The Conceptual Nature of Knowledge, Situations, and Activity

William J. Clancey

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

Clancey, W. J. 1997. The conceptual nature of knowledge, situations, and activity. In P. Feltovich, K. Ford, and R. Hoffman (Eds.), Human and Machine Expertise in Context (pp. 247–291). Menlo Park, CA: The AAAI Press.

(1). BACKGROUND: THE AUDIENCE AND THE PROBLEM

Research papers are ultimately personal statements, locating the author's developing thought along a path from what is now seen as naive toward what is viewed as a hopeful redirection. From 1974 to 1987, I was part of a community of AI researchers who devised computer programs that could diagnose diseases, engage in case-method discourse for teaching, and model students' problem solving strategies (Buchanan and Shortliffe, 1984). Following the rubric of "knowledge-based systems," we believed not only that knowledge could be represented in rules ("If there is evidence of bacterial meningitis and the patient is an alcoholic, then therapy should cover for diplococcus organisms"), but also that a body of such rules would be functionally equivalent to what an expert physician can do. We knew that the physician knew more, but we assumed his or her knowledge simply consisted of *more rules*.

The assumption that human knowledge consists exclusively of words organized into networks of rules and pattern descriptions ("frames") guided the creation of hundreds of computer programs, described in dozens of books such as *Building Expert Systems* (Hayes-Roth et al., 1983), *Intelligent Tutoring Systems* (Sleeman and Brown, 1982), and *The Logical Foundations of Artificial Intelligence* (Genesereth and Nilsson, 1987). Certainly, these researchers realized that processes of physical coordination and perception involved in motor skills couldn't easily be replicated by pattern and rule descriptions. But such aspects of cognition were viewed as "peripheral" or "implementation" concerns. According to this view, intelligence is *mental*, and the content of thought consists of networks of words, coordinated by an "architecture" for matching, search, and rule application. These representations, describing the world and how to behave, serve as the machine's knowledge, just as they are the basis for human reasoning and judgment. According to this "symbolic approach" to building an artificial intelligence, descriptive models not only *represent* human knowledge, they correspond in a maplike way to *structures stored* in human memory. By this view, a descriptive model is an explanation of human behavior because *the model is the person's knowledge*—stored inside, it directly controls what the person sees and does

The distinction between representations (knowledge) and implementation (biology or silicon), called the "functionalist" hypothesis (Edelman, 1992), claims that although AI engineers might learn more about biological processes of relevance to understanding the nature of knowledge, they ultimately will be able to develop a machine with human capability that is not biological or organic. This strategy has considerable support, but unfortunately, the thrust has been to *ignore the differences* between human knowledge and computer programs and instead to tout existing programs as "intelligent."

Emphasizing the similarities between people and computer models, rather than the differences, is an ironic strategy for AI researchers to adopt, given that one of the central accomplishments of AI has been the formalization of *means-ends analysis* as a problem-solving method: Progress in solving a problem can be made by describing the difference between the current state and a goal state and then making a move that attempts to bridge that gap.

Given the focus on symbolic inference, cognitive studies have appropriately focused on aspects of intelligence that rely on descriptive models, such as in mathematics, science, engineering, and medicine—the professional areas of human expertise. Focusing on professional expertise has supported the idea that "knowledge equals stored models" and hence has produced a dichotomy between physical and intellectual skills. That is, the distinction between physical skills and "knowledge" is based on an assumption, which was instilled in many professionals in school, that "real knowledge" consists of scientific facts and theories. By this view, intelligence is concerned only with articulated belief and reasoned hypothesis.

But understanding the nature of cognition requires considering more than the complex problem solving and learning of human experts and their tutees. Other subareas of psychology seek to understand more general aspects of cognition, such as the relation of primates to humans, neurological dysfunction, and the evolution of language. Each of these requires some consideration of how the brain works, and each provides some enlightening insights for robot builders.[1] In this respect, the means-ends approach I promote is a continuation of the original aim of cybernetics: to compare the mechanisms of biological and artificial systems.

By holding current computer programs up against the background of this other psychological research, cognitive scientists can articulate differences between human knowledge and the best cognitive models. Although questions about the relation of language, thought, and learning are very old, computational models provide an opportunity to test theories in a new way—by building a mechanism out of descriptions of the world and how to behave and seeing how well it performs. Gardner (1985) says this is precisely the opportunity afforded by the computational modeling approach:

> Only through scrupulous adherence to computational thinking could scientists discover the ways in which humans actually *differ* from the serial digital computer—the von Neumann computer, the model that dominated the thinking of the first generation of cognitive scientists. (p. 385)

---

[1] See especially the articles and commentary in *Behavioral and Brain Sciences*, e.g., Donald (1993).

Gardner concludes from such comparisons that cognitive scientists should substantially broaden their view of mental processes. This chapter is in the same spirit, stepping out from what AI programs do to inquire how such models of cognition relate to human knowledge and activity. I frame strategies for bridging the gap, as well as appropriately using the technology developed to date.

An exposition of the differences between people and computers necessarily requires examples of what computers cannot yet do. Such descriptions are to some extent poetic—a style of analysis promoted by Oliver Sacks in books such as *The Man Who Mistook His Wife for a Hat*—because they cannot yet be programmed. This analysis irks some AI researchers and has been characterized as "asking the tail of philosophy to wave the dog of cognitive science" (Vera and Simon, 1993). Through an interesting form of circularity, descriptive models of scientific discovery shape how some researchers view the advancement of their science: If aspects of cognition cannot be modeled satisfactorily as networks of words, then work on these areas of cognition is "vague," and comparative analysis is "nonoperational speculation." Here lies perhaps the ultimate difficulty in bridging different points of view: The scientific study of human knowledge only partially resembles the operation of machine learning programs. In people, nonverbal conceptualization can organize the search for new ideas. Being aware of and articulating this difference is pivotal in relating people and programs.

To understand people better, a broader view of conceptualizing is required, one which embraces the nonverbal, often called "tacit," aspects of knowledge (Subsections 1.1 and 1.2). "Situated action" can then be understood as a psychological theory (Subsection 1.3). To illustrate how knowledge, situations and activity are dynamically related to descriptions, I present the example of how the Seaside community developed from a central plan (Section 2). This example reveals how descriptions such as blueprints, rules of thumb, and policies are used in practice—they are not knowledge itself, but *means* of guiding activities and resolving disputes. In this analysis, I will distinguish between *concepts* (what people know), *descriptions* (representations people create and interpret to guide their work), and *social activity* (how work and points of view are coordinated). On this basis, I articulate the difference between information-processing *tasks*, as described in cognitive models of expertise (Chi et al., 1988), and *activities*, which are *conceptualizations* for choreographing how and where tasks are carried out (Section 3). The confusion between tasks and activities is rooted in the identification of descriptions with concepts, and accounts for the difficulty in understanding that *situations are conceptual constructs*, not places or problem descriptions (Section 4). Finally, from this perspective, having reconstellated knowledge, context, and representational artifacts, I consider specific suggestions for using tools such as expert systems and computer programs in general (Section 5).

(1.1). Concepts are More than Networks of Words

In this chapter, I explain the idea of situated cognition (e.g., see Gardner, 1985; Lakoff, 1987; Sacks, 1987; Bruner, 1990; Edelman, 1992), which I take as a broad approach for re-relating human knowledge and AI programs. In contrast to the dominant view of AI in the past decades, the theory of situated cognition claims that knowledge is not a *set of descriptions*, such as a collection of facts and rules. AI researchers and cognitive scientists model knowledge by descriptions in cognitive models and expert systems. But the map is not the territory (Korzybski, 1941): Human conceptualization has properties relating to learning and flexibility that make human knowledge different from procedures and semantic networks in a computer program. Situated cognition research explores the idea that conceptual knowledge, as a capacity to coordinate and sequence behavior, is inherently formed as part of and within physical performances.

The force of this claim is that a machine constructed from networks of words alone, which works in the manner of the production rule architecture described by Newell (1991), cannot learn or perform with the flexibility of a human (Dreyfus and Dreyfus, 1986). The hypothesis is that a theory of knowledge that equates meaning and concepts with networks of words fundamentally fails to distinguish between *conceptualization (*a form of physical coordination), *experience* (such as imagining a design), and *cultural artifacts* (such as documents and expert systems). Such distinctions are made by Dewey (1902), most obviously in his critique of Bertrand Russell's "devotion to discourse" (Dewey, 1939). Today Dewey's view is associated with the "contextualism" of ecological psychology (Barker, 1968; Turvey and Shaw, in press) and the sociology of knowledge (Berger and Luckman, 1966). Earlier in the century it was called "functionalism" (Harrison, 1977), meaning "activity-oriented,"[2] in the philosophy of James (1890), Dewey, and Mead (1934), and carried further into a theory of language as a tool by the social psychology of Bartlett (1932) and Vygotsky (Wertsch, 1991).

In contrast, the AI literature, exemplified by a collection (vanLehn, 1991) that presents the work of many distinguished researchers, equates the following terms:

---

[2]"Functionalism" is a theory about the evolutionary, developmental, instrumental content of knowledge, not to be confused with the "functionalist" hypothesis that a physical symbol system in a computer is "functionally equivalent" to the operation of the human brain.

"knowledge"
"knowledge representations"
"representations"
"mental models"
"knowledge base"
"concepts"

For example, the following recently appeared in the *AI Magazine*: "The situationalists are attacking the very *idea* of knowledge representation—the notion that cognitive agents think about their environments, in large part, by manipulating internal representations of the worlds they inhabit"(Hayes et al., 1994, p. 17). Here the idea of "knowledge representation" is equated with the idea that "knowledge is representational." By this view, a representation of knowledge is not just a description in a model, *a scientist's representation of a subject's knowledge*, but literally something manipulated internally in the subject's brain. The computational view hypothesizes that the scientist's model and subject's knowledge are equivalent in both notation (knowledge representation language) and architecture (the knowledge base interpreter and relation of sensation and models to motor processes). Zenon Pylyshyn stated this hypothesis explicitly in his commentary presented at an AI symposium in 1988:

> The choice of both notation and architecture are central empirical issues in cognitive science, and for reasons that go right to the heart of the computational view of mind. It's true that in the physical sciences, theoretical notation is not an empirical issue. But in cognitive science our choice of notation is critical precisely because the theories claim that representations are written in the mind in the postulated notation: that at least some of the knowledge is explicitly represented and encoded in the notation proposed by the theory. The architecture is likewise important because the claim is that these are literally the operations that are applied to the representations.... In cognitive science, theories claim that the mind works the way the model does, complete with notation and architecture. What is sometimes not appreciated is that computational models are models of what literally goes on in the mind (Pylyshyn, 1991, p. 219).

Sometimes human knowledge and descriptions in a model are equated quite deliberately, as in Zenon Pylyshyn's frank statement; other claims about concepts, mental models, and knowledge bases become so ingrained that scientists do not reflect upon them. George Lakoff (1987) provides perhaps the best historical review of the paradigm:

> 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 (pp. xii-xiii).

Since the late 1980s, with the airing of alternative points of view, some AI researchers have argued that claims by situated cognition adherents about the descriptive modeling approach were all straw men, or that only expert systems were based on the idea that human memory consisted of a storehouse of descriptions. Certainly, the idea that "knowledge equals representation of knowledge" is clear in the expert systems literature:

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

Knowledge engineers in the decade starting about 1975 viewed expert systems as just a straightforward application of Newell and Simon's physical symbol system hypothesis:

> A consequence of the prominence of the physical symbol system hypothesis is the recent emergence of the representation of knowledge as one of the most central enterprises of the field. Almost every AI project of recent vintage—from natural language understanding to visual perception to planning to expert systems—has employed an explicit symbolic representation of the information in its domain of concern. General languages for representing arbitrary knowledge are becoming a focus in this preoccupation with using symbols for facts and metainformation for a given domain.... One of the working hypotheses in this field is that knowledge is representational; that is, "knowing" consists in large part of representing symbolically facts about the world. This lends support to Newell's physical symbol system hypothesis.... (p. 45-6)

Notice how the claims go beyond saying that "knowledge is representational" to argue that knowledge is "explicit" and "symbolic," which in expert systems means that knowledge is represented as rules or other associational patterns in words. The symbols are not just arbitrary patterns, they are *meaningful encodings*:

> It is sufficient to think of symbols as strings of characters and of symbol structures as a type of data structure... The following are examples of symbols: Apple, Transistor-13, Running, Five, 3.14159. And the following are examples

> of symbol structures: (On Block1 Block2) (Plus 5 X) (Same-as (Father-of Pete) (Father-of (Brother-of Pete)))." (p. 61)

Although it is true that this point of view, equating knowledge with word networks, remained controversial among philosophers, it was the dominant means of modeling cognition throughout the 1980s. 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)

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, that knowledge consists of stored descriptions, and, by assumption, this is not the source of theoretical deficiencies:

> 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)

Although Hayes, Ford, and Agnew in particular have tried to associate the view that knowledge representations are knowledge exclusively with expert systems research, it is easy to find examples in the cognitive psychology literature, as the quote from Bresnan attests. For example, Rosenbloom recently described how Soar's architecture "supports knowledge": "Productions provide for the explicit storage of knowledge. The knowledge is stored in the actions of productions, while the conditions act as access paths to the knowledge." (Rosenbloom et al., 1991, p. 81) Again, by this view knowledge is something that can be stored. Soar's productions *are* knowledge.

But Hayes et al. are correct that one can find more balanced treatments. Michalski provides the following appraisal of machine learning:

> An intelligent system must be able to form concepts, that is classes of entities united by some principle. Such a principle might be a common use or goal, the same role in a structure forming a theory about something, or just similar perceptual characteristics....
>
> In research on concept learning, the term "concept" is usually viewed in a narrower sense...namely, as an equivalence class of entities, such that it can be *comprehensibly described* by no more than a small set of statements. This description must be sufficient for distinguishing this concept from other concepts. (Michalski, 1992, p. 248, emphasis added)

When knowledge is equated with descriptions comprehensible to human readers, a mental model is equated with the data structure manipulations of a computer program, and all representing in the mind is reduced to a vocabulary of symbols composed into texts. Consequently, when situated cognition researchers deny that mental representing is a process of manipulating text networks (e.g., Brooks, 1991; Suchman, 1987), some AI researchers interpret this as claiming that there are "no internal representations" *at all* (Hayes et al., 1994) or "no concepts in the mind" (Sandberg and Wielinga, 1991; Clancey, 1992b). Actually, the claim is that human concepts cannot be *equated* with descriptions, such as semantic networks. Put another way, manipulating symbolic expressions according to mathematical transformation rules and conceptualizing are different kinds of processes.

"Knowledge," as a technical term, is better viewed as an analytic abstraction. Like energy, knowledge is not a substance that can be in hand (Newell, 1982). Sometimes it is useful to view knowledge metaphorically as being a thing; describing it and measuring it as a "body of knowledge." For example, a teacher planning a course or writing a textbook adopts this point of view; few people argue that such forms of teaching should be abolished. But more broadly construed, human knowledge is dynamically forming as adaptations of past coordinations (Edelman, 1992). Therefore we cannot inventory what someone knows, in the sense that we can list the textual contents (facts, rules, and procedures) of a descriptive cognitive model or expert system.

AI researchers are often perplexed by these claims. One colleague wrote to me:

> There is something of great value that humans store in libraries—usually called knowledge—that helps us to interact successfully with our environment (not to mention entertaining us). There is something (knowledge seems like a good word for it) that is useful to various degrees across people, cultures, and time. It sounds like you are denying this.

Identifying knowledge with books in a library is identifying human memory with texts, diagrams, and other descriptions. This is indeed the folk psychology view. But just as we found that the brain is not a telephone switchboard, situated cognition claims that progress in understanding the brain is inhibited by continuing to identify knowledge with *artifacts* that knowledgeable people create, such as textbooks and expert systems. AI needs a better metaphor if it is to replicate what the brain accomplishes. In the next subsection, I introduce the epistemological implications of situated cognition. The suggested metaphor is not knowledge as a substance, but as a dynamically-developed coordination process.

(1.2). Relating Knowledge, Activity, and Situations

From a psychologist's perspective, the theory of situated action (Mills, 1940; Suchman, 1987) claims that knowledge is *dynamically constructed* as we conceive of what is happening to us, speak, and move (Thelen and Smith, 1994). Most importantly, social scientists emphasize that a person's *conception of activity* is with respect to social norms, and this is how we coordinate our experience and action. Action is thereby *situated in a person's role* as a member of a community. The common idea in the literature of AI that "situated" means "in a physical setting" or merely "interactive" (Vera and Simon, 1993) distorts the psychological nature of the theory and the social nature of activity.

An activity is not merely a movement or action, but a complex choreography of identity, sense of place, and participation, which *conceptually regulates* our behavior. Such conceptual constraints enable value judgments about how we use our time, how we dress and talk, the tools that we prefer, what we build, and our interpretations of our community's goals and policies. That is, our conception of what we are doing, and hence the *context* of our actions, is always social, even though we may be alone. Professional expertise is therefore "contextualized" in the sense that it reflects knowledge about a community's activities of inventing, valuing, and interpreting theories, designs, and policies (Nonaka, 1991; Collins, this volume). This conceptualization of context has been likened to the water in which a fish swims (Wynn, 1991); it is tacit, pervasive, and necessary.

The construction of the planned town of Seaside, Florida illustrates how a community's knowledge enables it to coordinate *scientific facts* about the world (such as hurricane and tide data), *designs* (such as architectural plans), and *policies* (social and legal constraints on behavior). Schön (1987, p. 14) claims that "professionalism"—"the replacement of artistry by systematic, preferably scientific knowledge"—ignores the distinction between science, design, and policy. Expertise is defined by professionalism as if it were scientific knowledge alone, in terms of what can be studied experimentally, written down, and taught in schools. Correspondingly, the nature of knowledge is narrowed to “truths about the world,” and facts for solving problems are viewed in terms of mathematical or naturally-occurring objects and properties. Professionalism thus equates *the work of creating designs* and *interpreting policies*, in which we construct a social reality of value judgments, artifacts and activities, with *the work of science*. Consequently, the “social construction of knowledge” (Berger and Luckman, 1966) is equated with the development of theories about *nature*, when its force should instead be directed at understanding the social origin and resolution of problems in *everyday work*. As a result of this confusion, claims about knowledge construction in design and policy interpretation are viewed as forms of “relativism” and hence “antiscientific.”

Identifying the application of theory *in* practice with the development of theory itself (science) has led to some unfortunate exchanges in print. For example, when Lave says “The fashioning of normative models of thinking from particular, ‘scientific’ culturally valued, named bodies of knowledge is a cultural act” (Lave, 1988, p. 172), she is referring to how cognitive researchers and schools apply mathematical theory to evaluate everyday judgments, such as using algebra to appraise grocery shoppers’ knowledge in making price comparisons. Shoppers may measure as mathematically incompetent in tests, but on the store floor be fully capable of making the qualitative judgments that fit their needs (by comparing packages proportionally, for example). Thus, human problem solving is seen primarily through the glasses of formal theories—a normative model of “how practice should be” is fashioned from the world view of science. This is essentially how professionalism characterizes expertise in general. But Hayes et al. read this as *an attack on science*: “RadNanny claims that ... science merely reflects the mythmaking tendencies of the human tribe, with one myth being no more reality-based than another.” (p. 23)

The two points of view are at cross purposes: Lave criticizes the application of scientific norms of measurement and calculation to understanding human behavior; Hayes et al. criticize the application of cultural studies to understanding science. Reconciling these points of view involves allowing that some human knowledge is judgmental, nonverbal, and contextually valued, that is, not reducible in principle to scientific facts or theories.[3] Attempts to replace or equate all knowledge with descriptions

[3]Lave places “scientific” in quotes to highlight that using mathematical theory to model certain cultural phenomena is itself a cultural activity. She wishes to distinguish scientific *theories* from the *application of theories* in studying human behavior. This application is often equated with science itself, hence the scare

leaves out the perceptual-conceptual coordination that makes describing and interpreting plans possible, as the Seaside example illustrates.

(2). CONSTRUCTING A COMMUNITY

Seaside is a planned community, a beach front tract of about 80 acres located in the Florida panhandle. Practice, pattern descriptions, and theory all play a part in the ongoing development of the town. A brochure describes how the town was originally designed:

> Careful study of small towns in the south and in Florida in particular provided the planning team with a set of planning and building standards that had withstood the test of time. Street widths, distances between structures, sidewalks, street trees and lighting, building forms and material in these towns were documented and distilled into a code which made explicit the unwritten rules which for generations had guided the making of building and towns in the region. In reviving this neglected tradition, buildings would be produced which were well-adapted to the local climate and which worked together to form coherent streets and squares and a community with a strong sense of place. (From a Seaside brochure)

Against the backdrop of freshly-painted pastel homes of varying sizes, one finds BMWs and Weber grills, the stuff of the late twentieth century. Amid this oddly familiar and Disneyesque superorganization, one finds adaptations for the place and time (Dunlop, 1989):

> Though it is a strong plan armed with an equally strong companion code, it is not so rigid that it can’t be modified. Along the way, Leon Krier [builder of a prominent home in Seaside] looked at the plan and suggested adding the pedestrian paths that now form a second network through the town. Homeowners chose to pave the streets in red brick; they originally were of crushed shell, more pleasing to the eye but less so to bare feet.

Given the scale and ownership by individuals, the results were not entirely intended or controlled by the organizing committee:

---

quotes: Cognitive “science” is partly about the design of schools; cognitive research embodies certain cultural assumptions about what should be learned in schools (by virtue of what school subjects are studied) and how to assess the adequacy of teaching (in applying rational models of optimality to evaluate practical choices). Using science to direct everyday affairs is a cultural value, not part of science itself.

> Duany [one of the architects hired to produce the town plan] notes that in the plan were certain inadvertent "errors of alignment and misinterpretations," which today he views as fortuitous. "We've found that it added vitality. Now we're less concerned about perfection. Urbanism thrives on a certain amount of irregularity."

The regularity can be jarring, but not all the patterns were dictated by the plan:

> Concoctions have become more elaborate, even excessive, as if a dozen or more Victorian ship captains had landed there at once. "The code," says Duany, "does not actually generate cute Victorian houses. That just happens to be the taste of America today." The code does call for a variety of housing types, ranging from the Charleston side-yard house to the antebellum manse, and mandates what goes where...

Certainly the deliberate patterning makes Seaside a curiosity, but what makes it of artistic interest is the unexpected juxtaposition:

> That is not to say that Seaside has become a neatly patterned patchwork quilt. It still retains small elements of surprise and serendipity.

The Seaside example primarily illustrates how prescriptive theories are reinterpreted in a changing context. Rules for "how to build a Victorian house" are adapted to the Florida climate in a planned community. This goes beyond claiming that plans must be modified in action. One view is that plans must be modified because the world is messy, so our ideals can't be realized—as if the forces of darkness work against our rational desires. The rubric, "reducing theory to practice" suggests not merely an application or change in form, but a *loss* of some kind.

But in practice, standardized methods and procedures are as much *a problem* as a resource. The plan calls for certain kinds of wind or rain protection, but the available wood is not the oak of the Carolinas, only weaker pine. Certainly, we turn to the plans to know what to do, but as often we are turning elsewhere to decide what to do *about* the plans. Design rules and policies create problems; they are the origin not only of guidance and control, but of discoordination and conflict. Generally, rules only strictly fit human behavior when we view a community over a short time period, ahistorically.

Furthermore, the Southern towns we see today, on which Seaside is patterned, weren't generated by single coherent plans, dictating all homeowners' choices—just as Seaside today isn't rotely generated from the code. So the pattern descriptions and code are abstractions, lying between past practice and what Seaside will become, neither a description of the past in detail, nor literally what will be built.

Variations are produced by errors in the plan (e.g., irregularities in how the code is applied to the produce the more detailed plans of blocks and streets), misinterpretations, and serendipitous juxtapositions. Neighboring builders on the street make independent decisions whose effect in combination is often harmonious. The pattern of six honeymoon cottages on the beach is intentional; the preponderance of sea captain Victorians is a reflection of taste; and some of the irregularity reflects personal wealth and different preferences for using a home. Patterns we perceive in physical juxtapositions are emergent effects that we as observers experience, describe, and explain. Without some freedom—choices not dictated by the central control of the code—the effect would seem artificial precisely because it was too regular, and hence predictable. Openness to negotiation will vary; some restrictions (e.g., number of floors and building materials) are relatively constrained.

Understanding the nature of expertise requires understanding the negotiation process in the context of the emerging practice—not as an appeal to literal meanings and codes, but a dynamic reexamination of what's been built so far. What patterns are developing? How do the patterns relate to previous interpretations and developing understanding of what we are trying to accomplish? That is, expertise is as much the participation within a community of other designers, an *inductive* process of constructing new perspectives, as a deductive process of applying previously codified rules and theoretical explanations. In the next section, I consider in more detail how human knowledge in using plans is different from descriptive models of deliberation and learning. Human action is not, as the descriptive modeling view suggests, an artful combination of *either* following plans *or* situated action; rather, attentively following a plan involves *reconceiving* what it means during activity itself, and *that* is situated action.

(2.1). The Future-Orientation of Prescriptive Models

The interplay by which practice (what people do) and theory (descriptions of behavior and the world) shape each other is *dialectic*. What people do and produce is not dictated by the theory. The value of a theory, by this view, is not simply in how it corresponds to past practice, but in how well it serves to guide future practice.

A simple view of science is that scientists formulate experiences and observations about the world in scientific models. The models are valued because they describe the observed patterns and predict future phenomena in detail. Furthermore, models are valued when they have engineering value; they enable us to build buildings that withstand a hurricane and tell us how far back and high off the sand to build the houses of Seaside. In this respect, the models of Seaside homes are intended to accurately describe the past, but exist for their value in predicting a harmonious effect in the new community.

Christopher Alexander (1977) conceived of architectural design in just this way (summarized by Figure 1). On the one hand, we have artifacts and activities in the world. We perceive the world and sense similarities, in our everyday process of making sense and acting in our world (Schön, 1979). In our practice as designers (architects, city planners, robot builders), we represent our experiences of similarity descriptively, using classifications, grammars, and rules to represent patterns. Examples include a botanical classification, a bug library in a student modeling program, and Alexander's "pattern language" of hundreds of design configurations found in homes and communities that people find satisfying. Next, as scientists we seek to understand why these patterns exist, so we can control or create them deliberately or predict what will happen next. We create causal stories about how the patterns developed, indicating how properties of objects influence each other over time. As engineers, we then use our theories to build and manipulate things in the world.

In this way, models of architecture (and social practice in general) have predictive and explanatory value: Proceeding deductively (from the right side of Figure 1), theories predict that certain patterns won't occur. Or perhaps, the incidences of a type will be rare. For example, the theories behind Alexander's architectural pattern language suggest that given a choice people will put their bedrooms on the east side of a house, and not in the basement.

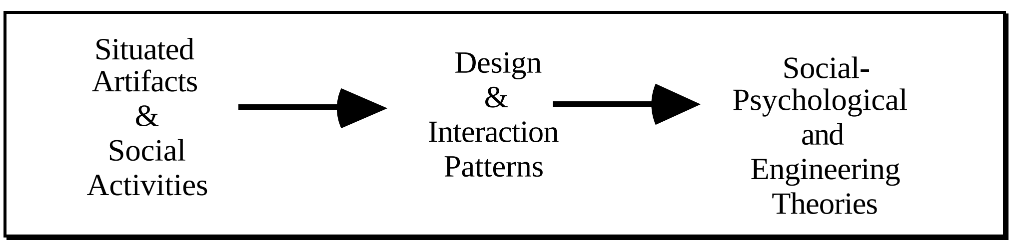

Figure 1. Relation of the experienced world, pattern descriptions, and theories.

This view of how patterns are discovered and explained fits Dewey's (1902, 1939) argument that representations are tools for inquiry. Dewey emphasized that such representations may be external (charts, diagrams, written policies) or internal imaginations (visualizations, silent speaking). The notion of representational accuracy is future-directed, in predicting success in making something. In traditional science, this "making" is experimental, in a laboratory. In business, procedures predict organizational success in efficiency and competitive effectiveness. The purpose of description is forward-looking, an orientation of control and/or change. The purpose of theorizing isn't accurate description of the past, per se, but to be knowledgeable of and in the future.[4]

Representational artifacts play a curious role in changing human behavior: On the one hand, the reflective process of observing, describing, and explaining promotes change by enabling invention of alternative designs. But these ways of seeing, talking, and organizing can become conservative forces, tending to rigidify practice. Understanding social change, particularly how to promote change in professions and business is a fundamental problem that is trivialized by the view that human behavior is driven by descriptions of fact and theory alone (or equivalently, that emotions and social relations are an unfortunate complication).

To understand change, we need to understand stability. This entails understanding the nature of interpretation by which theories are comprehended and used to guide activity. Agreement isn't reached by just sharing facts, but by sharing ways of coordinating different interests—a choreography, not a canon. Coherence and regularity are phenomena of a community of practice—a group of people with a shared language, tools, and ways of working together (Wenger, 1990). Theories (codes, plans, rules) are developed and interpreted within the ongoing development of values, orientations, and habits of the group.

Relating Figure 1 to the Seaside example, Southerners didn't literally apply a "code for building a Southern town" in their past activity. Instead, the descriptions created today are abstractions and idealized reorderings that tell a rationalized story about how building occurs. In contrast, the past activity itself was, to paraphrase Schön, "improvisation-in-

---

[4]Minsky's initial idea of "frames" considered them as schemas for directing future action. But such descriptions, as in the cognitive models that were based on this idea, were assumed to be mostly subconsciously manipulated. The situated cognition view is that descriptions and mental models are consciously generated and interpreted. By hypothesis, other forms of internal representing in the brain (e.g., the signal processing of primary vision) are different in kind—they are perceptual-conceptual "couplings" not inferences (e.g., see Edelman, 1992; Freeman, 1991; Thelen and Smith, 1994; Clancey, in press). Put another way, when descriptive models involve interpreting what rules and frames mean, they are modeling conscious human behavior. When they suppose that such descriptions are just "symbols" mechanically related, they are modeling subconscious neural processes. Situated cognition claims that these two forms of interpretation have dramatically different capabilities for reorienting behavior. As Bartlett emphasized, conscious story telling and interpretation provide a means of "turning around on our own schemata." Recently Polk and Newell (1995) have modeled how text manipulation in reading occurs consciously as a behavior; this begins to distinguish between texts, descriptive models, and neural processes.

action." This is depicted in Figure 2. Rules neither strictly describe the past nor control the future. Creating and interpreting (standards, exceptions, repairs) occur within activity. Vocabularies and rules fit human behavior only when viewed in narrow contexts, ahistorically. Similarly, now that people have a representation in hand, a plan for Seaside, they won't rotely apply it in the future. Within practical limits, they *improvise interpretations* that suit their activity as it develops, for example, changing the paths to suit barefoot walks to the beach.

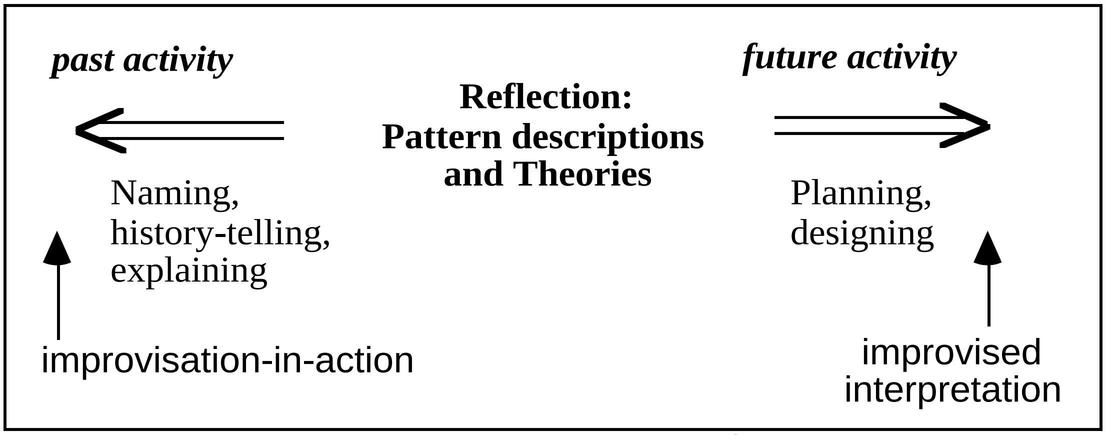

Figure 2. Descriptions lie between performances.

Describing the world and describing behavior occur in reflection, as actions that will at some level be automatic or immediate. Some interpretations of situated action miss this point, viewing action as *either* improvised *or* planned:

> The shooting-the-rapids example illustrates one important way in which experts make use of plans. The plan normally consists of an approximate path through rapids, as broad as feasible, taking advantage of the main currents and avoiding the obvious obstacles.... The most important property of such a plan is that it minimizes the number of occasions when an emergency calling for situated action will arise. (Vera and Simon, 1993, p. 41)

But situated action is not reverting to something more primitive or out of control. Rather, a dynamic adaptation is always generalizing our perceptions, our conceptions, and coordinations as we act. This reconceptualization occurs moment by moment and is a necessary part of abandoning a plan and looking more carefully to recategorize the situation: Rafting down a river, I might be reflecting on the rapids that just narrowly overturned my boat, telling myself that the water is higher than I expected, realizing that the deep holes have shifted from last season's run. I decide that I will have to portage around the next bend. A descriptive account abstracts my behavior and sees only smooth, "deliberate" execution or highly-reactive adjustments. Situated cognition claims that we are *always automatically adjusting* even as we follow a plan. That is, the relation is *both-and*: We are always recategorizing circumstances, even as we appear to proceed in lock-step with our predescribed actions. The claim is that descriptive cognitive models do not work this way, but the brain does. Descriptive models of "opportunistic planning" suppose a mixture of bottom-up and top-down processing, but this is again manipulating descriptions within a given ontology—a fixed language of objects, attributes, and events by which a model is indexed and matched against situation descriptions. In the brain, recoordination is dynamic, involving a mixture of perceptual recategorization and reconceptualization, sometimes on many levels at once (Varela, 1995; Edelman, 1992).

This example should make clear that situated action is not action without internal representations (leading to the claims that the baby has been thrown out with the bathwater)—indeed, the claim is starkly different! The claim is about the nature of the *representing process*. Our internal representing is *coupled* such that perception, movement, and conceptualization are changing with respect to each other moment-by-moment in a way that descriptive "perceive-deliberate-act" models do not capture. But identifying knowledge with text and all knowledgeable behavior with deliberation, a descriptive modeling theorist finds a dynamic model to be incomprehensible, a violation of the traditional engineering approach :

> Apparently one cannot talk about the brakes of a car without mentioning the engine, the steering wheel, the tires, the ignition system and the seat belts in one breadth. The need for "wholeness" denies the possibility of characterizing and explaining the mechanisms that govern components. (Vera and Simon, 1993, p. 124).

Indeed, the point is that the brain is not like a car in its linear, causal coupling of fixed entities, but operates by a kind of mechanism engineers have yet to replicate in an artificial device (Freeman, 1991). Structures in the brain form during action itself (Merzenich et al., 1983), like a car whose engine's parts and causal linkages change under the load of a steeper hill or colder wind. Such mechanisms are "situated" because the repertoire of actions becomes organized by what is perceived.

The rapids example is obviously on a different time scale than building a town. But the relation between conceptualization and planful descriptions I have characterized occurs throughout human life. In general, human conceptualization is far more flexible than the stored knowledge representations of a cognitive model suggest. Just as the river forces adjustments at nearly every moment, building the town involves flexibility and negotiation that expert system models of reasoning do not capture. Individual decisions and behaviors are in general shaped by an *a priori* mixture of personal and social descriptions, plans, and codes. The pace of surprise in the town is different from running rapids, but local adaptations are occurring when new buildings are proposed and blueprints are interpreted during construction. Knowledge base characterizations of expertise, even with respect to these relatively static diagrams, are impoverished precisely because they attempt to equate practice—what experts actually do in real settings over time—to a code, the rules and scripts of the knowledge base. A more appropriate understanding of the relation views knowledge base representations—insofar as they are part of the discourse of the expert or expressed in text books, policies, etc.—as something the expert refers to in his or her own practice, as a guide, a resource, something that must be interpreted.

Human activity, whether one is rafting down a river or managing a construction site, is broadly pre-conceived and usually pre-described in plans and schedules (even the rafting company uses computer reservations). But the details are always improvised (even when you are pretending to be a robot). At some level, all "actions" happen in a coordinated way without a preceding description of how they will appear. The grainsize of prior description depends on the time available, prior experience, and your intentions (which are also variably pre-described depending on circumstances).

This analysis raises questions about packaging theories and policies in a computer system and delivering it to some community as an expert system or instructional tool. Knowledge engineers could build an expert system that embodied the Seaside plan, but would such a tool address the practice of collaboration? Would it relate to the participants' problems in negotiating the points of view of different expertise? The "capture and disseminate" view of "reproducing knowledge" (cf. Hayes-Roth, cited in Section 1.1) does produce useful tools. But situated cognition suggests knowledge engineering hasn't considered the conceptual problems people have in reconciling different world views, which is what forces reconceptualization in conversations between the carpenter, the home owner, the town council, the county inspector, etc. That is, the original expert system approach ignores the fact that *there are many experts* and they would benefit from tools for working together. Ironically, our view in 1975 of "many experts" when building Mycin was that *physicians* might disagree about how to build the knowledge base, not that there were different professional *roles* to reconcile. Problems are "ill-structured" not just because there are many constraints and too much information, but also because *different participants are playing different roles and claiming different sets of "facts."*

To understand how situated cognition suggests new ways of using expert system technology in tools for collaborative work, we need to explore further what people are conceptualizing, which produces these different views of the world, and why these conceptualizations cannot be replaced by a program constructed exclusively from descriptions. In the next section, I contrast *how people conceive of activities* with the task analysis of knowledge engineering, which equates intentions with goals, context with data, and problems with symbolic puzzles.

(3). ACTIVITIES VS. TASKS

To understand the idea that *knowledge* is inherently social (as well as inherently neural), we must first understand that human *action* is inherently social.[5] The difficulty is understanding that "action" is meant in the broad sense of an "activity." The activity of being a construction site coordinator comprises many individual "tasks," such as "test hypothesis" or "combine conjunctive goals," as found in expert systems and cognitive models. In this section I will explicate how a person's conception of his or her activity is socially *oriented* and *shaped*. On this basis I will reformulate the meaning of "the social construction of knowledge" (Section 3.2), the relation of activities to tasks and goals in expert systems (Section 3.3), what constitutes a "problem" (Section 3.4), and the social nature of work (Section 3.5).

(3.1) The Conception of Activity

Our everyday way of talking about social activity places primacy on the individual and marks "social" as being a matter of choice or a *kind* of activity. For example, we socialize at a party; we engage in "social chat" before settling down to work; we may decide to join others for a drink after work in order to "be social." In common parlance, social activities are things that people do together: parties, meetings, tours. We may "opt out" of a social activity, and go our own way. By the cognitive "individualist" point of view, social activities par excellence are special occasions, such as weddings.

---

[5]I am trying to convey a "both-and" point view: Both social and neural. But "inherently" has different implications. Knowledge is inherently social in *content* and inherently "neural" in *form*. This doesn't mean that we won't be able to reproduce the functions of neurons by silicon or some other substrate. I am simply emphasizing that 1) there is of course an internal aspect, and 2) we need to reproduce the dynamic characteristics of neurons relative to categorization and coordination. One could say "inherently psychological," but this wouldn't convey that it is *how the neural mechanism works* that we need to understand and replicate if we are to create a machine capable of situated action.

Social activities in our everyday experience may also be things we do reluctantly, such as attending a meeting. In business settings, meetings appear to take time away from “the real work,” which is individually-directed. Work is having your nose to the grindstone; social activity is having fun talking to people about non-serious things.

In each of these examples, activity is viewed as being social or not: Individual activity is when I am alone, social activity is when I am interacting with other people. This is essentially the biological, either-or view of “activity”—a state of alertness, of being awake doing something. But the social scientist, in describing human *activities as social*, is not referring to kinds of activities per se. Rather what we do, the tools and materials we use, and how we conceive of what we are doing, are *culturally* constructed. Even though an individual may be alone, as in reading a book, there is always some larger social activity in which he or she is engaged. Indeed, as we will see, descriptive accounts provide an inadequate view of subjectivity and, hence, the attitude of “individualism,” because they do not emphasize the inherently social aspect of identity.

For example, suppose that I am in a hotel room, reading a journal article. The cognitive perspective puts on blinders and defines my task as comprehending text. From the social perspective, I am *on a business trip*, and I have thirty minutes before I must go by car to work with my colleagues at Nynex down the road. The information processing perspective sees only the symbols on the page and my reasoning about the author’s argument. The social scientist asks, “Why are you sitting in that chair in a hotel room? Why aren’t you at home?” That is, to the social scientist my activity is not merely reading—I am also on a business trip, working for IRL at Nynex in White Plains, NY.

We are always engaged in social activity, which is to say that our activity, as human beings, is always shaped, constrained, and given meaning by our *ongoing* interactions within a business, family, and community. Sitting in my hotel room, I am still nevertheless on a business trip. This ongoing activity shapes how I spend my time, how I dress, and what I think about.

Why has the social nature of all activity been so misunderstood in common parlance and cognitive science? Social activity is probably viewed in a commonsense way as being *opposed to* work because of the tension we feel when *we are engaged in an activity* and because of the time or place are obligated (by our commitments) to adopt another persona. In our culture, we engage in multiple, contrasting activities every day, which split our attention and loyalty. We leave the family in the morning in order to “go to work.” We end our pleasantries at the start of the meeting in order to “get down to business.” We end a discussion with a colleague in our office in order to “do something meaningful.” Ending one form of engagement, we experience a tension and conflict in the change of theme, which is marked conventionally as a shift from "socializing" to "working." Both activities are social constructions, but the limiting of our options leads us to view the more narrowly-defined style as being "work" and the freedom we have left behind as "being social."

Although a social scientist may cringe to have it put this way, a cognitive scientist might begin by thinking of activities as being forms of *subjugation*, or more neutrally, *constraint*. Human activities are always constrained by cultural norms. During a movie in a theater, I cannot yell out to a friend across the room whether he would get me some popcorn. I cannot in general talk very loudly to the person next to me. If my back hurts, I cannot stand by my seat (and block the view). Sitting in a hotel room on a business trip, I cannot decide to take a bath or go for a walk if I am expected to be over at Nynex in 10 minutes.

The standard examples of activities suggest a form of passivity: being on a tour in a museum, attending a religious service, listening quietly to a lecture. The trick in understanding activities is realizing that such *following* or *adherence to a norm* is inherent in all activities. Day by day, you make choices, identifying with a group of people and participating in social practices that limit (and hence give meaning to) your behavior.

Again, the individualist will object: But what about *after work*, when I am sitting at home alone in my easy chair, reading *Atlas Shrugged*? I'm in control of my time, I do what I want to do. And what about on the weekend, when I am gardening or taking a walk by myself. Surely these are not social!

But how we conceive of free time, the very notion of "after work," "weekend" and "gardening" are socially constructed. Again, "constructed" here means that what people tend to do occurs within a historically-developed and defined set of alternatives, that the tools and materials they use develop within a culture of manufacture and methods for use, and that how they *conceive* of this time is with respect to cultural norms.[6] Weekends in

---

[6]Again, saying that tools are cultural does not mean that they are not physical. I am trying to convey how a "both-and" causal process operates: Tools are both culturally conceived and physical at the same time. The either-or perspective suggests that we can "understand" a phenomenon from a single point of view at a time. This is true for the operation and use of mechanical devices; refer again to Vera and Simon's characterization of a car. The linear perspective prefers to say "*partly* cultural" and "*totally* physical." This attempt to divide up the causal influences is adequate for some purposes. But understanding historical change requires understanding how cultural, biological, and physical interactions interpenetrate. Gould's (1987) critique of linear views of evolution is based on the same notion of dialectic causality. The dialectic

Bali are not the same as weekends in Detroit. Our very understanding of "time alone" is co-determined with respect to our understanding of "being at work," "being on a business trip," and "being at a party." Although we are not literally confined in the same way we might be while sitting at a contract negotiation meeting, on the weekend we are nevertheless engaging and acting within an understanding of the realm of possible actions that our culture makes available. Sitting at home on Sunday morning reading the NY Times over coffee is a culturally-constructed event. The meaning of any activity depends on its context, which includes how we conceptually contrast one range of activities with another.

Of course, someone can "opt out" and go live in the Na Pali Headlands of Kauai out in the jungle. But if you have left Detroit and your desk job, you are now engaged in the activity of "opting out." You are "going back to nature," "seeking a simple life." Although free to romp around naked, you cannot escape the historical social reality that defines your activity in terms of making one choice and not another, of participating or not. The meaning of your activity of being in Kauai will be co-determined by your understanding of the activity of working in Detroit. Sitting in a chair in a lecture hall, we may be bored by a lecture and dream about the beaches of Kauai. But even then, we are still in the activity of "attending a lecture," though our activity might be best described as "not paying attention to the lecture."[7]

An activity is therefore not just something we do, but a *manner* of interacting. Viewing activities as a *form of engagement* emphasizes that the conception of activity constitutes a means of coordinating action, a manner of being engaged with other people and things in the environment, what we call a *choreography*. Every human actor is in some state of participation within a society, a business, a community. My activity within the Institute for Research on Learning is "working at home," an acceptable form of engagement, a way of participating in IRL's business. Within Portola Valley, where I am working, I am not participating in the activities of the schools or the town council. Like most people I am in the activity of "going about my own business." Like most of my neighbors, I don't know the names of the people who live around me; we are all in the activity of "minding our own business." Even *not* talking to my neighbor is a kind of choreography.

The idea of activity has been appropriately characterized in cognitive science as an *intentional state*, a mode of being. The social perspective emphasizes time, rhythm, and

---

view claims that *understanding the interactions between physical and social constraints* is essential for understanding what has evolved. This is perhaps most clear in the evolution of language (Donald, 1991).

[7]This example comes from Frake (1977).

place (Hall, 1976): An activity is a *process framework*, an encompassing fabric of ways of interacting that shapes what people do. Activities tend to have beginnings and ends; activities have the form “While I am doing this, I will only do these things and not others.” While I am working at home, I will stay in my office and write; I will not call my spouse and chat, I will not read my electronic mail, I will not make phone calls, I will not stop in mid-morning and go out for a walk (but I will swim at noon).

People understand *interruptions*, “*being on task,*” and *satisfaction* with respect to activities. For example, contrast your experience when interrupted by different people when you are reading: a stranger in a train, a colleague in your office, your spouse when you’re reading the paper in the morning. Your conceptual coordination of the interruption is shaped not just by your interest in what you are reading (and why you are reading it) but by the activity in which you are engaged. Activities provide the background for *constructing situations*; they make locations into events. Different activities allow me to walk casually down the middle of the street on the afternoon of Palo Alto’s Centennial Saturday, but run for my life when crossing on any other day.

With these considerations in mind, I return to the descriptive view of goals and knowledge, specifically as formulated in expert systems. How is the conception of activity related to what a knowledge engineer represents? What other kinds of tools would be useful?

(3.2). The Relation of Activities to Tasks and Goals in Expert Systems

Activities are broadly intentional, but not confined by the kinds of goals that define an expert system’s operation. People conceive of goals and articulate them *within* activities. From the descriptive modeling perspective, the encompassing and composed nature of activities is often missed because the modeler starts by choosing *one activity* to model and the point of view of one role fulfilling one predefined goal within this activity. With such a design approach, human action appears to be a relation between defined goals, data, and decisions. For example, in modeling medical diagnosis (Buchanan and Shortliffe, 1984), we chose the physician’s activity of examining a patient. We even viewed this narrowly, focusing on the interview of the patient, diagnosis, and treatment recommendation, ignoring the physical exam. But the physician is also in the activity of “working at the outpatient clinic.” We ignored the context of patients coming and going, nurses collecting the vital signs, nurses administering immunizations, parents asking questions about siblings or a spouse at home, etc. In designing medical expert systems like Mycin, we chose one activity and left out the life of the clinician. We ignored union meetings, discussions in the hallways about a lost chart, phone calls to specialists to get dosage recommendations, requests for the hospital to fax an x-ray, moonlighting in the Emergency Room. Indeed, when we viewed medical diagnosis as a *task* to be modeled, we ignored most of the *activity* of a health maintenance organization! Consequently, we developed a tool that neither fit into the physician's schedule, nor solved the everyday problems he encountered.

This is not to say that Mycin, if it had been placed in the clinic, would not have been useful. Rather, the thrust of this shift in perspective, broadening our point of view from goals to activities, partially explains *why Mycin was never used at all*, and second, reveals *opportunities for using the technology that we never considered*. Indeed, the design of Mycin reveals how our conception of activity shapes the design process. We viewed physicians as the center, knowledge as stored, knowledge acquisition as transfer, and the knowledge base as a body of universal truths (conditional on universal situation descriptions). Rather than developing tools for facilitating conversations, this *rationalist approach works to eliminate conversations* and replace human reasoning by automatic deductive programs (Winograd and Flores, 1986). In a similar analysis, Lincoln et al. (1993) suggest that medical practitioners need notational devices with spreadsheet-like operations that help manage the interactions of tentative inferences within volatile, uncertain situations. But if the expert performances of a physician are explained only in terms of knowledge stored in the head, tools for developing models interactively in a team are not considered.

To provide another example, Figure 3 represents the activities of a contractor who is sitting alone in his trailer on the Seaside site. The contractor has many identities: on his job that morning, he might be thinking about struggling with a county building code form which needs to be filed. In this activity, he is coordinating the contractors on the site, perhaps reconciling the work of the plumbers and electricians at a particular home. He wears the badge of the "Seaside Construction Co.," to which his decisions must also be accountable. But at the same time, he holds certain principles about how to be a supervisor, which he seeks to bring to his company to change their practices, and has a conception of how to advance his career by doing well in this company. Of course, we can include this worker's conception of how living in Florida influences his decisions about pursuing this career in the Florida Panhandle, rather than Miami, etc.

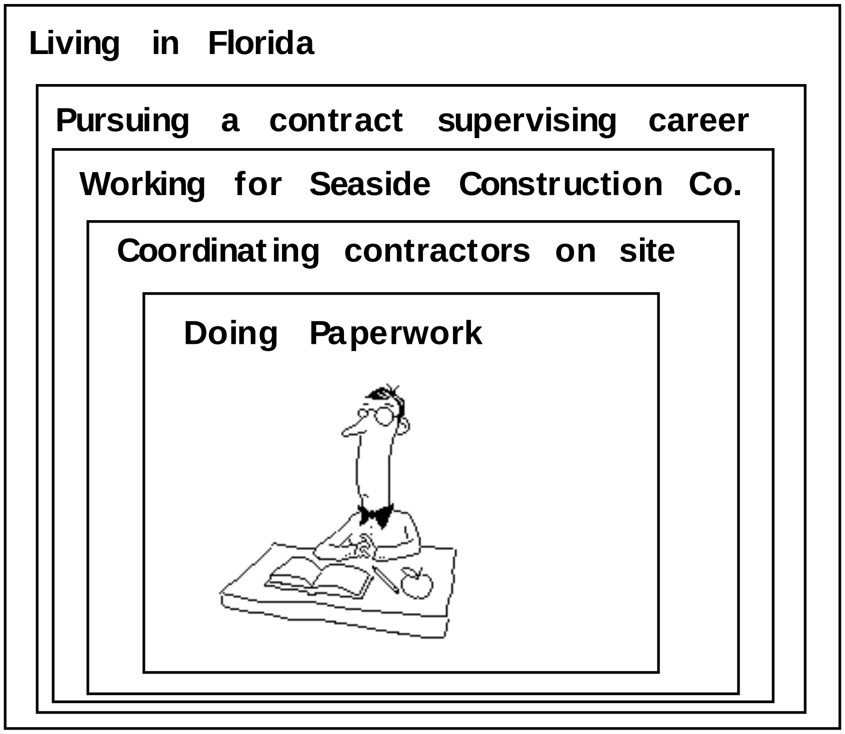

Figure 3. Ongoing activities of a Seaside contractor, “working alone.”

Although these activities are described as nested levels, other relatively disjoint identities and participation frameworks will not fit strictly into such a hierarchical ordering. For example, the worker is a sports fan, attends a certain church, supports a political party, volunteers for community service, etc. Any of these may or may not be relevant conceptions having a bearing on day-to-day contracting work. The important point is that this person has all of these identities and will experience a conflict if an activity suggests a way of coordinating his views, talk, and actions that is different from how he might behave in another context, which he perceives is also a relevant conception of what he is currently doing. For example, if the county requests an action that he believes is not in accord with the union’s principles, he will have a problem. One source of creativity lies in juggling multiple identities and carrying ideas from one community to another.

To recapitulate: All human action—deliberation, goal defining, theory application, information description, policy interpretation, planning—occurs within activities. Conception of activities is usually implicit, serving as the *background* against which problems arise and judgment is based. In Winograd and Flores’ (1986) analysis this background is the origin of our sense of trouble in a situation, which they called “breakdown.” Breakdown is not just a difficulty in interpreting text, a failure for a description to apply to a new situation (as Winograd and Flores emphasized), but more generally is *any conceptual discoordination* between perspectives within and between people, suggesting different ways of characterizing facts and evaluating judgments.

In contrast with this view of situated action, the idea of “rational action,” also called “cognitivism,” suggests that goals and prescriptive rules *control* actions. By this view, tasks are isolated things to do and context is just the given world of data to operate upon. That is, behavior is “conditional” on the facts. The contrary, “situated,” view is that defining goals, claiming what constitutes the facts, and following plans and policies all occur within nested activities. In this sense, all action is *situated* in the actors’ conceptions of what they are *supposed to be* doing, that is, norms, values, and roles—their identities. Articulation of goals, facts, and methods—that is, creation of descriptions and interpretations of representational artifacts—arises within this *conceptual* frame[8].

Activities are *the grounding of intentionality*—“what I am doing now” is defined with respect to my activities. Hence, my intention is not just to finish reading a journal article or to write a report, but “to make a contribution to the research project,” “to convince the client that I should continue to belong to this project,” “to convey what IRL means by participatory design.” Viewed narrowly, cognitive goals are *information* oriented—developing models and choosing actions. The goals of activities involve—conceptually and physically—reaffirming and developing forms of engagement, membership, and identity (Wenger, 1990; Lave and Wenger, 1991); they are *participation* oriented.

What I take to be information and see in a situation depends on my conception of my activity. When I walk into the medical clinic with the conception of a medical anthropologist, I see notes on paper and hear conversations that constitute data by how they are arranged in space and time. When I walk into the medical clinic as a cognitive scientist, I listen to the physician’s information requests and hear the names of symptoms and diseases. To understand the process by which people segment the world into objects and events, we need to understand that *perceiving is situated in activities.*

In these respects, activities are the context for all that we do. But by reducing activities to descriptions of goals, data, and actions in the descriptive approach (Section 1.1), knowledge engineering and cognitive modeling unwittingly reduced *conceptions* of context to descriptions. This research misunderstood the functional character of conceptions to *coordinate* what we perceive, our judgments, and how we interact.

---

[8]In this respect, the focus of the sociology of knowledge and social psychology on conceptual interactions and constructions is more inclusive than ecological psychology, which emphasizes perceptual interactions in physical niches (Turvey and Shaw, 1995).

Given the relation between knowledge and activities, we can better understand how knowledge is socially constructed.

(3.3). Social Construction of Knowledge

An individual's capacity to engage in an activity may be characterized as *knowledge*. Thus, "knowledge is socially constructed" means first, that *knowledge develops and has value within activity*, and second, *activities are socially constructed*. Descriptive modelers often interpreted the phrase "social construction of knowledge" in terms of their assumption that knowledge consists of statements about beliefs and theories. But activities are the primary construction: Ways of interacting and viewing the world, ways of structuring time and space, pre-date human language and descriptive theories. We can see this today in chimpanzees and other animals who conceive of social relations, build and maintain homesites, and distinguish between activities such as foodgathering and play—all without a descriptive language for modeling what they are doing[9]. Within a discourse, an activity mediated by language (Wertsch, 1991), statements about beliefs and theories provide a way of changing activities (Mills, 1940; Schön, 1979; cf. Figure 2). Again, situated action does not deny the role or importance of plans, but emphasizes that planning occurs within an *already conceptually-coordinated* activity.

To understand what "social construction of knowledge" means, you must first understand that *activities*, the choreographies of human action, *develop within ongoing activities*. Our capacity to plan what we will do, to design new methods and tools, and to formalize what we know, develops within and depends upon our pre-existing activities. For example, when caregivers at a health maintenance organization join to form a new outpatient clinic, they are already engaged in activities such as "being a technical clinical assistant at the County Hospital," "being a physician escaping from private practice," and "representing the union advocating greater responsibility for nurses." These pre-existing conceptions of interacting and these identities provided the context within which a new clinic was formed.

Knowledge, then, develops *within activities*. The knowledge required to accomplish goals, such as caring for a patient, is determined, in large part, by the scientific and health

[9]By a descriptive language, I mean a language like natural language or many AI knowledge representation and planning languages, with syntactically-distinguished subjects, actions, objects, and modal qualifiers, enabling historical, prescriptive, or explanatory statements to be made about objects in the world, their properties, and events. As AI research has demonstrated, such distinctions are necessary in order to *model* the world and behavior (Edelman, 1992).

care establishment to which the clinicians belong. The members of the local clinic help define what constitutes competence by their choices of whom they want to work with and how they talk about each other's work.

The idea that knowledge is a possession of an individual person is as limited as the idea that culture is going to the opera.[10] Culture is pervasive; we are participating in a culture and shaping it by everything we do (Hall, 1976). Knowledge is pervasive in all our capabilities to *participate* in our society; it is not merely beliefs and theories *describing* what we do.

The difficulty in understanding how knowledge is socially constructed partly stems from social scientists' lack of success in articulating the conceptual nature of knowledge. For descriptive modelers, social construction of knowledge is by default understood to mean the social construction of written facts about nature and equated with relativism (e.g., see Slezak, 1989; Lakoff, 1987). The practice of science is cultural (Gregory, 1988), but the effect is not so much on what someone sees when looking in a microscope and even less on how numbers are tallied to formulate an equation. Rather, social construction operates on and through concepts, activities, designs, and policies. Expertise is not just about scientific facts and laws, but about value-laden artifacts and conventions and how to coordinate them.

For example, the physics researchers at Stanford University's Linear Accelerator Center (SLAC) are constrained in their work by government funding, the politics of promoting international science competitively against other states such as Texas, the terrain that limits building around a major fault zone, etc. The "science" we see at SLAC in the next decades will not be a purely experimental effort, but will reflect the savvy of the Director in securing funding and building huge new devices within the context of engineering and political constraints. Again, the activities of the scientists at SLAC cannot be understood solely in terms of the scientific horizon of physics, but are framed by being a member of SLAC, being a Californian competing for federal dollars, and being an American scientist rushing in time to beat the efforts at CERN in Switzerland. Choices of what to explore next in the realm of physics are generated and evaluated in this context.

---

[10]Again, knowledge as a functional capacity is historically-personal (subjective), and knowledge as a coordination process is neurological. The problem with the "possession" metaphor is not that it denies the social character, but too quickly slips into viewing knowledge as a static thing and the cultural aspect as just a coloring or flavoring of objective truths. For example, it is common in AI to define knowledge as "true belief." By this view, culture is what we do on Saturday evenings or explains why some cooking pots are made from iron and others from clay.

Similarly, the medical profession is not just deductive application of physiological science. Drug dosages, diets, exercise programs, etc. must all be designed with respect to the practices of people, which constrain time, memory, and will to carry through procedures (Feltovich et al., 1992).

The social construction of knowledge includes the construction of written theories and facts about the world. But, as these examples illustrate, the force is not on what scientists find out about nature, but which facts are relevant and what designs are valued within the constraints of engineering and social affairs. In these choices, we are always involved in *constructing communities of practice*: SLAC, a school district, a seaside town, a subdiscipline of AI. Knowledge in this realm, particularly the professional knowledge called "expertise," concerns how to make interpretations of policy ("judgments") that generate successful, harmonious designs.

According to Schön (1987), the difficulty of formulating design knowledge lies in a combination of the rapidly-changing character of activities and the difficulty of codifying rules about "the ill-defined melange of topographical, financial, economic, environmental, and political factors" (p. 4). This is summed up by Kyle: "We know how to teach people how to build ships but not how to figure out what ships to build" (quoted by Schön, 1987, p. 11). In effect, it is difficult to formulate *a priori* what kinds of problems will arise and on what basis they should be resolved. The next section elaborates this distinction between problems in practice and the formal problem solving of "professionalism" studied in cognitive science of the 1970s and 80s.

(3.4). Problematic Situations, Not Puzzles

The scientific view of problem solving is that one starts with certain data, a goal, and certain theories about how goals and facts are related. But in practice the problem is often which *kinds* of facts are relevant and how to justify action within a matrix of conflicting regulations and competing judgments. Following Dewey, Schön emphasizes that a problem must be *constructed* to fit the methods and theories of a practice:

> When a practitioner sets a problem, he chooses and names the things he will notice. In his road-building situation, the civil engineer may see drainage, soil stability, and ease of maintenance; he may not see the differential effects of the road on the economies of the towns that lie along its route. Through complementary acts of naming and framing, the practitioner selects things for attention and organizes them, guided by an appreciation of the situation that gives it coherence and sets a direction for action. So problem setting is an

> ontological process—in Nelson Goodman's (1978) memorable word, a form of worldmaking. (Schön, 1987, p. 4)

Casting a situation in some language defines a "space" for reasoning about alternative designs, diagnoses, and plans. Schön calls this process "problem framing." In contrast with the classical AI view of problem solving (Newell and Simon, 1972), Schön argues that framing a problem is not a matter of searching and filtering through given facts, per se, but of *creating* information (cf. von Foerster's 1970 critique of information processing theories of memory). The descriptive view maintains, for the most part, that the world is encountered as objects with properties. The situated cognition view is that we segment the world perceptually, within the rubrics of our activities. In our interpretative process of qualifying and weighing experiences, we participate in such a way that our process of seeing and naming has created a world, the conceptual space in which we coordinate our thought and action. "Creating information" means that our *interpretations* claim which facts are meaningful to the problem at hand; by this we define the problem (reifying our conceptions into what Newell and Simon call a "problem space" description).

Schön goes on to describe how differences in opinion are not reducible to arguments about facts. Instead, different conceptualizations lead professionals with different backgrounds to perceive and name *different sets of facts* as being relevant. These differences derive from conceptualization of activities—not more names and facts—that underlie each professional's attending, valuing, and sense-making:

> Depending on our disciplinary backgrounds, organizational roles, past histories, interests, and political/economic perspectives, we frame problematic situations in different ways. A nutritionist, for example may convert a vague worry about malnourishment among children in developing countries into the problem of selecting an optimal diet. But agronomists may frame the problem in terms of food production; epidemiologist may frame it in terms of diseases that increase the need for nutrients or prevent their absorption; demographers tend to see it in terms of a rate of population growth that has outstripped agricultural activity; engineers, in terms of inadequate food storage and distribution; economists, in terms of insufficient purchasing power or the inequitable distribution of land or wealth. In the field of malnourishment, professional identities and political/economic perspectives determine how people see a problematic situation, and debates about malnourishment revolve around the construction of a problem to be solved. Debates involve conflicting frames, not easily resolvable—if resolvable at all—by appeal to data. Those who hold conflicting frames pay attention to different facts and make different sense of the facts they

> notice. It is not by technical problem solving that we convert problematic situations to well-formed problems; rather, it is through naming and framing that technical problem solving becomes possible. (Schön, 1987, p. 4-5)

In terms of the "practice <-> pattern <-> theory" framework (Figure 1), Schön is characterizing how a practitioner comes to perceive a difference, articulate a pattern, and frame a problem in technical terms. In contrast, the process of structuring the world, of perceiving and naming order, has been characterized in AI research, not so much as a conceptual capability of people, but as inherent in types of situations. The laboratory perspective of giving problems to a subject suggested that there were two kinds of problems: *well-structured problems*, which could be mapped directly into a known problem-solving language and procedure, and *ill-structured problems*, which were experienced as confusing in some way, requiring restatement and often more information before they could be resolved. Observing that this classification was relative to the subject's knowledge, Simon (1973) concluded that all problems are potentially ill-structured and there is nothing fundamentally different between playing chess and designing a ship—both can be described by categorizing states and operators in the General Problem Solver. Design problems and their solutions are thereby reduced to puzzles and mathematical operations.

Situated cognition argues that problem solving is a particular kind of activity occurring within other ongoing activities, which are the context that produces the troublesome situation and provides the framing values and goals for justifying our action. By this process, judgments are made objective—our justifications relate to the principles, methods, and practices of the communities to which we belong (Berger and Luckman, 1966).

In contrast, cryptarithmetic, theorem proving, chess, and other "problems" studied by Newell and Simon (1972) are merely puzzles existing in a mathematical world of well-defined rules of play. Carrying over these ideas to medical expert systems in the early 1970s, knowledge engineers viewed medical practice as the rote application of facts and causal relations between organisms and therapies. In viewing every patient encounter as a "problem," designers of consultation systems never understood either the conceptual nature of *trouble* as a discoordination, or how it was resolved in everyday situations.

In reducing medical knowledge to descriptive, scientific models of disease, knowledge engineers, as well as cognitive psychologists adopting the expertise perspective (Chi et al., 1988) lumped together written scientific facts, conceptions, experience with therapeutic designs (regimens). For the most part, this viewpoint ignored the political factors underlying the distribution of decision-making between the medical subspecialties and between nurses and physicians. For example, it is common to find in a module of caregivers three MDs, one physician's assistant, and four nurses of different varieties. Whose professional knowledge is coded in Mycin? From this perspective, Mycin models the knowledge of an infectious-disease specialist in a large, tertiary care hospital. That is, Mycin doesn't contain "medical knowledge," per se, but was intended to *model one role in a certain activity*. This role and activity were rarely reflected on or deliberately pursued in the design because the content of the knowledge base was viewed as "medical knowledge," *universal truths* about medicine. "Case-based reasoning" and "roles" are well-known ideas in AI research, but experience from the past and the choreography of roles tend to be reduced to *more descriptions*, which get thrown into the pot of scientific facts and rules.

The flattening of knowledge into facts about the world can be seen clearly in how the term "frame" was used in anthropology to refer specifically to activities (Frake, 1977), while in AI's knowledge representation research it meant any "unit" of knowledge, from a description of a political context to a description of a chair. By this "atomization" of knowledge into uniform pieces, the nesting of concepts and communities by which the world is segmented, facts are interpreted as relevant, and designs are invented, is represented as a hierarchy of graphs (called "contexts"). In this view, every surrounding "frame" becomes a "context," and the distinctions among scientific data, practical design constraints, and interpersonal choreography are lost. Any claim of "openness to interpretation" and "creating of information" suggested arbitrariness and scientific relativism, so the difference between nature and culture became muddled. "Cultural knowledge" becomes just more facts in long-term memory (Lave, 1988, p. 89).

To summarize, human knowledge comprises much more than written scientific facts and theories. Problems arise not in selecting facts, but in conceptualizing how we should view the activity we are currently engaged within: What *differences* (Bateson, 1972)—kinds of facts—make a difference? What perspective (economic, physical, political, medical) should be adopted? How are conflicting judgments to be reconciled in the practice of our conversations, design procedures, and regulating policies? Who should be invited to participate in this discourse and what rights should they be accorded? How will we reach a decision? On what time frame? How will we answer to the competing viewpoints, which suggest that our designs are unproven, that they have failed in the past, that they are too costly? Expertise consists of the ability to make value judgments for framing problems, which in turn establishes a reified "problem space," in which the technical methods of science and engineering may proceed (Schön, 1987).

Knowledge, context, and trouble are conceptual. Through the thought process of conceptualization—which is still poorly understood—we articulate problem descriptions, facts, and rules for guiding our action. Dewey called this process of thinking, describing, and manipulating descriptions "inquiry." He argued in 1939 that Bertrand Russell's rationalist view (which became the foundation of descriptive modeling) fundamentally confused the origin and role of statements ("propositions") in problem solving:

> The exclusive devotion of Mr. Russell to discourse is manifested in his assumption that *propositions* are the subject-matter of inquiry, a view assumed so unconsciously that it is taken for granted that Peirce and I likewise assume it. But according to our view—and according to any thoroughgoing empiricist—*things and events* are the material and objects of inquiry, and propositions are *means* in inquiry, so that as conclusions of a given inquiry they become means of carrying on further inquiries. Like other means they are modified and improved in the course of use. (p. 573)

Thus, Dewey saw statements as being out in the world of our conscious experience, lying between performances (Figure 2). Once we understand that conceptual coordinations in different modalities—including for example rhythm, imagery, accent, and gestures—are not equivalent to *descriptions* of the world and our behavior, we can better understand the activity of describing and comprehending in recoordinating our activity. But if we equate conceptualizations and descriptions, we will have little idea *how problems arise* and what resources we draw upon for generating and improving our descriptions of the situation, how we will decide, and what we will do.

As we have seen, the view of knowledge as true descriptions does not adequately explain what is problematic when professionals from different disciplines attempt to work together. Correspondingly, by the descriptive view, an individual is just a repository and applier of knowledge. By the social view, the individual's contribution is more dynamic and unique. The relationship of the individual to the group is different than is suggested by the "cognitive tasks" view of work.

(3.5) "Social Activity" vs. Individual Work

The reduction of problem solving to rote rule manipulation and parsing of text has distorted not only our view of expertise and context, but our understanding of how work actually gets done. Having decomposed knowledge in terms of technical calculi, knowledge engineers and cognitive modelers are left with a residue of "value judgments," which appear as immature forms of real (scientific) knowledge, and a residue of "social relations," which appear as messy, but necessary considerations for getting information and communicating decisions. The resulting dichotomization makes the core of knowledge individual cognition and the remainder "social factors." The core is "hard science" and the periphery (so-constrained) is "soft."

The dichotomization of individual and society is reflected in the Cold War drama between the forces of democracy and communism. The theory of situated action suggests that our identity as supporters of individual rights may have inhibited our scientific understanding that knowledge and work are inherently social (Hall, 1976). It is perhaps not a coincidence that Soviet psychologists were deeply affected by the relation between the state and the individual and sought to understand how the influences interpenetrated. For example, Lev Vygotsky emphasized that an individual's understanding develops within the pre-existing social fabric of activities—the conceptual segmentation and ordering of time, place, and events, which develops into and is manifest in habits, norms, means of labor, and roles. What is "socially shared" is not just language, tools, and expressed beliefs, but conceptual ways of choreographing action, by which descriptions and artifacts develop and are given meaning.

As individuals we participate in the process of constructing what will be the norm, how performance will be evaluated, what ideas will be valued, and what tools will be used. The experience of being in an activity is not usually that of subjugation, as I first introduced the notion, but what people usually call "being constructive." Norms of the group allow redirecting its path: Activities may include means of communication and negotiation by which individual ideas and preferences are heard and incorporated. Conflicts are foremost differing conceptions of activities in which *individuals* believe themselves to be engaged. If the permissible styles permit a tradeoff, a complainer may compromise; the activities of the group, or what constitutes a norm or acceptable variation, will change.

When work is identified with technical knowledge, interpersonal capability and knowledge about other people and their abilities are viewed as non-essential or simply nice add-ons for getting the job done. In this way, the view that knowledge is objective and technical obscures that knowledge is *about what people do during their lives.* Rather than being about things and properties, knowledge is first and foremost about how to belong, how to interact, what to do productively with your time. That is, knowledge is inherently *personal*, as Polanyi (1958) put it, because it has a tacit dimension and develops within cultural commitments. Again, the conception of activities provides a way of understanding subjectivity without relegating it to "uncertain belief" (opinion) or "misconception" (theoretical error). In this respect, the theory of the social construction of knowledge, because it emphasizes values, roles, and interpersonal choreography, provides a better accounting for the nature of individual knowledge than cognitivism. How ironic that raising the banner of “the social” appears to deny the importance of the individual!

To bring this back to “knowledge-level” descriptions appearing in cognitive models (Newell, 1982), knowledge is therefore not just about a social *world*, but about *activities* and the *social-political-physical reality* in which activities occur. Knowledge is not just about tasks, but about forms of *participation*—the who, what, where, and why of behavior. Reality for a human being is not just the facts of nature, but an *identity* as a person. The social construction point of view tries to show that identity is not just a collection of technical procedures and scientific beliefs, as for example in a student model of a teaching program. A knowledge-level description would, in its entirety, not merely characterize the information-processing (or model-building) behavior of a person, but the timing, the locations, the roles, and the identities of that person’s life (Wenger, 1990). For example, in the everyday workplace, knowledge about what other people know is essential for assigning jobs, getting assistance, and developing teams. This social-psychological viewpoint does not replace “goal” or “task” by “activity,” but places behavior in a broader analytic context.

In this section, I presented a theoretical perspective contrasting knowledge about activities and task descriptions. I will conclude by considering briefly two areas in which this perspective can be applied—in pursuing the goals of AI to replicate human intelligence (Section 4) and in applying knowledge engineering to produce useful tools (Section 5).

(4). WHAT IS CONCEPTUALIZING IF NOT MANIPULATING STORED DESCRIPTIONS?

The foundational assumption that “representations” in the brain and “representations” on paper can be treated isomorphically (Vera and Simon, 1993), although useful in some pedagogical respects, has limited value as a productive simplification for cognitive science. Dewey argued this point with Russell 55 years ago, so why hasn’t the idea taken hold? Bartlett (1932) posed the same question, and concluded:

> It is because the force of the rejection of associationism depends mainly upon the adoption of a functional point of view; but the attitude of analytic description is just as important within its own sphere. . . .
>
> In various senses, therefore, associationism is likely to remain, though its outlook is foreign to the demands of modern psychological science. It tells us something about the characteristics of associated details, when they are

> associated, but it explains nothing whatever of the activity of the conditions by which they are brought together. (p. 308)

Bartlett contrasts the functional point of view, an understanding of social and biological conditions, with *description of patterns of behavior* (e.g., associations expressed as rules or semantic networks). Again, AI seeks not only to provide useful explanations of behavior, but to replicate the associational capability of the brain. Modern psychological science has been slow to develop a mechanism other than stored descriptions of associations. However, through the work of Hebb (1949) and his followers in connectionism a different kind of architecture, not based on networks of words, has been explored. In this section I will briefly survey some considerations in developing a machine with human capabilities of speech and understanding, as framed by the situated cognition perspective.

Figure 4 summarizes how situated cognition relates human knowledge, practice, and representational artifacts. Broadly speaking, the box labeled “knowledge” corresponds to conceptualizing and other representing processes in the brain. Cognitive scientists describe these dynamic capabilities and processes in terms of static perceptual categories, named concepts and properties, habitual procedures, etc. The box labeled “practice” corresponds to human behavior, including conversations, ways of talking, turn-taking, posturing, gesturing, etc., studied by subfields of discourse analysis, interaction analysis, etc. The box labeled “descriptions” corresponds to documents of all kinds, standardized database record structures (vocabularies), dictionaries, knowledge bases of expert systems, cognitive models, corporate policies, etc. Although all speech could be placed in this box, it is useful to separate written and other codified representations that can be stored, transferred to other people, and later viewed and interpreted. This separation is useful because documents, as artifacts, play a special role historically in the development of a community, which speech (even memorized narratives) does not have (Donald, 1991).

In terms of the Seaside example, the contractor’s *knowledge* includes his conception of his activities; his *practice* includes how he spends his time, his manner of speaking to other workers, how he organizes mail and requests in his office, etc.; his *descriptions* include letters he writes, forms he fills out, regulations he posts on the company bulletin board, etc. By this view, conceptual coordinating occurs in all actions, including the formalizing and interpreting of descriptions. Conceptualizing is a dynamic process of reconstructing “global maps” relating perceptions, other conceptualizations, and motor actions (Edelman, 1992). Conceptualizing is inherently *multimodal* (even when verbal organizers are dominating), *adaptive* (Vygotsky: “Every thought is a generalization”), and *constitutes an interactive perceptual-motor feedback system.* Conceptualizing is itself a behavior in animals capable of imagery and inner speech. ("Hearing" a tune in one's head is also an example of conceptualizing.) "Concepts," as formalized in descriptive cognitive models, are *names for conceptualizations* of objects, events, and relations, which may or may not be articulated in the discourse of a practice.

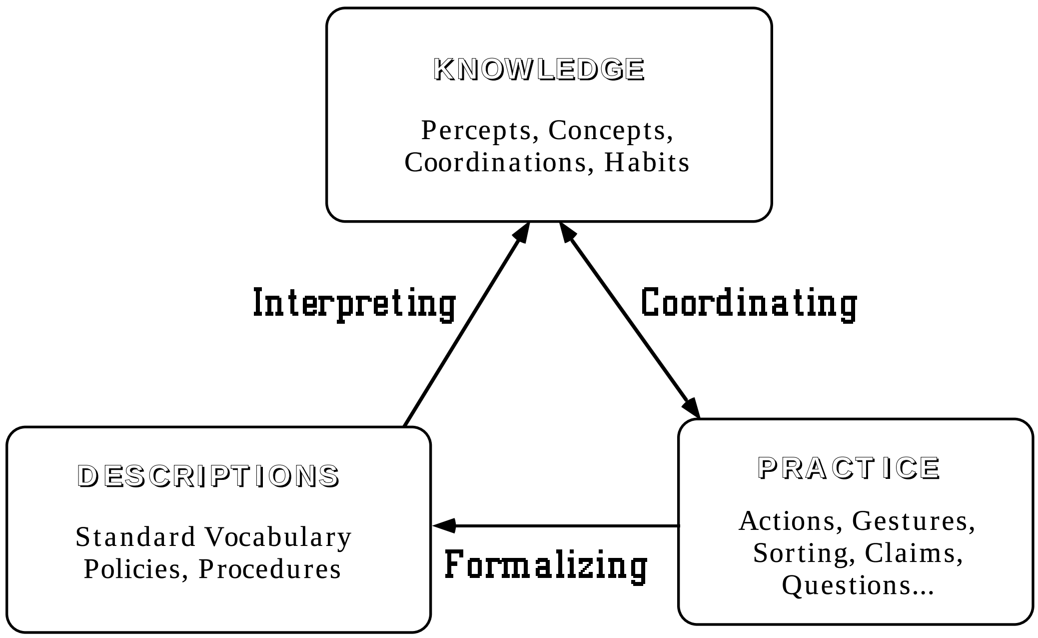

Figure 4. Simplified view distinguishing between conceptualization (knowledge), action in the world (practice), and text, diagrams, and computer programs (descriptions, commonly called "representations").

What physical recoordinating occurs as we speak and comprehend text? When cognitive modelers identify concepts with text networks, this scientific question does not arise—physical coordinating is viewed as an *effect* of comprehending text, not its basis. Examining the extremes of human experience, studies of creativity and dysfunction have discovered that conceptualization concerns much more than relating words; our knowledge includes conceptualization of scenes, rhythm, sequential ordering, identities, and values (e.g., see Gardner (1985), Sacks (1987), and Rosenfield (1992)).

Figure 4 is also intended to represent that knowledge and practice are co-determined: What we are doing is conceived as we are doing it. The effects are of course serial, as any conversation reveals. But the effects are also dialectic, in the sense that what we perceive and what we conceive, although separable categorizations in the brain, are postulated to co-determine each other. This is a kind of parallelism and interactivity different from modular architectures based on simultaneous, but independent formation. Chaos models probably come closest to characterizing how areas of the brain, functioning together, but generalized and “modularly” substitutable, can co-organize each other (Freeman, 1991).

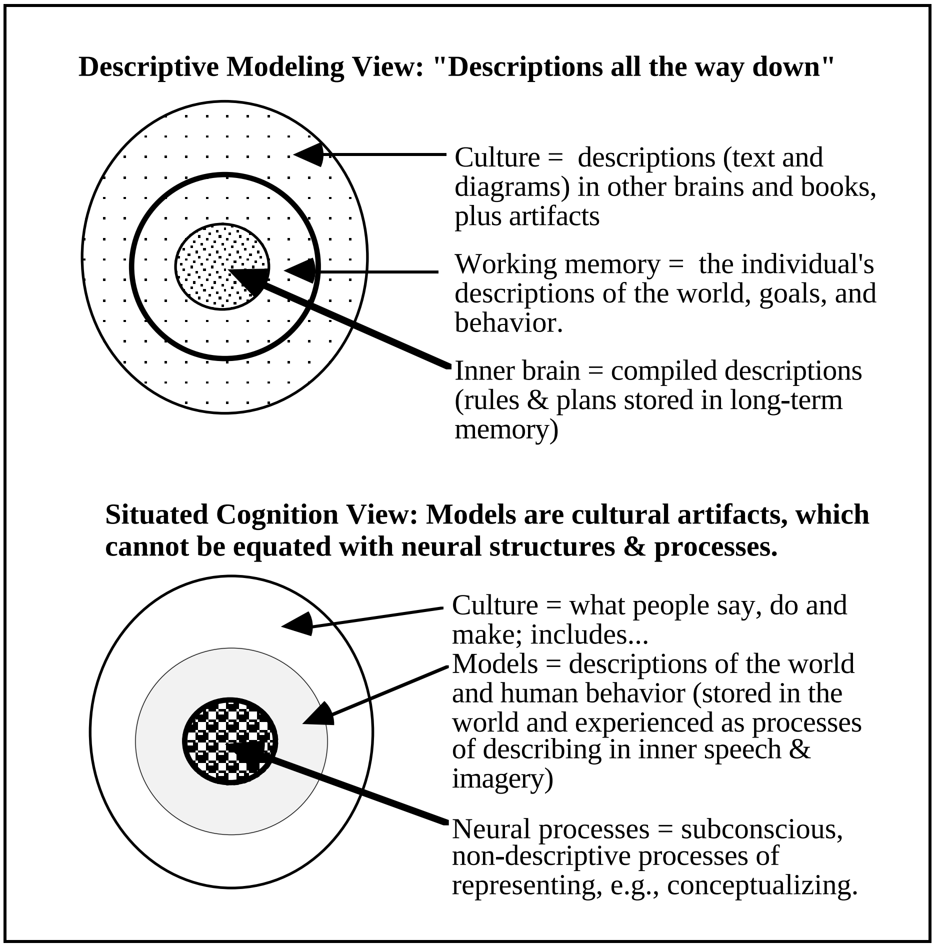

Figure 5. Philosophical positions relating models to the individual brain and culture.

Figure 5 illustrates the distinction between the descriptive and situated perspectives in another way. By the view that knowledge consists of text networks, called “representations,” representations are viewed as a single kind of thing, called “symbol structures,” located in the world, in working memory, and in long-term memory (e.g., see Simon, 1973; Vera and Simon, 1993). By the situated cognition view, a distinction is made between models on paper or in computer programs, imagined experiences, and conceptual processes. Human knowledge and culture can be described, *but are not reducible* to a body of descriptions. Put another way, we cannot understand how descriptions are created and given meaning unless we make a distinction between representations that are consciously manipulated and other forms of representing. That is, we cannot understand the nature of language if we call all forms of representing, inside and out, “symbol processing” (Lakoff, 1987; Edelman, 1992).

In the information-processing view (Newell and Simon, 1972; vanLehn, 1991), there is a sharp line between the individual and the environment, such that the relation is of input-output—taking in data and putting out what has been created inside. Neural processes are viewed as being similar in kind to conscious reasoning, involving storage, matching, assembly of descriptions (Vera and Simon, 1993).

In the situated cognition view, the functional distinction between conscious and subconscious is emphasized, and the line between culture and individual experience is less distinct. Individual experience is coupled through the conceptualization of activity to what other people say and do. Neural processes are not viewed as analogous to writing text, matching descriptions, or deducing the implications of rules. By this view, perception operates without a preconceived *description* of what is interesting in the world to perceive (Schön, 1979).

Figure 5 suggests that people are not information processors in the manner of expert systems. There are at least two ways of approaching this argument: Looking inwards to consider the nature of memory and perception, and looking outwards to consider the nature of social action. Both arguments relate to the idea of interpretation and meaning. Some philosophers and linguists argue that meaning cannot be equated with stored descriptions. Social scientists argue that descriptions must be creatively interpreted in practice. Both arguments claim that “creative interpretation” (however it may work) cannot be deduced from other descriptions alone. That is, the social and philosophical arguments of situated cognition ultimately make claims about the architecture of intelligence, and this is not the architecture of expert systems or most cognitive models. What are the implications of this shift in perspective for the design and use of tools in practice?

(5). ACTIVITY-BASED TOOL DESIGN

The Seaside example reveals how narrowly the task analysis of knowledge engineering views the knowledge, context, and problems of workers. When we view the construction manager in isolation, we describe his tasks as scheduling workers, planning construction, allocating space on site, and repairing equipment. Within these tasks, we describe the manager as modeling, gathering information, and explaining decisions. The expertise of being a construction manager is well-defined by such a framing of tasks, allowing us to describe problem spaces and rules that express his work in terms of solving a puzzle.

In developing a tool for the Seaside, knowledge engineers might have first thought to *deliver* a knowledge base that embodied the “Seaside Community Town Plan” and tools for *enforcing* that plan in the various decisions of the builders and homeowners. We now see the problem differently: Who interprets plans with whom, how can the plan be questioned and modified? How can we help people deal with difficult cases?

First, different personnel are engaged in different activities, suggesting that different tools might be needed. Indeed, it is perhaps more fruitful to view people who are collaborating as creating problems for each other (often inadvertently) and creating information (descriptions and evaluations of past designs, interpretations of policy) to help solve problems. The *work* of collaboration is partially caused by the need to coordinate different activities occurring in one place with shared resources.

The idea of collaboration as “reading from the same score” or “singing with one voice” is misleading; people are following *many scripts* because they are expected to fulfill different roles. In a medical clinic, for example, seven workers might be following five scripts based on their professional jobs which will have names like MD, PA, LVN, TCA, and RN. Software designers need to customize tools in terms of activities and not only tasks.

Kukla et al. (1992) describe how an activity-based analysis of conversations reveals what people need to see about other people’s work, when they need to see it, and where. On this basis, Kukla created collaborative tools for process control in a Monsanto chemical plant. The key design idea is to facilitate conversations by studying what people want to show each other when they are discussing different issues ranging from routine quality control and scheduling to troubleshooting. The study focused on unexpected problems that arise and that could not be easily automated; hence the tool provides assistance in situations of high risk and high cost. The program provides a means of accessing and comparing historical data from different viewpoints, video of current conditions, and instrument readouts. A surprising aspect of Kukla’s designs is that there is not “one collaboration tool” per se, but *different components and interface designs* for different roles in conversations about different topics.

Second, although some problems are common, tool design can go beyond automating what is routine to provide tools for helping people with non-routine situations. Studies of problem framing (Schön, 1979) indicate different phases that the puzzle view of cognition has poorly characterized:

- Framing, giving a name to a situation.
- Recounting the history of what you observed and did.
- Telling related stories from experience.

- Ordering events, claiming temporal relationships.
- Tentatively configuring a causal explanation of underlying processes.
- Reconceptualizing the meaning of observations, theories, and policies.

Lincoln et al. (1993) have described how the conception of a “computerized patient record” fails to relate to the representational work of clinicians who need a notational structure for describing the story of the patient, the patient’s disease, and encounters with the patient. Similarly, Zuboff (1988) applied situated cognition ideas to suggest how computers can be used to “informate” and not just automate the workplace. These investigators claim that providing access to historical information is the first problem. How strange that we proposed to give physicians Mycin in 1977, when the medical community didn’t even have databases!

Third, the idea that *learning is a process of conceiving an activity*, and activities are inherently social, puts emphasis on improving learning by addressing issues of membership, participation in a community, and identity. Participatory design of computer tools follows this approach (Greenbaum and Kyng, 1991). Participatory design deliberately brings together tool designers, workers, and researchers in a multidisciplinary collaboration (i.e., a choreographed conflict). Through scenarios, role-playing, and other ways of projecting design implications, the design of the computer tools is oriented to the practice of the users.[11]

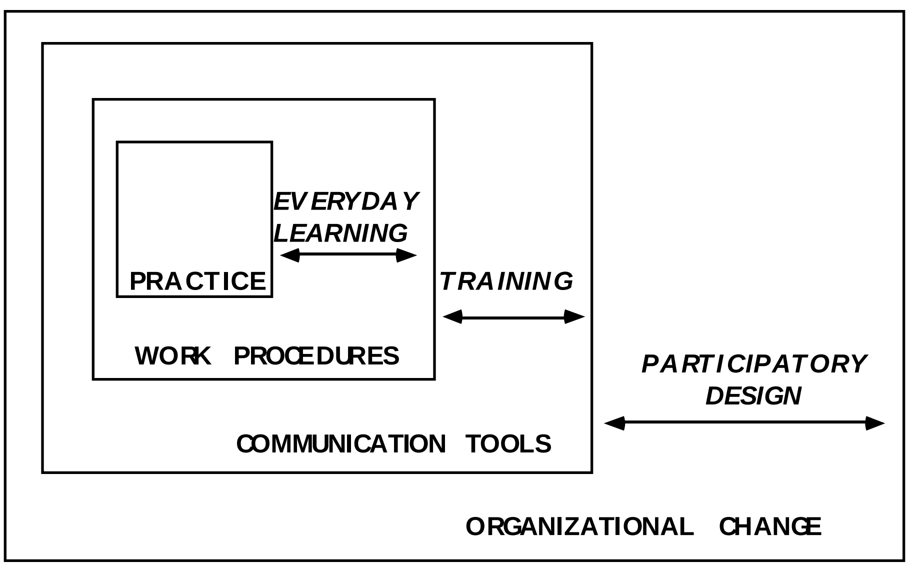

Figure 6. Aspects of organizational learning: Coordinating everyday practice, procedures, tools, and organizational change.

---

[11]See the special issue of the *Communications of the ACM*, “Representations of Work,” Volume 38, Number 9, September 1995, for further examples and discussion.

Figure 6 puts these ideas together to show how learning can be conceived as recoordinating occurring on different time scales:

- *everyday practice* (relating conception of practice to procedures and regulations),
- *training* (relating conception of procedures and regulations to workplace tools), and
- *participatory design* (relating workplace tools to the redesign process).

In contrast, a theory of knowledge based on descriptions would view behavior as constrained by designs and plans: procedures coordinate practice; tools coordinate the application of procedures; the design process coordinates the invention of tools. Both views are useful: One imposes order by design, the other explains how order actually develops and is sustained.

To summarize, an activity-based perspective suggests how to design computer tools to promote learning: Facilitate conversations and hence participation, allowing multiple perspectives to be viewed and compared. Focus especially on how problematic situations arise and are resolved.

(6). CONCLUSIONS

In this chapter, I considered the conceptual relation of knowledge, situations, and activities. I emphasized that knowledge is more than written scientific facts and theories. Professional expertise in particular is framed by a person's conceptualization of multiple, ongoing activities, which are essentially identities, comprising intentions, norms, and choreographies. These conceptions are an important source of conflicts, judgments, and values in human action. Expertise consists of generally useful concepts for coordinating activities, especially ways of framing problematic situations so that technical, descriptive methods can be applied in a routine manner (Schön, 1987).

By equating knowledge with scientific theory, designs, and policies, the "symbolic approach" of AI equated knowledge with text networks such as classifications, grammars, dictionaries, and causal networks, and other forms of scientific representations such as equations and diagrams. This view serves us well for describing some of the patterns of human reasoning. But machines based on this "architecture of intelligence" inadequately replicate the human ability to coordinate conceptions of meaning, physical skills, intention, visual scenes, timing, and attitude (Sacks, 1987).

By equating knowledge, situations, and activities with descriptions, the descriptive approach reduced "context" in problem solving models to goals and data. Consequently, the conceptual nature of context was oversimplified and its *content*, the aspect of identity in making judgments and choreographing action, was often ignored.

Ironically, by treating social conception as merely context (and context as just more data), the individual nature of cognition was glossed by descriptive modeling—personal differences were viewed as matters of opinion, different knowledge (stored in the head), or just differing awareness of the facts. Knowledge was viewed as universal because it consists mainly of scientific theories. Knowledge representation research consequently focused narrowly on topics such as “qualitative models of physics.” Indeed, the lingering worry that “individual differences” were not yet accounted for in descriptive models of learning reflects not understanding and relating *kinds of conceptualizing* and hence forms of representing other than describing (Gardner, 1985). Even studies of diagrams reduce pictures to words (Larkin and Simon, 1987).

With so many ideas equated or conflated in the descriptive modeling approach—knowledge, science, descriptions, and context—the rail against “decontextualized knowledge” (Brown et al., 1988) was interpreted as an attack on “abstract description” (Sandberg and Wielinga, 1991) and even an attack on science itself (Slezak, 1989). Instead, the force of the criticism was not upon the formalization of descriptive models, but upon their adequacy in directing everyday human affairs. Descriptive modeling has made solid contributions to scientific and engineering modeling (Clancey, 1992a). But building expert systems and intelligent tutoring systems on these principles alone produces tools that don’t address the dilemmas of everyday design in arguments about which facts are relevant and how to interpret policies (Greenbaum and Kyng, 1991; Kukla et al., 1993; Sachs, 1995). That is, the epistemology of expert systems is inadequate for understanding or enhancing collaboration.

The descriptive view of knowledge suggested that knowledge engineers should just codify theories and plans in tools and deliver them to workers. More recent efforts go beyond packaging expertise, to help people converse about designs and policy in unanticipated situations. These new tools use the same representational and automated modeling techniques originally developed for expert systems. But the design process starts with a better understanding of interpersonal aspects of interpreting, questioning, and modifying theories and rules—and this perspective is applied to the software design process itself. On this basis, we focus not just on delivery of preordained plans, but on construction of new conceptions, helping reconcile inherent conflicts in resources, timing, and values that arise as people with different expertise work together.

ACKNOWLEDGMENTS

This research has been supported in part by Nynex Science and Technology, Inc., a major HMO in Southern California, and a gift from the Xerox Foundation. I would like to thank my colleagues at IRL and our partners for sharing their insights: Gitti Jordan, Charlotte Linde, Judith Gregory from IRL, Jean Gilbert at Kaiser-Permanente, and at Nynex, Pat Sachs, Dave Torok, Maarten Sierhuis, and David Moore. But especially I have benefited from my friends in AI, who have encouraged me to keep trying to make sense of social theories: John McDermott, Luc Steels (Brussels), Rolf Pfeifer (Zurich), Kurt vanLehn, and the editors of this volume.

References

Alexander, C. (1977). *A Pattern Language.* New York: Oxford University Press.

Barker, R.G. (1968). *Ecological Psychology: Concepts and Methods for Studying the Environment of Human Behavior*. Stanford, CA: Stanford University Press.

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.

Berger P.L and Luckmann, T. (1966). *The Social Construction of Reality: A Treatise in the Sociology of Knowledge.* Garden City, NY: Anchor Books.

Brooks, R.A. (1991). Intelligence without reason. In *Proceedings of the 12th International Conference on Artificial Intelligence* (pp. 569-595). San Mateo, CA: Morgan-Kaufmann Publishers.

Brown, J. S., Collins, A., and Duguid, P. 1988. *Situated cognition and the culture of learning.* Institute for Research on Learning, Technical Report No. 88-0008. Palo Alto, CA. Shorter version appears in *Educational Researcher*, **18**(1), February, 1989.

Bruner, J. (1990). *Acts of Meaning.* Cambridge, MA: Harvard University Press.

Buchanan, B.G. and Shortliffe, E.H. (editors) (1984). *Rule Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*. Reading, MA: Addison-Wesley.

Chi, M.T.H., Glaser, R., and Farr, M.J. (editors) (1988). *The Nature of Expertise*. Hillsdale: Lawrence Erlbaum Associates.

Clancey, W.J. (1992a). Model construction operators. *Artificial Intelligence*, **53**(1),1-124.

Clancey, W.J. (1992b). Representations of knowing: In defense of cognitive apprenticeship. *Journal of Artificial Intelligence in Education*, **3(**2),139-168.

Clancey, W.J. (1993). Situated action: A neuropsychological interpretation (Response to Vera and Simon). *Cognitive Science*. **17**(1), 87-116.

Clancey, W.J. (in press). *Situated Cognition: On Human Knowledge and Computer Representations.* New York: Cambridge University Press.

Collingwood, R. G. (1938). *The Principles of Art.* London: Oxford University Press.

Dewey, J. (1939). Experience, knowledge, and value: A rejoinder. In P.A. Schilpp (editor), *The Philosophy of John Dewey*, Evanston, IL: Northwestern University, pp. 515-608.

Dewey, J. (1902/1981). *The Child and the Curriculum.* Chicago: University of Chicago Press. Reprinted in J.J. McDermott (editor), *The Philosophy of John Dewey*, Chicago: University of Chicago Press, pp. 511-523.

Donald, M. (1991) *Origins of the Modern Mind: Three stages in the evolution of culture and cognition.* Cambridge, MA: Harvard University Press.

Donald, M. (1993). Multiple book reviews of *Origins of the Modern Mind*. *Behavioral and Brain Sciences*, **16**(4)737-791.

Dreyfus, H.L and Dreyfus, S.E. (1986). *Mind Over Machine.* New York: The Free Press.

Dunlop, B. (1989). Coming of age. *Architectural Record*, July, 96-103.

Edelman, G.M. (1992). *Bright Air, Brilliant Fire: On the Matter of the Mind.* New York: Basic Books.

Effken, J.A. and Shaw, R.E. (1992). Ecological perspectives on the new artificial intelligence. *Ecological Psychology*, **4**(4), 247-70.

Feltovich, P.J. Coulson, R.L., Spiro, R.J., and Dawson-Saunders, B.K. (1992). Knowledge application and transfer for complex tasks in ill-structured domains: Implications for instruction and testing in biomedicine. In D. Evans and V. Patel (editors) *Advanced Models of Cognition for Medical Training and Practice*, Berlin: Springer-Verlag, pp. 213-244.

Fodor, J.A. (1975). *The Language of Thought.* Cambridge, MA: Harvard University Press.

Fodor, J.A. (1981). *RePresentations: Philosophical Essays on the Foundations of Cognitive Science*. Cambridge, MA: The MIT Press.

Frake, C.O. (1977). Plying frames can be dangerous: Some reflections on methodology in cognitive anthropology. *The Quarterly Newsletter of the Institute for Comparative Human Development*, **1**(3), 1-7. Also in Frake, C.O. (editor), *Language and Cultural Description.* Stanford, CA: Stanford University Press, 1980, pp. 45-60.

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.

Genesereth, M.R. and Nilsson, N.J. (1987). *The Logical Foundations of Artificial Intelligence*. Los Altos, CA: Morgan Kaufmann.

Gould, S.J. (1987). *An Urchin in the Storm: Essays about Books and Ideas.* New York: W.W. Norton and Company.

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.

Hall, E.T. (1976). *Beyond Culture.* New York: Doubleday.

Harrison, R. (1977). Functionalism. In B.B. Wolman (editor), *International Encyclopedia of Psychiatry, Psychology, Psychoanalysis, and Neurology,* Volume 5. New York: Aesculapius Publishers, pp. 144-6.

Hayes, P.J., Ford, K.M., and Agnew, N. (1994). On babies and bathwater: A cautionary tale. *AI Magazine*, **15**(4):14-26, Winter.

Hayes-Roth, F., Waterman, D.A., and Lenat, D.B. (1983). *Building Expert Systems*. Reading, MA: Addison-Wesley.

Hebb, D.O. (1949/1961). *The Organization of Behavior, a Neuropsychological Theory*. New York: Wiley.

Korzybski, A. (1941). *Science and Sanity*. New York: Science Press.

Kukla, C.D., Clemens, E.A., Morse, R.S., and Cash, D. (1992). Designing effective systems: A tool approach. In P.S. Adler and T.A. Winograd (editors), *Usability: Turning Technologies into Tools.* New York: Oxford University Press. pp. 41-65.

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

Larkin, J. H., & Simon, H. A. 1987. Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science*, **11**(1), 65-100.

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.

Lincoln, T.L., Essin, D.J., and Ware, W.H. (1993). The Electronic Medical Record: A challenge for computer science to develop clinically and socially relevant computer systems to coordinate information for patient care and analysis. *The Information Society* **9**: 157-88.

Mead, G. H. (1934). *Mind, Self, and Society.* Chicago: University of Chicago Press.

Merzenich M.M., Kaas J.H., Wall, J., Nelson R.J., Sur M., and Felleman D. (1983). Topographic reorganization of somatosensory cortical areas 3b and 1 in adult monkeys following restricted deafferentation. *Neuroscience*, **8**(1), 33-55.

Mills, C.W. (1940). Situated actions and vocabularies of motive. *American Sociological Review*, **5**:904-913.

Michalski, R.S. (1992). Concept learning. In S.C. Shapiro (editor), *Encyclopedia of Artificial Intelligence*, Volume 1, New York: John Wiley & Sons, Inc. pp. 249-259.

Nonaka, I. (1991). The knowledge-creating company. *Harvard Business Review*. November-December. 96-104.

Newell, A. and Simon, H.A. (1972). *Human Problem Solving*. Englewood Cliff, NJ: Prentice-Hall, Inc.

Newell, A. (1982). The knowledge level. *Artificial Intelligence*, **18**(1), 87-127, January.

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

Polk, T. A. and Newell, A. (1995). Deduction as verbal reasoning. *Psychological Review*, **102**(3):533-566.

Polanyi, M. (1958/1964). *Personal Knowledge: Towards a Post-Critical Philosophy*. New York: Harper and Row.

Rosenbloom, P.S., Newell, A., and Laird, J.E. (1991). Toward the knowledge level in Soar: The role of the architecture in the use of knowledge. In K. vanLehn (editor), *Architectures for Intelligence: The Twenty-Second Carnegie Symposium on Cognition*, Hillsdale: Lawrence Erlbaum Associates. pp. 75-112.

Rosenfield, I. (1992). *The Strange, Familiar, and Forgotten* . New York: Vintage Books.

Sachs, P. (1995). Transforming work: Collaboration, learning, and design. *Communications of the ACM*, **38**(9), 36-44, September. Special issue on “Representations of Work.”

Sacks, O. (1987). *The Man Who Mistook His Wife for a Hat.* New York: Harper & Row.

Scribner, S. and Sachs, P. (1991). Knowledge acquisition at work. *IEEE Brief*. No. 2. New York: Institute on Education and the Economy, Teachers College, Columbia University. December.

Sandberg, J.A.C. and Wielinga, B.J. (1991). How situated is cognition? In *Proceedings of the 12th International Conference on Artificial Intelligence*. San Mateo, CA: Morgan-Kaufmann Publishers. (pp. 341-346).

Schön, D.A. (1979). Generative metaphor: A perspective on problem-setting in social policy. In A. Ortony (editor), *Metaphor and Thought*. Cambridge: Cambridge University Press. 254-283.

Schön, D.A. (1987). *Educating the Reflective Practitioner*. San Francisco: Jossey-Bass Publishers.

Simon, H.A. (1973). Ill-structured problems. *Artificial Intelligence*, **4**, 181-201.

Sleeman, D. and Brown, J.S. (1982). *Intelligent Tutoring Systems.* London: Academic Press.

Slezak, P. (1989). Scientific discovery by computer as empirical refutation of the Strong Programme. *Social Studies of Science*, Volume 19. London: Sage. pp. 563-600.

Suchman, L.A. (1987). *Plans and Situated Actions: The Problem of Human-Machine Communication.* Cambridge: Cambridge Press.

Thelen, E. and Smith, L.B. (1994). *A Dynamic Systems Approach to the Development of Cognition and Action.* Cambridge, MA: Bradford Books.

Turvey, M.T. and Shaw, R.E. (1995). Toward an Ecological Physics and a Physical Psychology. In R. L. Solso and D. W. Massaro (eds.), *The Science of the Mind.* pp. 144-172. New York: Oxford University Press.

vanLehn, K. (1991). *Architectures for Intelligence: The Twenty-Second Carnegie Symposium on Cognition*, Hillsdale: Lawrence Erlbaum Associates.

Varela, F.J. (1995). The re-enchantment of the concrete. In Steels, L., & Brooks, R. (editors), *The "Artificial Life" Route to "Artificial Intelligence": Building Situated Embodied Agents*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Vera, A.H. and Simon, H.A. (1993). Situated action: A symbolic interpretation. *Cognitive Science*. **17**(1), 7-48.

Von Foerster, H. (1970). Thoughts and notes on cognition. In P.L. Garvin (editor), *Cognition: A Multiple View.* New York: Spartan Books, pp. 25-48.

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.

Wertsch, J.V. (1991). *Voices of the Mind: A Sociocultural Approach to Mediated Action.* Cambridge, MA: Harvard University Press.

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 (editors), *Design at Work: Cooperative design of computer systems.* p. 45-64.

Zuboff, S. (1988). *In the Age of the Smart Machine: The Future of Work and Power.* New York: Basic Books.