather general symbol manipulators. This idea has become so widely adopted that few today realize that it was a breakthrough at the time it was introduced.

As Newell and Simon (1976) state clearly in their physical symbol system hypothesis, intelligent behavior appears to require the ability to encode, interpret, and manipulate symbolic structures. This was a major assumption of early AI research, and the most impressive successes in the field's 55 year history have relied on this fundamental capability. Computational systems that carry out complex reasoning in mathematics (Gelernter, 1959; Slagle, 1962) and physics (Novak, 1977), programs that understand the meanings of natural-language sentences (Winograd, 1972), and mechanisms for generating novel plans (Sacerdoti, 1974) and designs (Gero & Radford, 1985) all represent, create, and operate over symbolic structures.

This emphasis runs counter to recent trends in many branches of AI, which, over the past few decades, have retreated from this position. Some subfields, such as machine learning, have abandoned almost entirely the use of interpretable symbolic formalisms, caring only about performance, however achieved. Other subfields, like knowledge representation and constraint satisfaction, have retained a focus on symbols but limit their formalisms for reasons of efficiency or analytical tractability. Both developments constitute steps backward from the physical symbol system hypothesis, and they bode ill for our efforts to fathom the complex nature of intelligence.

Again, this narrowing of scope has been associated with successful applications, convincing many that all research should follow a similar trajectory. Yet the arguments for powerful structured representations, and processes that operate on them, remain as powerful as they were in AI's earliest days. Focusing on efficiency, accuracy, or formal analyzability will distract us from developing systems with the same breadth and flexibility associated with human intelligence. For these reasons, the cognitive systems paradigm encourages the use of rich, structured representations to support these capabilities.

However, note that this commitment does not imply a reliance on any particular class of structured formalisms. Relational logic is a well-known example, but many other notations have similar representational abilities. Nor does any stigma attach to frameworks that include numerical information provided these serve mainly as annotations on relational structures. The key issue concerns whether an approach can represent, and operate over, the complex relations that arise in reasoning, problem solving, language understanding, and other high-level cognitive processes.

4. System-Level Research

A third feature that characterized much early AI work was an emphasis on *system-level* accounts of intelligence. Because researchers envisioned comprehensive theories of the mind, they naturally recognized the need for their programs to comprise a number of interacting components. The argument for this approach was compelling: because intelligence is clearly such a complex and

multifaceted phenomenon, even partial accounts should incorporate multiple capabilities and aim to explain how these different processes can work together to support high-level mental activities of the sort observed regularly in humans.

This systems perspective was a recurring theme that held for many research projects. Early programs for natural language processing (e.g., Winograd, 1972) often included not only modules for syntactic analysis, but also mechanisms for semantic processing and for multi-step reasoning. Similarly, many learning systems combined multiple induction techniques and, in some cases, augmented them with reasoning mechanisms (e.g., Pazzani, 1988). The phrase *intelligent systems*, which has seen wide use, reflects the realization that understanding the nature of mental abilities depends on artifacts that comprise interacting components.

Many AI systems developed during this period were given distinctive names that served as shorthand for a constellation of mutually supportive mechanisms. Although these programs may have introduced new algorithms, papers often emphasized interactions among their components. A related trend was the introduction of high-level programming languages, such as OPS (Forgy, 1979) and Prolog (Clocksin & Mellish, 1981), designed to aid the development of intelligent systems. Each such framework came with a distinctive syntax that reflected its theoretical assumptions about intelligence. These two ideas came together in Newell's (1990) notion of a *cognitive architecture* that provides both a set of interacting mechanisms and a high-level language for intelligent agents.

Despite these promising beginnings, by the 1990s many researchers had come to focus on their energies on component algorithms rather than integrated systems. This resulted partly from AI finding its primary home in computer science departments, which typically gave higher status to the study of algorithms. A second influence was the emphasis on conference publications, which provided sufficient space to describe algorithms but not enough for system-level accounts. Another factor was the relative ease of evaluating algorithms, both formally and experimentally, which made it easier to both produce and publish papers on such topics. Finally, university professors found it far simpler to teach AI courses as a set of unrelated techniques than to present coherent frameworks for complex intelligent systems.

In recent decades, these trends have led to greatly decreased interest in system-level research and to the fragmentation of AI into a set of disconnected subfields. Yet the need for such work has not diminished, and the reasons for pursuing it are as sound as they were half a century ago. Like any complex phenomenon, we cannot understand intelligence merely by studying its parts; we must decipher how these components interact. An abiding concern with issues that arise during such integration is a third distinguishing feature of the cognitive systems paradigm, which reflects this focus in its name.

This does not mean that cognitive systems has no room for work on new component mechanisms, which of course are required to make system-level research possible. However, they are best introduced and studied in terms of their interactions with other processes. For instance, a researcher might replace one module in a cognitive architecture with a new variant that better supports other aspects of the framework. Thus, the paradigm encourages research on component processes in the context of larger-scale cognitive systems, even though it does not favor work on isolated algorithms for their own sake.

5. Heuristics and Satisficing

Another central assumption of initial AI research was that intelligence involves *heuristic search* (Newell & Simon, 1976). Although not the only field to adopt the search metaphor, AI was distinctive in its use of heuristics that, although not guaranteed to produce results, often made problems tractable which could not be solved otherwise. In fact, the field was sometimes referred to as *heuristic programming* to distinguish it from mainstream computer science, which emphasized ‘algorithms’ that provided guarantees. On this dimension, AI differed from fields like operations research, which limited their attention to tasks for which one could find optimal solutions efficiently.

Instead, many AI researchers had the audacity to tackle more difficult problems to which such techniques did not apply. Their approach involved developing search methods that relied on heuristics to guide search down promising avenues and away from fruitless ones. The resulting systems often *satisfied* (Simon, 1956) by finding acceptable rather than optimal solutions, and even these were not guaranteed. However, in practice, heuristic methods could often solve problems that the more constrained approaches could not.

In some cases, heuristics were stated as symbolic rules that suggested candidates or eliminated alternatives. This approach was common in research that emphasized structured representations, as in work on planning (e.g., Sacerdoti, 1974) and reasoning (e.g., Newell et al., 1959; Slagle, 1963). In others cases, heuristics were encoded as numeric evaluation functions. This idea was widely used in research that adopted simpler representations, as in work on game playing (e.g., Greenblatt et al., 1967). Despite their differences, both paradigms assumed that complex tasks involved search through combinatorial spaces and that exhaustive approaches were intractable, making heuristic methods the natural option.

Unfortunately, recent decades have seen many AI researchers turn away from this practical attitude and adopt other fields’ obsession with formal guarantees. For example, much recent work in knowledge representation has focused on constrained formalisms that promise efficient reasoning, even though this restricts the reasoning tasks they can handle. Research on reinforcement learning often limits itself to methods that provably converge to an optimal control policy, even if the time required for convergence makes them completely impractical. And the recent popularity of statistical approaches has resulted partly from the belief, often mistaken, that techniques with mathematical formulations provide guarantees about their behavior.

Although it makes sense to use nonheuristic methods when they apply to a given problem, this does not mean that researchers should study only those tasks that such approaches can handle. The original charter of AI was to address the same broad class of tasks as humans, and experiments have repeatedly established that people rely on heuristic methods in all but the simplest contexts. Moreover, studies have shown that, in many cases, this strategy is remarkably effective. Thus, it seems clear heuristics should occupy a central place in the cognitive systems agenda.

Nevertheless, this emphasis on heuristics does not mean that research on cognitive systems should never utilize methods that offer guarantees like optimality. This is common for low-level processing that provides the infrastructure for high-level cognition, since tasks at this level often do not have the combinatorial character of more complex tasks. Nonheuristic techniques even have

a role at higher levels, say when used to handle special cases in the context of a more general framework. However, the cognitive systems paradigm recognizes the importance of heuristics for many complex tasks and it does not glorify formal guarantees for their own sake.

6. Links to Human Cognition

The notion of heuristics bears on another assumption prevalent in early AI research – that the design and construction of intelligent systems has much to learn from the study of *human* cognition. Many central ideas in knowledge representation, planning, natural language, and learning (including the importance of heuristics) were originally motivated by insights from cognitive psychology and linguistics, and many early, influential AI systems doubled as computational models of human behavior (e.g., Newell, Shaw, & Simon, 1959). Few researchers attempted to model the details of human behavior, but many AI scientists exhibited a genuine interest in cognitive psychology and the ideas that it offered.

The field also looked to human activities for likely problems that would challenge existing capabilities. Research on expert medical diagnosis (e.g., Shortliffe & Buchanan, 1975), intelligent tutoring systems (Sleeman & Brown, 1982), artistic composition (e.g., Cohen, 1975), scientific discovery (e.g., Langley, 1981), and many other topics were all motivated by a desire to support activities considered difficult for humans. This approach produced a wealth of ideas about how to construct cognitive artifacts, many of which were demonstrated in the context of innovative, impressive, and interesting intelligent systems.

However, as time passed, fewer researchers adopted this perspective, preferring instead to draw their inspirations and concerns from more formal fields. Still worse, fewer chose to work on challenging intellectual tasks that humans can handle only with considerable effort or advanced training. Attention moved instead to problems on which computers can excel using simple techniques combined with rapid computing and large memories, like data mining and information retrieval. Even for challenging problems like playing chess that require heuristic search, the vast majority of recent work has emphasized fast processing and large storage.

Although such efforts have resulted in clear successes, I claim that they reveal little about the nature of intelligence in either humans or machines. There remains a need for research that is motivated by the cognitive abilities that humans exhibit so effortlessly, with limited computational power and limited working memories, both as a guide for selecting tasks to drive research and as a source of ideas about how to tackle them. For these reasons, links to human thinking should be an important aspect of the cognitive systems paradigm.

Still, connections to human behavior are not a requirement for valuable research on cognitive systems. One can develop accounts of intelligence that are uninfluenced by results in psychology or that even diverge explicitly from knowledge about human cognition. This is true especially for lower levels of processing that hold less relevance to the paradigm. What remains important is that these systems treat seriously the high-level constraints that confront humans, such as the need to process dialogue as it arrives or the need to learn incrementally over time. In contrast, work that ignores these constraints, such as methods for nonincremental data mining, have far less relevance to understanding cognitive systems.

7. Exploratory Research

A final characterization of early AI was its commitment to exploratory research. Because there existed few examples of intelligent artifacts, a common strategy was to identify some intellectual ability of humans, design and implement a system that exhibited it, and demonstrate its behavior on a set of convincing examples. This typically required multi-year projects, with most dissertations following the model. The results from such efforts were compelling because they showed clearly that some new facet of intelligence had been explained in computational terms. The first decades of AI was a heady period, when demonstrations of new abilities proceeded rapidly despite the relatively small number of researchers.

Some examples will help communicate the spirit of the times. The first implemented AI system, the Logic Theorist (Newell et al., 1957), successfully proved a number of theorems in propositional logic, in some cases with novel reasoning chains. Another early program proved a variety of theorems in plane geometry (Gelernter, 1959), while SAINT solved symbolic integration problems from an MIT calculus examination (Slagle, 1962). SHRDLU (Winograd, 1972) interacted with users in written English about a simulated environment that it manipulated in response, while ISAAC (Novak, 1977) understood and solved physics word problems with diagrams. AARON (Cohen, 1975) composed and drew creative works of art, while Bacon (Langley, 1981) rediscovered empirical laws from the history of physics. Each of these systems, along with many others, demonstrated ability on a new class of problems and offered insights into the nature of intelligence.

However, the 1980s and 1990s saw a shift in research style. As people began to develop new approaches to established problems, it became natural to compare the behaviors of different methods. Experimental comparisons were favored in machine learning (Kibler & Langley, 1988) and natural language, while formal analyses were more common in the reasoning community. Explicit performance metrics were adopted for such comparisons, sometimes under pressure from funding agencies, and these became central to many conference-sponsored competitions. Increasingly, publication required showing that one's approach fared better on such metrics than previous work, which left little room for papers on more innovative but less optimized approaches. This situation now holds across most AI subfields, although their metrics and other details differ.

Such a framework for measuring progress might be justifiable if AI had already achieved broad coverage of human intellectual abilities. But existing systems for planning, reasoning, language, learning, and other areas clearly lack the breadth or flexibility seen in people. As a result, there remains a need for exploratory research on cognitive systems that demonstrate a wider range of capabilities, even if they are not as efficient or accurate as current techniques. The field also has room for innovative approaches to established tasks that provide theoretical insights, even if they are not competitive in the narrow sense.

Of course, this does not mean that careful empirical studies or formal analyses have no place in the cognitive systems paradigm. Controlled experiments can still reveal the sources of power, and mathematical analyses can clarify representational capacity or ability to scale. Intelligence involves a complex set of phenomena that we can only hope to understand through a collection of complementary methods that include demonstrations of new capabilities, high-level comparisons to human cognition, systematic experiments, and formal analyses.

8. Borders of the Paradigm

Although the research emphases just described partially characterize the paradigm of cognitive systems, I can clarify it further with some illustrative examples. My purpose is not to provide a hard specification of the field's boundaries, which would be undesirable for our developing community, but rather to offer a set of contrasts between work that clearly fits within the field's charter and other efforts that fall outside.

For instance, the focus on high-level cognition and structured representations suggests that research on methods for deep understanding of language clearly are within these borders, whereas shallow methods for information retrieval and text mining are not. Similarly, the emphasis on system-level accounts indicates that work on standalone algorithms, such as ones for classification learning, also falls outside its boundaries. The concern with heuristics suggests that progress on flexible but possibly inefficient reasoning mechanisms offers positive examples, whereas highly constrained techniques for efficient theorem proving or constraint satisfaction are less relevant.

As we have seen, progress in cognitive systems is often motivated by efforts to reproduce aspects of human intelligence. However, the paradigm emphasizes systems that exhibit the same qualitative capabilities as humans under the same qualitative conditions and constraints, as opposed to ones that fit quantitative results from psychological experiments. Neither must these systems incorporate mechanisms that are consistent with knowledge of human cognition; not all approaches to achieving broadly intelligent systems need operate in the same manner as humans. Research on cognitive architectures has aims that place it firmly within the paradigm, but work on new models stated within an existing architecture should emphasize novel insights (e.g., unexpected abilities or limitations) that have resulted from the exercise.

As noted earlier, cognitive systems supports a broader range of evaluation styles than many AI branches. Although experimental studies with quantitative performance metrics can be informative, the paradigm welcomes any research that makes explicit claims and provides evidence to support them. Such evidence may take many forms. For instance, demonstrating that a system supports an entirely new functionality is perfectly acceptable. New approaches to aspects of intelligence that have already been studied are also legitimate, provided that claims are backed by well-reasoned, convincing arguments. Experimental comparisons to existing methods are welcome but not required, and research that extends existing techniques should do more than show minor improvements on standard performance metrics.

Although typical work on cognitive systems involves an implemented computer program, this is certainly not required. Research may instead analyze a challenging problem that current methods cannot handle, especially if they describe the problem clearly and propose ways to evaluate progress. Surveys of existing approaches to a class of problems and their limitations can also be valuable, especially if they propose alternative ways to address the problem class that would not suffer from the same drawbacks. Proposals for new representational formalisms, insightful characterizations of human intelligence, or relevant theoretical analyses all have merit, especially if they clarify how to construct computational artifacts.

Importantly, the cognitive systems movement takes no stance about which representational formalisms and associated mechanisms are appropriate. Research projects in the logical, case-based, and probabilistic frameworks all fit within the paradigm, provided they deal with rich cognitive

structures and present material in a way accessible to the broader community. Progress on computational learning methods is also relevant, especially if they acquire rich cognitive structures at roughly the same rates as humans and take advantage of prior knowledge to constrain learning.

Generally speaking, the cognitive systems movement pursues research carried out in the original spirit of AI, which aimed to design and implement computer programs that exhibited the breadth, generality, and flexibility often observed in human intelligence. This contrasts with typical examples of recent work that aim to develop highly efficient or very powerful techniques for a narrow class of constrained tasks. Such research has its place, but I claim that the cognitive systems paradigm offers the most promising path to achieving computational artifacts that exhibit human-level intelligence in a broad and flexible manner.

9. Fostering the Cognitive Systems Movement

The six assumptions summarized earlier remain as valid now as they were over 50 years ago, in the first days of AI. They hold our best hope for achieving that field's original goals – to reproduce the full range of human intelligence in computational artifacts – and they deserve substantially more attention than they have received in recent years. Thus, it seems natural to ask how we can generate interest in the cognitive systems approach to understanding the mind and encourage its wider adoption within the research community.

One important response involves education. Most AI courses ignore the cognitive systems perspective, and few graduate students read papers that are not available on the Web, which means they are often unfamiliar with the relevant older literature. Instead, we must provide a broad education in AI that cuts across different topics to cover all the field's branches and their role in intelligent systems. The curriculum should incorporate ideas from cognitive psychology, linguistics, and logic, which are far more important to the cognitive systems agenda than ones from mainstream computer science. One example comes from a course on artificial intelligence and cognitive systems (<http://www.cogsys.org/courses/langley/aicogsys11/>) that I have offered at Arizona State University, but we need many more.

We should also encourage expanded research activity within the cognitive systems paradigm. Funding agencies can have a major effect here, and the past decade has seen encouraging developments on this front. During this period, DARPA in the USA supported a number of large-scale programs with this emphasis (Brachman & Lemnios, 2002), and the US Office of Naval Research has long shown a commitment to the paradigm. The European Union has also funded some projects (e.g., Christiansen, Sloman, Krujiff, & Wyatt, 2009) in the area. Continued government support of relevant research will aid progress, but this will require committed and informed people to join funding agencies as program officers.

The field would also benefit from more audacious and visionary goals to spur the field toward greater efforts on cognitive systems. For instance, the General Game Playing competition (<http://games.stanford.edu>) has fostered research on general intelligent systems, and Project Halo's (www.projecthalo.com) vision for a 'digital Aristotle' has inspired exciting progress on knowledge-based systems. But we also need demonstrations of flexible, high-level cognition in less constrained settings that require the combination of inference, problem solving, and language into more com-

plete intelligent systems. The Turing test has many drawbacks but the right spirit, and we would benefit from more efforts toward integrated systems that exhibit the breadth and flexibility of humans. Challenging tasks will reexcite junior and senior researchers about the original vision of artificial intelligence.

Of course, we also require venues to publish the results of research on cognitive systems. From 2006 to 2011, the annual AAAI conference included a special track on ‘integrated intelligence’ that encouraged submissions on system-level results. Last year, the AAAI Fall Symposia on Advances in Cognitive Systems (<http://www.cogsys.org/acs/2011/>) attracted over 75 participants, and the current journal, along with an associated conference of the same name, is an offshoot of that event. We need more alternatives along these lines to help counter the mainstream bias that favors papers which report on narrow tasks, standalone algorithms, and incremental improvements. Broader criteria for scientific progress are necessary to advance the field, making room for papers that analyze challenging problems, demonstrate new functionality, and reproduce distinctive human capabilities.

In summary, the original vision of AI was to understand the principles that support high-level cognitive processing and to use them to construct computational systems with the same breadth of abilities as humans. The cognitive systems paradigm attempts to continue in that spirit by utilizing structured representations and heuristic methods to support complex reasoning and problem solving. The movement focuses on integrated systems rather than component algorithms, and human cognition provides a source of ideas for many of these programs. These ideas have their origins in the earliest days of AI, yet they have lost none of their power or potential, and we stand to benefit from their readoption by researchers and educators. Without them, AI seems likely to become a set of narrow, specialized subfields that have little to tell us about intelligence. Instead, we should use the assumptions of the cognitive systems paradigm as heuristics to direct our search toward true theories of the mind. This seems the only intelligent path.

Acknowledgements

This essay expands on my presentation at the 2011 AAAI Fall Symposium on Advances in Cognitive Systems, as well as on material that appeared in *AI Magazine* and *AISB Quarterly*. This work was supported in part by Grant N00014-10-1-0487 from ONR, which is not responsible for its contents.

References

- Brachman, R., & Lemnios, Z. (2002). DARPA’s new cognitive systems vision. *Computing Research News*, 14, 1.
- Christiansen, H. I., Sloman, A., Krujiff, G.-J., & Wyatt, J. (2009). *Cognitive systems*. Berlin: Springer-Verlag.
- Clocksinn, W. F., & Mellish, C. S. (1981). *Programming in Prolog*. Berlin: Springer-Verlag.
- Cohen, H. (1975). Getting a clear picture. *Bulletin of the American Society for Information Science*, December, 10–12.
- Forgy, C. L. (1979). *The OPS4 reference manual* (Technical Report). Department of Computer Science, Carnegie-Mellon University, Pittsburgh, PA.

- Gelernter, H. (1959). Realization of a geometry theorem proving machine. *Proceedings of the International Conference on Information Processing* (pp. 273–282). Paris: UNESCO House.
- Gero, J. S., & Radford, A. D. (1985). Towards generative expert systems for architectural detailing. *Computer-Aided Design*, 17, 428–435.
- Greenblatt, R. D., Eastlake, D. E., & Crocker, S. D. (1967). The Greenblatt chess program. *Proceedings of the 1967 Fall Joint Computer Conference* (pp. 801–810). Anaheim, CA: AFIPS.
- Kibler, D., & Langley, P. (1988). Machine learning as an experimental science. *Proceedings of the Third European Working Session on Learning* (pp. 81–92). Glasgow: Pittman.
- Langley, P. (1981). Data-driven discovery of physical laws. *Cognitive Science*, 5, 31–54.
- Langley, P. (2011). The changing science of machine learning. *Machine Learning*, 82, 275–279.
- Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
- Newell, A., Shaw, J. C., & Simon, H. A. (1959). Report on a general problem-solving program. *Proceedings of the International Conference on Information Processing* (pp. 256–264). Paris: UNESCO House.
- Newell, A., & Simon, H. A. (1976). Computer science as empirical enquiry: Symbols and search. *Communications of the ACM*, 19, 113–126.
- Novak, G. (1977). Representations of knowledge in a program for solving physics problems. *Proceedings of the Fifth International Joint Conference on Artificial Intelligence* (pp. 286–291). Cambridge, MA: Morgan Kaufmann.
- Pazzani, M. (1988). Integrating empirical and explanation-based learning methods in OCCAM. *Proceedings of the Third European Working Session on Learning* (pp. 147–166). Glasgow, Scotland: Pittman.
- Sacerdoti, E. D. (1974). Planning in a hierarchy of abstraction spaces. *Artificial Intelligence*, 5, 115–135.
- Shortliffe, E. H., & Buchanan, B. G. (1975). A model of inexact reasoning in medicine. *Mathematical Biosciences*, 23, 351–379.
- Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, 63, 129–138.
- Slagle, J. R. (1963). A heuristic program that solves symbolic integration problems in freshman calculus. In E. A. Feigenbaum & J. Feldman (Eds.), *Computers and thought*. New York: McGraw-Hill.
- Sleeman, D., & Brown, J. S. (Eds.) (1982). *Intelligent tutoring systems*. New York: Academic Press.
- Winograd, T. (1972). Understanding natural language. *Cognitive Psychology*, 3, 1–191.