

COGNITIVE SCIENCE

William J. Rapaport

Department of Computer Science, Department of Philosophy,
and Center for Cognitive Science
State University of New York at Buffalo, Buffalo, NY 14260

rapaport@cs.buffalo.edu

<http://www.cs.buffalo.edu/pub/WWW/faculty/rapaport/>

October 29, 1996

Abstract

This is a draft of the long version of an article to appear in Anthony Ralston, Edwin D. Reilly, and David Hemmendinger (eds.), *Encyclopedia of Computer Science, 4th edition* (New York: Van Nostrand Reinhold, 1998). It is a revised and updated version of Rapaport, William J. (1993), “Cognitive Science”, in Anthony Ralston & Edwin D. Reilly (eds.), *Encyclopedia of Computer Science, 3rd edition* (New York: Van Nostrand Reinhold): 185–189, which is an abridged version of Rapaport, William J. (1990), “Cognitive Science”, *Technical Report 90-12* (Buffalo: SUNY Buffalo Department of Computer Science). Comments on this draft are eagerly solicited. This document is *Technical Report 96-19* (Buffalo: SUNY Buffalo Department of Computer Science) and *Technical Report 96-4* (Buffalo: SUNY Buffalo Center for Cognitive Science).

1 DEFINITION

Cognitive science is the interdisciplinary study of cognition. Cognition includes mental states and processes such as thinking, reasoning, remembering, language understanding and generation, visual and auditory perception, learning, consciousness, emotions, etc. Some cognitive scientists limit their study to *human* cognition; others consider cognition independently of its implementation in humans or computers: “cognition, be it real or abstract, human or machine” (Norman 1981b: 1). Some cognitive scientists study cognition independently of the cognitive agent’s environment (cf. Jerry Fodor’s “methodological solipsism” (1980)); others study it “within the context of the person, the society, the culture” (Norman 1981b: 1; cf. Hutchins 1995ab).

Cognitive science can also be defined as, roughly, the (hopefully non-empty) intersection of the disciplines of computer science (especially artificial intelligence), linguistics, philosophy (especially philosophy of mind and philosophy of language), and psychology (especially cognitive psychology). Some writers on cognitive science add cognitive anthropology to this list, which deals in part with the human societal and cultural context mentioned above. Most would add the cognitive neurosciences, which are concerned with the “implementation” of mind in human physiology. In most other academic disciplines, (usually) a common methodology is brought to bear on a multitude of problems. By contrast, in cognitive science, many *different* methodologies—those of the several cognitive sciences—are brought to bear on a *common* problem: the nature of cognition, i.e., “intelligent cognitive behavior” (Norman 1981b: 1).

Cognitive science’s approach to the study of mind is often contrasted with that of behaviorism. The behaviorist approach to psychology seeks to describe and predict human behavior in terms of stimulus–

response correlations, with no mention of unobservable (hence, “unscientific”) mental states (including mental constructs such as symbols, ideas, or schemata) or mental processes (such as thinking, planning, etc.) that might mediate these correlations (cf. Gardner 1987: 11). A behaviorist who would be willing even to talk about the “mind” would view it as a “black box” that could only be understood—to use computational jargon—in terms of its input–output behavior. Cognitive science in general (and cognitive psychology in particular) seeks to understand human cognitive functions in terms of mental states and processes, i.e., (to use computational jargon again) in terms of the algorithms that mediate between input and output. (Nonetheless, insofar as behaviorism is concerned with the “intelligent cognitive” behaviors listed above, it, too, is a cognitive science.)

2 COGNITION AND COMPUTATION.

The notion that mental states and processes intervene between stimuli and responses sometimes takes the form of a “computational” metaphor or analogy, which is often used as the identifying mark of contemporary cognitive science: The mind is to the brain as software is to hardware; mental states and processes are (like) computer programs implemented (in the case of humans) in brain states and processes. Some cognitive scientists make the stronger claim that mental states and processes *are* (expressible as) algorithms: “cognition *is* a type of computation” (Pylyshyn 1985: xiii). Others make a weaker, but more general, claim that cognition is *computable*, i.e., that there are algorithms that have the same input–output behavior as cognitive processes (cf. Rapaport, forthcoming).

Thus, according to the computational view of cognitive science, (1) there are mental states and processes intervening between input stimuli and output responses, (2) these mental states and processes either *are* computations or else are *computable*, and—hence—(3) in contrast to behaviorism, mental states and processes are capable of being investigated scientifically (even if they are not capable of being directly observed).

Insofar as the methods of investigation are taken to be computational in nature, computer science in general and artificial intelligence in particular have come to play a central role in cognitive science. It is, however, a role not without controversial philosophical implications: For if mental states and processes can be expressed as algorithms, then they are capable of being implemented in non-human computers. The philosophical issue is simply this: Are computers executing such algorithms merely simulating mental states and processes, or are they actually exhibiting them? Do such computers think?

Those who hold to the computational view of the mind as a unifying methodology reject the definition of cognitive science as “nothing more than six disciplines in search of a grant-giving agency” (Johnson-Laird 1981: 147) or “like the proverbial blind men trying to understand the elephant . . . [or] simply a political union” (Pylyshyn 1985: xi). In contrast, others (including some computer scientists) hold the more catholic view that “‘cognitive science’ is a broad rubric, intended to include anyone who is concerned with phenomena related to mind” (Winograd 1981: 261). They “believe in the value of multiple philosophies, multiple viewpoints, multiple approaches to common issues. . . . [A] virtue of Cognitive Science is that it brings together heretofore disparate disciplines to work on common themes” (Norman 1981c: 275–276).

Even cognitive scientists who disagree about the weak vs. the strong computational view of the mind are usually willing to agree that computer programs force cognitive scientists “to make intuitions explicit and to translate vague terminology into concrete proposals; they provide a secure test of the consistency of a theory . . . ; they are ‘working models’ whose behavior can be directly compared with human performance” (Johnson-Laird 1981: 185–186; cf. Pylyshyn 1985: 76). That is, the proper methodology of cognitive science is to express one’s theories about (human) cognition in a computer program (rather than, say, in English or in the languages of mathematics, logic, or statistics, as other sciences do) and then to compare the program’s behavior with (human) cognitive behavior. Although this methodology is consistent with the denial of the strong computational view—i.e., human cognitive behavior might be *simulable* by a computer program without itself *being* computational—it *accepts* the weak form of the computational view of the mind

as at least a working hypothesis.

3 VARIETIES OF COGNITIVE SCIENCE.

Currently, there are two major “paradigms” of computational cognitive science. To lead up to these, several dichotomies (albeit overly simplified ones) can be made:

(1) Researchers who study (human) cognitive behavior either (a) believe that there *are* mental states and processes that mediate input stimuli and output responses (this position may be called “cognitivism”) or else (b) believe that there are no such mediating states or processes (or that it is unscientific to talk about any unobservable such states or processes—the position of behaviorism).

(2) Cognitivists believe either (a) that all mental states and processes are computational in nature (and here there is a further dichotomy between (i) the weak and (ii) the strong computational views) or else (b) that at least some (and perhaps all) such processes are not computational. Position (2b) is held by a number of researchers who believe that there are inherent limitations on the ability of computers to simulate or produce mental phenomena (e.g., Searle 1980, 1990; Penrose 1989; Dreyfus 1992; Edelman 1992; cf. Johnson-Laird 1988: 26). It is certainly a position that provides many of the most interesting and hardest challenges to the computational cognitivists. One such challenge is the problem of the nature of consciousness. Another is the problem of subjective qualitative experiences (“qualia”)—e.g., what kind of computational theory can account for our experience of pain or of the color green? But (2b) is also a position that is often ridiculed as “mysticism” or as a contemporary version of vitalism.

(3) The dichotomy between the two major paradigms is between (a) those computational cognitivists who believe that cognitive computations are “symbolic” and (b) those who believe that they are, rather, “connectionist”.

3.1 Symbolic Computational Cognitive Science.

The foundations of symbolic computational cognitivism may be found in the “Physical Symbol System Hypothesis” and the “Representational Theory of the Mind”. The Physical Symbol System Hypothesis, due to Allen Newell and Herbert Simon (1976), is offered as a solution to the problem of “how it is possible for mind to exist in this physical universe” (Newell 1981: 84; cf. Pylyshyn 1985: 75): Mind exists as a physically implemented “symbol system”. The concept of a physical symbol system is “the most fundamental contribution . . . of artificial intelligence and computer science to” cognitive science (Newell 1981: 38). A *symbol system* is any effectively computable procedure, i.e., a universal machine (which, by Church’s Thesis, could be a Turing machine, a recursive function, a general-purpose digital computer, etc.). A *physical symbol system* is a physical implementation of such a symbol system. The Physical Symbol System Hypothesis states that a physical system is capable of exhibiting intelligent behavior (where intelligence is defined in terms of *human* intelligence) if and only if it is a physical symbol system (cf. Newell 1981: 72). This is taken to be an empirical hypothesis, whose evidence comes from work in symbolic, i.e., non-connectionist, artificial intelligence (Newell 1981: 73). Newell argues that *intelligent* physical systems are physical *symbol* systems since intelligence requires *representations* of a wide variety of goals and states, and since such flexible representations require symbols (hence the Representational Theory of the Mind; cf. Newell 1981: 58, 62; Pylyshyn 1985: xii, 24). It is the first of these reasons—the requirement of representations—that is empirical; the second—that the representations must be symbolic—is challenged by connectionism. The converse claim, that physical *symbol* systems are capable of being *intelligent* physical systems, has been challenged by the non-computationalists of position (2b), above.

The Representational Theory of the Mind, which can be seen as a consequence of the Physical Symbol System Hypothesis, says that cognition is best understood as computations over mental representations. One particularly strong form of the Representational Theory of the Mind is Fodor’s “language of thought”

theory (1975), which says that the mental representations are a language (sometimes called “mentalese”). Fodor’s theory of *methodological solipsism* (1980) holds that the syntax of the language of thought is all that cognitive science needs to deal with, i.e., that the cognitive agent’s environment (and the input–output transducers)—while important for understanding how information gets into and out of the mind—are irrelevant for understanding how the mind works.

(4) There is, perhaps, a fourth dichotomy among the symbolic computational cognitivists, between (a) those who are satisfied with symbolic algorithms whose input–output behavior is the same as human cognitive behavior and (b) those who are only satisfied with symbolic algorithms that not only are input–output equivalent to human cognitive behavior but also are equivalent in all but the details of physical implementation, i.e., equivalent in terms of subroutines and abstract data types. A particularly strong form of (4b) also requires the algorithms to be equivalent to human cognitive behavior at the level of space and time complexity (cf. Pylyshyn 1985: xvi).

According to the Physical Symbol System Hypothesis and the Representational Theory of the Mind, when a physical system—be it computer or human—executes a “cognitive” algorithm, the representations are brought to life, so to speak, and made to behave according to the rules of the symbol system; the symbol system becomes dynamic, rather than static. If cognition is representational and rule-based in this way—i.e., if (or, more conservatively, to the extent that) cognitive behavior consists of transformations of representations according to rules—then a computer that behaves according to (physical implementations of) these rules causally applied to (physical implementations of) these representations *is* behaving cognitively and is not merely simulating cognitive behavior (cf. Sokolowski 1988).

Although the Physical Symbol System Hypothesis and the Representational Theory of the Mind offer an answer, which is satisfying to most computer scientists, to Descartes’s question of how mind and body can interact (namely, mind can be implemented in body), they are not without their detractors. Of particular note are the objections of Terry Winograd, who did pioneering work in the symbolic paradigm of artificial intelligence. Winograd cites a biologist, Humberto Maturana, who straightforwardly denies the Representational Theory of the Mind: “cognition is not based on the manipulation of mental models or representations of the world” (Winograd 1981: 248). Instead, according to Winograd and Maturana, there are cognitive “phenomena that *for an observer* can be described in terms of representation, but that can also be understood as the activity of a structure-determined system with no mechanism corresponding to a representation” (Winograd 1981: 249). This view echoes the “intentional stance” theory of the philosopher Daniel Dennett (who has many more sympathies with computational cognitivism). According to Dennett, it makes sense to treat certain complex systems (e.g., chess-playing computers) *as if* they had beliefs and acted intentionally even though there might not be anything in their structure that corresponded in any way to beliefs or intentions (Dennett 1978). Recent work in non-representational artificial intelligence and “situated cognition” is consistent with this approach (cf. Brooks 1991a, Kirsh 1991, *Cognitive Science* 1993).

3.2 Connectionist Computational Cognitive Science.

The “connectionist” (or “neural network”, or “parallel distributed processing”) approach to artificial intelligence and computational cognitive science can be seen as one way for a system to (appear to) behave intelligently without being a “symbol system” and yet be computational. On this approach, large numbers of very simple processors (“nodes”) are connected in multiple ways by communication links of varying strengths. Input nodes receive information from the external world. The information is propagated along the links to and among intermediate (or “hidden”) nodes, finally reaching output nodes. If the output is not what was expected (e.g., if it does not match a “training set” of sample input–output pairs), the strengths of the links are adjusted (by a variety of automatic techniques, such as “back propagation” and “simulated annealing”). This process is repeated until the system “settles down” into a stable configuration that exhibits the desired (cognitive) behavior. Connectionist systems and techniques have been developed for learning features of natural language, for aspects of visual perception, and for a number of other cognitive (as well as non-cognitive) phenomena. There is a wide range of types of connectionist methods, many of which are, in fact, highly representational, but most of which are “distributively representational”, by which is meant

that the kind of information that a *symbolic* artificial-intelligence program would represent using various symbolic knowledge-representation techniques is, instead, “represented” by the strengths and connectivity patterns of the links. Rather than having intelligence “programmed” into the system using explicit rules and representations, intelligence is sometimes held to “emerge” from the organization of the nodes and links. (Useful tutorials are Knight 1989, 1990; good surveys of connectionism are *Cognitive Science* 1985; Graubard 1988; and—from a critical standpoint—Pinker & Mehler 1988 and the debate between Fodor and Zenon Pylyshyn (1988) and Smolensky (1991).)

3.3 Compromises between These Approaches.

There is a tendency for advocates of one of these two computational cognitive paradigms to assert that it is better than the other, usually on the grounds that some cognitive behavior has been implemented more or less successfully in the one but not (or not yet) in the other. There is also a growing tendency for some researchers to seek a “meeting of the minds”, usually on the grounds that the cognitive behaviors that have been implemented more or less successfully using the one paradigm are precisely the ones that have not (yet) been successfully implemented using the other. In particular, cognitive processes that are easy to implement symbolically (e.g., problem solving, reasoning, game playing, certain aspects of linguistic competence) tend to be ones that are relatively difficult for humans or that have to be explicitly taught, while those that have proven difficult to implement symbolically (e.g., certain aspects of visual perception and learning) tend to be those that “come naturally” to humans. This paradox of (symbolic) artificial intelligence has its counterpart in the debate over connectionism. The processes that have proven difficult to implement symbolically appear to be susceptible to connectionist techniques. The construction of “hybrid” symbolic/connectionist systems based on this complementarity may prove to be a major advance in our understanding of cognition.

A second way of merging the two approaches is to view connectionism as a lower level of cognitive processing, i.e., as a way of implementing symbolic processes. Thus, e.g., logical reasoning, which is well suited for symbolic computing, might be implemented in a connectionist system in such a way that certain of its “connectivity patterns” reliably represent precisely the things that would be explicitly represented symbolically. Some proponents of the Physical Symbol System Hypothesis would put this slightly differently. They would say that at some stage in the sequence of levels that describes a computer (beginning with the “device level”—the description in electronic terms), there must be a level that implements a symbol system. (Cf. Newell 1981: 76, Pylyshyn 1985: 96, Clark 1990.) Several connectionist implementations of “symbolic” algorithms have been investigated, but as yet there is no general theory of how this might be accomplished.

4 COGNITIVE SCIENCE RESEARCH.

One way of exploring the content of cognitive science is to look at research in the individual cognitive-science disciplines that could equally well be considered research in cognitive science *per se*.

4.1 Artificial Intelligence.

Given the computational view of cognitive science, it is arguable that all research in artificial intelligence is also research in cognitive science. Nonetheless, certain applications of artificial intelligence techniques to problems in engineering, management, etc., and, perhaps, the development of expert systems do not seem to fall within the scope of cognitive science. Certainly, however, those aspects of artificial intelligence research that might be considered to be “computational psychology” or “computational philosophy” are also in the domain of cognitive science (cf. Shapiro 1993). Among these are the following: the early work by Newell and Simon on problem solving (the Logic Theorist, the General Problem Solver; cf. Newell et al. 1957, 1963), as well as recent work on the SOAR project (Newell 1990); aspects of knowledge representation that attempt to reflect cognition (e.g., some uses of semantic networks, such as M. Ross Quillian’s original theory (1968)

and more recent systems such as Stuart C. Shapiro’s SNePS (Shapiro 1979; Shapiro & Rapaport 1992), and John R. Anderson’s ACT* systems (cf. Anderson 1995); Marvin Minsky’s theory of frames (cf. Minsky 1981); Roger Schank’s theory of scripts and conceptual dependency (Schank & Abelson 1977); work on “naive” or “qualitative” physics, which attempts to develop systems that can reason about physics in ways that humans do on an everyday basis (rather than in ways that professional physicists do) (cf. Hayes 1985); machine learning; planning; reasoning; natural-language understanding and generation; and computational vision. (For surveys of these and other topics, cf. Shapiro 1992.)

4.2 Linguistics.

Linguistics is another discipline that is arguably wholly subsumed by cognitive science (at least, to the extent that computationalism is de-emphasized). After all, language is often held to be the “mirror of the mind”—the (physical) means for one mind to communicate its thoughts to another. But it was with the development of transformational grammar by Chomsky that cognitivism replaced behaviorism in linguistics. Subsequent work on a variety of computationally tractable “successors” to transformational grammar (e.g., head-driven phrase-structure grammar; Pollard & Sag 1994) and other work in computational linguistics is clearly a part of cognitive science (cf. Winograd 1972, Allen 1995), as is work by such cognitive linguists as George Lakoff (e.g., his work on metaphors (Lakoff & Johnson 1980) and on natural categories (Lakoff 1987)) and Leonard Talmy (e.g., his work on “fictive motion”, Talmy 1996). (For surveys of recent contributions by linguistics to cognitive science, see Lakoff 1987, Pinker 1994.)

4.3 Philosophy.

Philosophers have long studied the nature of mind and language, and much recent work in the philosophy of mind, the philosophy of language, epistemology, and the philosophy of consciousness¹ has been informed by research in the other cognitive-science disciplines (cf. Beakley & Ludlow 1992). In addition to the work by Dennett and Fodor mentioned earlier, and Hilary Putnam’s (1960) theory of Turing-machine functionalism as a solution to the classic philosophical problem of the relationship of mind and body, the objections to the nature and possibility of success of artificial intelligence that have been raised by philosophers have served as research goals for artificial intelligence researchers—such criticisms must also be considered as part of cognitive science. The two major lines of criticism are those due to Hubert Dreyfus and John Searle. Dreyfus (1992) argued, on the basis of the phenomenological school of philosophy, that since computers do not have a (human) body, do not have (human) purposes or needs, and do not share the human cultural milieu, they will never be truly intelligent. Searle (1980) has argued that the Turing Test, although perhaps an indicator of the presence of intelligent *behavior*, fails as an indicator of the presence of intelligence. His “Chinese Room Argument” purports to show that a computer cannot understand natural language: Suppose that an English-speaking human who knew no Chinese was locked in a room and equipped with a program (written in English) for manipulating Chinese ideographs in such a way as to convince native Chinese speakers that they were communicating with another native speaker of Chinese. Such a person, according to Searle, would pass the Turing Test, yet (by hypothesis) would not understand Chinese. (For an excellent, annotated bibliography on the philosophy of mind, see Chalmers 1996.)

4.4 Psychology.

Cognitive psychology, of course, is a central cognitive-science discipline. Among recent work by cognitive psychologists that has been influential in cognitive science are: Eleanor Rosch’s theory of categorization (1978), Philip N. Johnson-Laird’s theory of reasoning based on mental models (Johnson-Laird & Byrne

¹Recent sources of information on the philosophy of consciousness include: Jackendoff 1987; Lycan 1987, 1996; Dennett 1991; Flanagan 1992; Searle 1992; Tye 1995; and Chalmers 1996).

1991), Stephen Kosslyn’s theory of mental imagery as a non-propositional kind of mental representation (1994), and John R. Anderson’s theory of human associative memory (1995).

4.5 Cognitive Science.

But perhaps the most important research topics in cognitive science are those that are truly interdisciplinary, i.e., in which researchers from the several cognitive sciences apply their differing methodologies to a common problem and, conversely, inform their own studies with results of investigations from the complementary disciplines. Prime examples of these would be: (1) research in visual perception, which has been investigated in psychology, in artificial intelligence (cf. Marr 1982), and in robotics, not to mention in physiology and biophysics (cf. Leibovic 1990); (2) research into mental imagery, which, in addition to the work in psychology mentioned above, has received critical philosophical attention from Pylyshyn and Dennett (cf. Block 1981) and has also been investigated using neuroscientific techniques; (3) research on categorization, where results from psychology (notably the work of Rosch (1978))—influenced by philosophical studies of family resemblance (Wittgenstein 1953) and natural kinds (cf. Wolterstorff 1970)—have largely overturned the “classical” philosophical view going back to Aristotle of there being necessary and sufficient conditions for membership in a category; (4) research on the logic of belief and knowledge, in which people from artificial intelligence have not only adapted, for use in artificial intelligence programs, systems of epistemic and doxastic logics developed by philosophers, but have also offered solutions to many open problems in these logics that philosophers have largely ignored (cf. Fagin et al. 1995); (5) research on cognition and emotion, which has been studied by psychologists and AI researchers (cf. Ortony 1988; Bates 1994; Sloman & Poli 1996); and (6) a research project on indexicality and narrative understanding that has involved input from AI, communicative disorders, literary theory, geography, linguistics, philosophy, and psychology (Duchan et al. 1995).

5 THE FUTURE OF COGNITIVE SCIENCE.

If cognitive science is to become a discipline in its own right, and not just a congeries of parts of other disciplines, perhaps its best hope lies not only in such multi-pronged attacks on common problems, as just discussed, but in single research groups whose members come from different disciplines yet who work together on common problems of cognition. The range of disciplines, and the levels of analysis, are by no means settled and, indeed, are widening in scope. As Donald Norman (1981b: 3) observes, “To some, the very essence of a cognitive system is that of a symbol processing system”, while to others the essence is that of connectionist neural networks. To yet others, the essence of cognitive science is more holistic, viewing the mind as an integral component of the larger world—of society, of culture—*not* (solely) understandable in terms of symbol manipulation (syntax) but in need of a semantics—an understanding of the relations of the mental symbols to the external world. Thus, two major open issues for a complete understanding of cognition are—looking inwards—how the mind is implemented and how the very fact of its implementation in particular kinds of physical or biological mechanisms influences the nature of cognition, and—looking outwards—how and to what extent the nature of cognition is shaped by the socio-cultural world that minds find themselves in.

6 HISTORY OF COGNITIVE SCIENCE.

Both the symbolic and connectionist approaches to computational cognitive science can trace their origins to two major lines of investigation. First, there was the development of symbolic logic at the turn of the century and Warren McCulloch and Walter Pitts’s application of logic to the analysis of the behavior of neural networks (1943). Second, there were Alan Turing’s analyses of computation (1936) and—using the Imitation Game (now known as the Turing Test)—of whether computers could think (1950). Cognitivism

burst upon the scene in 1956. In (or very near) that year, the following cognitive theories appeared: George Miller’s theory of human short-term memory (Miller 1956), Noam Chomsky’s analysis of formal grammars (Chomsky 1957), Jerome Bruner and colleagues’ study of thinking (Bruner et al. 1956), and Newell and Simon’s Logic Theorist—the first artificial intelligence program (presented at the first artificial intelligence conference, at Dartmouth, in 1956, organized by Minsky and John McCarthy; cf. Newell et al. 1957). In 1979, the journal *Cognitive Science* appeared; two years later, the first annual meeting of the Cognitive Science Society was held (its proceedings are now published by Lawrence Erlbaum Associates). Other major cognitive-science journals include: *Behavioral and Brain Sciences*, *Cognition*, *Linguistics and Philosophy*, *Mind and Language*, *Minds and Machines*, and *Philosophical Psychology*; in addition, most journals in the specific cognitive-science disciplines also have articles on cognitive science. Finally, there has been a surge recently of research centers and institutes of cognitive science, as well as graduate and undergraduate degree programs, including university departments of cognitive science. (For a more detailed history, see Gardner 1987.)

Central sources of information on cognitive science include Norman 1981, Pylyshyn 1985, Gardner 1987, Collins & Smith 1988, Johnson-Laird 1988, Posner 1989, Leiber 1991, Goldman 1993, von Eckardt 1993, Luger et al. 1994, Osherson 1995, Stillings et al. 1995, and Thagard 1996. On the World Wide Web, see, e.g., “Cognitive and Psychological Sciences on the Internet” [<http://www-psych.stanford.edu/cogsci/>] and “The World-Wide Web Virtual Library: Cognitive Science” [<http://www.cog.brown.edu/pointers/cognitive.html>].

7 ACKNOWLEDGMENTS.

I am grateful to the following of my cognitive-science colleagues at UB for comments on earlier drafts: Jean-Pierre Koenig, K. Nicholas Leibovic, Stuart C. Shapiro, Jim Swan, Deborah K. W. Walters, and David Zubin.

8 REFERENCES.

1. Allen, James (1995), *Natural Language Understanding, 2nd edition* (Redwood City, CA: Benjamin-Cummings).
2. Anderson, John R. (1995), *Learning and Memory: An Integrated Approach* (New York: John Wiley & Sons).
3. Bates, Joseph (1994), “The Role of Emotion in Believable Agents”, *Communications of the ACM* 37.7 (July): 122–125.
4. Beakley, Brian, & Ludlow, Peter (eds.), *The Philosophy of Mind: Classical Problems/Contemporary Issues* (Cambridge, MA: MIT Press).
5. Block, Ned (ed.) (1981), *Imagery* (Cambridge, MA: MIT Press).
6. Brooks, Rodney A. (1991a), “Intelligence without Representation,” *Artificial Intelligence* 47: 139–159.
7. Brooks, Rodney A. (1991b), “Intelligence without Reason,” *Proceedings of the 12th International Joint Conference on Artificial Intelligence (IJCAI-91; Sydney, Australia)* (San Mateo, CA: Morgan Kaufmann): 569–595.
8. Bruner, Jerome S.; Goodnow, Jacqueline J.; & Austin, George A. (1956), *A Study of Thinking* (New York: John Wiley).
9. Chomsky, Noam (1957), *Syntactic Structures* (The Hague: Mouton).

10. Chalmers, David J. (1996), "Contemporary Philosophy of Mind: An Annotated Bibliography," [<http://ling.ucsc.edu/~chalmers/biblio.html>].
11. Chalmers, David J. (1996), *The Conscious Mind: In Search of a Fundamental Theory* (New York: Oxford University Press).
12. Clark, Andy (1989), "Beyond Eliminativism," *Mind and Language* 4: 251–279.
13. *Cognitive Science* (1985), Special Issue on Connectionist Models and Their Applications, Vol. 9, No. 1 (January–March).
14. *Cognitive Science* (1993), Special Issue on Situated Action, Vol. 17, No. 1 (January–March).
15. Collins, Allan, & Smith, Edward E. (eds.) (1988), *Readings in Cognitive Science: A Perspective from Psychology and Artificial Intelligence* (San Mateo, CA: Morgan Kaufmann Publishers).
16. Dennett, Daniel C. (1978), "Intentional Systems," in Daniel C. Dennett, *Brainstorms* (Montgomery, VT: Bradford Books): 3–22.
17. Dennett, Daniel C. (1991), *Consciousness Explained* (Boston: Little, Brown and Co.).
18. Dreyfus, Hubert L. (1992), *What Computers Still Can't Do: A Critique of Artificial Reason* (Cambridge, MA: MIT Press).
19. Duchan, Judith F.; Bruder, Gail A.; & Hewitt, Lynne E. (eds.) (1995), *Deixis in Narrative: A Cognitive Science Perspective* (Hillsdale, NJ: Lawrence Erlbaum Associates).
20. Edelman, Gerald M. (1992), *Bright Air, Brilliant Fire: On the Matter of the Mind* (New York: Basic Books).
21. Fagin, Ronald; Halpern, Joseph Y.; Moses, Yoram; & Vardi, Moshe Y. (1995), *Reasoning about Knowledge* (Cambridge, MA: MIT Press).
22. Flanagan, Owen (1992), *Consciousness Reconsidered* (Cambridge, MA: MIT Press).
23. Fodor, Jerry A. (1975), *The Language of Thought* (New York: Thomas Y. Crowell Co.).
24. Fodor, Jerry A. (1980), "Methodological Solipsism Considered as a Research Strategy in Cognitive Psychology," *Behavioral and Brain Sciences* 3: 63–109.
25. Fodor, Jerry A., & Pylyshyn, Zenon W. (1988), "Connectionism and Cognitive Architecture", *Cognition* 28: 3–71; reprinted in Steven Pinker & Jacques Mehler (eds.), *Connections and Symbols* (Cambridge, MA: MIT Press, 1988).
26. Gardner, Howard (1987), *The Mind's New Science: A History of the Cognitive Revolution, with a New Epilogue by the Author: Cognitive Science after 1984* (New York: Basic Books).
27. Goldman, Alvin I. (1993), *Readings in Philosophy and Cognitive Science* (Cambridge, MA: MIT Press).
28. Graubard, Stephen R. (ed.) (1988), "Artificial Intelligence," special issue of *Daedalus*, Vol. 117, No. 1 (Winter); reprinted as *The Artificial Intelligence Debate: False Starts, Real Foundations* (Cambridge, MA: MIT Press, 1988).
29. Hayes, Patrick J. (1985), "The Second Naive Physics Manifesto," in Jerry R. Hobbs & Robert C. Moore (eds.), *Formal Theories of the Commonsense World* (Norwood, NJ: Ablex Publishing Corp.): 1–36.
30. Hutchins, Edwin (1995a), *Cognition in the Wild* (Cambridge, MA: MIT Press).
31. Hutchins, Edwin (1995b), "How a Cockpit Remembers Its Speeds," *Cognitive Science* 19: 265–288.
32. Jackendoff, Ray (1987), *Consciousness and the Computational Mind* (Cambridge, MA: MIT Press).

33. Johnson-Laird, Philip N. (1981), "Mental Models in Cognitive Science," in Donald A. Norman (ed.), *Perspectives on Cognitive Science* (Norwood, NJ: Ablex Publishing Corp.): 147–191.
34. Johnson-Laird, Philip N. (1988), *The Computer and the Mind: An Introduction to Cognitive Science* (Cambridge, MA: Harvard University Press).
35. Johnson-Laird, Philip N., & Byrne, Ruth M. J. (1991), *Deduction* (Hillsdale, NJ: Lawrence Erlbaum Associates).
36. Kintsch, Walter; Miller, James R.; & Polson, Peter G. (eds.) (1984), *Methods and Tactics in Cognitive Science* (Hillsdale, NJ: Lawrence Erlbaum Associates).
37. Kirsh, David (1991), "Today the Earwig, Tomorrow Man?," *Artificial Intelligence* 47: 161–184.
38. Knight, Kevin (1989), "A Gentle Introduction to Subsymbolic Computation: Connectionism for the A.I. Researcher," *Technical Report CMU-CS-89-150* (Pittsburgh: Carnegie Mellon University, School of Computer Science).
39. Knight, Kevin (1990), "Connectionist Ideas and Algorithms," *Communications of the ACM* Vol. 33, No. 11, pp. 59–74.
40. Kosslyn, Stephen M. (1994), *Image and Brain: The Resolution of the Imagery Debate* (Cambridge, MA: MIT Press).
41. Lakoff, George (1987), *Women, Fire, and Dangerous Things: What Categories Reveal about the Mind* (Chicago: University of Chicago Press).
42. Lakoff, George, & Johnson, Mark (1980), *Metaphors We Live By* (Chicago: University of Chicago Press).
43. Leiber, Justin (1991), *An Invitation to Cognitive Science* (Cambridge, MA: Basil Blackwell).
44. Leibovic, K. Nicholas (ed.) (1990) *Science of Vision* (New York: Springer-Verlag).
45. Luger, George F.; with Johnson, Peder; Stern, Carl; Newman, Jean E.; & Yeo, Ronald (1994), *Cognitive Science: The Science of Intelligent Systems* (San Diego: Academic Press).
46. Lycan, William G. (1987), *Consciousness* (Cambridge, MA: MIT Press).
47. Lycan, William G. (1996), *Consciousness and Experience* (Cambridge, MA: MIT Press).
48. Marr, David (1982), *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information* (New York: W. H. Freeman).
49. McCulloch, Warren S., & Pitts, Walter H. (1943), "A Logical Calculus of the Ideas Immanent in Nervous Activity," *Bulletin of Mathematical Biophysics* 7: 115–133; reprinted in Warren S. McCulloch, *Embodiments of Mind* (Cambridge, MA: MIT Press, 1965): 19–39.
50. Miller, George A. (1956), "The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information," *Psychological Review* 63: 81–97.
51. Minsky, Marvin (1981), "A Framework for Representing Knowledge," in John Haugeland (ed.), *Mind Design: Philosophy, Psychology, Artificial Intelligence* (Montgomery, VT: Bradford Books): 95–128.
52. Newell, Allen (1981), "Physical Symbol Systems," in Donald A. Norman (ed.), *Perspectives on Cognitive Science* (Norwood, NJ: Ablex Publishing Corp.): 37–85.
53. Newell, Allen (1990), *Unified Theories of Cognition* (Cambridge, MA: Harvard University Press).
54. Newell, Allen; Shaw, J. C.; & Simon, Herbert A. (1957), "Empirical Explorations with the Logic Theory Machine," *Proceedings of the Western Joint Computer Conference* (New York: IEEE); reprinted in Edwin A. Feigenbaum & Jerome Feldman (eds.), *Computers and Thought* (New York: McGraw-Hill).

55. Newell, Allen; Shaw, J. C.; & Simon, Herbert A. (1963), "GPS: A Program That Simulates Human Thought," in Edwin A. Feigenbaum & Jerome Feldman (eds.), *Computers and Thought* (New York: McGraw-Hill).
56. Newell, Allen, & Simon, Herbert A. (1976), "Computer Science as Empirical Inquiry: Symbols and Search," *Communications of the ACM* 19: 113–126.
57. Norman, Donald A. (ed.) (1981a), *Perspectives on Cognitive Science* (Norwood, NJ: Ablex Publishing Corp.).
58. Norman, Donald A. (1981b), "What Is Cognitive Science?," in Donald A. Norman (ed.), *Perspectives on Cognitive Science* (Norwood, NJ: Ablex Publishing Corp.): 1–11.
59. Norman, Donald A. (1981c), "Twelve Issues for Cognitive Science," in Donald A. Norman (ed.), *Perspectives on Cognitive Science* (Norwood, NJ: Ablex Publishing Corp.): 265–295.
60. Ortony, Andrew (ed.) (1988), *The Cognitive Structure of Emotions* (Cambridge, UK: Cambridge University Press).
61. Osherson, Daniel N. (general ed.) (1995), *An Invitation to Cognitive Science, 2nd edition*, 4 volumes (Cambridge, MA: MIT Press).
62. Penrose, Roger (1989), *The Emperor's New Mind: Concerning Computers, Minds, and the Laws of Physics* (Oxford: Oxford University Press).
63. Pinker, Steven (1994), *The Language Instinct: How the Mind Creates Language* (New York: William Morrow).
64. Pinker, Steven, & Mehler, Jacques (eds.) (1988), *Connections and Symbols* (Cambridge, MA: MIT Press).
65. Pollard, Carl, & Sag, Ivan A. (1994), *Head-Driven Phrase Structure Grammar* (Chicago: University of Chicago Press).
66. Posner, Michael I. (ed.) (1989), *Foundations of Cognitive Science* (Cambridge, MA: MIT Press).
67. *Proceedings of the Annual Conference of the Cognitive Science Society* (1981–) (Hillsdale, NJ: Lawrence Erlbaum Associates).
68. Putnam, Hilary (1960), "Minds and Machines," in Sidney Hook (ed.), *Dimensions of Mind* (New York: New York University Press): 148–179; reprinted in Hilary Putnam, *Mind, Language and Reality* (Cambridge, U.K.: Cambridge University Press): 362–385.
69. Pylyshyn, Zenon (1985), *Computation and Cognition: Toward a Foundation for Cognitive Science; 2nd edition*, (Cambridge, MA: MIT Press).
70. Quillian, M. Ross (1968), "Semantic Memory," in Marvin Minsky (ed.), *Semantic Information Processing* (Cambridge, MA: MIT Press): 227–270.
71. Rapaport, William J. (forthcoming), "How Minds Can Be Computational Systems", *Journal of Experimental and Theoretical Artificial Intelligence*; preprinted as *Technical Report 96-10* (Buffalo: SUNY Buffalo Department of Computer Science) and *Technical Report 96-1* (Buffalo: SUNY Buffalo Center for Cognitive Science).
72. Rosch, Eleanor (1978), "Principles of Categorization," in Eleanor Rosch & Barbara B. Lloyd (eds.), *Cognition and Categorization* (Hillsdale, NJ: Lawrence Erlbaum Associates): 27–48.
73. Schank, Roger C., & Abelson, Robert P. (1977), *Scripts, Plans, Goals, and Understanding* (Hillsdale, NJ: Lawrence Erlbaum Associates).

74. Searle, John R. (1980), "Minds, Brains, and Programs," *Behavioral and Brain Sciences* 3: 417–457.
75. Searle, John R. (1990), "Is the Brain's Mind a Computer Program?" *Scientific American* (January 1990): 26–31.
76. Searle, John R. (1992), *Rediscovery of the Mind* (Cambridge, MA: MIT Press).
77. Shapiro, Stuart C. (1979), "The SNePS Semantic Network Processing System," in Nicholas Findler (ed.), *Associative Networks: Representation and Use of Knowledge by Computers* (New York: Academic Press): 179–203.
78. Shapiro, Stuart C. (ed.) (1992), *Encyclopedia of Artificial Intelligence, 2nd edition* (New York: John Wiley & Sons).
79. Shapiro, Stuart C. (1993), "Artificial Intelligence", in Anthony Ralston & Edwin D. Reilly (eds.), *Encyclopedia of Computer Science, 3rd edition* (New York: Van Nostrand Reinhold): 87–90.
80. Shapiro, Stuart C., & Rapaport, William J. (1992), "The SNePS Family," *Computers and Mathematics with Applications* 23: 243–275; reprinted in Fritz Lehmann (ed.), *Semantic Networks in Artificial Intelligence* (Oxford: Pergamon Press, 1992): 243–275.
81. Sloman, Aaron, & Poli, Riccardo (1996), "SIM-AGENT: A Toolkit for Exploring Agent Designs", in Mike Wooldridge, Joerg Mueller, & Milind Tambe (eds.), *Intelligent Agents, Vol. II (ATAL-95)* (Berlin: Springer-Verlag): 392–407.
82. Smolensky, Paul (1991), "Connectionism, Constituency, and the Language of Thought", in Barry Loewer & Georges Rey (eds.), *Meaning in Mind: Fodor and His Critics* (Cambridge, MA: Basil Blackwell).
83. Sokolowski, Robert (1988), "Natural and Artificial Intelligence," *Daedalus*, Vol. 117, No. 1 (Winter 1988): 45–64.
84. Stillings, Neil A.; Weisler, Steven E.; Chase, Christopher H.; Feinstein, Mark H.; Garfield, Jay L.; & Rissland, Edwina L. (1995), *Cognitive Science: An Introduction; 2nd edition* (Cambridge, MA: MIT Press).
85. Talmy, Leonard (1996), "Fictive Motion in Language and 'Ception'", in Paul Bloom, Mary A. Pearson, Lynn Nadel, & Merrill F. Garrett (eds.), *Language and Space* (Cambridge, MA: MIT Press): 211–276.
86. Thagard, Paul (1996), *Mind: Introduction to Cognitive Science* (Cambridge, MA: MIT Press).
87. Turing, Alan M. (1936), "On Computable Numbers, with an Application to the Entscheidungsproblem," *Proceedings of the London Mathematical Society*, Ser. 2, Vol. 42, pp. 230–265; reprinted in Martin Davis (ed.), *The Undecidable: Basic Papers on Undecidable Propositions, Unsolvable Problems, and Computable Functions* (New York: Raven Press, 1965): 116–154.
88. Turing, Alan M. (1950), "Computing Machinery and Intelligence," *Mind* 59; reprinted in Alan Ross Anderson (ed.), *Minds and Machines* (Englewood Cliffs, NJ: Prentice-Hall, 1964): 4–30.
89. Tye, Michael (1995), *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind* (Cambridge, MA: MIT Press).
90. von Eckardt, Barbara (1993), *What is Cognitive Science?* (Cambridge, MA: MIT Press).
91. Winograd, Terry (1981), "What Does It Mean to Understand Language?", in Donald A. Norman (ed.), *Perspectives on Cognitive Science* (Norwood, NJ: Ablex Publishing Corp.): 231–263.
92. Winograd, Terry (1983), *Language as a Cognitive Process* (Reading, MA: Addison-Wesley).
93. Wittgenstein, Ludwig (1953), *Philosophical Investigations* (New York: Macmillan).
94. Wolterstorff, Nicholas (1970), *On Universals: An Essay in Ontology* (Chicago: University of Chicago Press).