Situated Cognition: Empirical Issue, 'Paradigm Shift' or Conceptual Confusion?

Peter Slezak
Program in Cognitive Science
University of New South Wales
Kensington, Sydney NSW 2052
peters@huon.cs.unsw.edu.au

Abstract

The self-advertising, at least, suggests that 'situated cognition' involves the most fundamental conceptual re-organization in AI and cognitive science, even appearing to deny that cognition is to be explained by mental representations. A. Vera and H. Simon have rebutted many of these claims, but they overlook an important reading of situated arguments which may, after all, involve a genuinely revolutionary insight.

Introduction: Paradigm Shift?

The recent special issue of *Cognitive Science* (Vol. 17, No. 1, Jan. 1993) is evidence of the growing importance of the radical claims of 'situated action', but it is also evidence of the peculiarity of the debate surrounding these developments. The exchanges reveal an unusual recalcitrance in which both sides charge each other with misrepresentations and confusions. For example, defending the orthodox symbol system approach to cognition against Clancey's (1993) critique, Vera and Simon show a certain impatience with philosophizing and insist that "Cognitive psychology is an empirical science", recommending that we "get out of the armchair" and "get on with the job" (1993, p. 132). However, this seemingly unexceptionable advice will be to no avail when the problems arise from conceptual confusions rather than straightforward empirical issues. In this sense, Vera and Simon are begging the question against their critics by complaining of assertions "almost entirely unbuttressed by empirical evidence" (1993, 118) and by simply insisting that "cognitive science proceeds exactly as biology, chemistry, geology or physics do" (119). As Kuhn (1970) has famously shown, such sciences proceed in the way that Vera and Simon suggest only during periods of "normal science" when fundamental principles are not themselves in question. Rightly or wrongly, situated cognition, in some of its many guises, professes to be challenging the very foundational assumptions or Kuhnian "paradigm" of cognitive science.

While it is undeniable that there is much justice in the substantive counter-criticisms of Vera and Simon, I believe that they have missed the very point and character of a central situated criticism. In the nature of the case, the available evidence will look the same on either approach - just as the available observations at the time of Copernicus's *De Revolutionibus* did not decisively distinguish it from the Ptolemaic system. To this extent, though not articulated in detail, there is some justification for Agre's (1993) talk of a clash of 'worldviews'. Accordingly, it is worth briefly indicating the warrant for seeing the dispute in these, admittedly abused and overused, Kuhnian terms.

Physical Symbol System Hypothesis

In their classic articulation of the 'Physical Symbol System Hypothesis', Newell and Simon (1976) characterize it as a "Law of Qualitative Structure" whose other examples include the cell doctrine in biology, plate tectonics in geology, the doctrine of atomism, and Darwinian evolution by natural selection. This statement and the practice of cognitive science itself, leave no doubt the foundational status of the Hypothesis, - a fact which helps to explain the characteristic talking past one another among its critics and defenders. It is striking to notice the aptness of Kuhn's portrayal of such disputes in which "two scientific schools disagree about what is a problem and what is a solution" and therefore "they will inevitably talk through each other when debating the relative merits of their respective paradigms" (1970, 109-10).

Thus, Vera and Simon strenuously reject Clancey's charge of conflating different notions of symbols, but there is a certain construal of this criticism which Vera and Simon miss. The five categories of symbol distinguished by Vera and Simon do not suffice to absolve them of the charge because what is at stake is the precise *conception* of these categories and not their mere distinctness as Vera and Simon assume. Clancey's (admittedly somewhat unclear) objection concerns a conflation of the special properties of symbols in one category with those of another. The clearest evidence for this inadequacy of Vera and Simon's response is in their treatment of the very central properties of symbols which they identify as those of being patterned and having denotation (1993, 125). But "being patterned and denoting" are hardly unproblematic. It is precisely this issue which is the vexed problem of intentionality or "psychosemantics" which has preoccupied philosophers for the past decade or so. Jerry Fodor's (1980) well-known paper on 'Methodological Solipsism' and the ensuing literature have been concerned with explaining how semantics or meaning can be explained in terms of the purely syntactic patterns of physical symbols. The burgeoning philosophical discussions are ample evidence of the fact that there is a remarkable consensus, at least, on Fodor's judgement that "of the semanticity of mental representations we have, as things now stand, no adequate account" (1985, p. 28).

Typically, Stalnaker (1991, 229) says "There is little agreement about how to do semantics, or even about the questions that define the subject of semantics". Brian Smith, too, confesses "It should be admitted that how this all works - how symbols 'reach out and touch someone' - remains an almost total mystery" (1987, 215). This specific problem has also arisen as a critique of 'logicism' within AI, where critics have been pointing not merely to specific shortcomings of current approaches within the broad knowledge representation paradigm, but to inherent flaws in the very conception of knowledge representation.

Frege's Formal Symbol Game

Briefly, Fodor's problem (like Searle's) is that of finding a way of reconciling the meaningless, syntactic formality of computational symbols and their causal interactions (Vera and Simon's patterns) with the fact that psychological generalizations are couched in semantic, intentional terms which refer to the real world outside the system or organism (Vera and Simon's denotation). In particular, it seems that different thoughts, in the sense of denoting different things and having different truth conditions, can nonetheless have identical internal substrates. It is perhaps no accident that within AI, a debate has been proceeding about the "logicist" approach which is committed to a classical logical conception according to which an abstract formal system gets its meaning from a model theory. Woods (1987), Smith (1987), Rosenschein (1985) and others have been insisting that the classical logical view is fundamentally misguided in its conception of the way in which a system gets to relate to the external world. Although Fodor (1991, p. 279) has recently declared that "Procedural semantics seems to have disappeared", the obituary may be a little premature (see Woods 1987, Hadley 1989, 1990), and there seem to be important affinities between this approach and the "embedded" or "situated" approaches. Both offer a way out of the seemingly intractable problem of reconciling the formal, meaningless character of symbol systems with their intentional, semantic contents.

In order to appreciate the bearing of these critiques on the symbol system hypothesis, it is important to look more closely at its origins in order to discern its underlying conception. Significantly, Newell and Simon note

> The roots of the hypothesis go back to the program of Frege and of Whitehead and Russell for formalizing logic ... a characteristic view, often referred to as the "symbol game". Logic ... was a game played with meaningless tokens according to certain purely syntactic rules. All meaning had been purged. ... We could call this the stage of formal symbol manipulation. (1976, 43)

As we have now seen, this is the clearest statement of just what Fodor has called the "formality condition" which gives rise to "methodological solipsism" and the problems of intentionality. This is the puzzle of explaining how the intentional or semantic, representational properties of mental states can arise from the purely meaningless syntactic properties of their tokens. At the very least the properties of patterning and denotation cited by Vera and Simon are hardly as unproblematic as their discussion suggests.

Nils Nilsson (1987, 1991) is explicit in embracing a certain conception of knowledge representation in AI which brings into relief the sources of concern. Central to his appropriately named 'logicist' position is the commitment to the representation of knowledge in the form of explicit, declarative sentences encoded in some formal language like the predicate calculus. In keeping with such a model-theoretic conception, Nilsson asks the key question:

> In what sense can we say that a collection of predicate calculus sentences represents knowledge about the world? Our answer to this question involves the notions of interpretations and models of sentences in the predicate calculus. ...
>
> We will call the world objects, relations, and functions invented by the designer, the intended model of the sentences the designer uses to describe the world. (1991, 36,7)

We see here that Nilsson is explicit in his requirement that the intended model is ascribed by the designer and that the representational symbols must be intelligible to the designer. This is to embrace a specific conception of "observer attribution" which has been one focus of 'situated' critiques. An indication of the problem at the heart of the very conception of a symbol is also seen in Newell's explanation of the Physical Symbol System Hypothesis:

> The idea is that there is a class of systems which manipulate symbols, and the definition of these systems is what's behind the programs in AI. The argument is very simple. We see humans using symbols all the time. They use symbol systems like books, they use fish as a symbol for Christianity, so there is a whole range of symbolic activity, and that clearly appears to be essential to the exercise of mind" (1986, p.33)

This explanation is striking for the explicitness with which it assimilates internal, mental representations with our external communicative symbols. Books and pictograms are taken as exemplars for the kind of symbol also to be posited to explain cognition. This is to assimilate the 'intrinsic' intentionality of mental representations with the 'derived' intentionality of our symbolic artifacts. This assimilation has been among the foundational assumptions of cognitive science but it is open to question. It is, indeed, the sense of "conflation" intended in Clancey's critique though not addressed by Vera and Simon in their reply. We see the same "conflation" clearly and self-consciously articulated in Dennett's remarks in an earlier review of Fodor's (1975) *Language of Thought* which was the philosophical manifesto for the Newell and Simon enterprise:

> What is needed is nothing less than a completely general theory of representation, with which we can explain how words, thoughts, thinkers, pictures, computers, animals, sentences, mechanisms, states, functions, nerve impulses, and formal models (inter alia) can be said to represent one thing or another. (1978, 91)

The hoped-for unification in cognitive science was to be effected precisely by showing that these heterogeneous items are all species of the same genus. Dennett explains further:

> It will not do to divide and conquer here - by saying that these various things do not represent in the same sense. Of course that is true, but what is important is that there is something that binds them all together, and we need a theory that can unify the variety. (1978, 91)

Minds in Machines or Friendly Photocopiers?

Before attempting to further clarify what I take to be the insight in the situated critique, it is important to acknowledge the justice and force of the rebuttals by Vera and Simon. The problem is complicated by the fact that some of the radical propaganda for situated cognition is unwarranted and made on behalf of claims which are entirely unremarkable, much less revolutionary, and readily subsumed under the traditional symbol system approach as Vera and Simon have argued. Thus, it is essential to give a preliminary indication of these claims before turning to those which may not be so easily accommodated within the traditional approach. For example, despite purporting to be a fundamental challenge to the very principles on which AI research has been based, many recommendations from the new 'situated' perspective appear to involve little more than encouraging an appreciation of the setting in which machines are to be used by people. That is, Agre (1990) says, "designing computers that are to operate in isolation is one thing, but designing computers that are to occupy an important place in the lives of real people is something else" (1990, 376). This seems to be explicitly downplaying, if not abandoning, the enterprise of AI as it has been conceived since Turing's work - that is the enterprise of designing autonomous intelligent machines. Instead, we see here an explicit contrast with this project and the emphasis instead upon designing artifacts which are intrinsically dependent on interaction with humans. Lucy Suchman (1987, p. 170), too, seems primarily concerned with human-machine communication. However interesting and important this may be, the talk of "real social settings", "reciprocal interpretation", "a field of socially organized meaning", "mutual intelligibility in the collaborative pursuit of a goal" (Stefik & Bobrow 1987, 221) - all this clearly envisages human-computer interaction rather than AI in the classical sense of minds in machines.

The "Canonical" Cottage Cheese Case

James G. Greeno (1989, 285), too, challenges the adequacy of the entire information processing framework on which AI and cognitive science rests, asking "whether our science of cognition is fundamentally limited because it deals only with symbolic computation". Despite the radicalism of Greeno's challenge to the entire information processing paradigm and even "our current epistemological beliefs" (1989, 286), the situation may be more banal. Greeno suggests that the grounds for questioning the orthodox symbolic approach are empirical considerations which suggest that people "use resources in the situations, *rather than symbolic computation*" (1989, p. 286; emphasis added). Greeno cites the Weight Watcher who had studied calculus but nevertheless answers a question about a daily allotment of cottage cheese by means of a simple, directly physical, geometrical, operation dividing up a portion of cheese, rather than by any symbolic operation such as a multiplication on fractions. *Ergo*, reasoning is not symbolic but "situated" (ibid. p. 286).

However, such illustrative examples fail to establish the radical claims since the talk of computation and symbol processing here equivocates between the use of symbols by a person at the level of *conscious awareness*, and the use of symbols as a substrate underlying the psychological phenomena. It is a non-sequitur to conclude from the absence of symbols and computation in *conscious* reasoning that there are no representations *at all* (Greeno, 1989, p. 290), and that the whole information processing paradigm is wrong. There is certainly nothing in the 'physical symbol system' hypothesis of Newell and Simon which requires that the symbols in question are those which are *consciously* applied to the problem, though this is the view that Greeno seems to be attacking. That is, typical symbolic explanations in cognitive science apply to the Weight Watcher precisely in the "situated" mode and not in Greeno's alternative portrayal of "symbolic" reasoning. Thus, Hobbes's remark that "reasoning is but reckoning" or Descartes's explanation of stereo vision as calculation of parallax, were claims about something *underlying* the contents of our conscious awareness. Helmholtz, too, made a similar point about the unconscious processes of inference. The work of Kahneman and Tversky is precisely analogous to the Cottage Cheese Case: even experts in formal statistics and probability theory are poor in their ability to make probability judgements in certain everyday "situated" contexts, but the postulated 'biases and heuristics' are not conscious calculations. On Wason's (1966) selection task, inferences made under realistic, concrete circumstances are better than when a logically identical problem is posed in a more abstract, formal manner. However, the "direct" reasoning in these cases does not mean that they are unmediated by symbols altogether. Vera and Simon are right to complain that all such illustrations of "situated" cognition are precisely the kind of phenomena which the symbol hypothesis was designed to explain.

The Direct Objects of Perception

The Weight Watchers' case is supposed to illustrate the thesis that the person's actions are somehow *unmediated* by mental representations. Citing Gibson and Heidegger, Greeno complains about conceiving internal representations of the environment as part of the individual's knowledge, since these do not capture the way in which the organism interacts *directly* with the structure of the world (ibid, 290). This is a notorious worry in philosophy and arose for Locke's theory of perception and also for Ayer's 'sense data' theory. In both these cases, the "ideas" or "sense-data", respectively, were taken to be the *direct* objects of perception and, in this sense, proxies mediating between our minds and

the world. Aware of the difficulties, Berkeley's solution was to get rid of the external world. Greeno's solution is to get rid of the representations. On one reading, it is a radical anti-mentalism. But Greeno's complaint against the classical symbolic approach is misplaced since the symbols are not properly thought of as somehow intruding between the world and the person. Nevertheless, on a quite different reading, Greeno's point can be taken as affirming this very point about the directness of our connection with the world notwithstanding our cognitive representations which may be seen to *constitute* such a direct relationship. Greeno and Moore (1993) in their rejoinder to Vera and Simon (1993) can be seen not to deny internal mental representations as such, but to be complaining of a certain *conception* of such representations which takes them to be analogous to our externally observable, communicative symbols. This criticism, while not supported by the Weight Watcher examples, nevertheless amounts to a radically alternative view of cognition.

Greeno's treatment of failures in action serves to confirm my diagnosis of the problem. Greeno endorses the Heideggerian view that representations are only constructed or used in cases where "direct interaction with the situation is unsuccessful" (290) and "... when normal connected activity fails in some way" (289). On one reading, this is a bizarre notion, for it makes it seem that mental representations somehow pop into being only when there is some kind of failure. Again, such a view comes to make more sense provided that we confine the conception of representation to *the contents of conscious thought*. In this case it does make sense to say that skilled behaviours will be attended to, and in this sense, (explicitly) "represented", only when the usual smooth performance is disturbed in some way. However this familiar feature of skilled actions can hardly constitute grounds for challenging the very foundations of symbolic, information processing AI or the role of representations in cognitive science.

Situatedness and Symbol Grounding

In Greeno's view "an individual is considered as interacting with the structures of situations directly, rather than constructing representations of structural features and interacting with the representations" (1989, p. 290). It is not difficult to discern the same kind of worry which motivates Searle: For Greeno, as for Searle, there seems to be a stark difference between the *direct* understanding of things in the world, and that which is mediated by abstract, meaningless symbolic representations. However, Greeno and Searle appear to be mistakenly taking the perceptual apprehension of an *external* object as a criterion for meaningfulness. Greeno takes the case in which we really do, in fact, deal with external objects (such as the cottage cheese) and notices, rightly, that this seems not to be mediated (consciously) by symbolic reasoning; and, conversely, when we *do* (consciously) use symbolic operations, Greeno notes the disconnection of the symbols from real events in the world (p. 293). It is in the concern with this disconnection of symbols - their meaninglessness - that we see the affinity with Searle and his analogous contrast between the Chinese symbols and those of English.

The differences between Greeno and Searle are as illuminating as the similarities: Searle looks at the putative mental symbols and complains that they are not meaningful, while Greeno looks at the meaningful objects of thought and complains that they are not symbolic. I think that these are complementary ways of making the same mistake and also pointing to an important insight. Searle rightly protests that formal symbols are without meaning, and Greeno rightly protests that meanings are not symbolic. But they are also both mistaken in blaming the symbols, since these may be the wrong place to look for the semantic content of mental representations. At least, this appears to be the important insight of the "situated" theorists - among whom I include 'Procedural' semanticists. Greeno's point may perhaps be reformulated so as not to deny the existence of a symbolic substrate to reasoning, but to argue that, like other possible levels of analysis such as the chemical or quantum levels, it is not a relevant level for the explanation of cognitive phenomena.

Observers' Attributions

Critics of 'logicism' in AI from 'Procedural' approaches have also drawn attention to the problem of 'observer attribution' as a source of deep problems in representing knowledge (see Hadley 1989). The puzzles of semantics may be due to a compelling equivocation on profoundly different meanings of 'meaning'. It is the failure of symbols to appear meaningful when apprehended by conscious awareness which Searle mistakenly takes to be relevant their role in intentionality. As we have seen, it appears that the same requirement has been explicitly placed on the knowledge representations in AI by logicists such as Nils Nilsson (1987, 203), and has been the focus of attention by recent critics who complain of the inadequacies of traditional model theoretic semantics (Woods, 1987, 233). The parallels between the Chinese Room puzzle and the independent debates in AI help us to understand the force of the "situated" critique. For example, Searle dismisses the relevance of connections between symbols and the external world (the 'Robot Reply') as begging the question, but the "situated automata" of Rosenschein and Brooks may be seen to be offering a significantly different kind of 'Robot Reply'. In particular, they may be understood, not as showing how the substrate symbols are "meaningful" in Searle's inappropriate sense, but rather how *the relevant notion of meaningfulness is quite different*, having to do with the "embededdness" of reactive automata in a context. The appearance of question-begging arises from the fact that this is not so much to answer Searle's question (or Fodor's), but rather to suggest its irrelevance to the real problem of the semantic content of mental representations. Searle's notorious argument gains its force from taking the Newell and Simon symbol hypothesis seriously, and therefore may be regarded as a *reductio ad absurdum* of this conception of symbol rather than as a refutation of AI. As Hadley notes, the problem of classically conceived symbols being "parasitic upon externally supplied interpretations" by the system builder is not one to be

solved, but rather it is one to be circumvented through an entirely different conception of the source of meaningfulness - namely in the actual procedures which connect representations to external objects in the real world. Smith (1986. 45-48), too, like Clancey also emphasizes that "this external attribution of significance is at the heart of the notion of computation. ... the use of the term "symbol" in computational contexts ... is a matter fraught with complexity, subtlety and tacit attribution" (see also Woods, 1987 and Rosenschein 1985). Woods (1987) complains against accounts of belief which are based on model-theoretic semantics rather than the *more abstract reasoning procedures of perception and action*. These latter provide a specific account of the manner in which the symbols of the system are linked to their referents in the real world. Woods' approach can perhaps be encapsulated by saying that logic is not psychology. Rosenschein summarizes and contrasts:

> In the old approach, the machine is viewed as manipulating data structures that encode logical assertions as linguistic objects. In the new approach, logical assertions are not part of the machine's knowledge base, nor are they formally manipulated by the machine in any way. Instead, these assertions are framed in the metalanguage of the designer, who uses them only to express (to himself!) the information content of the states of the machine he is designing. (1985, p. 355)

No Representations in the Mind?

Clancey appears to be denying the existence of mental representations, but this can perhaps be understood as something more subtle and interesting than a born-again Behaviourism - namely, denying the assumption that internal and external representations are relevantly similar for the purposes of theoretical understanding. If representations, properly speaking, are paradigmatically our external artefacts, then, in *this* sense, there are no mental representations.

A consequence of the standard assimilation of external and internal representational symbols (see Rumelhart and Norman 1983) is the incorporation of an interpreter as observer-analogue or homunculus into our cognitive theories. According to Rumelhart and Norman, "An interpreter, therefore, must be capable of examining symbols and executing the actions that they specify" (ibid, 79). This particular homunculus is ineliminable because it is intrinsic to the very conception of the symbols on which substantive models rest. Unlike the homuculi in cognitive theories which are permissible because they can be successively "discharged" (Dennett, 1978), the present one is an inherent feature of the symbolic, representational conception.

Smith (1987) points out that the phrase "the representation of knowledge" is ambiguous in a certain way. In one sense knowledge is being represented as in a statement, theory or other formalism. In another sense, the expression is referring to the representations which are the object of such formalism (See also Cummins, 1986 p.126). Chomsky has drawn attention to the same ambiguity in the usage of the term "grammar" according to which we may be referring to the linguists' theory or else to the internal psychological, mental entities in the speaker's head which are described by the theory. Clancey's point that knowledge-level descriptions are "not structures or mechanisms inside the agent" is to be understood essentially like Chomsky's competence-performance distinction which was intended to capture the fact that the highly abstract, idealized formalisms of a grammar need have no direct relation to any internal process or mechanisms though they must, nonetheless, be seen as "psychologically real".

Logicism

As we saw, Nilsson (1987, 1991) is explicit in embracing the traditional account of symbols and their meaning which brings into relief the sources of concern. Against the proceduralist position that representations must be such as to be utilized by the system rather than undertood by the designer, Nilsson protests that "the very existence of libraries full of books written in declarative form would seem to argue against the position of the proceduralists" (1987, p. 203). But citing our own symbolic artifacts such as books whose intentionality is "derived", is precisely to beg the question at issue and to confirm the diagnosis of Clancey and others about the ills of the traditional approach. It is for essentially the same reasons that Smith (1991, p. 282) has aptly termed the CYC project of Lenat and Feigenbaum the "Electric Encyclopedia" for it involves merely automating the kind of symbol to be found in a book.

The traditional, logicist conception of symbols and their meaning provides a framework from within which the puzzles of "psychosemantics" are insoluble. If "meaning" or semantic interpretation is inherently a matter of some intended model in the logicians sense (Nilsson 1991, p.37), then there can be no way of explaining why the actual external world has any special status as the "right" one (See Birnbaum 1991). The agonizing over this problem in philosophy might be a symptom of the fact that the semantic content or intentionality of mental representations must be conceived on an entirely different basis. This seems to be clearly one of the many things that "situatedness" is offering. As Hadley (1989) notes, "For a machine to possess (or tacitly know) the semantics of its machine language in the sense required by PS [procedural semantics], is just for the machine to stand in the right causal relationship to that machine language" (1989 p. 110, 111). Of course, what "the right causal relationship" might be is the big question.

Conclusion: Eliminating the Agent

The insight of situated cognition is that cognition cannot be mediated by meaningless symbols and must be somehow directly connected with the world. However, given a problematic three-part relation,

$$\text{agent} \leftrightarrow \text{representation} \leftrightarrow \text{world}$$

getting rid of the intermediate, intervening relatum, as suggested by situationists, is only one possibility to achieve directness, and only marginally more plausible than getting rid of the third relatum, the world (e.g. Berkeley). The aim of securing a direct connection between cognition and the

world can also be achieved by the seemingly paradoxical move of dispensing with the *agent* and, though not expressed in these terms, it is this approach which is perhaps the insight implicit in the situated critique of symbols. The rejection of representations may be best understood as recommending the rejection of a certain *particular conception* of them in favour of one which is not based on their being intelligible in the manner of external artefacts such as words and pictures etc. A crucial equivocation on distinct meanings of 'meaning' has led to the postulation of symbols having meaning in this observer-relative sense in which a representation is necessarily apprehended and understood by someone. Undeniably, computational symbols are meaningless in *this* sense, but this is no more problematic than the meaninglessness and mediation of action potentials or synaptic activations. Rejecting this inappropriate criterion of meaning actually amounts to rejecting the *agent* as homunculus in the system. The directness of cognition with the world is thereby achieved by eliminating this notion of meaningful symbol and its accompanying agent-homunculus.

The alternative conception of symbols and their meaning is one in which the representations have appropriate functional relations to other representations and to the world in the sense of permitting intelligent, adaptive behaviour. Thus, for example, it is not literally the absence of representations as such which makes Brooks' automata exemplifications of 'situated cognition', but rather the lack of *observer-ascribed* meaningful representations.

Long before the emergence of 'situated cognition', Jerry Fodor (1980) quipped that he didn't know what *Dasein* was, but he was sure that there was plenty of it around. Fodor's article was concerned with the vexed question of the formality of computational symbols, and his allusion to *Dasein* supports my diagnosis that the hei-falutin', hei-phenated, Heideggerian turn of situated cognition is essentially the philosophical problem of intentionality and 'symbol grounding' in a different guise.

References

Agre, P.E. (1990). Book Review of Lucy A. Suchman, Plans and Situated Actions, *Artificial Intelligence*, **43**, 369-384.

Agre, P.E. (1993). The Symbolic Worldview, *Cognitive Science*, **17**, 1, 61-70.

Birnbaum, L. (1991). Rigor Mortis: A response to Nilsson's 'Logic and Artificial Intelligence', *Artificial Intelligence*, **47**, 57-77.

Clancey, W.J. (1993). Situated Action: A Neuropsychological Interpretation, *Cognitive Science*, **17**, 1, 87-116.

Cummins, R. (1986). Inexplicit Information, in M. Brand and R.M. Harnish (eds.), *The Representation of Knowledge and Belief*, Tucson: The University of Arizona Press.

Dennett, D.C. (1978). *Brainstorms*, Montgomery, Vermont: Bradford Books.

Fodor, J.A. (1975). *The Language of Thought*, New York: Crowell.

Fodor, J.A. (1980). Methodological Solipsism Considered as a Research Strategy in Cognitive Psychology. *Behavioral and Brain Science*, 3, 63-109.

Fodor, J.A. (1985). Fodor's Guide to Mental Representation. in *A Theory of Content and Other Essays*, Bradford/MIT, 1990, 3-29.

Fodor, J.A. (1991). Replies. B. Loewer and G. Rey eds., *Meaning in Mind*, Oxford: Basil Blackwell.

Greeno, J.G. (1989). Situations, Mental Models and Generative Knowledge, in D. Klahr and K. Kotovsky (eds.) *Complex Information Processing: The Impact of Herbert A. Simon*. Hillsdale, N.J.: Lawrence Erlbaum.

Greeno, J.G. & Moore, J.L. (1993). Situativity and Symbols, *Cognitive Science*, **17**, 1, 49-60.

Hadley, R.F. (1989). A Default-Oriented Theory of Procedural Semantics, *Cognitive Science* **13**, 107-137.

Hadley, R.F. (1990). Truth Conditions and Procedural Semantics. In Philip P. Hanson ed., *Information, Language and Cognition*. Vancouver: University of British Columbia Press.

Kuhn, T.S. 1970. *The Structure of Scientific Revolutions*, Second Edition. Chicago: University of Chicago Press.

Newell, A, (1986). The Symbol Level and the Knowledge Level. In Z. Pylyshyn and W. Demopoulos eds., *Meaning and Cognitive Structure*. Norwood: Ablex.

Newell, A. & Simon, H.A. (1976). Computer Science as Empirical Inquiry: Symbols and Search. *Communications of the ACM*, 19, 113-126.

Nilsson, N.J. (1987). Commentary on McDermott. *Computational Intelligence*, 3, 202-203.

Nilsson, N.J. (1991). Logic and Artificial Intelligence, *Artificial Intelligence*, **47**, 31-56.

Rosenschein, S.J. (1985), Formal Theories of Knowledge in AI and Robotics, *New Generation Computing*, **3**, 345-357.

Rumelhart and Norman (1983). Representation in Memory, Center for Human Information Processing, *Technical Report* CHIP 116, University of California, San Diego.

Smith, B.C. (1986). The Link from Symbols to Knowledge. In Z. Pylyshyn and W. Demopoulos eds., *Meaning and Cognitive Structure*. Norwood: Ablex.

Smith, B.C. (1987). The Correspondence Continuum. *CSLI Report 87-71*.

Smith, B.C. (1991). The Owl and the Electric Encyclopedia, *Artificial Intelligence*, **47**, 251-288.

Stalnaker, R. (1991). How to do Semantics for the Language of Thought. B. Loewer and G. Rey eds., *Meaning in Mind*, Oxford: Basil Blackwell.

Stefik, M.J. & Bobrow, D.G. (1987) Review of Winograd and Flores, Understanding Computers and Cognition, *Artificial Intelligence*, **31**, 2, 220-226.

Suchman, L.A. (1987) *Plans and Situated Actions*, Cambridge: Cambridge University Press.

Vera, A. & H. Simon (1993). Situated Action: A Symbolic Interpretation, *Cognitive Science*, **17**, 1, 7-48.

Wason, P.C. (1966). Reasoning. Foss, B. (ed) *New Horizons in Psychology, I.* Harmondsworth: Penguin, 135-151.

Woods, W.A. (1987). Don't blame the tools. *Computational Intelligence*, **3**, 228-237.