Symbol Systems

Michael L. Anderson[1] and Don Perlis[1,2]
(1) Institute for Advanced Computer Studies
(2) Department of Computer Science
University of Maryland, College Park
College Park, MD 20742
ph: (301) 405-1746, fax: (301) 314-1353
mikeoda@cs.umd.edu, perlis@cs.umd.edu

July 23, 2002

Abstract

A symbol is a pattern (of physical marks, electromagnetic energy, etc.) which denotes, designates, or otherwise has meaning. The notion that intelligence requires the use and manipulation of symbols, and that humans are therefore symbol systems, has been extremely influential in artificial intelligence.

keywords: symbol, reasoning, intelligence, representation, logic

This article has two parts. We begin with a presentation and discussion of the idea of a physical symbol system (PSS), as formulated by Newell and Simon. This notion, and the associated physical symbol system hypothesis (PSSH), were first presented—under somewhat different names—in (Newell and Simon 1972), with later fuller formulations (Newell and Simon 1976, Newell 1980) and still later elaborations (Newell 1990). The second part consists of a discussion of various objections to PSS/PSSH and replies, especially with reference to the themes of symbol grounding, situated cognition, embodiment, and situated robotics.

PART ONE

In 1972 Allen Newell and Herbert Simon published a classic book on Human Problem Solving, in which among many other things they described information processing systems (IPSs). Then in 1975 Newell and Simon were jointly awarded the 1975 ACM Turing Award, and on that occasion they presented their paper (later published in CACM, 1976): “Computer Science as Empirical Inquiry: Symbols and Search”. In this paper they offer a more succinct description of these systems, under the now-standard name of physical symbol systems:

> One of the fundamental contributions of computer science has been to explain, at a rather basic level, what symbols are... Symbols lie at the root of intelligent action... One [structural] requirement [for intelligence] is the ability to store and manipulate symbols. (p. 114)
>
> ...
>
> A physical symbol system consists of a set of entities, called symbols, which are physical patterns that can occur as components of another type of entity called an expression (or symbol structure). Thus, a symbol structure is composed of a number of instances (or tokens) of symbols related in some physical way (such as one token being next to another)... A physical symbol system is a machine that produces through time an evolving collection of symbol structures. (p. 116)

So defined, PSSs are indeed very broadly conceived, largely because they conceive of symbols in very broad terms. Indeed it would be difficult to specify any physical entity that would not count as a symbol on their definition.

Associated with a PSS are various further notions, such as: (i) an expression in a PSS *designates* an object if the PSS's behavior depends on the object (for instance by affecting the object); and (ii) a PSS can *interpret* one of its expressions E, if E designates a process that the system can carry out.

Designation and interpretation are intended to connect the PSS to the world. That is, the PSS must be able, somehow, to ground its symbols (see "The Symbol Grounding Problem" this volume) in real referents, and to act upon them. Moreover, in allowing for the double role of object (designatee) and process in a single expression (e.g. "pull the chain" is both a sentence and a process) a PSS appears able to play the role of a self-modeling agent. This plays a key role in the development of Newell and Simon's criteria for general intelligence; see discussion below on universal machines.

With this as background, Newell and Simon assert the Physical Symbol System Hypothesis (PSSH):

> A physical symbol system has the necessary and sufficient means for general intelligent action... This is an empirical hypothesis. (p. 116)

Given how broadly a PSS is defined, human brains would seem to count as PSSs. Our brains appear to manipulate symbol structures and carry out processes on the basis of (some of) those structures, which in turn affect objects in the world (as well as other symbol structures). Thus the PSSH might seem obviously true, but Newell and Simon do not see it as obvious. Their hope is that ongoing research in artificial intelligence will succeed in "bringing forth empirical evidence" in favor of PSSH. In the remainder of their 1976 paper they discuss evidence for PSSH, especially in terms of heuristic search (see "SEARCH" in this volume), as they describe via the Heuristic Search Hypothesis:

> The solutions to problems are represented as symbol structures. A physical symbol system exercises its intelligence in problem solving by search—that is,

> by generating and progressively modifying symbol structures until it produces a solution structure.
>
> ...
>
> Physical symbol systems must use heuristic search to solve problems because such systems have limited processing resources. (p. 120)

In 1980 Newell published a paper titled "Physical symbol systems", which gives much more detail about the nature of PSSs, relating them to standard theoretical machinery. For instance, in this paper PSSs are defined simply as universal machines (in the sense of Turing machines): a machine that can be so programmed as to simulate any computationally-possible procedure whatsoever. This in turn requires that a machine be encodable as an expression that can serve as data, hence universality entails the necessity of symbolic expressions.

Although seemingly very different from the 1976 definition, Newell argues convincingly that universal machines are nothing more nor less than the PSSs he and Simon had defined earlier. He further argues that universal machines provide the first satisfactory definition of what constitutes a symbol, namely, harking back to what we saw above, it can designate something, via a given machine (or symbol system). Thus Newell turns the definitional process backwards, defining symbol in terms of symbol system. A symbol, then, is directly tied to its use in a physical context, rather than having a prior existence. As a corollary, almost anything whatsoever can be a symbol. This in turn can lead to objections, which we take up below.

Given the equivalence between a PSS and a universal machine, the PSSH implies that universality is essential for intelligent behavior. Implicit in this, and made clearer in other work (on production systems and especially SOAR) is the idea that it is the ability to represent—with symbolic expressions—one's own behavior that allows for behavioral changes and hence learning.

There are many other capacities of PSSs besides designation and interpretation, and these are spelled out by Newell in some detail. They all have to do with internal manipulation of expressions as well as input and output, to make the two basic capacities of designation and interpretation as powerful as possible.

Newell also outlines a series of levels of description of a physical computer: the device level (the electronics); the circuit level (electrical processes); the logic level (memory values and operations on them); the program level (the PSS level for a computer); and the processor-memory-switch level (the level of description of the various large-scale computer units such as memory devices, processors, I/O devices, etc). He argues that there must exist a neural (biological) level of organization that supports a symbol structure, an organization he calls an architecture; this he regards as an empirical hypothesis on a par with the PSSH.

This hypothesis is a forerunner to his definition of the Knowledge Level (Newell 1982), which is in addition to the above standard hierarchy for computers as normally envisioned. Suppose, at the program (PSS) level, one were to implement an additional level of description via a special "intelligent" program. This program would have stored knowledge (at the symbol level) and would bring that knowledge to bear on whatever problem it encountered. It would thus be a kind of reasoning engine; Newell calls this new level of behavior the knowledge level, and he formulates a principle of rationality for it: If an agent knows that one of its possible actions will achieve one of its goals, then the agent will perform that action. This is clearly an idealization, since it typically is not possible to bring (all) available knowledge to bear (due to resource limitations); Newell suggests that human cognition is at best an approximation to a knowledge-level system.

PART TWO

The computational model of human intelligence, for example as expressed in PSSH, has been extremely influential in the development of Artificial Intelligence, resulting in many techniques for simulating—and many systems which display—intelligent, if limited, behavior. These include techniques for problem solving and planning, especially various search techniques (Russell and Norvig 1995), architectures like SNePS (Shapiro 1979) and systems like SOAR (Newell 1990, Laird *et al.* 1987, "SOAR" this volume)) and Hilare II (Giralt, *et al.* 1991), as well as logic-based production systems which are time-situated and able to handle uncertainty and contradictory information (Bhatia, *et al.* 2001).

However, this approach to understanding and reproducing intelligent behavior has been increasingly criticized from a perspective which foregrounds the importance of interaction with and utilization of the external environment, and the practical orientation of real-world agents, in shaping and guiding not just particular instances of cognition but in determining the nature of cognition in general. We will provide just a brief account of the debate, below, primarily with an eye to better understanding PSSH. For a fuller account of situated and embodied cognition, and the challenge it poses to the more traditional approaches to AI, see the entries for "Embodiment", "Situated Cognition", "Situated Robotics" and "The Symbol Grounding Problem" in this volume, and (Anderson forthcoming).

To proponents of the situated or embodied approach to understanding intelligence, PSSH suggests a picture of cognition along the following lines:

1. A symbol system is characterized by a distinctive kind of decomposition of cognitive functions, such that the sensorimotor system is independent of, and functions primarily as the source of inputs to, and the target of outputs from, the reasoning system. Problem-solving proceeds in terms of temporally and conceptually distinct steps: the world is sensed (input is received from the sensory system); a model of the world is built; a plan of action is formulated via computation on the model; an action is taken (output is sent to the motor system).

2. The symbols, in terms of which the world is modeled, and by the transformations of which cognition proceeds, are meaningful primarily in terms of their internal relations to other symbols, and not in virtue of their physical relation to external objects, the behavioral dispositions of the cognitive system, or the particularities of their physical instantiation. Symbols "denote" objects or aspects of the environment, and their representative function rests on this relation of denotation.

Proponents of embodied or situated cognition (henceforth combined as simply SC) question whether intelligent organisms, humans included, satisfy the above descriptions. They do not deny that cognition involves abstract reasoning and planning, nor that humans employ symbols when engaging in this sort of reasoning. Rather, they suggest that abstract reason is only the tip of the cognitive iceberg; that it rests on and requires a great number of substantive cognitive capacities; and that these other capacities are not symbolic in nature, but rather involve states and processes that are tightly coupled to, and proceed via interaction with, the environment of the agent in question. Thus, for particular examples:

1. Whereas symbolic representation suggests an abstract relation of denotation, many internal representational states are in fact directly causally coupled to objects in, or aspects of, the environment. Thus a visual representation may cause a very particular pattern in the sight centers of the brain, which changes with changes in environment or in the relationship between the environment and the perceiver. This inner state is representational only in virtue of this continuing causal coupling, which suggests that the information about the world it contains should be understood not on a grammatical model, which involves abstract denotation of objects and their relations, but rather in terms analogous to the Watt governor, which carries information about the speed of an engine (in the angle of its arms) only in virtue of its direct coupling with that engine.

2. Whereas PSSH suggests that the functioning of the sensorimotor and reasoning components of an intelligent system can be understood largely in isolation from one another, SC maintains that the process of sensing and representing in fact involves the continual cooperation of these components, which should perhaps not therefore be presented as functionally decoupled. In general, what is worth paying perceptual attention to, and what concepts and categories are appropriate to bring to bear in representing the world, depend upon what one is doing (Clancey 1993). Further (and partly for this reason), representations tend to be cast in functional terms, "the-bee-that-is-chasing-me" rather than "bee12" (Agre and Chapman 1987, Bickhard 1993), and the world is understood, in part, in terms of the actions it invites or "affords"; a chair is perceived not in terms of abstract qualitative descriptions, but as affording sitting (Gibson 1979). This suggests that the contents and meaning of inner mental states, and the processing they undergo, cannot be understood in isolation from the ongoing activity of the representing agent. Likewise, whereas PSSH suggests that cognition should be understood in terms of the four step sense-model-plan-act cycle, and that the calculation on symbols which occurs after sensing and before acting is the meat of thinking, SC claims that in fact cognition can (and often does) involve interaction with the environment at any stage—for instance, rotating a puzzle piece to make it easier to visualize, or writing down the intermediate results in a complex mathematical calculation. Thus rather than being just the result of cognition, a given action can be *part* of the cognitive process.

3. Whereas PSSH suggests that cognition should be understood in terms of abstract rules for manipulating abstract symbols, SC maintains that cognition is in fact rooted in basic coping strategies and the embodied experience of thinking agents. Thus, even the

apparently abstract rules of logic are best understood in terms of more basic experience. (Lakoff 1987), for instance, suggests that the *exclusive or* (xor: $(p \vee q)$ & $\neg(p$ & $q))$, is rooted in and derived from our basic experience with physical objects and containers: item x can be in one box or the other, but not both.

One important response to this challenge can be found in (Vera and Simon 1993). There the authors argue that proponents of SC have interpreted the notion of a "symbol" in a restricted sense, much more narrow than originally intended, and that when a symbol system is understood more broadly, there is no necessary antithesis or tension between SC and PSSH. To get the clearest idea of what they have in mind, we quote their characterization of symbol systems at length:

> A physical symbol system is built from a set of elements, called symbols, which may be formed into symbol structures by means of a set of relations. A symbol system has a memory capable of storing and retaining symbols and symbol structures, and has a set of information processes that form symbol structures as a function of sensory stimuli, which produce symbol structures that cause motor actions and modify symbol structures in memory in a variety of ways.
>
> A physical symbol system interacts with its external environment in two ways: (1) It receives sensory stimuli from the environment that it converts into symbol structures in memory; and (2) it acts upon the environment in ways determined by symbol structures (motor symbols) that it produces. Its behavior can be influenced both by its current environment through its sensory inputs, and by previous environments through the information it has stored in memory from its experiences.
>
> Henceforth, we will usually refer to both symbols and symbol structures simply as "symbols". Symbols are patterns. In a computer, they are typically patterns of electromagnetism, but their physical nature is radically different in different computers (compare the vacuum tubes of the 1940's with integrated circuits of today). And, in any event, their physical nature is irrelevant to their role in behavior. The way in which symbols are represented in the brain is not known; presumably, they are patterns of neuronal arrangement of some kind.
>
> When we say that symbols are patterns, we mean that pairs of them can be compared (by one of the system's processes) and pronounced alike or different, and that the system can behave differently, depending on the same/different decision.
>
> We call patterns symbols when they can designate or denote. An information system can take a symbol token as input and use it to gain access to a referenced object in order to affect it or be affected by it in some way. Symbols may designate other symbols, but they may also designate patterns of sensory stimuli, and they may designate motor actions. Thus, the receipt of certain patterns of sensory stimulation may cause the creation in memory of the symbol (say, CAT) that designates a cat (not the word "cat," but the animal). Of course, this does not guarantee that there is really a cat out there: That depends on the

> veridicality of the processes that encode the stimulus into the symbol designating a cat. Similarly, a motor symbol may designate the act of "petting" (with some parameters to assure that the cat will be the object of the petting). (pp. 8-9)

It is not clear that this characterization of PSS's is very distant from that offered in the name of SC, above. In particular, Vera and Simon explicitly maintain that symbols are "patterns that designate or denote," and define such symbols separately from the patterns of sensory stimuli, in response to which symbols may be generated. It also appears to have the same general "flavor," suggesting as it does the sense-model-plan-act cycle to which SC objects; a physical symbol system operates as follows: stimuli are received, symbols which denote the environment are generated, these symbols are processed to produce more symbols, some of which designate actions, which are then sent to the motor system to cause a given motor response. However, their discussion of SC (which they are calling situated action, or SA) suggests that they may have in mind a somewhat broader definition of symbol:

> In some situations, an actor's internal representations can be extremely simple, but no one has described a system capable of intelligent action that does not employ at least rudimentary representations. Perhaps the barest representation encompasses only goals and some symbolization of a relation between goal and situation, on the one hand, and action in the other. But some representation of these is unavoidable if action is to be purposive. ...In systems like Pengi (Agre and Chapman 1987) and the creatures of Brooks (Brooks 1999), often taken as paradigmatic examples of applied SA, there are substantial internal representations, some of them used to symbolize the current focus of attention and the locations of relevant nearby objects, others used to characterize the objects themselves in terms of their current functions. (pp. 38-9)

Or, as they write on the previous page: "That the symbols in question are both goal-dependent and situation-dependent does not change their status. They are genuine symbols in the traditional information-processing sense. ...'The-bee-that-is-chasing-me' is a perfectly good symbol; it denotes a distinct class of object in the world (i.e., any bee that is engaging in the activity of chasing me)." (p. 37)

This claim seems convincing in reference to Pengi. Pengi's inner states are symbolic because they do in fact denote, albeit in terms of a functional characterization. It is much less clear that the inner states in Brooks' creatures are of this sort, nor that we should accept this apparent broadening of the term "symbol" to include *any* internal state. Genghis, Brooks' six-legged robot, achieves walking behavior with a very simple set of controllers (Brooks 1999, ch. 2). The position of each leg is represented in terms of 2 numbers, α and β, which give, respectively, its horizontal and vertical position. An α-balance machine continually sums the six α values and sends the sum to each leg, so that if one leg is forward, all legs will be sent a signal casuing them to move back slightly to compensate. There is also a simple controller attached to each individual leg, such that if the β is positive (the leg is up), it increases α by suppressing the signal from the α-balance machine. In addition there is a controller which decreases β (puts the leg down) whenever β is positive, and an up-leg trigger which can cause the leg to go up by suppressing the leg-down command. Finally, there is a walk trigger which sends a timed signal to the up-leg triggers, e.g. to cause three of the legs to go up every 2.4 seconds. When the legs go up, they move forward and down (because of the leg-forward and leg-down controllers); while the three legs on the ground move backwards (because of the α-balance machine). Then the other three legs go up, etc. This walk trigger can, in turn, be connected to sensors to cause it to run only under certain conditions, e.g. when the IR sensors register heat. Although there is certainly *representation* here, it is difficult to see this activity in terms of denotation and symbol processing; it seems much more natural to understand the representations of leg position as instances of causal coupling, and Genghis' walking behavior in reactive terms. Although one can stretch the definitions of "symbol" and "denotation" to cover this case, doing so would not thereby erase the difference between the "symbols" in Genghis, and their relation to the things they "denote," and those employed in Pengi. Every environmentally reactive system has internal states; surely some of these states are not symbolic—they neither denote nor can they be systematically combined with other inner states to form symbol structures.

Vera and Simon approach the issue of affordances—the perceived invitation by an environment to take certain actions—in a similar fashion, accepting the importance of affordances to guiding agency, but insisting that these affordances be understood symbolically, i.e., in terms of internally encoded states.

> We have already seen that when people are dealing with familiar situations, using habitual actions, their internal representations, at the conscious level, may be almost wholly functional, without any details of the mechanisms that carry out these functions. The "affordances" of the environment, represented internally, trigger actions.
>
> Of course, the absence of consciousness of mechanisms implies neither that mechanisms are absent nor that they are non-symbolic. To acquire an internal representation of an affordance, a person must carry out a complex encoding of sensory stimuli that impinge upon eye and ear. And to take the corresponding action, he or she must decode the encoded symbol representing the action into signals to the muscles. (p. 41)

Vera and Simon may be on more solid ground here, for affordances are not the same as causes; they represent the behavioral *options* offered by an environment, and only act to trigger a given action under particular cognitive circumstances. A chair affords sitting, but triggers sitting only if I want to sit down. Thus, it seems that it must be possible for these representations to be related to other complex cognitive states like desires, which may in turn play roles in longer term plans and activities. Still, it is worth pointing out that there are important distinctions to be drawn between affordances, which are often unconscious, highly action-oriented and perceptually active (in that they directly inform the content of perception) inner states, serving primarily allow substantial and inexpensive coordination between sensory input and motor responses, and abstract concepts like CAT, which although certainly grounded in sensory experience, also have substantial logical, hierarchical and lexical relations, important to their meaning, the full extent of which is *not* perceptually active or involved in sensory-motor coordination. Thus, while it *may* be plausible in some circumstances to understand the role of the latter symbol in the linear terms suggested by PSSH—whereby one senses something, which triggers the symbol CAT, which symbol is processed within a network of beliefs and desires (I like cats and like to pet them), thereby producing the symbol structure $Pet(CAT)$, which, when sent to the motor system causes my petting of the cat—it is not immediately obvious that these simple coordinations between the sensory, motor, and conceptual systems are sufficient, or of the right form, to support the deployment of affordances.

This is relevant to the larger question of how best to understand human problem solving behavior, which, as Vera and Simon accept, often involves interactions with the environment whose role is epistemic (Kirsh and Maglio 1995), i.e. part of the problem-solving itself (aimed for instance at simplifying the problem at hand, as when a puzzle piece is rotated to better "see" where it might fit) rather than an implementation of a solution reached by cognition alone. They are certainly right to insist that continual interaction with the environment is fully compatible with PSSH:

> [S]equences of actions can be executed with constant interchange among (a) receipt of information about the current state of the environment (perception), (b) internal processing of information (thinking), and (c) response to the environment (motor activity). These sequences may or may not be guided by long-term plans (or strategies that adapt to feedback of perceptual information). (p. 10)

This way of modeling environmental interaction in problem-solving treats our interaction with the puzzle piece— sense-model-look_for_fit, rotate-model-look_for_fit, rotate-model-see_fit-place—as just a variation of sense-symbolize-plan-act with more frequent attention to incoming sensory information and its changes. At a certain level of abstraction, this is no doubt the case, for surely epistemic actions involve internal processing just as do pragmatic actions. Yet it seems that, when the object of an action is to help bound a calculation, simplify a search, or otherwise change the epistemic parameters of a task, the coordinations required between the sensory, symbolic, and motor systems are somewhat different from those required in the case of pragmatic action. This is certainly not to say that the execution of epistemic actions does not require symbols and symbol processing. Indeed, it may be that epistemic actions are even more symbol-centric than pragmatic ones, for they may involve modeling not just the environment but also the self and its cognitive capacities; alternately it could be that a different decomposition of cognitive functions would diminish the need for symbolic coordination between the parts, even in the case of epistemic actions. Proponents of SC hope for the latter, while PSSH expects the former; but no one knows the answer yet.

When the debate between SC and PSSH is cast in terms of *whether* cognition depends upon symbols, or *whether* it is primarily reactive and interactive, it has a tendency to devolve into semantic questions: what is the meaning of "symbol"?; is a sense-symbolize-plan-act system really interactive? This is not the most productive impasse; for in point of fact both positions have room for symbols *and* interaction. There are nevertheless substantive differences between the two positions, which we suggest come down to three main issues:

1. The nature and role of concepts: We define a concept as a structured, contentful inner object (constituent of thought), which is semantically evaluable, re-deployable and largely stable, and which has hierarchical and logical connections to *other* concepts. Given this definition, there is a disagreement about when such inner structures are needed to explain intelligent behavior. This definition of a concept is similar to, although somewhat narrower than, the PSS definition of a symbol, suggesting that PSSH expects to discover a very central role for concepts at nearly all stages of cognition. In contrast, SC expects a larger role for non-conceptual content, which is an inner state the contents of which are best specified in terms of the abilities, skills and dispositions of the agent, or in terms of significant, non-articulable bodily or perceptual experience (Chrisley 1995).

2. The details of cognitive decomposition, and of the relations and coordinations between the parts: Although PSSH is not by definition committed to a particular implementation of intelligent agency, it has tended in practice to isolate perception and action from reasoning components, and to handle coordination between sub-systems or agents in terms of the distribution and interpretation of symbolically encoded information. In contrast, SC is committed to systems which have much more interpenetration between perception, action and reasoning systems (as in the case of affordances), and in which coordinations between sub-systems are indirect (as when all systems have access to the same sensor stream, but do not directly exchange information), mediated by the environment (as when Genghis, in lifting one leg, naturally increases the weight on the other 5 legs, allowing for the "information" that a leg has gone up to be transmitted to the other leg controllers with no direct information exchange), or sub-symbolic (Braitenberg 1984, Edelman 1992).

3. The nature and origin of higher-order cognition: Although Vera and Simon do not directly address the SC claim that the contents and rules of higher-order cognition is rooted in more basic experience (as with the claim that *xor* is derived from experience with objects and containers), it seems that such claims, insofar as they primarily pertain to the question of how such laws can be understood or learned by humans, are fully compatible with PSSH. Still, the overall project of naturalism, of which SC is one instance, faces difficulties in accounting for certain formal properties of systems like logic and arithmetic (e.g. completeness), and the apparent necessity of the truths they express (one would expect a system based on contingent experience to be likewise contingent, yet $2 + 2 = 4$ does not appear to be a contingent truth). And questions remain about the nature of abstraction (and how to implement an abstracting agent), and self-modeling, which seem necessary to high-level cognition no matter how one accounts for its origin.

REFERENCES

Anderson ML (forthcoming) Embodied cognition: A field guide. *Artificial Intelligence.*

Agre PE and Chapman D (1987) Pengi: An implementation of a theory of activity. *The proceedings of the Sixth National Conference on Artificial Intelligence (AAAI-87)* 268-272.

Bhatia M, Chi P, Chong W, Josyula D, O'Donovan-Anderson M, Okamoto Y, Perlis D and Purang K (2001) Handling uncertainty with active logic. *Proceedings of the AAAI Fall Symposium on Uncertainty in Computation.* Menlo Park, CA: AAAI Press.

Bickhard HM (1993) Representational content in humans and machines. *Journal of Experimental and Theoretical Artificial Intelligence* 5:285-333.

Brooks RA (1999) *Cambrian Intelligence: The Early History of the New AI.* Cambridge, MA: MIT Press.

Clancey WJ (1993) Situated action: A neuropsychological interpretation (Response to Vera and Simon). *Cognitive Science* 17(1):87-107.

Gibson JJ (1979) *The Ecological Approach to Visual Perception.* New York: Houghton Mifflin.

Giralt G, Alami R, Chatila R and Freedman P (1991) Remote operated autonomous robots. In *Intelligent Robotics: Proceedings of the International Symposium* 1471:416-427. Bangalore, India: International Society for Optical Engineering (SPIE).

Kirsh D, and Maglio P (1995) On distinguishing epistemic from pragmatic actions. *Cognitive Science* 18:513-549.

Laird JE, Newell A and Rosenbloom PS (1987) SOAR: An architecture for general intelligence. *Artificial Intelligence* 33(1-2):113-129.

Lakoff G (1987) *Women, Fire and Dangerous Things: What Categories Reveal About the Mind.* Chicago: University of Chicago Press.

Russell S and Norvig P (1995) *Artificial Intelligence: A Modern Approach.* Upper Saddle River, NJ: Prentice-Hall.

Newell A (1990) *Unified Theories of Cognition.* Cambridge, MA: Harvard University Press.

Newell A (1982) The knowledge level. *Artificial Intelligence* 18:87-127.

Newell A (1980) Physical symbol systems. *Cognitive Science* 4:135-183.

Newell A and Simon HA (1976) Computer science as empirical enquiry: Symbols and search. *Communications of the Association for Computing Machinery* 19(3):113-126.

Newell A and Simon HA (1972) *Human Problem Solving.* Englewood Cliffs, NJ: Prentice-Hall.

Shapiro SC (1979) The SNePS semantic network processing system. In Findler NV, ed., *Associative Networks: Representation and Use of KNowledge by Computers* 179-203. New York: Academic Press.

Vera AH and Simon HA (1993) Situated action: A symbolic interpretation. *Cognitive Science* 17:7-48.