

Representations of Knowing: In Defense of Cognitive Apprenticeship

A response to Sandberg & Wielinga

William J. Clancey

**Institute for Research on Learning
66 Willow Place
Menlo Park, CA 94025**

Journal of AI in Education, 3(2): 139-168, 1992, in response to the paper by J. Sandberg and B. Wielinga, "Situating Cognition: A paradigm shift?" a revision of their IJCAI-91 paper, "How situated is cognition?"

Abstract

Sandberg and Wielinga argue in their paper, “Situated Cognition: A paradigm shift?” that “there are no strong reasons to leave the traditional paradigm of cognitive science and AI.” They are certainly correct that we should not “disregard evidence and achievements of Cognitive and Instructional Sciences.” But they fail to appreciate the implications of the storehouse view of knowledge, which suggests that learning is like putting tools in a shed. Situated cognition arguments against traditional views of learning transfer suggest that human memory does not consist of stored facts and procedures. Perhaps because of the difficulty of this connection—and imagining what the alternative could be—Sandberg and Wielinga also misconstrue cognitive apprenticeship (“formal education should not just be replaced by ‘cognitive apprenticeship’”). They misunderstand the idea of effectively relating formalized subject material to everyday practice, believing Collins et al. to be against the teaching of theories and generalities altogether, when in fact they favor AI applications to education. The practical implication of cognitive apprenticeship is to refocus instructional research on the design process itself: We should design computer systems in partnership with students, teachers, and practitioners in the context of use, so we can produce programs that people can afford and want to use, that promote creativity, and that relate in an honest, pragmatic way to everyday life.

1. Introduction: Career changes and new designs

Why did Winograd, Brooks, Wenger, Brown, et al. experience a major change in their understanding of cognition? Does this change have practical implications in their work? Obviously, these scientists insist that a major shift has happened in their careers, which causes them to see computer programs in a new way. Winograd subtitled his book, “A new foundation for design.” We can see for ourselves that roboticists, computational ethologists, and workplace analysts have invented new designs on the basis of their new conceptions of the mind (Maes, 1990; Greenbaum and Kyng, 1991).

Concerning instructional design, Collins, Brown, and Newman (1986) place new emphasis on intelligent tutoring systems that help students reflect on reasoning processes (e.g., Guidon-Manage (Rodolitz and Clancey, 1990)). They argue that situated cognition helps us better relate theoretical models like Neomycin’s disease taxonomy to medical practice, so that learning and using such models occurs within the context of a community of practice (as opposed to being handed over as objective facts, existing independently of human modelers and practitioners). Yet, in their review of situated cognition, Sandberg and Wielinga say on one hand that the conceptual emphasis of cognitive apprenticeship is “high brow,” and on the other hand accuse Collins et al. of dismissing conceptual learning entirely! A serious miscommunication has occurred.

Can we understand situated cognition in a way that respects the experience and insights of our colleagues? Can we understand the implications for improving the credibility and value of AI applications to education, rather than feeling threatened that the old way of thinking was inadequate? My conclusions to these questions may startle readers who believed Sandberg and Wielinga to be raising important, valid points. Most of the difficulty arises because they confuse decontextualized learning (*how* concepts are taught) with formal learning (*what* is taught). They fail to see that cognitive apprenticeship is a contextualized way of teaching abstractions. This confusion, in turn, is apparently caused by conflating different situated cognition perspectives into one universal, stereotyped view: Sandberg and Wielinga confuse design ideas for robots (maps of the world aren't necessary for insect-like navigation) with arguments about the nature of human knowing (concepts aren't static things stored in the brain). Furthermore, Sandberg and Wielinga don't realize that arguments against "transfer" are arguments against *this metaphor for describing learning*, not against the possibility of abstracting or improving performance from one experience to the next. All of these difficulties appear to rest on Sandberg and Wielinga's belief that today's cognitive models are "functionally" adequate, and that nobody assumes that models of knowledge are practically describing stored structures in the brain.

I will begin by refuting the claim that situated cognition is attacking a straw man, focusing on how it overturns stored-schema cognitive models. I will then address the paper point by point, explaining how cognitive apprenticeship supports continued research in AI applications to education. Finally, I explain why Lave et al. believe that describing learning in terms of "transfer" conflates the activity of *knowing* with knowledge representations.

2. Sandberg and Wielinga misrepresent cognitive science

In large part, Sandberg and Wielinga (hereafter S&W) base their dismissal of situated cognition views on the claim that nobody ever really believed the contrary. They claim that cognitive models were always assumed to be descriptions of behavior, not literal descriptions of things in the mind:

The identity hypothesis not only implies an identity relation at the level of observable behavior, but also at the level of the structures that are being manipulated by certain processes, and at the level of the mechanisms that make it possible to operate on these structures. We think that the identity hypothesis was never widely held in this strict sense.

Thus, of three levels—behavior, representational structures, and mechanisms—S&W think that cognitive models are believed to be identical to only human behavior. They cite Dennett as supporting their view: "No one supposes that the model maps onto the process of psychology and biology all the way down." But Dennett is saying that nobody supposed computational models to map onto neural cell biology or the chemistry of human physiology. Contrary to what S&W claim,

many, if not most, cognitive scientists have claimed that cognitive models map onto an internal level of representational structures and processes, not just behavior.

To show this, I will focus on the prevalent cognitive science view that human memory is a place where symbolic structures are stored and problem solving consists of accessing and manipulating representations subconsciously. As Dennett suggests, most people hold that the mechanism—how these structures and operations are *implemented* in hardware—is assumed to be irrelevant. Gardner makes this point clearly:

If there was to be an identity, it obviously could not reside in the hardware, but, as Putnam pointed out, might well occur in the software: that is, both human beings and machines—and any other form of intelligent life, from anteaters to Antipodeans—could be capable of realizing the same kinds of program. Thus, the equation occurs at a much higher level of abstraction—a level that focuses on the goals of cognitive activity, the means of processing at one's disposal, the steps that would have to be taken, the evaluation of steps, and the kindred features. (Gardner, 1985, p.78)

Researchers sometimes point this out, emphasizing that they are only committed to the brain being a physical symbol system, not that the physical structures are identical:

Human thinking (cognition) can be regarded as a *computational* process. The basic notion is that human thought (including perception, understanding, and even perhaps emotion) is *the result* of the manipulation of information—homologous to data structures that might be employed in computers. [Footnote: Clearly there is no assumption of literal similarity in data structures, only a representational equivalence. (There are many obvious differences between computers and brains, e.g., computers are composed of silicon-based memory...)] (Evans and Patel, 1990, p. 10)

Lakoff, in his critique of the predominant view (which S&W say never existed) spells out the assumptions:

The traditional view is a philosophical one. It has come out of two thousand years of philosophizing about the nature of reason. It is still widely believed despite overwhelming empirical evidence against it.... We have all been educated to think in those terms....

We will be calling the traditional view *objectivism* for the following reason: Modern attempts to make it work assume that rational thought consists of the manipulation of abstract symbols and that these symbols get their meaning via a correspondence with the world, objectively construed, that is, independent of the understanding of any organism...

A collection of symbols placed in correspondence with an objectively structured world is viewed as a *representation* of reality.... Thought is the mechanical manipulation of abstract symbols. The mind is an abstract machine, manipulating symbols essentially in the way a computer does, that is, by algorithmic computation. Symbols that correspond to the external world are internal representations of an external reality...

Though such views are by no means shared by all cognitive scientists, they are nevertheless widespread, and in fact so common that many of them are often assumed to be true without question or comment. Many, perhaps even most, contemporary discussions of the mind as a computing machine take such views for granted. (Lakoff, 1987, pp. xii-xiii)

This view equates human knowledge with stored symbols and reasoning with symbol manipulation:

Thus, knowledge is seen to be of paramount importance, and AI research has shifted its focus from an inference-based paradigm to a knowledge-based paradigm. Knowledge is viewed as consisting of facts and heuristics. (Davis and Lenat, 1982, p. xvi)

This view isn't held just by AI researchers, seeking whatever mechanism will work, but by psychologists dedicated to understanding human memory. Chi et al. state in their introduction to "The Nature of Expertise":

There is now a cognitive science related to the representation and execution of expert performance. This science has developed a technology in the form of programs for performing tasks formerly done only by experts. Although this technology is still primitive, it represents an important contribution of fundamental research on the nature of representation in memory. Behind this technology is a better understanding of what it means to be an expert. Expertness lies more in an elaborated semantic memory than in a general reasoning process. Such knowledge is present not only in the performance of unusual people, but in a skill like reading which is widely distributed in most of us. We are beginning to understand the nature of the propositional network underlying such representation. The expert has available access to a complex network without any conscious representation of the search processes that go on its retrieval. (Chi, et al., 1988, p. xxxv)

Clearly, Chi believes that knowledge consists of representations that are literally stored in a "semantic memory" consisting of a "complex network" that is subconsciously searched. This is precisely what the identity hypothesis claims, that there is a correspondence between a knowledge base, searching, and matching performed by a computer, and comparable structures and processes in the human. This is the hallmark of the information-processing approach:

The comprehension model is an information-processing model: It identifies certain processes and mechanisms (a short-term memory buffer, cyclical processing, memory retrieval) that interact to produce comprehension. (Miller, et al., 1984, p. 5)

The theory posits a set of processes or mechanisms that produce the behavior of the thinking human. Thus the theory is reductionistic; it does not simply provide a set of relations or laws about behavior from which one can often conclude what behavior must be. (The elementary processes and their organization, of course, are not explained: reduction is always relative.) Thus, the theory purports to explain behavior—and not just to describe it, however parsimoniously. (We are aware that some would dispute such a distinction, viewing all causal explanations as simply descriptions.) (Newell and Simon, 1972, p. 9)

The processes posited by the theory presumably exist in the central nervous system; they are internal to the organism. As far as the great debates about the empty organism, behaviorism, intervening variables, and hypothetical constructs are concerned, we take these simply as a phase in the historical development of psychology. Our theory posits internal mechanisms of great extent and complexity, and endeavors to make contact between them and visible evidences of problem solving. That is all there is to it. (Newell and Simon, 1972, p. 9-10)

We confess to a strong premonition that the actual organization of human programs closely resembles the production system organization.... (Newell and Simon, 1972, p. 803)

By the mid 1970s, most cognitive modeling was based on the idea that human long-term memory was literally a semantic store:

Here is the essence of the frame theory: When one encounters a new situation (or makes a substantial change in one's view of a problem), one selects from memory a structure called a *frame*. (Minsky, 1977, p. 355)

For many theorists who use it, the term *schema* has come to be synonymous with the term *long-term memory structure*. The schema-is-a-structure assumption is clearly evident in the cognitive scientific literature and needs no elaboration.... Schemata are generally claimed to be pre-existing knowledge structures stored in some location in the head. (Iran-Nejad, 1987, p. 111)

Conceptual dependency and MOPS research is based on the idea that representations are literally stored in memory. Faced with difficulties, researchers produce only variations of a storage model:

Elizabeth Loftus has established a "rewriting" effect during question answering.... But these phenomena cannot be simulated in a system that treats question answering as a purely passive retrieval process. Memory alterations can only occur if the retrieval process somehow acts on memory, to refine or alter its previous contents. (Lehnert, 1984, p. 36)

Of course, some researchers, stopping to reflect on the assumptions of the field, were surprised to see how far the theories had gone:

More interesting, and perhaps more serious, is the confusion between purposive and mechanistic language that characterizes much of the writing in cognitive science. As if it were the most natural thing in the world, purposive terminology has been imported into an information-processing framework: subgoals are stored in short-term memory; unconscious expectations are processed in parallel; opinions are represented propositionally; the mind contains schemata. (Miller, et al., 1984, p. 6)

Nevertheless, as Mandler points out, the basic assumption of cognitive science that computational models map onto a representation and process level inside the head is rarely questioned:

The central themes that emerged during those 5 years [1955-60] and that mark the cognitive sciences are the concepts of representation and process. They are the primary foci of all the relevant disciplines, and it is symptomatic of our acceptance and their importance that we rarely hear anybody question these two foundations.... We are more concerned about distinctions between analogic and propositional representations or between declarative and procedural knowledge. (Mandler, 1984, p. 306)

Perhaps nowhere are the assumptions more clear and the difficulties more severe than in models of language (Winograd and Flores, 1986). Bresnan even reminds her colleagues that they all operate within the paradigm of the identity hypothesis, and it is, by assumption, not the source of their difficulties:

The cognitive psychologists, computer scientists, and linguists who have questioned the psychological reality of grammars have not doubted that a speaker's knowledge of language is mentally represented in the form of stored knowledge structures of some kind. All theories of mental representation of language presuppose this. What has been doubted is that these internal knowledge structures are adequately characterized by transformational theory... (Bresnan, 1984, p. 106)

Of special relevance to education is how these assumptions, especially that knowledge consists of stored representations, was picked up by AI researchers and used to justify the design of intelligent tutoring systems as a means of transferring expertise to students:

Much of what constitutes domain-specific problem-solving expertise has never been articulated. It resides in the heads of tutors, getting there through experience, abstracted but not necessarily accessible in an articulatable form. (Sleeman and Brown, 1982, p. 9)

They [ITS programs] approach teaching from a subset viewpoint: expertise consists of a set of facts or rules. The student's knowledge is modelled as subset of this knowledge. (Goldstein, 1982, p. 51)

The processes of the student are divided into two homunculi—a problem solving specialist and a learning specialist—with the [genetic] graph serving as the student's basic memory structure for procedural knowledge. (Goldstein, 1982, p. 71)

I have gone to great length to reveal what should be obvious. Sandberg and Wielinga may never have believed that human memory was a place where representations were stored. They may not have realized that Newell and Simon intended the production system model to describe symbolic programs that were actually running in the brain. But S&W's historical understanding of cognitive science is simply wrong. And therefore they do a great injustice to the situated cognition view to say it is arguing against something nobody ever believed.

The challenges brought by situated cognition are not easy to understand. They require a major shift in world view, perhaps no easier to comprehend than the shift from an earth-centered to heliocentric solar system. The challenge of Jean Lave, for example, is serious and difficult to comprehend, but we must first begin by acknowledging that it is in fact a challenge to what it is commonly believed:

I wish to avoid, respectively, functionalist and phenomenological reductions of the constitutive order and lived-in world to internal representations and inter-subjectively constructed ones. (Lave, 1988, p. 194).

To explain what this means, I will clarify (and sometimes correct) what S&W say about situated cognition. I will then focus on the issues of transfer and classroom learning raised by Lave, Brown, et al.

3. Specific comments on Sandberg and Wielinga's paper

One good thing about Sandberg and Wielinga's paper is their simple statement of situated cognition points of view¹. This provides the opportunity for clarification, which I will do in an enumerated fashion, using the paper's headings.

Introduction

1. "Knowledge is always relative, and open to interpretation."

In my first paper about situated cognition (Clancey, 1991c), I attempted to build on Newell's analysis of the knowledge level. Therefore, I tended to adopt his language and didn't always distinguish between knowledge, as a capacity to behave, and representations of knowledge (e.g.,

¹Note that remarks 1 - 5 and 29 are Sandberg and Wielinga's paraphrases of the situated cognition position; the remainder of the remarks are their own questions and counter claims.

expert system rules). S&W's reading is therefore understandable. My claim is that knowledge *representations* are open to interpretation. This distinction is clear enough in Figure 13.4 (Clancey, 1991c, p. 394).

2. "Memory can no longer be seen as a storage place for representations—there are no representations left to store."

But we are constantly generating representations that we must store or they will be lost. Off on a hike, I might be thinking to myself about a talk I plan to give. As I generate ideas, I must represent and store them on a sheet of paper I keep in my pocket, or I will forget them:

Writing was an important innovation precisely because it was the first technology of symbolic representation that permitted...inscribing, passing about, inspecting, storing, and destroying.... Activity of these forms has to go on outside our heads because it's not practical for them to go on inside our heads. (Agre, 1988)

3. "We as observers are inclined to confuse patterns that characterize behavior over time with mechanisms that determine behavior moment to moment."

Braitenberg (1984) deserves credit for saying this, especially as justification for inventing new kinds of robot mechanisms, well prior to Brooks' work.

4. "...representations we equip, for example, our expert systems with do not reflect structures that cause expert behavior."

"Reflect" might be ambiguous here. I prefer to say, "Representations do not *correspond to stored structures* inside the expert's head."

5. "We as observers interpret them [representations] semantically, the system however cannot."

Yes, when we comment on what a representation means, we are conceptualizing, not retrieving definitions and linking them by stored grammatical rules. Of course, a person can simulate a machine by speaking rotely (e.g., reciting a classroom definition).

The identity hypothesis

6. "This assumption is what we call the *identity hypothesis*."

As the quote by Gardner indicates, this is a common term in writing about the foundations of cognitive science. I use the term prominently on page 382, Figure 13.3 (Clancey, 1991c).

7. "An expert system that solves problems through behaviour similar to that of a human expert can be viewed as a theory of problem solving behavior of that expert."

That's not how many knowledge engineers, with little interest in psychology, view expert systems. Mycin wasn't intended to be a model of human problem solving, but rather to be functionally equivalent in recommending therapy for infectious diseases.

The performance-oriented expert systems, on the other hand, start with productions as a representation of knowledge about a task or domain and attempt to build a program that displays competent behavior in that domain. These efforts are not concerned with similarities between resulting systems and human performance.... They are intended simply to perform the task without errors of any sort, humanlike or otherwise. (Buchanan and Shortliffe, 1984, p. 26)

8. "The knowledge structures we describe, the representations we form, are assumed to be functionally equivalent to whatever devices humans use to solve problems, in the sense that they give rise to similar behavior."

Yes, in the sense that the input/output behavior is what matters—producing the correct answer from a given set of data. [In contrast, a "process model" of human problem solving (Newell and Simon, 1972; Kintsch, et al., 1984) attempts to replicate the ordering of steps and intermediate results.] The Mycin gang believed that human knowledge consists of stored facts and rules; therefore a huge knowledge base could replicate human capabilities. Why this hypothesis is wrong is the essence of situated cognition. S&W evidently don't feel the impact of situated cognition because they still believe functional equivalence is possible by manipulating representations syntactically.

In effect, situated cognition rejects the physical symbol system hypothesis of Newell and Simon. The claim is that cognition does not consist solely of manipulating stored and matched symbols that have a physical reality in the brain. There is something more. Nobody is denying that creating and using representations is the hallmark of human intelligent behavior. We are simply saying that a symbol-manipulation model fails to account for how representations are created, and fundamentally misconstrue what representations are (cf. Lakoff). This does not say that existing cognitive models are useless. But it does discredit their adequacy for understanding learning processes, and calls into question their use for justifying the design of computer expert systems to replace people (Winograd and Flores, 1986).

Beliefs about how our models relate to the mechanism of thought have major implications for how we use computers and how we frame the difficulties in our research. Many knowledge engineers in the 1980s believed that just storing more structures in the program will enable programs to eventually reach human levels of performance. The knowledge acquisition bottleneck is after all based on the idea that the structures we need for expert systems are already encoded in the heads of experts and we need to extract them:

Traditionally the transmission of knowledge from human expert to trainee has required education and internship years long. Extracting knowledge from humans and putting it in computable forms can greatly reduce the costs of knowledge reproduction and exploitation. (Hayes-Roth, et al., 1983, p. 5)

The real question is whether the memory-as-stored-structures model is functionally equivalent to human capability. Functionalism says it is:

No one is likely to confuse a program embodying a piece of physics with the actual physical process that is being simulated. A program that represents a wave breaking on a shore is manifestly different from a real wave, and it would be absurd to criticize the program on the grounds that it was not wet. No sane person is likely to assume that the real wave is controlled by a computer program: it is governed by physical forces that are simulated by the program.

All theories are abstractions, of course, but there is a more intimate relation between a program modeling the mind and the process that is modeled. Functionalism implies that our understanding of the mind will not be further improved by going beyond the level of mental processes. The functional organization of mental processes can be characterized in terms of effective procedures, since the mind's ability to construct working models is a computational process. If functionalism is correct, it follows not only that scientific theories of mentality can be simulated by computer programs, but also that in principle mentality can be embodied within an appropriately programmed computer: computers can think because thinking is a computational process. (Johnson-Laird, 1983, p 8-9)

In my view, situated cognition hypothesizes that the workings of the brain are not computational, in the sense of being equivalent to a stored-program processor—we will never get the full range of human learning and creativity out of such a mechanism (Clancey, 1991a; 1991b; 1991e). Again, this is not to deny the value of expressing cognitive theories as effective procedures; but it does deny the identity hypothesis of Newell and Simon as well as the weaker “brain as computer” functionalist hypothesis: Information-processing theories are merely descriptions, they do not explain learning, problem-framing, or creativity (Schön, 1987) because as mechanisms they cannot replicate it. If situated cognition is right, then expert systems can never achieve what people can do, and there are important limitations in our current designs. This is what Winograd and Flores (1986) were warning us about; it is why Brooks and others are rallying around new mechanisms.

9. “Models cannot be tested under the assumption of the full identity hypothesis.”

Presumably “full identity hypothesis” refers to mapping all the way down to chemical structures of the human body. But intermediate correlations can be sought and experimentally tested. For example, Bransford, et al. (1977) and Jenkins (1974) provide experiments that refute the schema-as-stored-structure hypothesis.

10. “The assumption that traditional instruction is solely based on the identity hypothesis is an oversimplification.”

Yes, of course, there are many other arguments against decontextualized instruction. For example, Lave says that people misunderstand the relation between representations and human culture. Just as knowledge is not to be equated with representations of knowledge, contexts of use are not to be equated with representations of contexts. We are admonished not to “relegate culture, acquisition of knowledge and memory to an internalized past.” (Lave, 1988, p. 18)

The view that knowledge is stored suggests that interactions between people are structured and controlled by pre-existing structures stored in the head. The opposing view is that neural and social structures coordinating our behavior come into being during our interactions. They are dialectically related, so we can say my social interactions and my neural interactions constrain each other. They are both ongoing. They have their own emergent structure. They are coherent in their own terms (on different levels of organization). In other words, situated cognition calls on us to reconceive the idea of culture, learning, and memory as not reducible to (fully explained or caused by) stored descriptions of experience.

Knowledge and memory

11. "...knowledge is not something that we store as structures in memory. Clancey calls this the *externalisation move*. Through this externalisation Clancey places knowledge (as understood at the knowledge level (Newell, 1982)) as well as symbolic representations *outside* the thinking agent."

First, note that the word "knowledge" commonly conflates 1) representations of knowledge created by an observer (e.g., rules, goals, plans, strategies attributed to an agent) and 2) the *agent's* articulated models of the world and behavior plans. Both are representations, not knowledge.

To disentangle the terms knowledge and representation, we must be clearer about how representations are created and used. An agent is representing when he or she imagines perceptual experiences (silent speaking, visualizing, dreaming) or creates external symbol structures (e.g., writing rules in a knowledge base). The agent treats such expressions (imagined perceptual experiences or external notations) as *symbolic* by virtue of subsequent operations of representing that comment upon them (Clancey, 1991c). The externalization move (Clancey, 1991c, p. 382, 400-401) claims that "all symbol manipulation is going on above the line, in the agent's behavior."² This contrasts with the view presented by Newell (1982, p. 99) that physical symbol manipulation occurs *internal* to the agent, in a hidden and inarticulable fashion³.

Second, I place knowledge, as a capacity, outside the thinking agent in the sense that we can't explain intelligent behavior *exclusively* in terms of neural processes, but must include the structure

²More precisely, symbolic *structures* only exist and are manipulated in the external environment, outside the agent. By hypothesis, this means that, in carrying on an internal conversation, the agent is representing without creating or storing physical structures that *represent* his or her speech. Speaking to myself is unlike listening to someone else because comprehending and speaking to oneself occur dialectically—no mediating physical structures need to be perceived. Of course, I can reflect at another level on the occurrence or meaning of the internal dialogue itself (e.g., I can be surprised that the sequence of thoughts occurred at all).

³Referring to the distinction of mechanism, representational structure, and behavior, the externalization move places representing, perception of representations, and manipulation of representations in the realm of behavior. The cognitively impenetrable processes remaining inside (the mechanisms) are by hypothesis not constituted of stored and secretly-manipulated representations.

of the environment. For example, if we want to predict my behavior when packing for a camping trip, we should study the structure of the boxes of materials I store in my closet (Clancey, 1991d). This doesn't mean that materials in the closet are knowledge any more than data structures are knowledge. It just means that my capacity to behave intelligently is always an interaction of my neural processes and structures in my environment (e.g., how I have organized my desk, how I have represented my plans in a "to do" list) and what I have represented about my experiences and ideas (e.g., in notebooks). As Lave, Suchman, et al. emphasize, these environmental structures are in large part provided by other people (e.g., encyclopedia, office layout and furniture, word-processor tools).

12. "Knowledge and its symbolic representations are the results of a sense-making process in which an observer describes patterns of behaviour of an intelligent agent."

This refers to theorizing about agents, a particular kind of representation activity. Other sense-making concerns patterns in physical, informational, and organizational processes, including what we usually call scientific analysis.

13. "Constructing a representation means seeing something in a new light."

More generally, we are perceiving in a new way (Bamberger and Schön, 1983). We cannot easily distinguish between perceiving, conceiving, representing, and coordinating activity, except for making qualitative distinctions between the form, intentions, and levels of representational activity (e.g., verbal representations, framing and history-telling representations, jazz improvisation).

14. "People can recall things. We can recall names, numbers, sentences from a poem, laws of physics and so on."

But rote behavior is hardly the mark of intelligence. Tape recorders, databases, and cameras "recall" better than people could ever hope. Recalling is a misleading term, as is "reconstructing." We re-experience, we re-enact, we re-present. As James (1892) says, when the clock strikes twelve bells at noon every day, we find it unnecessary to say that it is recalling anything. The mere act of recurrence (perceptually attributed by an observer) says nothing about the mechanism. Furthermore, the idealized view of human recall ignores changes in intonation, rhythm, loudness, and any contextual aspect of what the utterance means to the speaker.

This is typical of cognitive theorizing: We order events into uniform categories and then puzzle over why the categories are so nicely patterned, as if they were inherent in the events of the world. The real question of interest is how did we create such representations and why are our ways of seeing and coordinating behavior stable? Rather than placing our representations inside the heads of our subjects, and assuming that is the end of the matter (cf. Newell and Simon: "That is all there is to it"), we must turn around and ask, where did these representations, of our own making, come from?

15. “So, how does Clancey's functional architecture explain recall of decontextualised information? We have to assume that some process is replayed that generates the symbolic representation of Newton's law.”

Replaying suggests a structure that is reinterpreted, like a needle moving over a phonograph or a series of commands in a program. Replaying is the wrong metaphor. A process is constructed that bears an analogical resemblance (in behavior) to what previously occurred because it is built out of perception-action maps that were previously active together (Bartlett, 1932; Edelman, 1987; Clancey, 1991b, Section 1.3). By hypothesis, the maps so-constructed are always newly manufactured (or selected, according to Edelman), not merely reactivated:

Suppose I am making a stroke in a quick game, such as tennis or cricket.... I do not, as a matter of fact produce something absolutely new, and I never merely repeat something old. The stroke is literally manufactured out of the living visual and postural “schemata” of the moment and their interrelations. (Bartlett, 1932, p. 202)

When we recite “ $F=MA$ ” we are reCOORDINATING our speaking, just as in hitting a tennis ball for the thousandth time. The symbolic representation is only generated when we move our pens or our mouths.

Crucially, new internal neural constructions apparently integrate (compose) already active neural processes at that moment, perhaps accounting for our learning of new behavior sequences in a context-sensitive, but always generalized way (Clancey, in preparation, a). This rejects the idea that “learning is a second-order effect” (Newell and Simon, 1972, p. 7). Learning doesn't merely follow performance (REPLAY + REFLECT), but rather is inherent in every movement and perception.

16. “If such a memory capability exists, why not call it a piece of knowledge?”

A capability is not a thing. Pieces have permanence, they are localized. Such properties are true of knowledge representations, but not knowledge itself. Knowledge is like energy; it is a capacity. It cannot be in hand (Newell, 1982). S&W are confusing the *product*, what people say, with the *process* of speaking (cf. Collingwood, Section 4).

17. “Why can we not create such a memory phrase through explicit communication (i.e., teaching decontextualised subject matter)?”

Of course we can get children to memorize things. Nobody is denying the possibility of teaching theories. Rather, Brown et al. argue against teaching *formal subject matter* in a *decontextualized way*. The question is how do *representations of subject matter* relate to *knowledge*? What are the implications of the claim that knowledge cannot be exhaustively represented (inventoried), that culture cannot be reduced to a list of facts and laws, that practice is emergent and reproduced in activity itself, not driven by rules of behavior? How does reflection (theorizing about goals, information, and causal relations) naturally occur in the course of activity? These are questions that

arise when you change your view that knowledge consists of facts and heuristics that can be exhaustively represented and transmitted in textbooks, computer tutors, or classroom exercises.

Of course, few people believe that learning to be a physician, for example, consists entirely of memorizing facts and rules. But we haven't had a good theoretic framework for integrating formal learning with practice, especially because we viewed stored theoretic knowledge as the essence of "real understanding." The difference may appear to be subtle, but it is powerful and important for how technology gets developed and incorporated in practice.

18. "Even if in the neural machinery that implements the functional architecture there is no recognisable *place* where a behavioural phrase is stored, at least functionally it can be viewed as a unit that can be activated as such."

This may appear to be a good metaphor if you believe that remembering is producing *identical* behaviors. In doing this, we focus on the words and laws, the representations as things, which then become mapped onto hypothesized units being activated in the speaker's head. Lave et al. want us to focus on attitudes, orientations, ways of interacting. These are not recalled things, but behaviors. When you shift to this point of view, you stop talking about knowledge as units that are activated or taught in a curriculum. Recalling laws and manipulating symbols becomes a means, not the ends. In effect, Brown et al. see schooling (e.g., algebra, physics, history classes) as providing tools to children with only the most meager sense of what they are for.

19. "It is possible to teach new strategies to experts...changes the expert's behaviour.... In terms of Clancey's functional architecture, this means that process memory is extended with new phrases through the perception of external representations (symbols explaining a new procedure), not through actual acting."

No, you have to act in order to perceive (comprehend) instructions. Reading is an interactive process; new strategies aren't simply absorbed into the experts' heads. If comprehension of instructions is to lead to changed behavior, by hypothesis it must re-coordinate perceptual-action processes that occur during activity. This means for example, that you must be activating and recomposing neural processes (perhaps perceptual categories or details) that are active when you perform the actions you are being instructed about. This description is vague, but it's good enough for orienting the search for new mechanisms (see Freeman, Edelman, Iran-Nejad).

20. "It is not clear to us whether this corresponds to the reflection process postulated by Clancey, but in any case the functionality of the traditional memory model can explain this phenomenon just as well as the process memory architecture."

I am referring to Donald Schön's use of the term "reflection," in which we represent our activity (Schön, 1987; 1990). Reflection is not something that occurs internally, in a hidden way, separate from our activity. It is always part of an ongoing activity, of a set of concerns, attitudes, and

orientation towards what is important, what we are trying to do, what we are paying attention to. By comparison, the AI view of reflection as inspecting internal structures (“meta-knowledge”) is sterile and syntactic. For a person, reflecting is conceiving new ideas.

21. “Admittedly such [perceptual] patterns are primarily indexical and functional. They impose structure on the observed reality and through that structure make a focussed problem solving process possible. However, there appears to be no good reason why such patterns cannot be viewed as a psychological reality.”

It is important to remember that these patterns in behavior are always generated by a generalizing, learning process. They are not stored descriptions that are matched like frames. They are interactive recompositions, albeit routine. We don’t know how habits are formed neurologically.

Patterns of expertise do have psychological reality, in two ways: First, they are real for us, the perceivers of these patterns. They are statements of our observations. Second, as scientific descriptions, they are useful psychological models (e.g., novice vs. expert studies, MOPS, Neomycin’s diagnostic procedures).

However, psychological reality is not to be confused with the claim that psychological models are functionally equivalent to human performance. A basic claim of situated cognition is that such patterns in themselves are inadequate as a mechanism for generating the range of human behaviors. Models based on stored patterns produce *behaviors that appear similar*. They replicate behavior patterns, routines. In themselves, patterns and patterns of pattern modifications cannot generate the novel behaviors we observe in people. As I said at the Banff Knowledge Acquisition Workshop in 1986: “To be following a pattern (observer claim) is not necessarily to be following a pattern-thing (mechanism).” After all, we are not automatons.

22. “Again there does not appear to be a compelling reason why a process memory should give a better account of these phenomena than a traditional theory of memory.”

Maybe not from a layman’s point of view. But see Bartlett, Bransford, Jenkins, Rosenfield, Edelman, et al. for scientific evidence that the traditional storage theory of memory is inferior to a constructive model.

My theory of process memory (Clancey, in preparation, a) is an attempt to solve the combinatoric problem of symbolic reasoning. The promise of a self-organizing system is to avoid the problem of indexing and search, as well as the “frame problem” of knowing when internal, stored models are out-of-date.

23. “Again, the process memory does not seem to give us much more than conventional theories of memory storage, memory retrieval and depth of processing.”

It depends on your goals. I doubt my grandmother will care to change her point of view. But connectionists have been greatly inspired by alternative views of memory. For psychologists trying to understand how the brain works; the shift from information-processing to constructivist theories is radical (Rosenfield, 1988). If you are only concerned about education, I believe you’ll understand the implications more at the level of reconceiving the nature of culture, practice, and collaboration. When you see that people’s ability to interact and coordinate behaviors is the *foundation* for their ability to theorize in similar ways (rather than the other way around), you will begin to feel the force of the argument.

24. “There appears to be no compelling need for a radical change in the paradigm of AI with respect to the notions of knowledge and memory. As Marr (1977) has pointed out, we are studying information processing *problems* and not in the first place mechanisms.”

But inventing architectures for producing intelligent behavior remains a central concern of AI (vanLehn, 1991). Mechanisms may appear unimportant to S&W because most alternatives over the past 35 years have been based on one idea: A stored memory of representations. The issue has been *what kind* of representations: Declarative or procedural? Propositional or analog? Rules or frames? Distributed in networks or localized? Case-based experience memories or abstract theories? Bottom-up matching or top-down matching? Driven by goals or driven by data? (Cf. Mandler, Section 2)

From the robot-design perspective, we need to know what to put inside the robot. How we view knowledge will have a big effect on what mechanisms we try to invent. Brooks showed that navigation without built-in maps is possible, where previously people were using cameras and complicated matching of representations. To AI researchers, the mechanism is of interest, not just how we *view* or talk about what the mechanism does.

Furthermore, the very idea of “information-processing problems” fails to account for how situations become problematic and are framed by ways of seeing and talking; the information-processing approach itself distorts the nature of intelligent behavior:

The puzzles or problems are assumed to be objective and factual. They are constructed “off-stage” by the experimenters, for, not by, the problem solvers. The process of their construction is not relevant to problem-solving activity and not accessible to inspection. Problem solvers have no choice but to solve problems. (Lave, 1988, p. 35)

Problems of the closed, “truth or consequences” variety are a specialized cultural product, and indeed, a distorted representation of activity in everyday life, in both sense of the term—that is, they are neither common nor do they capture a good likeness of the dilemmas addressed in everyday activity. (Lave, 1988, p. 43)

25. “Knowledge can still be viewed at the functional level as ‘mental substance’ more or less of which can be available to a thinker.”

It is true that this view is often still useful. But we should be suspicious of teachers who *only* present the information-processing view that everyday practice consists of solving problems that are given, with well-defined information and goals:

If we see professional knowledge in terms of facts, rules, and procedures applied non-problematically to instrumental problems, we will see the practicum in its entirety as a form of technical training...

If we focus on the kinds of reflection-in-action through which practitioners sometimes make new sense of uncertain, unique, or conflicted situations of practice, then we will assume neither that existing professional knowledge fits every case nor that every problem has a right answer. (Schön, 1987, p. 39)

26. “Moreover, this substance has internal structure which can be made explicit in empirical studies...knowledge can at least functionally be viewed as a substance that can be communicated in such a way that it can be used.”

Again, this continues the view that knowledge is inherently structured, that these structures are shared among people by being representationally-equivalent and individually stored, and this sharing of goals, beliefs, and heuristics and other representations is what constitutes and provides the basis for a culture of practice:

The enterprise [of learning-transfer research] rests on the assumption of cultural uniformity which is entailed in the concept of knowledge domains. ‘Knowledge’ consists of coherent islands whose boundaries and internal structure exist, putatively, independently of individuals. So conceived, culture is uniform with respect to individuals, except that they may have more or less of it. (Lave, 1988, p. 43)

Lave is saying that abstracting cultural domains leads us to assume that our represented categories of knowledge have a transcendental standing (“independently of individuals”). What we call culture resides and exists in the activities of people, not just processes in the brain. Conventions are not things (like grammars) stored in the head. What is shared especially is the capability to interact together. The inherent complexity of interpersonal experience is violated by the idea that behavior is reducible to internal structures in the head, which by being uniform from person to person supposedly enable us to interact together.

In summarizing their argument by saying “this substance has internal structure,” S&W fail to accept that human memory is not a place where things are stored. The internal structure of neural processes is constantly changing and always freshly created (albeit out of previous activations). S&W don’t feel the force of situated cognition claims because they hold to their view that knowledge can be well modeled by the stored-schema view. To keep repeating that stored-schema models have value misses the point. It is the inadequacies of this view that Schön, Lave, et al., are calling us to understand.

Educational implications—the role of concepts in intelligence

27. “Central to Brooks' view (1990a), and to the 'situationist position' in general is that internal (symbolic) representations and thus concepts do not play a necessary part in intelligent behaviour.... Much human activity calls for the use of concepts—apart from the question of how they are represented.... Brown et al. and Clancey...seem to adopt Brooks' denial of concepts by agreeing that there are no internal representations.”

Brooks is proceeding in the bottom-up way advocated by Braitenberg: What can be done without plans and maps of the world? How can a dog play ball with me, “knowing” that it is the game we play every afternoon, without being able to represent ball-playing in some language? We believe the dog knows the ball as a thing and knows how to interact with it. We don't give much away from human uniqueness if we admit that the dog is at least perceptually categorizing the ball as a recognizable entity (and, indeed, he might respond to the word). But this is not equivalent to the claim that conceptualizing is having symbolic representations inside the head. A dog can know the concept “ball” without being able to speak a word, which is contrary to most views of human knowledge, in which stored classifications and networks of words are the basis of conceptual knowing.

We must not confuse *the experience* of seeing or hearing something with *physical symbol structures* (e.g., words on paper). We must not confuse “knowing a concept” (or being knowledgeable) with linguistic names and definitions. We typically view concepts as things, equivalent to the names of categories and defined by their properties. It may be more productive to view concepts as recurrent (but always re-adapted) ways of perceiving and coordinating behavior (cf. Bartlett on playing tennis):

Thus all concepts—even apparently well-defined, abstract technical concepts—are always under construction. They are never wholly definable, and they defy the sort of categorical descriptions that are conventionally used in teaching: part of their meaning is always inherited from the contexts in which they are used. (Brown, et al., 1988, p. 5)

Using a pencil means having it in hand, but “using a concept” doesn't mean possessing some *thing*. S&W belie their continued adherence to the idea that *knowledge is a collection of representations* when they suggest that using a concept means *having* a representation inside. My dog and I can be representing in activity (e.g., gesturing, advancing forward in a playful way) without mentally manipulating representations (cf. Bamberger and Schön (1983)).

Also, notice that representing internally (e.g., visualizing or talking to ourselves) is not the same as having *representations* inside. I am perceiving when I visualize, but this is quite different from creating some physical representational structure and looking at it (as I am doing in my typing now). Put another way, I can be representing privately without creating or storing physical symbol structures.

Finally, this quote reveals S&W's tendency to discuss a "situationist position" as if we all—anthropologists, robot designers, educators, and philosophers alike—shared a common vocabulary and set of beliefs. Conflating the views of Brooks, Lave, Brown, Collins, and myself leads S&W to misunderstand what each of us have to say. We don't agree all the time, nor do we have common goals in developing this theory. We all go through stages of initial confusion and resolution⁴. For example, after reading the first draft of Brown, et al. (1988), like S&W, I wrongly equated formal learning with decontextualized learning, not realizing that the argument was against teaching abstractions out of context, not against teaching them at all. I wrote to Brown:

You have some distance to go in explaining the value of formal learning. You appear to completely dismiss "decontextualized learning," by definition claiming that it can only go on after experience.... The fact is that an apprentice doesn't and can't begin by becoming chummy with the masters. It isn't worth the master's time. (Is this person going to work in my field? Is he capable?) Nor is it necessary that everybody start with the analog of pressing clothes in order to get introduced to the culture and content of a field.... I'm just trying to anticipate reactions by people who will think you've gone too far or haven't considered the practical implications. (Personal communication, October 1988)

Today, I am struggling like anyone else to understand some of Lave's positions; I am prodding Brooks to be more careful to contrast his robot design claims with human psychology; and I am attempting to articulate my own expansion of the ideas so they are comprehensible to an AI and education audience. The fact is that S&W don't have a single foe to slay and I don't have a single position to defend. We're all in this together, trying to make sense of each other's representations, pointing and cajoling each other to attend to experience, to be more scholarly, to work out new models and their implications.

28. "Here again, concepts or no concepts is the question."

No. The question is whether knowledge consists of a network of concepts stored in memory or whether concepts are names observers use for recurrent ways of seeing and acting. The question is whether we act by referring to representations in a hidden way, or whether to *use* a concept (as in reading, in contrast to exhibiting knowledge of a concept) means to be awake, alert, and perceiving a representation. The question is whether it's representations all the way down, or whether the core—what enables us to talk in recurrent ways, to define, and describe laws—is a non-representational, non-computational mechanism.

⁴Other commentary on Brown, et al. (1988) includes some of S&W's objections: that knowledge has a static character like knives and pots, that Brown et al. are against formal learning, that the cottage cheese story is an example of ineptitude, and that experts have a greater "storehouse of usable knowledge" than just plain folks (Palincsar, 1989). See also Wineburg's (1989) claim that situated cognition ideas are not new and are well known by education theorists.

Educational implications—How situated should our education be?

29. “[Situated cognition claims that] to learn a subject means to take part in a particular community of practice. This entails what is termed ‘cognitive apprenticeship.’”

Cognitive apprenticeship is a term applied by Collins, Brown, and Newman (1986) to learning formal theories, in a specific kind of community of practice:

The computer enables us to go back to a resource-intensive mode of education, in a form we call cognitive apprenticeship (Collins, Brown, Newman, 1986)... Cognitive apprenticeship employs the modelling, coaching, and fading paradigm of traditional apprenticeship, but with emphasis on cognitive rather than physical skills. (Collins, 1988, p. 1)

Where conceptual and factual knowledge is addressed, cognitive apprenticeship emphasizes its use in solving problems and carrying out tasks. That is, in cognitive apprenticeship, conceptual and factual knowledge is exemplified and situated in the contexts of its use... It is this dual focus on expert processes and situated learning that we expect to help solve the educational problems of brittle skills and inert knowledge. (Collins, et al., 1986, p. 4)

Cognitive emphasizes that apprenticeship techniques actually reach well beyond the physical skills usually associated with apprenticeship to the kinds of cognitive skills more normally associated with conventional schooling. (Brown, et al., 1988, p. 25).

As examples of computer programs that promote cognitive apprenticeship, Collins (1988) cites WEST and Anderson’s geometry tutor. Collins, et al. (1986) emphasize the importance of representations for making processes visible, particularly for reflecting on one’s own performance. A key idea is that concepts are revealed and used in contexts similar to how they will be used in real life (Collins, 1988, p. 2). This is to be contrasted with either rote learning in a classroom without the opportunity to perform, be coached, and reflect. It is also contrasted with learning by doing in the real world, without the opportunity to use multiple representations, to explore without risk, or to replay what you have done (e.g., see Roschelle, 1991).

In effect, S&W believe that Brown, Collins, Duguid, and Newman are arguing against formal learning, when they are actually arguing how to make formal learning effective. S&W act as if cognitive apprenticeship is an argument against AI applications to education, but Collins (1988) uses AI programs to illustrate his points!

30. “Educational researchers advocating this view on learning provide a lot of anecdotal evidence to support their claims (Resnick, 1988; Schoenfeld, 1985; Lave, 1990; Pea, in press). What can we say about the promises of cognitive apprenticeship? What empirical findings do or do not support its claims?”

These researchers do tell many stories to exemplify their claims. But these stories are lifted from extensive, well-documented research programs. S&W’s rhetorical questions criticize Collins et al., for having no evidence or experimentation to back up their claims. In fact, their papers are full of human and computer teaching examples, including the work of Lampert, Schoenfeld, Scardamalia & Bereiter, Palincsar & Brown, Fischer, White & Frederiksen, vanLehn, etc. It is unclear why S&W fail to realize that the AI and education research of the past decade advances the idea of cognitive

apprenticeship, since Collins et al. clearly cite that work, even in the first statement of their thesis: “Current work on developing explicit, cognitive theories of domain skills, metacognitive skills, and tutoring skills is making the crucial first steps in the right direction.” (Collins, et al., 1986, p. 31)

31. “One of the consequences of the situated stance, in our view, is to refrain from decontextualising, from abstracting the particularities of a problem situation, to construct concepts.... Wouldn’t cognitive apprenticeship have the risk that decontextualisation will never take place, so the student would become an able practitioner in a limited set of situations...?”

But cognitive apprenticeship is *based on the idea of teaching abstractions in context*. In listing the benefits of situated learning, Collins et al. say:

Learning in multiple contexts induces the abstraction of knowledge so that students acquire knowledge in a dual form, both tied to the contexts of its uses and independent of any particular context. This unbinding of knowledge from a specific context fosters its transfer to new problems and new domains. (Collins, et al., 1986, p. 28)

What is at issue is how concepts are introduced and applied, how behavior is shaped by use of representations, and how reflection is facilitated by multimedia presentations. S&W are right that cognitive apprenticeship seeks to refrain from decontextualizing, but only in the sense of presenting abstractions in isolation, *not in the sense of ruling out abstraction*. Cognitive apprenticeship shows how abstracting should be related to practice. Indeed, Collins, et al. (1986) emphasize helping students articulate theories, compare strategies, and form and test hypotheses. Collins, et al. (1986) clearly articulate principles for combining theory and practice: increasing complexity, increasing diversity, and global before local skills (p. 25).

Educational implications—A situated example

32. “One could ask whether the solution described is indeed adequate. What if the weight-watcher’s material had been some liquid?”

The cottage cheese example has been the subject of considerable derision. But Brown, et al. (1988) are making a crucial point about the need for schooling to legitimize experiential knowledge and relate it to formal techniques:

Though schooling seeks to encourage problem solving, it disregards most of the inventive heuristics that students bring to the classroom. Instead of deploying such inventiveness to good effect, schools tend to dismiss it out of hand. It thus implicitly devalues not just individual heuristics, which may be fragile, but the whole process of inventively structuring cognition and solving problems. (Brown, et al., 1988, p. 14)

Bamberger’s (1991) work with musical notation exemplifies the point. Rather than teaching students conventional music notation from the start, she gives them bells, blocks, and pen and paper and encourages them to invent their own notation for well-known songs. The result is an amazing diversity of methods, many of which capture nuances of musical experience that standard notation cannot express. Students learn alternative ways of representing experience, how a given

representation can be meaningfully interpreted in different ways, and the advantages of having a standard notation.

Brown, Bamberger, Schön, Lave, and others are suggesting that we work with children's experience more closely, and give them a chance to create their own representations. We must reward inventiveness, rather than force-feeding society's preferred ways of thinking. To do this, we must value individual perspectives and non-formal, non-standard ways that students naturally build from their own experience.

The cottage cheese example celebrates inventiveness (Brown, et al., 1989). Nobody argues that the weight-watcher's methods could replace a formal coordinate system. But his method exemplifies the tricks that people invent all the time in their work (Lave, 1988). Representations evolve in our practice; adequacy is measured by our purposes and materials. If the material had been liquid other techniques might have been invented (as Scribner (1984) found among the milk loaders).

Brown et al. are not dismissing formal notations; but they want us to find ways of teaching conventions without destroying creativity, without leading kids to believe there is only one way to think, only one "correct" way of working problems. In particular, students are encouraged to use whatever materials are around them, to attend to their own experience and values, and to work in ways they find to be efficient and aesthetically pleasing. Underlying these ideas is the belief that the innovations of future workers, scientists, and engineers depends on inventively restructuring conventional views of problems (Schön, 1979; 1987; Zuboff, 1988), which in turn is inseparable from restructuring social interactions (Kukla, et al., 1990; Jordan and Alpert, 1991; Wenger, 1990).

Educational implications—More on cognitive apprenticeship

33. "Clancey is rather optimistic about the questions that students in cognitive apprenticeship would be able to pose. Instead of simply taking in what is being put forward as the expert view, they would become aware of differences of opinion that characterize the field.... Rather high-brow questions for a student to ask?"

I suggested that instead of just teaching Neomycin's disease taxonomy, we should get students to consider its origin, why taxonomies are important, and how disagreements arise in practice and are resolved (Clancey, 1991d). I am perplexed that such considerations would appear to be "high-brow" for a medical student. On the contrary, I have worried that the diagnostic procedure of Neomycin, with its esoteric terminology (e.g., "group and differentiate") is too theoretical.

I am claiming that intelligent tutoring systems offer the opportunity and advantage of teaching metacognitive skills (Clancey, 1988). For example, Guidon-Manage (Rodolitz and Clancey, 1990) is designed to give students practice in abstracting data requests in terms of a strategic language. By hypothesis, this provides a means of understanding a more-experienced physician's behavior in the clinic, as well as a means for the student to articulate what he or she doesn't know. In suggesting

that we reveal where such models come from, their limitations, and how to keep them up-to-date, I am proposing further metacognitive reflection to relate problem-solving theories to medical life. The ideas of cognitive apprenticeship stimulate me to make such connections.

34. “Are there not just as many sound arguments against cognitive apprenticeship as there are in favor?”

In arguing against cognitive apprenticeship, S&W are arguing against conceptual learning, metacognition, and the development of intelligent tutoring systems, which by every other remark, we would believe them to support. Situated learning does not mean “no abstractions,” but rather reconnecting formal education to everyday life:

Cognitive apprenticeship attempts to develop densely textured concepts out of and through continuing authentic activity. The term is closely allied to our image of knowledge as tool. Cognitive apprenticeship supports learning in a domain by enabling students to acquire, develop, and use conceptual tools in authentic domain activity, just as craft apprenticeship enables apprentices to acquire and develop the tools and skills of their craft through authentic work and membership of their trade. Through this process, apprentices enter the culture of practice. (Brown, et al., 1988, p. 24)

Attempting to realize this idea, researchers at the Institute for Research on Learning, for example, are working with scientists from Lawrence Livermore Labs to develop a mathematics curriculum on computers with neighboring high school teachers and Stanford University education students. The idea is to incorporate real-world examples of scientific modeling, to involve experts in classroom activity, and to invent new ways of teaching representations. This new beginning is motivated by the idea of cognitive apprenticeship; it is to be contrasted with the earlier approach of developing tools in the isolation of our laboratories and delivering them to students and teachers (e.g., Sleeman and Brown, 1982). In effect, cognitive apprenticeship ties to new ideas of software design (Greenbaum and Kyng, 1991)—not just theories of cognition, but new ways of developing tools in the context of practice.

35. “One should compare apprenticeship in tailoring to the way tailoring skill is acquired in western society, and not to a totally different type of skill like mathematics.”

Lave and Wenger (1991) claim that apprenticeship in West Africa fits patterns found the world-round. Of special interest is how the “organization of production” relates to “increasing participation and knowledgeable.” Contrasting tailoring with mathematics skills misses the point: Tailoring involves mathematics skills, but mathematics concepts are being articulated, learned, and enacted within a coordinated framework of culturally meaningful activity. Lave, Scribner, and others have focused on mathematics specifically in order to relate decontextualized symbol pushing to situated learning. The difference is between learning calculus and becoming a member of a society that uses calculus. Obviously, decontextualization occurs on a spectrum; few people are arguing that classroom learning be abolished. Cognitive apprenticeship is simply an argument for relating models

of the world and problem solving to everyday life, bolstered by new views of culture, cognition, and knowledge.

Conclusion

36. “Clancey arrives at the situationist stance by attacking the identity hypothesis.”

I wish the argument were that simple. In the final paragraphs of “Frame of reference” (Clancey, 1991c), I cite a dozen authors who influenced me over twenty years. I believe that attacking the identity hypothesis—that is, arguing against stored structures in memory—is a good way of summarizing the ideas, in a way that clearly has impact for how intelligent machines are designed. Furthermore, I believe that the stored-schema model of memory is at the heart of most cognitive theories of knowledge of the past few decades. Starting here helps us understand why being a member of a community of practice cannot be reduced to learning theories about the world or behavior rules.

37. “Situated cognition work today appears to throw away the baby with the bath water. It is all very well to postulate new architectures and mechanisms for cognition, but disregarding existing, viable theories on the grounds that they do not hold for the full range of complex phenomena is too drastic...”

First, situated cognition does not dismiss schema models as invalid or useless. The effect is to relate models to human behavior in a new way. Crucially, this involves recognizing the contributions of AI programming techniques so we don’t confuse limitations in the models with the tools themselves (Clancey, in press). For example, I said in the DELTA talk,

I will start by saying something about Artificial Intelligence, because we are not just going to throw away the old ways of building programs and the old ways of thinking. Instead, I believe we can generalize what AI-programming is in terms of a modeling methodology. (Clancey, 1991d, p. 4)

Furthermore, I have repeatedly said that our cognitive science models now serve as *specifications* for the mechanism we must construct (Clancey, 1991c, p. 372). For example, I have suggested to Brooks et al. that they use knowledge-level descriptions to formalize and compare the capabilities of robot designs (Clancey, in preparation, b).

Second, as the quotes from Collins, et al., show, cognitive apprenticeship has been inspired in part by AI applications to education. We are called to build on ITS models by relating them to practice.

38. “There is a danger of reductionism, i.e., reducing the mind to a simple organism interacting with its environment and producing complex behavior through the applications of behavioural rules.”

Ironically, this is precisely what stored-schema models are attempting to do (Newell and Simon, 1972). Lave et al. attack the approach of reducing culture, knowledge, and learning to manipulation of representations by a cognitive processor. Situated cognition research claims that *representations are the product of activity*, not the inner mechanism. All coordinated activity is an ongoing product

of interactions, not generated by applying stored procedures: “Knowledge is not independent but, rather, fundamentally ‘situated,’ being in part a product of the activity, context, and culture in which it is developed.” (Brown, et al., 1988, p. i)

39. “What we argue against is the complete shift of paradigm that situationists claim. Disregarding evidence and achievements of Cognitive and Instructional Sciences, and AI, is in our view overstating the issues.”

Cognitive apprenticeship is based on using expert systems and qualitative process models in new ways, not disregarding them (Collins, 1988). Furthermore, schema models provide an excellent starting point for reconceiving the nature of mental processes. For example, can we account for Neomycin’s procedural model of diagnostic interactions in terms of a constructive process that recoordinates perception and action along previous next-next-next sequences (Clancey, in preparation, a)? Looking back at the original protocols, can we now make better sense of the “noise” in the original physician interviews? Watching students using computer tutors, can we view their interaction as having dialectic, multiple social, psychological, and neural levels of coherent structure (Roschelle and Clancey, in preparation)?

40. “There is as yet enough room within the traditional theories, to advance towards a full theory of intelligent behaviors... Stating that AI and indeed educational science, have been built on faulty assumptions is too extreme a view for us to accept.”

In the end, we reach opposite conclusions, though we share the same respect for previous work. We must not forget that situated cognition incorporates a huge body of previous research that the stored-schema model suggests is wrong-headed, notably the work of Dewey, Vygotsky, Bartlett, Collingwood, Mead, Wittgenstein, Bateson, Gibson, and Maturana (Clancey, 1991c). A vast amount of psychological and philosophical research is ignored and contradicted by the stored-schema model. For example, Schank writes books about memory without relating his models to Bartlett’s.

Can the stored-schema model “advance towards a full theory of intelligent behaviors?” S&W may be convinced, but other researchers are marching ahead in a different direction.

4. On Transferability

Brown, et al. (1988, p. 1) argue, following Lave (1988), that “a situated theory of knowledge challenges the widely held belief that the abstraction of knowledge from situations is the key to transferability.” In saying this, they are surely not arguing that abstracting from multiple contexts is to be avoided, for their theory of cognitive apprenticeship is fully concerned with the development of conceptual models through feedback and reflection (Collins, et al., 1986, p. 4). They are surely not arguing that learning doesn’t generalize, for they emphasize the importance of metacognitive skills,

that is, patterns of reasoning that transcend particular problem situations. What they do want to emphasize is how abstractions are taught:

Where conceptual and factual knowledge is addressed, cognitive apprenticeship emphasizes its uses in solving problems and carrying out tasks. That is, in cognitive apprenticeship, conceptual and factual knowledge is exemplified and situated in the contexts of use. (Collins, et al., 1986, p. 4)

Again, situated learning focuses on a new theory of knowledge. Brown, et al. want us to realize that *knowing* is not possessing representations (facts, rules, and algorithms). “Transfer” is possible not because the student has memorized abstractions, but because these have become ways of seeing and coordinating activity (cf. Bamberger and Schön). The idea was well expressed by Collingwood more than 50 years ago:

Language is an activity; it is expressing oneself, or speaking. But this activity is not what the grammarian analyses. He analyses a product of this activity, ‘speech’ or ‘discourse’ not in the sense of a speaking or a discoursing, but in the sense of something brought into existence by that activity. (Collingwood, 1938, p. 254)

The shift may appear subtle at first; to understand Collingwood’s point, I suggest focusing on the idea that human memory is not a place where concepts are stored and that speaking is conceiving:

Knowledge cannot be inventoried. Knowing something is not having a thing, some substance in hand. The same is true of representations of meaning or context. Having in hand a representation of what a word means or of a situation is not understanding or being in a situation. Comprehending is not storing away a representation. Knowledge, like energy, is not a substance. The idea that a knowledge base could be functionally equivalent to human capability fundamentally misconstrues the relation between *processes* and *pattern descriptions*. (Clancey, 1991e)

Lave et al. object to the use of the term “transfer” because it suggests that knowing is a matter of mechanically reapplying inert concepts in different situations. “Transfer” is an inappropriate word because it suggests that some substance or tool is simply being *carried over*, like static objects, from one situation to the next. Lave et al. do not deny that we improve performance, that there is an effect of practice. But they deny the model of learning that assumes that generalizing from experience is inherently a representational activity, in which what carries over from one situation to the next is just a set of stored abstractions:

When “tool” is used as a metaphor for knowledge-in-use across settings, there is assumed to be an interaction between tool and situation, but only an *application* of a tool on different occasions. Since situations are not assumed to impinge on the tool itself, a theory of learning transfer does not require an account of situations, much less of relations among them. (Lave, p. 41)

In my view, to say that “situations impinge on the tool” is to emphasize that neural and social processes that coordinate behavior come into being interactively, during activity itself (again, contra Newell and Simon, learning is not a secondary process).

To understand that this is an argument about the nature of human memory, it is helpful when reading Brown, et al., to substitute the term *knowing* wherever they say “knowledge”:

Much current cognitive science...assumes that knowing is a process that can be separated from the activities and situations in which it is used...assumptions that knowledge [*knowing*] can be usefully regarded as self-contained and discrete....

Knowledge [*knowing*] is fundamentally a production of the mind and the world, which like woof and warp need each other to produce texture and to complete an otherwise incoherent pattern. It is impossible to capture the densely interwoven nature of conceptual knowledge [*knowing*] in explicit, abstract accounts. (Brown, et al., 1988, p. 1)

Knowledge [*knowing*] we claim, is partially embedded in the social and physical world. (Brown, et al., 1988, p. 2)

Authenticity in activity is paramount for learning if conceptual knowledge [*knowing*] is not self contained but, rather, if it is the product of and structured by the activity in which it is developed and deployed; if, in short, not just learning but knowledge [*knowing*] itself is situated. (Brown, et al., 1988, p. 15)

That is, Brown et al. (and Collingwood) are getting us to focus on the activity of being knowledgeable, of knowing, as opposed to focusing on knowledge as representations that are stored and mechanically applied.

To repeat, activity is inherently *situated* because all human action is constructed on the spot, as an interaction of neural and environmental processes, structured on multiple levels of emergent social and neural organization (Clancey, in preparation, a). At any moment, what we see and hear, how we move, what we say—in short how we coordinate our activity—is the result of ongoing interactions, built partly out of our sensory interactions, partly out of our ongoing neural compositions (corresponding to what we call focus of attention, strategies, intentions, and attitude), and partly out of the neural and environmental organizations that have occurred in the past. The effect is dialectic: We cannot causally attribute patterns of behavior to either neural or environmental forces alone, but to their ongoing interaction. Hutchins says this well:

When the context of cognition is ignored, it is impossible to see the contribution of the structure in the environment, in artifacts, and in other people to the organization of mental processes. (Hutchins, in press)

It should now be obvious why I believe the stored-schema model of memory must be overturned in order for AI researchers, psychologists, and cognitive scientists to understand the situated theory of knowing. The argument is not merely a prescription—that you should teach concepts in their situations of use. The argument is not merely a claim that people are reactive or respond without constantly referring to plans and rules or stored facts. Rather, for psychologists the real force of situated cognition is that structures in the brain are always new and constantly coming into being through an interactive process. The details of this theory are certainly unclear, but evidence suggests a compositional process by which neural maps are constructed out of previously selected perception-action coordinations (Clancey, in preparation, a).

Every perception and every coordination is a generalization of what we have done before (Vygotsky, 1934). Every perception and action is new at some level. We don't retrieve old categories and features. We don't have a past available to us to compare the present to, except in our

representations (Gibson's main point). Perception and action and our sense of what we are doing now emerge together, dialectically. Our behavior is inherently interactive with our environment (and the environment includes what we imagine inside our head). This idea is not entirely new—see Dewey (1902) on “situation, interaction, and continuity.”

Representations in themselves cannot be the “key to transferability” because they must be interpreted to be used. Crucially, interpreting is also a situated act (Suchman, 1987). I perceive a recipe and say what it means. This means that using representations (whether stored in a book or reconstructed in my speaking) involves conceiving something new. Though I may appear to speak and move in ways similar to how I have coordinated my activity in the past, my habits and routines are always improvised at some level, always novel, always generalized out of my previous neural constructions (Bartlett, 1932). It is this ability to re-coordinate and to generalize that a theory of knowledge based on representations alone cannot explain.

The fundamental claim is that our capability to interpret abstractions, to reuse a concept, and in general to be innovative rests on this non-representational process, by which perceiving and interacting with our environment are possible:

Even in cases where a fixed doctrine is transmitted, the ability of a community to reproduce itself through the training process derives not from the doctrine, but from the maintenance of certain modes of coparticipation in which it is embedded. (Lave and Wenger, 1991, p. 16)

In effect, Lave and Wenger are saying that *practice cannot be reduced to theory*. Inherently, knowing is a complex, on-going phenomenon, tacit in your perception, in your coordination, in your activity. Regardless of how we model human behavior and the world in terms of concepts, rules, and laws, there is always something in the on-going activity that is not captured by descriptions (Alexander, 1979; Wynn, 1991). Ultimately, pattern descriptions must demean the processes of interaction being described, for they suggest that the representation is the phenomenon itself.

From the anti-objectivist perspective (refer to Lakoff quotes, Section 2), this means that the world can be modeled, but what we perceive is coming into being as a result of our interactions with the world and interactions within it (Gregory, 1988): Our pattern descriptions, our scientific laws, don't generate what we observe. From the situated perspective, this means that beliefs, design rationales, goals, and what constitutes information are all coming into being during our activity; they do not lie behind the scenes as platonic descriptions that generate our behavior by a retrieval and matching process:

In the classical structural analysis, aspects of behavior are explained by, and serve as empirical evidence for, preexisting, “underlying” systems. It is these systems that provide the object of which an analysis is a model. To the extent that actual processes are analyzed, they are “structuralized”—made to follow from, or instantiate, structures. The activity of understanding, in such a view, comes down to recognizing and implementing instances of structure, filling them in with an overlay of situational particulars, and relating them to a “context” (which is in turn structured)... [In the work of Lave and

Wenger] it is not merely that the structure issue is transposed from the level of mental representations to that of participation frames. Rather this transposition is compounded by a more subtle and potentially radical shift from invariant structures to ones that are less rigid and more deeply adaptive. One way of phrasing this is to say that structure is more the variable outcome of action than its invariant precondition.... (W.F. Hanks, foreword to Lave and Wenger, 1991, p. 16)

It involves a prereflective grasp of complex situations, which might be reported as a propositional disposition, but is not one itself. (p. 20)

5. Conclusion

Sandberg and Wielinga state that “there are no strong reasons to leave the traditional paradigm of cognitive science and AI.” But when we relate the analysis of situated cognition to theories of human memory, we find fundamental shifts from the cognitive science model of knowledge and learning. Bamberger, Brown, Collins, Iran-Nejad, Lave, Schön, Suchman, Winograd, et al. adopt this new view because as scientists they base their theories of instruction and tool design on what they believe to be the best available understanding of human reasoning.

Adopting a new view of knowledge and representations does not mean rejecting past descriptions as being worthless. Instead, just as Newton’s laws are perfectly fine for planning airline schedules, we should certainly continue to use models of problem-solving strategies, taxonomies of diseases, and rules of discourse in instructional programs. The theory of cognitive apprenticeship gives us a new understanding of how such representations should be related to practice, how they promote reflection, and why reflection is important. We have only meager theories of how reflection reorients behavior (how talk in general affects behavior) (Bamberger and Schön, 1983), but we have a better appreciation of the questions we should be asking (Clancey, in preparation, a; b).

In conclusion, S&W’s claim that “formal education should not just be replaced by ‘cognitive apprenticeship’” is based on a misunderstanding of the term. Cognitive apprenticeship is proposed as an attempt to make formal education useful for society. It is inspired and exemplified by AI programs like Guidon-Manage. But we need to relate these programs better to everyday practice. The situated theory of knowing explains why this is so. Cognitive apprenticeship tells us what to do: We must reveal to students how such models are created, convey models in a way that brings students into the community of model builders, and just as important, bring instructional design into the community of practitioners (Greenbaum and Kyng, 1991; Clancey, in preparation, c). As this new view emerges, we can appreciate how other work in AI and software engineering calls us to better relate our representational artifacts to human society (e.g., Ehn, 1988; Floyd, 1987; Hughes, et al., 1991; Pollack, et al., 1982; Stefik and Conway, 1988).

Is this a paradigm shift? To many researchers who sense a new vitality and humanism in the process of technology design, who are suddenly overwhelmed by new ideas for relating technology to everyday life, it certainly feels that way.

References

- Agre, P. 1988. *Writing and representation*. Unpublished MIT Technical Report.
- Alexander, C. 1979. *A Timeless Way of Building*. New York: Oxford University Press.
- Bamberger, J. 1991. *The mind behind the musical ear*. Cambridge, MA: Harvard University Press.
- Bamberger, J. and Schön, D.A. 1983. Learning as reflective conversation with materials: Notes from work in progress. *Art Education*, March.
- Bartlett, F.C. [1932] 1977. *Remembering--A Study in Experimental and Social Psychology*. Cambridge: Cambridge University Press. Reprint.
- Bateson, G. 1972. *Steps to an Ecology of Mind*. New York: Ballentine Books.
- Braitenberg, V. 1984. *Vehicles: Experiments in Synthetic Psychology*. Cambridge: The MIT Press.
- Bransford, J.D., McCarrell, N.S., Franks, J.J., and Nitsch, K.E. 1977. Toward unexplaining memory. In R.E. Shaw and J.D. Bransford (eds), *Perceiving, Acting, and Knowing: Toward an Ecological Psychology*, Hillsdale, NJ: Lawrence Erlbaum Associates, pp. 431-466.
- Bresnan, J. and Kaplan, R.M. 1984. Grammars as mental representations of language. In W. Kintsch, J. R. Miller, and P. Polson (eds), *Method and Tactics in Cognitive Science*. Hillsdale, New Jersey: Lawrence Erlbaum Associates, pp. 103-136.
- Brown, J. S., Collins, A., and Duguid, P. 1988. *Situated cognition and the culture of learning*. IRL Report No. 88-0008. Shorter version appears in *Educational Researcher*, **18**(1), February, 1989.
- Brown, J. S., Collins, A., and Duguid, P. 1989. Debating the situation: A rejoinder to Palincsar and Wineburg. *Educational Researcher*, **18**(2), 10-12.
- Buchanan, B.G., & Shortliffe, E. H. 1984. *Rule-Based Expert Systems: The MYCIN Experiments of the Heuristic Programming Project*. Reading: Addison Wesley,
- Chi, M.T.H., Glaser, R., and M.J. Farr (eds) 1988. *The Nature of Expertise*. Hillsdale: Lawrence Erlbaum Associates.
- Clancey, W.J. 1988. The knowledge engineer as student: Metacognitive bases for asking good questions. In H. Mandl and A. Lesgold (eds), *Learning Issues for Intelligent Tutoring Systems*, New York: Springer-Verlag, pp. 80-113.
- Clancey, W.J. 1991a. Why today's computers don't learn the way people do. In P. Flasch and R. Meersman (editors), *Future Directions in Artificial Intelligence*. Amsterdam: Elsevier, pp. 53-62.
- Clancey, W.J. 1991b. Review of Rosenfield's "The Invention of Memory," *Artificial Intelligence*, **50**(2):241-284, 1991.
- Clancey, W.J. 1991c. The frame of reference problem in the design of intelligent machines. In K. vanLehn (ed), *Architectures for Intelligence: The Twenty-Second Carnegie Symposium on Cognition*, Hillsdale: Lawrence Erlbaum Associates, pp. 357-424.
- Clancey, W.J. 1991d. Invited talk. *AI Communications—The European Journal on Artificial Intelligence* **4**(1):4-10.
- Clancey, W.J. 1991e. Situated Cognition: Stepping out of Representational Flatland. *AI Communications—The European Journal on Artificial Intelligence* **4**(2/3):109-112.
- Clancey, W.J. 1992. Model construction operators. *Artificial Intelligence*, **53**(1):1-115.
- Clancey, W.J. (in preparation a). *Interactive control structures: Evidence for a compositional neural architecture*. Submitted for publication.
- Clancey, W.J. (in preparation b). A Boy Scout, Toto, and a bird: How situated cognition is different from situated robotics. A position paper prepared for the NATO Workshop on *Emergence, Situatedness, Subsumption, and Symbol Grounding*. To appear in a special issue of the *AI Magazine*, Brooks and Steels (eds).

- Clancey, W.J. (in preparation c). The knowledge level reconsidered: Modeling socio-technical systems. To appear in *The International Journal of Intelligent Systems*, special issue on knowledge acquisition, edited by Ken Ford.
- Collingwood, R.G. 1938. *The Principles of Art*, London: Oxford University Press.
- Collins, A., Brown, J.S., and Newman, S.E. [1986] 1989. Cognitive apprenticeship: Teaching the craft of reading, writing, and mathematics. In L.B. Resnick (ed), *Cognition and Instruction: Issues and Agendas*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Collins, A. 1988. Cognitive apprenticeship and instructional technology. In B.F. Jones and L. Idol (eds), *Dimension of Thinking and Cognitive Instruction*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Davis, R. and Lenat D.B. 1982. *Knowledge-Based Systems in Artificial Intelligence*. New York: McGraw-Hill.
- Dewey, J. 1902. *The Child and the Curriculum*, Chicago: University of Chicago Press.
- Edelman, G.M. 1987. *Neural Darwinism: The Theory of Neuronal Group Selection*. New York: Basic Books.
- Ehn, P. 1988. *Work-Oriented Design of Computer Artifacts*. Stockholm: Arbeslivscentrum.
- Evans, D. and Patel, V. (eds) 1990. *Medical Cognitive Science*. Cambridge: Bradford Books.
- Floyd, C. 1987. Outline of a paradigm shift in software engineering. In Bjercknes, et al., (eds) *Computers and Democracy—A Scandinavian Challenge*, p. 197.
- Freeman, W. J. 1991. The Physiology of Perception. *Scientific American*, (February), 78-85.
- Gardner, H. 1985. *The Mind's New Science: A History of the Cognitive Revolution*. New York: Basic Books.
- Goldstein, I.P., 1982. The genetic graph: a representation for the evolution of procedural knowledge. In D.Sleeman and J.S. Brown, *Intelligent Tutoring Systems*. London: Academic Press, pp. 51-78.
- Greenbaum J. and Kyng, M. 1991. *Design at Work: Cooperative design of computer systems*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Gregory, B. 1988. *Inventing Reality: Physics as Language*. New York: John Wiley & Sons, Inc.
- Hayes-Roth, F., Waterman, D.A., and Lenat, D.B. 1983. *Building Expert Systems*. Reading, MA: Addison-Wesley.
- Hughes, J. Randall, D., and Shapiro, D. 1991. CSCW: Discipline or Paradigm? A sociological perspective. In L. Bannon, M. Robinson, and K. Schmidt (eds), *Proceedings of the Second European Conference on Computer-Supported Cooperative Work*. Amsterdam, pp. 309-323.
- Hutchins, E. in press. Learning to Navigate. In S.Chalkin and J. Lave (eds). *Understanding Practice*. New York: Cambridge University Press.
- Iran-Nejad, A. 1987. The schema: A long-term memory structure or a transient functional pattern. In R. J. Tierney, Anders, P.L., and J.N. Mitchell (eds), *Understanding Readers' Understanding: Theory and Practice*, (Hillsdale, Lawrence Erlbaum Associates)
- James, W. [1892] 1984. *Psychology: Briefer Course*. Cambridge, MA: Harvard University Press. Reprinted with annotations.
- Jenkins, J.J. 1974. Remember that old theory of memory? *Well, forget it! American Psychologist*, November, pps. 785-795.
- Johnson-Laird, P.N. 1983. *Mental Models: Towards a cognitive science of language, inference, and consciousness*. Cambridge, MA: Harvard University Press.
- Jordan, J. and Alpert, B. 1991. *Technology and Social Interaction*, Xerox-PARC Technical Report.
- Kintsch, W., Polson P.G., and Miller, J.R. (eds). 1984. *Method and Tactics in Cognitive Science*. Hillsdale, NJ: Lawrence Erlbaum Associates.

- Kukla, C.D., Clemens, E.A., Morse, R.S., and Cash, D. 1990. An approach to designing effective manufacturing systems. To appear in *Technology and the Future of Work*.
- Lakoff, G. 1987. *Women, Fire, and Dangerous Things: What Categories Reveal about the Mind*. Chicago: University of Chicago Press.
- Lave, J. 1988. *Cognition in Practice*. Cambridge: Cambridge University Press.
- Lave, J. and Wenger, E. 1991. *Situated Learning: Legitimate Peripheral Participation*. Cambridge: Cambridge University Press.
- Lehnert, W.G. 1984. Paradigmatic issues in cognitive science. In W. Kintsch, J. R. Miller, & P. Polson (eds), *Method and Tactics in Cognitive Science*. Hillsdale, New Jersey: Lawrence Erlbaum Associates, pp. 21-50.
- Maes, P. 1990. Designing Autonomous Agents, Guest Editor. *Robotics and Autonomous Systems* 6(1,2) 1-196.
- Mandler, G. 1984. Cohabitation in the cognitive sciences. In W. Kintsch, J. R. Miller, & P. Polson (ed), *Method and Tactics in Cognitive Science*. Hillsdale, New Jersey: Lawrence Erlbaum Associates, pp. 305-316.
- Maturana, H. R. 1983. What is it to see? ¿*Qué es ver?* 16:255-269. Printed in Chile.
- Miller, R., Polson, P.G., and Kintsch, W. 1984. Problems of methodology in cognitive science. In W. Kintsch, J. R. Miller, & P. Polson (eds), *Method and Tactics in Cognitive Science*. Hillsdale, New Jersey: Lawrence Erlbaum Associates, pp. 1-18.
- Minsky, M. 1977. Frame theory. In P.N. Johnson-Laird and P.C.Wason, *Thinking: Reasings in Cognitive Science*, Cambridge: Cambridge University Press, pp. 355-376.
- Newell, A. 1982. The knowledge level. *Artificial Intelligence*, 18(1):87-127, January.
- Newell, A. and Simon, H.A. 1972. *Human Problem Solving*. Engelwood Cliffs, NJ: Prentice-Hall.
- Palincsar, A.S. 1989. Less charted waters. *Educational Researcher*, 18(2), 5-7.
- Pollack, M.E. Hirshberg, J., and Webber, B. 1982. User participation in the reasoning processes of expert systems. *Proceedings of the National Conference on Artificial Intelligence*, Pittsburgh. pp. 358-361.
- Rodolitz, N.S., & Clancey, W. J. 1989. GUIDON-MANAGE: teaching the process of medical Diagnosis. In D. Evans, & V. Patel (eds), *Medical Cognitive Science*. Cambridge: Bradford Books, pp. 313-348.
- Roschelle, J. 1991. *Students' construction of qualitative physics knowledge: Learning about velocity and acceleration in a computer microworld*. Unpublished doctoral dissertation, University of California, Berkeley.
- Roschelle, J. and Clancey, W. J. in preparation. *Learning as Neural and Social*. Presented at AERA91, Chicago. To appear in a special issue of the *Educational Psychologist*.
- Rosenfield, I. 1988. *The Invention of Memory: A new view of the brain*. New York: Basic Books, Inc.
- Schön, D.A. 1979. Generative metaphor: A perspective on problem-setting in social policy. In A. Ortony (ed), *Metaphor and Thought*. Cambridge: Cambridge University Press. pp. 254-283.
- Schön, D.A. 1987. *Educating the Reflective Practitioner*. San Francisco: Jossey-Bass Publishers.
- Schön, D.A. 1990. *The theory of inquiry: Dewey's legacy to education*. Annual meeting of the American Educational Research Association, San Francisco.
- Scribner, S. 1984. Studying working intelligence. In B. Rogoff and J. Lave (eds), *Everyday Cognition: Its Development in Social Context*. Cambridge, MA: Harvard University Press, pp. 9-40.
- Sleeman, D. and Brown, J.S. 1982. *Intelligent Tutoring Systems*. London: Academic Press.
- Steels, L. 1989. Cooperation through self-organisation. In Y. De Mazeau and J.P. Muller (eds), *Multi-agent Systems*. Amsterdam: North-Holland.

- Stefik, M. and Conway, L. 1988. Towards the principled engineering of knowledge. In R. Englemore (ed), *Readings From the AI Magazine, Volumes 1-5, 1980-85*. Menlo Park, CA: AAAI Press: pp.135-147.
- Suchman, L.A. 1987. *Plans and Situated Actions: The Problem of Human-Machine Communication*. Cambridge: Cambridge University Press.
- Tyler, S. 1978. *The Said and the Unsaid: Mind, Meaning, and Culture*. New York: Academic Press.
- vanLehn, K. (ed) 1991. *Architectures for Intelligence: The Twenty-Second Carnegie Symposium on Cognition*, Hillsdale: Lawrence Erlbaum Associates.
- Vygotsky, L. (1934) 1986. *Thought and Language*. Cambridge: The MIT Press. Edited by A. Kozulin.
- Wenger, E. 1990. *Toward a theory of cultural transparency: Elements of a social discourse of the visible and the invisible*. PhD. Dissertation in Information and Computer Science, University of California, Irvine.
- Wineburg, S.S. 1989. Remembrance of theories past. *Educational Researcher*, **18**(2), 7-10.
- Winograd, T. and Flores, F. 1986. *Understanding Computers and Cognition: A New Foundation for Design*. Norwood: Ablex.
- Wynn, E. 1991. Taking Practice Seriously. In J. Greenbaum and M. Kyng (eds), *Design at Work: Cooperative design of computer systems*. Hillsdale, NJ: Lawrence Erlbaum Associates, pp. 45-64.
- Zuboff, S. 1988. *In the Age of the Smart Machine: The future of work and power*. New York: Basic Books, Inc.