Practice Cannot be Reduced to Theory:

Knowledge, Representations, and Change in the Workplace

William J. Clancey

Institute for Research on Learning
2550 Hanover Street
Palo Alto, CA 94304

Revised draft: January 15, 1993

(Appeared in S. Bagnara, C. Zuccermaglio, and S. Stucky (Editors),
Organizational Learning and Technological Change, pp. 16-46.
Papers from the NATO Workshop held September 22-26, 1992 in Siena, Italy)

Practice Cannot be Reduced to Theory: Knowledge, Representations, and Change in the Workplace

William J. Clancey

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

Abstract. Changing views of the nature of human knowledge change how we design organizations, facilities, and technology to promote learning: Learning is not transfer; using a plan is not executing a program; explanation is not reciting rules from memory. Such rationalist views of knowledge inhibit change and stifle innovate uses of technology. Representations of work (plans, policies, procedures) and their meaning develop in work itself. Representations guide, but do not strictly control human behavior. Every perception and action involves new, nonlinguistic conceptualizations that reground organizational goals and values. This essay explores how the epistemology of situated cognition guides business redesign.

Keywords. Organizational learning, situated cognition, situated learning, rational agent, business process design, reengineering, ethnography of work

1 Introduction

Change is always occurring in the workplace, often without a theoretical understanding of why one technology or organizational design works and another doesn't. Many programs for corporate change seek to impose policies, failing to acknowledge how people make sense of plans and standards in everyday work (Kling 1991, Suchman 1987, Wynn 1991, Zuboff 1988). On the other hand, focusing on the social nature of cognition may devalue how individual experiences shape new organizations and technologies.

This essay explores how new views of knowledge can be applied more systematically to organizational learning and use of technology. How can we help businesses create policies and standards that will guide creativity rather than stifle it? The central idea is to avoid equating knowledge with representations of knowledge. We want to shift from the "capture and disseminate" view of managing work to designing workplace processes that facilitate learning. In particular, explanations of delays and exceptions to standards can be exploited as opportunities for articulating new patterns and theories, instead of making an employee feel guilty for not rotely adhering to bureaucratic procedures.

Building on the idea of situated cognition, I broadly articulate the nature of knowledge, representations, and change in the workplace. I set the stage by giving examples of misconceptions about knowledge (Section 2) and summarize the tenets of the rationalist view of learning (Section 3). I illustrate the relevance of situated cognition to business process redesign by an example of a recent organizational change in a large corporation (Section 4). I then elaborate on the relation of practice, plans, and justifications, giving examples from software engineering and expert systems (Section 5). I present a model of workplace change in terms of intersecting communities of practice, using as an example the development of a medical information system (Section 6). Finally, I examine how representational change occurs in the course of reflecting on experience and reusing models in different settings over time (Section 7).

2 Common Misconceptions

At the NATO workshop at which this paper was first presented, I heard several statements about knowledge, perception, information, learning, memory, representations, consistent with the dominant views in cognitive science, sometimes called the "rationalist position" (Winograd and Flores 1986). Such views have been increasingly criticized since the mid-80s (e.g., Bickhard and Terveen in preparation, Clancey in press b, Dreyfus and Dreyfus 1986, Edelman 1992, Gasser, 1991, Lave 1988, Smith 1991, Suchman 1987), signaling a shift within cognitive science, sometimes called the "situated cognition position." The controversy is in fact quite old; most of the situated cognition ideas were clearly articulated by Dewey,

Collingwood, Bartlett, Bateson, Ryle, among others (Clancey in press b). The point of this paper is to develop implications for organizational learning, rather than redevelop the situated cognition argument. Nevertheless, it may be useful to illustrate the situated cognition point of view, particularly since it is controversial to many participants of the workshop itself. I quote seven statements made during the workshop, explain them according to what the participant probably intended, and use the opportunity to summarize situated cognition ideas.

"They have knowledge but cannot act."

What is meant is that a person can have a relevant, but incomplete model in hand, know that it is incomplete, and still not know what to do. This remark suggests an apparent paradox: A person can "have knowledge" but not be able to use it. This equates knowledge with representations and acting with applying representations. Knowledge is not a thing but a capacity to interact (Dewey 1938, Ryle 1949). Knowledge is more analogous to energy than a substance. Knowledge and representations are not interchangeable: We can represent what someone knows, but the representations are not the knowledge itself (the map is not the territory) (Clancey 1992b). Conversely, we can act without manipulating linguistic representations at all (but indeed, we often use maps and models to orient new behaviors).

"We have too much information."

What is meant is that we have too much data, more descriptions than we know what to do with. We lack abstractions, trend analyses, ordering, causal concepts, and theory. That is, we have too much *data* and not enough information. As Bateson (1972) would put it, we don't know what differences make a difference.

"Knowledge should be stored before it gets lost."

What is meant is that a representation should be recorded before it gets lost. In equating knowledge and representations, this remark views human memory as a place, so learning is like storing representations in a knowledge base (Rosenfield 1988, Clancey 1991b). This also suggests that using representations involves merely accessing them and applying them, leaving out the reconceptualization that always occurs when we read and comprehend text (Tyler 1978). The idea of getting "lost" focuses on how to retrieve representations when they are relevant. In contrast, human reminding is always relevant when it occurs (though the memory might not be useful) (Bartlett 1932). In effect, human learning and modeling are conflated in this remark, distorting how knowledge bases are actually created and used.

"Expert systems don't learn from experience."

What is meant is that the program is not revising its representations (knowledge base) as a result of feedback from incorrect behavior. But the deficiency is more stark: People learn with every motion, for example, in the process of conceiving questions, in the process of formulating rules—within a problematic situation, not just afterwards. The problem is that expert system reasoning is timeless and disembodied. Programs generate questions by instantiating plans, which are effectively grammatical descriptions of interactional strategies. In computers, every perception and action is linked by deliberation, rather than automatically adapted by processes of recategorization and recoordination at a prelinguistic level, as in people (Dewey 1896, Clancey in press). The learning deficiency of computer programs exists because human knowledge, memory, and perception-action coordination have been inadequately characterized.

"Mental models are in the head of the user of the computer system."

What is meant is that the person may have an experience of visualizing something or silently saying something. In this sense the experience is internal. But such experiences are not in the head in the same sense that representations are *in* a computer system. Relative to scientific models in general and knowledge bases in particular, mental models are not models at all—they are experiences, not substances. We talk about "running a mental model." But what is the "it" that is being run? A metaphor is taken literally here, as if a mental model is something stored and applied like a computer program. Experiences of creating and interpreting representations must not be equated with manipulation of a symbolic calculus (as in qualitative models of expert systems (Clancey 1989)).

"Learning is an individual, knowledge-based process and a social communication process."

What is meant is that learning has both individual and social aspects. To preserve an objective, given world, the remark suggests that an individual determines the truth and communicates it. But viewing learning as inherently individual misses how it occurs *in activity*, in cycles of perceiving and acting over time (Bamberger and Schön 1983). Articulating representations is inherently interactional, involving speaking or writing, or imagining. Furthermore, in social activity, individual perceptions and statements are occurring as part of coordinated discourse whose (social) forms shape goals, interests, and argumentation (Roschelle and Clancey, in press). Put another way, creating models inherently involves communication (always with oneself, often with others); communication is not something that occurs after the models are in hand.

"Perception is controlled by theory"

What is meant is that representations (maps, instructions, goal statements, beliefs, causal relations) are interpreted to influence where we look next and what we see. But interest, a sense of similarity or discord, and a sense of value in general precede creation of theoretical representations (Schön 1979). Consider how you find items of interest in an museum before you know what you are looking at. Perception is always value-laden and occurs at a pre-linguistic level, grounded in prior coordinations and perceptions (Dewey, 1896). How descriptions (e.g., reading instructions) change behavior becomes an important area to reconsider; theory does not control human perceptions and behavior in the sense that rules and programs control a computer's data processing (Suchman 1987).

3 The Rationalist Argument

Understanding the situated cognition perspective may be difficult at first because our concepts of knowledge, information, memory, perception, representations, etc. are intertwined. Change one and you must change them all. Fig. 1 summarizes the rationalist line of argumentation, which I claim is implicit in the seven statements we just analyzed.

The strong relation between normal science and the idea that knowledge consists of representations is partly responsible for prevalent views of memory and learning in cognitive science. Changing this view sometimes appears to threaten the scientific enterprise itself, or it may appear contradictory, suggesting that all theories are relative and it doesn't matter what we say (Slezak 1989. 1992). The key idea is to not equate phenomena being studied with models or equivalently, in the realm of human endeavor, not equate practice (what people do) with theory (what people say is true). The view that natural laws create regularities in the physical domain has been translated to the view that corporate policies and procedures create regularities in the workplace. Such causality between representations and behavior holds in computer systems, but not in human beings. Indeed, the presumption is that violation of business policies can only mean anarchy, when in practice anarchy is what would result if policies were literally followed.

It may appear at first strange to bring up issues of human memory in discourse about organizational learning. But we must root out fundamental assumptions about the nature of knowledge and representations if we are to understand how organizations actually generate and use policies, and what designs for organizations and technology would be more productive (Winograd and Flores 1986, Ehn 1988, Hirschorn 1984, Kukla, et al.. in press, Nonaka 1991). As the next example illustrates, change can occur without such theorizing, but at the risk of being ineffective or nonsystematic.

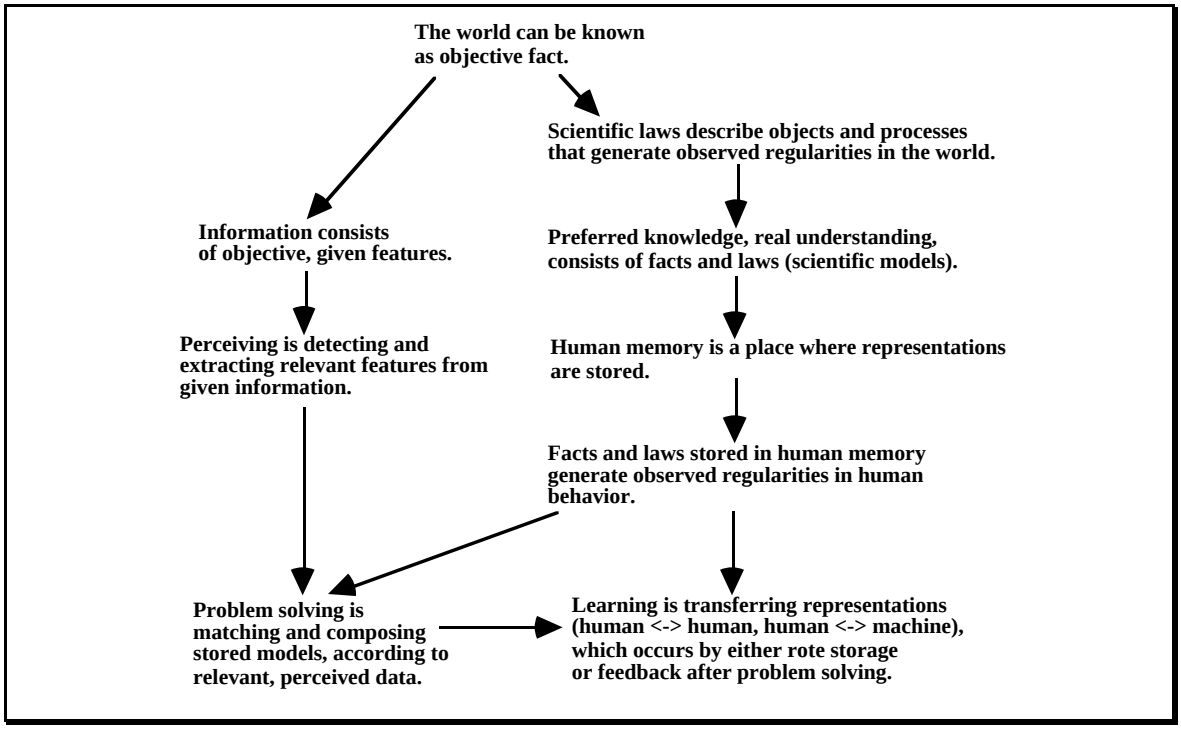

Fig. 1. Rationalist view of knowledge and representations

4 Organizational Design Example

Table 1 illustrates a shift that is occurring today in an American corporation. The shift is described in terms of changes to the organization:

- The changed organization is intended to build on internal diversity among existing employees, enhance teamwork and communication, promote learning through story telling, etc.
- The changed organization is intended to provide stable interactions with customers, so that client needs are "known deeply," while providing a single face or view of the corporation,

These changes appear innovative and intuitively appealing. But corporate management doesn't articulate the view of knowledge and learning implicit in this design. Hence the corporation is a poor position to compare alternatives and guide concomitant design of facilities and technologies in a principled way. Within the corporation, the lack of a principled theory might impede change by making this design appear arbitrary and potentially no better than any competing approach.

In effect, the corporation is shifting from a transfer view of learning to an interactive view. The old view suggests that knowledge resides in individuals and involves technical details; furthermore, the company's products and services are conceived internally, a priori, and disseminated to customers. One sales person serves many customers. Because sales people are specialists, each customer must potentially deal with many employees. Each sales person works for him or herself, seeking to move to the best geographic areas and to sell the most profitable products, in order to maximize personal income and promote his or her career.

In the new view, knowledge is about customer interactions. The view is inherently interpersonal and relational. Employees are generalists responsible for many products and services. Employees necessarily need to interact with each other, working as a team, to sell coherent systems specialized for customer needs. That is, each employee has an inherent need to get technical help from his or her colleagues. Communication within the corporation is encouraged and formalized by formation of functional workgroups and customer-liaison teams. Knowledge, as a capacity to interact with customers to define products and services, develops during interactions with customers (as opposed to being transferred in classes). The anticipated result is a more stable, single-face interaction with customers. Internally, there

should be less churning of sales representatives into more lucrative markets, and promoting loyalty instead to the customer-team relation.

Table 1. Epistemological shift implicit in an organizational redesign.

	Knowledge is about	**Knowledge resides in**	**Knowledge is developed by**
Individual View: Reify the individual employee, a constant player who moves in the corporation	Technical details of products and services (internal capacity)	Specialized employees (stored in individual heads)	Training given to individual
Interactional View: Reify company-customer relations as stable & responsive	Customer relations (interactive capacity)	Cross-functional team (manifest in activity)	Project activity of functional workgroup and teams

My point is not to argue specifically for this organization, although it does fit the theories of knowledge and learning I support. Rather, this example illustrates how explicating underlying principles of knowledge, memory, and learning can help us understand organizational alternatives. Furthermore, the example illustrates that theories of knowledge that might appear esoteric to business managers have direct implication for organizational design. In particular, the example illustrates that the jargon of today's organizations (e.g., "cross-functional workgroups") can be usefully related to situated cognition theories.

5 Practice, Plans, and Justification

In this section, I will consider in more depth the nature of plans from the perspective of software engineering and expert systems. We will examine how plans are created and used, and what happens when a plan is violated. In particular, we must change our understanding of how people justify their work. I will argue that organizational change is impaired if plans are not appropriately understood with respect to the inherent adaptiveness of human knowledge and activity. The rational view of stored knowledge and planned action inhibits organizational learning: It distorts descriptions of how work actually gets done, inhibiting reflection and preventing learning from past actions.

5.1 The rational view of work and plans

Many software managers believe that the problem of software development is to plan, control, and predict software programming throughout the development and maintenance cycle (Fig. 2). Design is framed as an activity that is rational, that is, an activity whose *every step is a deliberate choice* among alternative actions, justified by appeal to scientifically-derived standards and metrics. The problem of improving software development is conceived as defining appropriate standards for the development process and metrics for evaluating deviation from the ideal. Learning is conceived as a feedback process by which deviations are translated into issues that guide reorganization, development of better tools, and new standards of behavior. From this perspective, organizational learning is often conceived as a process of continuous improvement and *dissemination*, by which successes and failures are continuously analyzed and formalized as methods and tools to be transferred throughout the company.

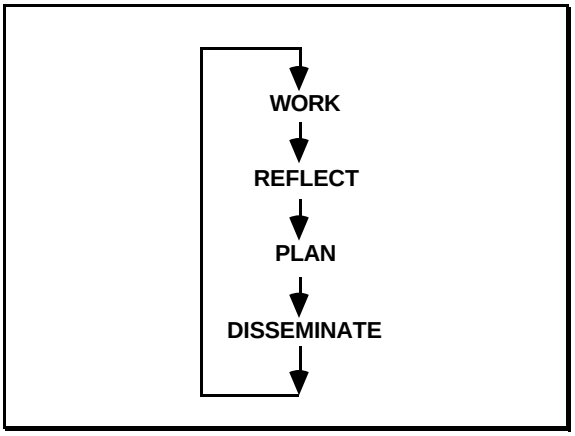

Fig. 2. Rational ordering of work, learning, and policies.

This rational view of work has merits, for it is true that progress depends to some extent on reflecting on past performance and formalizing procedures to improve future outcomes. But this view of human activity over emphasizes the planful nature of cognition. When people behave—for example when programmers write a line of code—they are not merely executing procedures or applying rules in the manner of a computer. Except when people are deliberately behaving like automatons, as in reciting a poem or reading a list of numbers, no human behavior has such a mechanical, template-driven nature. As we behave, the brain is not merely applying a stored plan, but is *constructing new ways of coordinating* what we see and do. Our moment-by-moment processes of interacting with our environment always result in new, adapted behaviors. Learning is occurring with every perception and movement.

In contrast, by the rationalist view intelligent behavior is always justifiable. Everything we do must have a reason, *which we have considered before acting*. This supposedly distinguishes us from animals, who don't represent their world and don't plan what to do. Supposedly people don't behave instinctively; they don't simply react. As civilized human beings, we are always supposed to follow rules, to behave according to scientific principles, to be methodical, to be lawful.

As an attitude for socially orienting behavior, this view has merits. But it is a poor description of how the human brain works. There are no stored plans, no procedures put away that we simply execute. Knowledge is not stored away like tools in a shed, which remain unchanged between uses. Learning is not just something that occurs on reflection, after we have acted. Every human action is unique, adapted, and hence process of learning. When I speak, I am not translating from an internal, hidden description of what I planned to say. I am conceiving. I may tell myself silently what I plan to say, but this silent telling is itself a form of reconceiving.

According to formal models of work, each step is viewed as manipulation of representations according to other representations (facts, rules, policies). In practice, each step involves action grounded on non-linguistic, value-oriented conceptualization and perceptual categorization. That is, action is not grounded on representations like written policies and standards. Following a recipe or implementing a corporate policy involves improvising and creating new causal stories and theories. Hence, rationalization is not reciting a plan that was followed in the manner a computer program is applied. Reflection is a time for reconceiving and theorizing, looking toward the future, not merely justifying with respect to rules.

5.2 The symbolic view of explanation

The example of "consultation program" explanation systems (Scott, et al. 1984, p. 338) reveals how our view of explanation becomes embodied in technologies and interactions between people and machines. Consider this typical example from the Mycin program, developed circa 1975 (Fig. 3).

```
** How did you determine that the aerobicity of ORGANISM-1 is
facultative?

The following were used:
RULE027 indicated there is strongly suggestive evidence (.8) that the
      aerobicity of ORGANISM-1 is facultative
RULE027 indicated there is weakly suggestive evidence (.2) that the
      aerobicity of ORGANISM-1 is anaerobic
Since this gave a cumulative CF of (.8) for facultative, and (.2) for
      anaerobic, it has been established that the aerobicity of
      ORGANISM-1 is facultative.
```

Fig. 3. Example of MYCIN's explanations (from Scott, et al., 1984, p. 347)

Table 2 summarizes the symbolic view manifest in the design of Mycin's explanation system. The second column summarizes how we view the relation of knowledge and representations today.

As pointed out by Lave (1988, Lave and Wenger 1991), the view of knowledge as consisting of stable, stored representations suggests a transfer view of teaching, advice giving, and explanation. Indeed, in the heyday of the Knowledge Systems Laboratory at Stanford, some of us wrote a manifesto summarizing our discovery that creating computer programs, giving advice to people in the workplace, and teaching students could all be accomplished by transferring representations from computer memory to user memory to student memory (Barr, Bennett, and Clancey 1979). Of course, nobody in the 1970s would have argued that MYCIN's advice was intended to be followed by rote. But aside from providing access to the program's deductions—an important innovation at the time—nothing in the design of the system related to the user's sense-making capabilities (see Clancey in press b, for further discussion).

Sometimes the deficiency of the transfer view of learning is described in terms of the context dependence of knowledge. But the issue is subtle. AI researchers know that knowledge base representations are not context independent. But they generally believe that the contexts to which knowledge base rules apply can be represented as conditions (features of the environment) that are themselves non-problematically given (objective) and stable. Indeed, the basic components of expert knowledge are called "situation-action rules." Similarly, the justifications for rules themselves are assumed to rest on a stable, objective, representational bedrock (e.g., see Wallace and Shortliffe 1984). In this way, all information about the world, all knowledge of the world, and all justifications for actions rest on representations—what I have characterized as "representational flatland" (Clancey 1991a). Hence, all action can be effectively controlled by "acquiring" the right representations. To understand what is wrong with this position, we need to consider in more detail the nature of human creativity and emergent aspects of behavior.

Table 2. Shift in view of explanation

	Exclusively symbolic view of knowledge	**Situated Cognition View**
Human knowledge	Facts and rules	Capability to interact; inherently unformalized; cannot be inventoried, grounded in non-linguistic perceptual categorizations, conceptualizations, and coordinations.
Knowledge base	Equivalent to human knowledge	A qualitative model of causal, temporal, spatial and subtype relations between objects and events in some domain of inquiry.
Explanation	Reciting facts and rules previously used to solve problems.	Reperceiving, reconceiving situations, theorizing freshly in a way that relates previously articulated models to non-theoretically controlled human actions.
	Looks backward to justify what happened.	Looks forward to improve future interactions.
Use of a consultation program	A mechanism to which people supply data and receive advice to be followed by rote.	A tool that facilitates perception of relevant data, reinterpretation of data categories, and reconception of previously articulated rules in each unique situation.

5.3 Emergent effects: How representations are created and used

A dynamic, *interactional view of behavior* can be difficult to understand, because the view that knowledge is stored and that speaking is a process of translating ideas into words is so ingrained in our society. A better psychological understanding requires turning the rationalist view inside out.

Representations—such as plans of what to do and reflections on our own behavior or what other people have said—are created in our activity, as we speak. Representations do not exist until we say something, or write something down, or imagine something in our head. Representations are created in the course of activity; they are not part of a hidden mechanism that drives our moving and speaking and seeing. Representations play an essential role in human behavior, but only in *cycles of behavior*, as we perceive what someone said, and say something else in return. As we perceive representations, we reorient our activity: We look somewhere else, we reperceive our partially completed work, and we organize our activity in a new way (Bamberger and Schön 1983).

An architect designing a building, a scientist writing an article, and a computer programmer all have in common the broad pattern of putting something into the environment—a line, a statement—and periodically stepping back to reconceive the nature of the whole. We are always interacting, putting some *thing* into the environment, reperceiving what we have done, and replanning (and this interacting can go on privately, in our imagination). Plans orient this process, but how the plans are interpreted is itself an interactive, non-predictable process. We cannot fully predict what the completed work product will look like because it is not generated by mechanically applying formulas. The architect, the writer, and the programmer don't know what they will produce until after they put their representations down on paper. Periodically, they will reperceive the whole and observe unanticipated effects. They observe trends, flows,

and directions that they didn't expect. New opportunities arise for reshaping the project. New plans are formulated to account for unanticipated difficulties or interesting new directions. Every behavior is an improvisation.

How can this interactional view improve organizational learning? We must begin by realizing that holding people accountable to representations, such as standard plans and procedures, must *take into account how people use representations*. In particular, we must remember that the very essence of human behavior is novelty, flexibility, and adjustment. We are always, at the root level, innovating. Reperceiving, readjusting, and replanning are integral to intelligent behavior.

Reflecting is necessary because we are creating something new when we design a building, an article, or a computer program. Each statement is a creative act. We must regularly stop and look at what we are doing. But also interactions between the parts can produce unanticipated effects. Both positive and negative properties can emerge, which are perceived as we reflect on what we have done so far. The building may become too large for its site, the article may become rambling and unevenly developed, the computer program may bog down against resource limits. Allowing for refocusing and replanning is therefore necessary for several reasons: creating something new can't be predicted, we want to benefit from serendipitous effects, and we must cope with undesirable interactions.

5.4 Coping with rationalist standards in a chaotic world

Where does this leave the programmer and software manager? They have been told that their activity must be planned, predictable, and controlled. But unless they are engaged in some rote, mechanical process, they cannot strictly predict or control what they will create. They won't know what they have or even what they can do until something of substance is on paper and they can reflect on where they are headed.

Pity the manager and worker who are held accountable to plans, but are expected to be creative and flexible. Programming manuals describe a highly-structured process full of documentation, obligatory liaisons between teams, approval meetings, checkpoints, and product specifications. The organization, in putting out these standards and procedures, seeks predictability and uniformity. Yet, in all this planning and proceduralization, there is no sense of what the manager or programmer should do first and what after that. How does the programmer know when he or she is making progress? New interactions are continuously being discovered, new opportunities and difficulties are emerging. We can see the date and where we stand in the schedule. But who knows what shortcuts or detours will occur tomorrow?

Human behavior is inherently *ad hoc*, inventive, and unique. But Western culture has biased us to believe that intelligent behavior is planned and under control. Every action is supposedly governed by representations. As adults we are supposed to be conscious of why we are doing what we are doing; we must have reasons. All of us, as members of this society, are influenced by this view of human nature. Unfortunately, since human behavior is not governed by representations in this manner, we find ourselves in a double bind.

Examining our own behavior, we find that we are simply doing things. We are always at some level being intuitive. Again, at some point, even the plans we generate can't be accounted for or strictly justified. But we must respect organizational authority, through our tacit agreement as participants of this society to be rational, to follow the regulations. We are caught: We don't reveal to other people or even to ourselves the true nature of how we do our work. Instead, we put up good appearances. We learn how to justify our behavior, and such representations of what we have done are always *post hoc*. Like mathematicians, we learn to tidy up our reports into nicely formatted proofs. It always appears as if we did things according to the rules: We made observations, we gathered evidence, we deliberated, we drew conclusions. We followed a plan, we were scientific. We were intelligent. We were good members of the organization.

Endorsing this rationalist view of human behavior inhibits learning. We don't reveal to ourselves or others how our work actually gets done. Indeed, working within the rationalist conception of science, we do not even have the concepts for describing how our mind works and how ideas develop in a community. We feel obligated to say that we followed the template, that our knowledge really did consist of schemas and rules, and we did what we were supposed to do. Rather than reflecting and seeing how we are really interacting with people and our environment, we (usually unknowingly) say what people want to hear: We endorse not only the rationalist view of knowledge, but the rationalist genre of justification. We get good at telling the rationalist story.

With this vested interest in a command and control mentality, neither the workers nor the managers are able to see how work really gets done. Indeed, we may be momentarily stopped in our tracks when "called to task" for not following a plan. Confusion becomes acute when we realize that no thing, no representation, underlies our behavior. There is ultimately no deeper plan or better understanding to refer to. Even when we realize that influences are beyond our control, we make up names like "tacit knowledge,"

as if to say that there is something written down and stored inside that properly justified our behavior, we just weren't thinking about it explicitly.

It takes courage to put forth the interactional view, to realize that there are no well-thought-out rules or tidy laws inside that make us the rational machines we hoped to be. We must accept that story-telling is all we have ever done or ever could do in justifying our behavior. Indeed, our integrity is called into question as we are called to justify ourselves according to the rationalist standard and cannot find the plans or schemas inside. There are no deeper facts to appeal to, we simply behaved.

5.5 Giving people the right to reflect

Improving learning, particularly within business, requires granting people the uniqueness and openness of experience. People must have the freedom, when appropriate, of saying, "I am not a machine, I am always doing something new. I am creating. I can't tell you what I will produce for sure until it is done. I can't say when the whole project will be done because I don't know what it will look like or how it will interact with other parts until it is well along." We must *give people a different view of the self and what a justification is*:

- What we do is always coming into being in our interactions, it can't be strictly predicted or controlled;
- Representations, such as business standards, orient, but don't control human behavior in the manner of computer programs;
- Reflection is itself a creative interaction, a reperception of emergent properties of the whole and interactions beyond the context of our immediate focus.

We inhibit learning when we view people as machine-like, suggesting that they follow instructions like a machine, and force them to justify behavior exclusively in terms of previously articulated plans. In one sense we have a poor view of what standards are for, how people use them, and how to relate standards to behavior. We have a poor view of accountability. We don't properly evaluate human work because we have obscured the inherent unplanfulness and inventiveness of every action.

Of course, in practice explanation does focus on understanding the relation of what we said we would do and what happened. What I am suggesting is an attitude and an understanding between employees and managers that *acknowledges the inevitability of diversions and discoveries*, and seeks to work with them, rather than apologizing or obscuring what is happening.

We must give people the right to acknowledge that they are constantly reperceiving and reorganizing their work. We must emphasize their crucial power to resee and reappraise what they have done so far. We must give them the right to be assessors of their own work, giving them time and tools to reflect and reconceive what they are accomplishing. We will do this not just because it is good, democratic policy, but because it is inherently the nature of human activity. Without the right to reassess and redirect, we are flying blind or idly, mechanically, doing what someone else conceived before current possibilities and problems emerged.

People do not simply plan and do. They continuously adjust and invent. Managing this process means *managing learning*, not managing application of a plan. Managing a creative process means *orienting the process of inventing new orientations*.

5.6 Disseminating practice, not theories

We do not deny the value of a technological view of work and transfer view of learning. Technological tools and methods exist. Certainly we desire other groups to know about them. But this means more than giving lectures, printing more manuals, or even seeding organizations with experienced players.

We must begin by better understanding organizational successes and failures. We must study what people do, understanding better how they interact with each other, their tools, and their physical environment. We must fundamentally respect the unplanned nature of behavior, the informal network in the community, and the novel, local contributions to standard company procedures. By looking more carefully, we can collect new ideas, respecting improvements, but especially respecting the climate of collaboration and exploration that allowed new contributions and adjustments to take hold and to spread.

We must view knowledge not as a storehouse of facts and procedures that can be inventoried, but as a capacity to interact, to reflect, to innovate. We must shift our view of assessment of productivity from "What have you done for me lately?" to "What surprises might you produce for me tomorrow?" Any investor knows that past performance is no guarantee of future results. The manager must be concerned not only with today's numbers, but tomorrow's promise. This requires sensitivity to trends and orientation, to what the organization is becoming.

Practice, what people actually do and how their communities evolve, is not reducible to descriptions of what they believe or what they do. The patterns we find are always a step removed from the continuously adapting and reconforming interactions. We might characterize a group's activity in terms of a certain vocabulary, a certain orientation, or a plan. But in the very next statement, the next moment of stopping to see what has been done so far, any member of the group is likely to surprise us, producing a new conceptualization, a new way of viewing the goals, and a new value for appraising work. Indeed, every theoretical description of the group is not just a snapshot, but a reflection of an observer's own interactions with the group, and the group itself is a moving target. What is needed is not better or more accurate theories, but an appreciation for the dynamics of the group's ongoing development, an understanding of where new contributions are manifest, and how they change the group's behavior.

6 Change

In this section, I consider broadly the implications of situated cognition for deliberately changing organizations and technology. The idea of participating in a workplace, as researchers, is considered from several perspectives, including the notions of participant observation and participatory design. Studying an ongoing software development effort, I illustrate how change can be understood in terms of interacting communities of practice. In this example, people with diverse roles and goals are brought together to construct formal descriptions of work that will be used as tools in the workplace. In Section 7, I consider what this reveals about tools and representational change in general.

6.1 Two spaces, two commitments

The laboratory research of computer scientists can be characterized as exploring what computers can potentially do. Computer scientists involved in action-oriented research consider the problem of designing and changing social organizations, facilities, and technology in everyday life (Fig. 4) (Brown 1991, Clancey in press c). This effort naturally requires collaboration between people with diverse backgrounds, including workers, managers, graphic designers, organizational specialists, psychologists, programmers, facilities specialists, etc. Research focuses on inventing design processes that facilitate fruitful collaboration (Greenbaum and Kyng 1991). Hence, we need to understand better how learning occurs in such interactions. Consistent with the situated cognition view of knowledge, note that our theories about work, training, and computer tools will be grounded in our experience in changing everyday life (Ehn 1988, Bannon 1991). Our knowledge about learning and tools for learning will first of all be an ability to participate in redesign communities, and only secondarily theories about that redesign (although we may easily create theories without having relevant knowledge and experience).

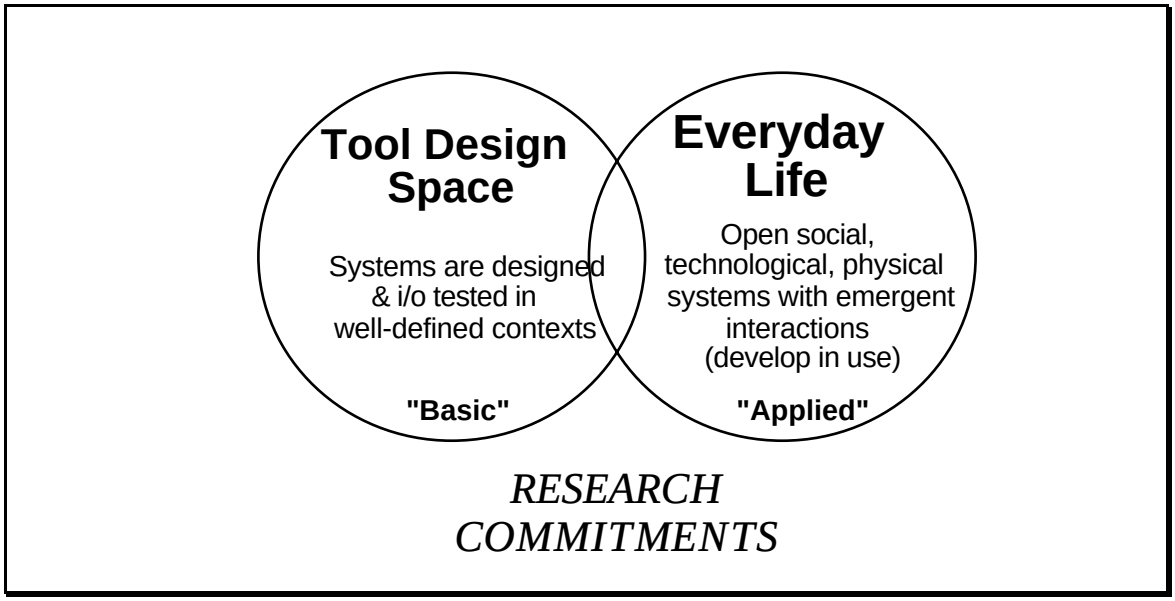

Fig. 4. "Applied" research concerns how to design systems that are part of everyday life.

An overriding theme of our focus on everyday life is making social processes visible in the workplace. This concern appears at two levels: in our models of business processes and in our metatheories of how the modeling and change processes occur. At each of these two levels, we advance our understanding by analyzing current models and contrasting them with what actually occurs in practice. We promote social and technological designs in which we find that collaboration and learning have been successful in the past. In this perspective, we are taking the idea of "participant observation" from ethnography and applying it to business redesign. We are recognizing that understanding how successful change occurs is a research problem, and it is just as basic as what occurs in the laboratory (hence the use of quote marks in Fig. 4). To develop this idea further, I first present my perspective on the ethnographic approach, and then illustrate what we have learned about the social processes of "knowledge acquisition" in constructing expert system software.

6.2 The ethnographic method of reflection

Ethnography is a method of observation based on participating in a community, so one can experience first-hand the habits, new ways of seeing, and evolving sense of values of the group (Jordan 1992). Ethnography brings us closer to experience, to practice, so we can partially avoid the conflict between what people do and what they say they do.

Again, the disparity between saying and doing is not caused by people's inability to remember what representations they used when acting. It makes no sense to speak of accuracy of recall, *per se*, during rationalization because there are no stored recipes, rules, or past decisions to appeal to. What is usually meant is the adequacy of the model of what occurred. Indeed, people can lie, and they may forget what representations influenced their behavior (e.g., what directions they were given). But we must remember that every human action is, at a certain level, new and direct. Saying what we did is a primary activity of theorizing, *describing for the first time*, which is inherently apart from the activity it is about. As Schön (Schön 1979) tells us, "To read the later model back onto the beginning of the process would be to engage in a kind of historical revisionism."

In ethnographic observation, we aren't seeking better descriptions, and hence procedures that we want everyone in the company to rotely follow. Believing that work can be managed better by better describing the work and writing more accurate job descriptions obscures the nature of the work, how such descriptions are actually used, and indeed the kind of work process we should be aiming to identify, reward, and disseminate. We seek instead to appreciate the dynamics by which a given community is continuously reperceiving what they are doing, theorizing about what they have accomplished, and planning what they might do next. We learn to see that the workplace is not a strictly organized and controlled place. Neither knowledge nor *work* is strictly decomposable into pieces. Components of work products are not simply

produced in a linear fashion and arranged on a table and assembled. Pieces of work (e.g., software modules) are coming into being along with the concepts for describing what the group is doing, values for appraising work, and theories that justify what is produced. We focus on how a community reflects, how it adjusts to changing conditions, how it develops day by day. That is, *we focus on how the community learns*.

In an ethnographic study, we might begin by considering how a community of practice interprets standard procedures—how their attitudes, experience, and sense of belonging to a larger community affect what they see in the procedures and how these perceptions organize their future work. Second, we might consider how the complex of interactions in the social, physical, and information processing environments allows for reflection, how individual reflections are composed into a group assessment of progress, and then again, how these representations influence their future work. Third, we consider how the group's ongoing assessments and restructuring interact with similar processes in other communities with which this group interacts. And then again, we consider how these intergroup interactions become manifest in further reflections and reorganizations of future work.

6.3 Example: Knowledge acquisition as a social process

As an example of how we are making social processes of change visible, I present what we have learned about the process of constructing models to be used in a tool for medical clinics[1]. Our study focuses on the representation of a vocabulary of patient symptoms, diagnoses, and treatments that form the basis of an online patient database. In subsequent sections, I draw on this example to again contrast situated cognition with rational theories of change.

Our initial studies of the modeling process indicate that a multidisciplinary group draws from many sources and projects different uses of their model of work (Fig. 5). A clinician may be viewing the model through the advantages for better patient care; a public policy expert may be viewing the model as a means to standardize patient care; a research MD may be focusing on gathering more accurate data in order to test alternative treatments; and a programmer may be focusing on the completeness and consistency of a patient data base. The negotiation of design sessions is therefore a complex interplay of previous and imagined future experiences, of conflicting constraints drawn from different communities of practice, of competing evidence, goals, and future uses.

This model of how representations are created and used can be compared first of all to the original "knowledge acquisition" process of interviewing an expert and codifying his or her knowledge. Of course, AI researchers realized early on that experts have different opinions; it was even proposed as early as 1975 to use MYCIN as a means of bringing national experts together and developing a unified model of infectious disease diagnosis and therapy. However, this diagram illustrates a more profound cause of disagreement: *The participants may agree on the need to develop a common language and models, but they have different uses in mind.* It's not just a matter of having different theories (as if theories are just individual opinions), but theories that develop and function within different communities of practice. Again, models are not be understood with respect to a single, objectively correct view, but within the social processes in which they are created and used. This requires a new analysis of what the participants in this process share: If they share a "common goal" it is at the level of spending time together to develop a representation that has multiple uses and is grounded (perhaps) in some empirical phenomenon (e.g., a population of patients) or tied to a common legal standard.

[1] This study is the work of Kären Wieckert, in collaboration with the T-HELPER project at Stanford University, in the Department of Medical Computer Science.

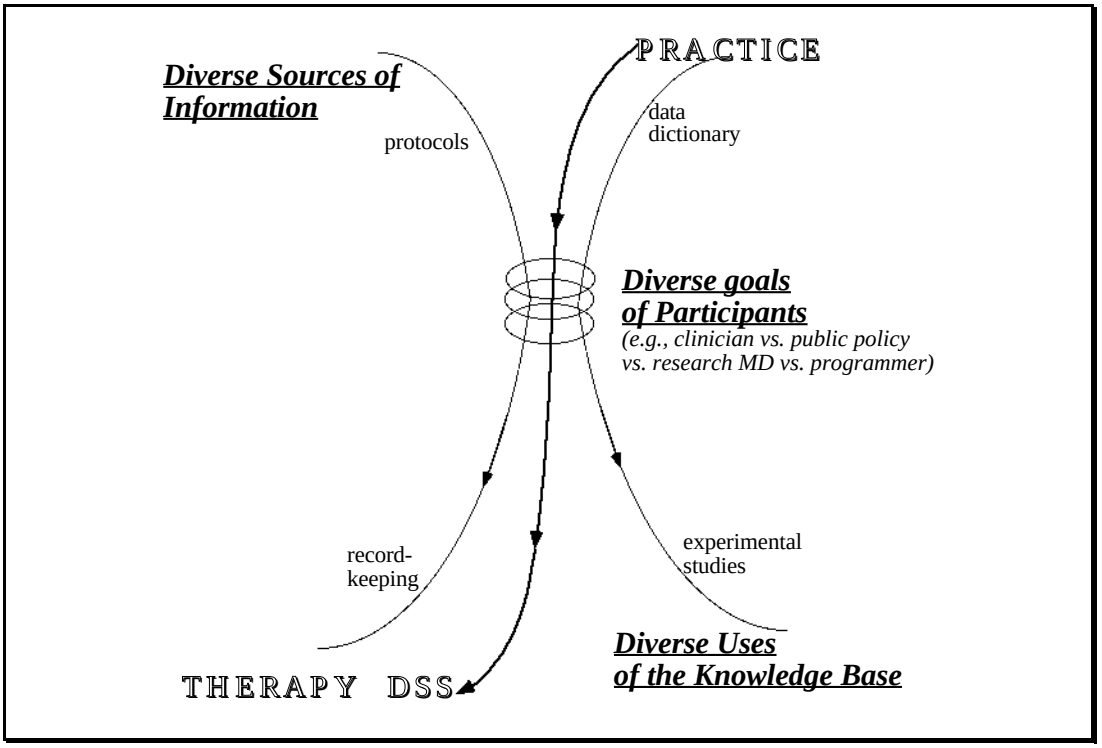

Fig. 5. Divergent view of participants from multiple communities of practice working together on a design project.

6.4 The interplay of practice and theory, workers and researchers

Our current perspective on representation creation and use is distinctly non-linear. It should be contrasted with the idea that a project has a start point, shared involvement, and a common end point (Fig. 6). For example, the idealistic view of participatory design is that of lockstep advancement, with everyone marching together toward a common goal—probably an idea that only those who have never worked in a group on a large project would imagine.

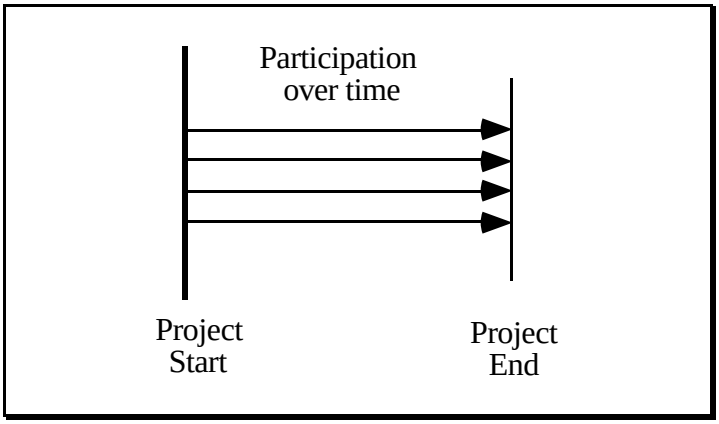

Fig. 6. Idealized, lock-step view of participatory design

The divergent view (Fig. 5) emphasizes that the design process is not linear and cannot be captured from a single perspective. That is, the practice of design itself, as a part of everyday life, remains as open to

change and interpretation as the world of the workplace that we seek to study and influence. This points out a weakness in a simple iterative view, cycling between design and practice (Fig. 7), what has been called "reciprocal evolution" (Allen 1991). In fact, project participants do not remain together in a single community of practice during design cycles, nor are influences and design effects contained within one community, as a cycling line might suggest. Furthermore, theorizing occurs as part of the everyday process of making decisions in work.

The idea of continuous learning emphasizes that the redesign process occurs within everyday work. Goals, policies and values are changing within every act in the workplace (Wynn 1991). That is, learning is occurring not just when we reflect, articulating models about work, but with every perception about what is happening in negotiating with customers and meeting with fellow employees. Judgment is required in at least minor ways even in determining that a situation is "routine." Redesign is in some sense occurring all the time. A key insight of workplace redesign efforts is putting into practice a process by which change that occurs as a matter of course everyday is appropriately valued and related to formal policies.

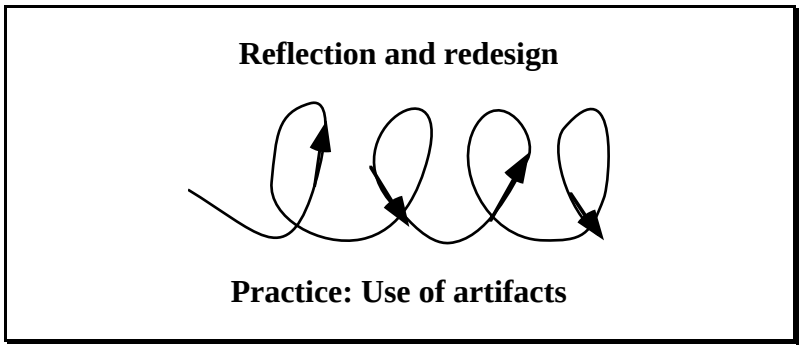

Fig. 7. Iterative model of reciprocal evolution, cycling between use of models and tools in practice and theorization with new design.

A better view of the design process combines the "influences and effects" diagram (Fig. 5) with the iterative process of reciprocal evolution. This revised diagram (Fig. 8) shows time moving from left to right, with new theoretical influences coming in over time from outside (arrows coming in from above), as well as from the workplace (arrows pointing upwards). In addition, effects diverge from initial intentions, including serendipitous uses (arrows coming from below) and effects in outside communities (arrows pointing downwards).

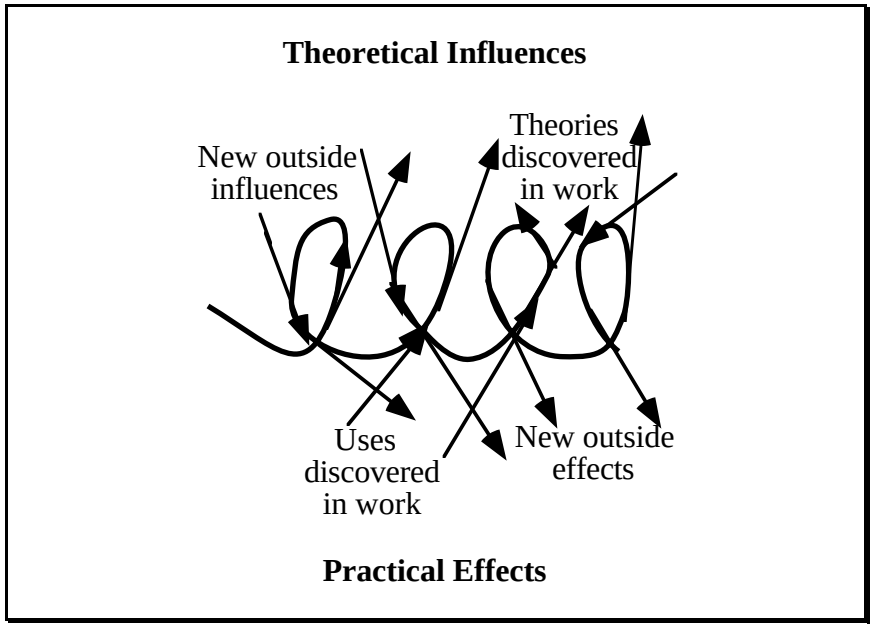

Fig. 8. Process map of participatory design involving multiple communities of practice.

Stepping back, we need to acknowledge that there are different gradations of change. At a certain grainsize, we are not merely "doing our work," but changing the business. Despite acknowledging that change occurs all the time, we view business redesign as something deliberate and special that involves other players in new activities, using specialized modeling methods and tools. One way to understand this is to consider how workers and researcher-designers are deliberately brought together in a new community of practice through participant observation and participatory design.

The social sciences are currently contributing two complementary ideas about the redesign process: participant observation (designers participate in users' world) and participatory design (users participate in designers' world). The practice of constructing models of work and theories of business reengineering (Hammer 1990) mediates these two worlds of practice and redesign (Fig. 9). Crucially, there are two practices: the practice of the workplace we seek to change and the practice of redesign. The redesign practice is a metapractice in the sense that it seeks to describe the practice of work and change it. So we want to acknowledge that the worlds of workers and designers are different, and there is need to develop new processes for working together, so that the workplace is included in the world of redesign.

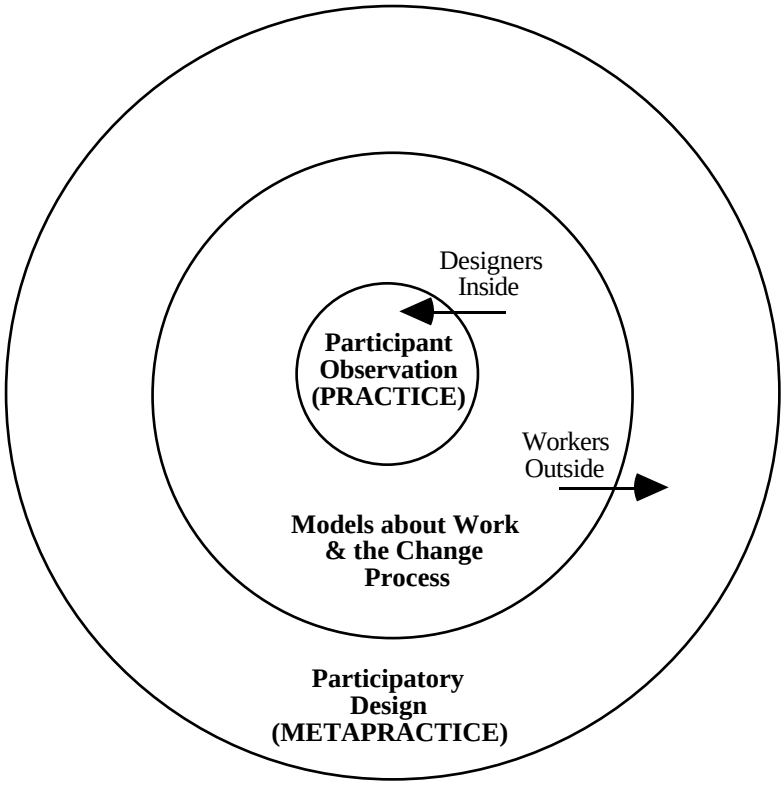

Fig. 9. Relation of participatory design to participant observation.

6.5 Individuals in changing communities of practice

To recap, we need to better understand work so we can understand how to bring about effective change. This understanding involves better understanding the nature and role of models, how identity is constructed through participation in communities of practice (Wenger in preparation), and how emergent understandings and interpersonal interactions can be encouraged and recognized. A key idea is that individuals contribute to group efforts by reflecting on what they say and do in *other communities* in which they participate. The experience of difference—which starts with a feeling of discord—is the source for articulating theories about past activities, as well as understanding emergent goals and values in the new group.

The example of vocabulary development from the T-Helper project, suggesting the divergent view of representational change (Fig. 5), is broadly applicable. First, we see that it fits different "capture and disseminate" activities:

- **standardization across communities of practice** (e.g., software reuse within divisions of a company, between research and manufacturing, and internationally);
- **curriculum design for school and workplace training** (e.g., collaboration between scientists, teachers, and educational researchers);
- **preserving corporate memory and intellectual property** (e.g., corporate efforts to "capture expertise" in materials and designs before people retire).

In short, this divergent picture describes what happens whenever an effort is made to bring together people with diverse roles to create a common vocabulary and models (e.g., curricula, knowledge bases, policies and procedures, national standards). Although project participants may create a single model as a result of a collaboration, the use of the model and its meaning will be prone to divergent interpretations in different communities of practice over time. Indeed, the modelling effort focuses articulation of past experience, as participants come to understand how constraints from different domains interact and

generate new ideas for synergistic interactions between different communities. Ideally, such reflections will be facilitated by researcher-designers in a business reengineering effort.

Furthermore, the diagrams (Figures 5 and 7) show what happens *within* the workplace itself. Besides describing what is happening in Stanford's laboratories as clinicians, research MDs, and computer scientists collaborate, the diagrams show what is happening inside the medical clinic itself between nurses, physicians, interns, and patients. That is, it is helpful to realize that the home turf of a single participant in a multidisciplinary team appears to him or her like another multidisciplinary team with divergent experiences and divergent future-oriented intentions.

The thrust of Fig. 5 is to view individuals as the locus of overlapping interests, which have their origins and ultimate meanings in participation in diverse communities of practice. Consequently, the researcher community itself is multivocal, with diverse influences and effects. For example, we are well aware that the Institute for Research on Learning (IRL) is a multidisciplinary group (Clancey 1992a). Each IRL researcher is always participating in many communities of practice (in academia, the schools, and businesses). New ideas at IRL develop from individual experiences of "reseeing" their previous experiences and future activities through the evolving language and world models of the IRL community. Thus, another way to express the divergent view of Fig. 5 is that a community of practice involves individuals interacting together, but *each participant looks backwards and forwards into the realm of personal experience.*

To continue the application of this idea to IRL, the effects we produce will occur in different communities *outside*, through the articulations and participation of IRL individuals. Even this picture gets complicated when we realize that an IRL anthropologist and computer scientist may attend a single-discipline workshop together, changing the experience of everyone all around.

Finally, the diagram of divergent sources and effects is useful for understanding strategic partnerships between research and business organizations. Here the inner ring of interaction (the loops at the center of Fig. 5) can involve workshops, research papers, and focused redesign projects. In such partnerships, we share the experience of a need for change; we negotiate what tools and methodologies will be used during our activities together. We may share a vision for what kind of artifacts and social organizations we wish to see develop from our work together. But always the redesign process will be facilitated by becoming aware of our diverse past experiences, values, and imagined results of our joint work. This analysis shows why scenario depiction is especially valuable. Story telling is seen as a way of grounding design ideas in individual experiences—to convey what life is like in another place as well as the impact of social and technological change.

7 Representational Tools

To illustrate some additional, subtle aspects about representation creation and use, consider how Fig. 5 developed. I drew this first as a way of "holding in place" the many players and constraints that Kären Weickert observed in her study of T-Helper design sessions. (This idea of "holding in place" relations via an arrangement of physical things is developed by Jeanne Bamberger (1991) in her analysis of children's representations of tunes.) I later used the figure for a research project presentation at IRL. In the course of that meeting, a video about reciprocal evolution was shown (Fig. 7). The curving lines of Fig. 7 reminded me of my own Fig. 5, and I realized that they could be overlaid (Fig. 8), and then contrasted with other views (Fig. 6). In this respect, the ideas of Fig. 5 were developed as commentary on the shortcomings of other depictions of the design process. A participant at the IRL research presentation commented that Fig. 8 was a "process map," leading me to understand that it showed Fig. 5 over time. That is, the diverse sources and purposes of a representation aren't all known at once, but articulated over time, as the representational artifact develops and is used. The process of representational change is shown here *as an inherently social feedback process*. Notably, the evolving system is open, as new constraints are imposed from above by outside communities of practice (e.g., by funding or political agencies) and opportunities for new uses below become known.

In effect, the diagram serves as a way for focusing and reperceiving relations that were not originally represented in the diagram itself. The idea of outside communities of practice, serendipitous use, and feedback were added later. By bringing in the reciprocal evolution idea, I realized that the IRL project in which that idea was developed could be described by Figures 5 and 8. This suggested that I apply Fig. 5 to other IRL projects, yielding analogies with curriculum development, software reuse, and corporate memory. After presenting an earlier, transparency version of this figure at an AI conference, I decided to drop an inner caption ("A knowledge base is not objective"), because it seemed to be a side issue for that

audience. This reminding, adaptation, and theorization process subsumes diverse ideas and projects under a common way of seeing (Fig. 5), thus generalizing my understanding of the competing forces and processes involved in many settings. I came to see each project I reconsidered as involving multidisciplinary participants with diverse goals, but nevertheless constituting a community of practice. Thus the idea of a community of practice was changed. A similar experience occurred in the creation of Fig. 10.

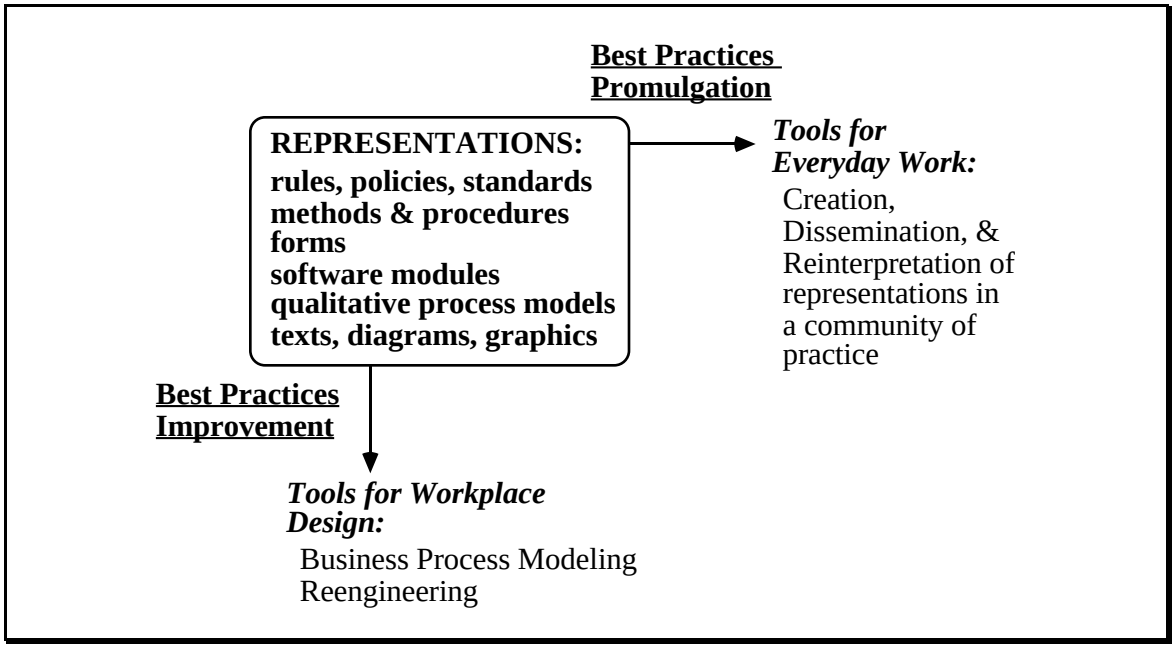

Fig. 10. Representations as tools for inquiry

This figure was conceived at a whiteboard during another IRL presentation for a corporate client. I wanted to convey our interest in how representations were created and used, showing that we were focusing on two domains of concern: tools for workers and tools for work redesign. The list of representations was generated on the fly, to illustrate how broadly we view the idea. The labels "best practices improvement" and "promulgation" actually came from my listener, as he described what he heard me saying. Schön (1987) calls this process of reconceptualization "backtalk" when it occurs within an individual's experience in reperceiving a previously created representation; it can obviously occur as well between individuals.

This example illustrates again how a picture can be used to "hold in place" different views and different activities. The "best practices" titles and the figure caption were added the next time I presented the figure, at a research conference. In effect, multiple interpretations and purposes are "held in place" by the representation itself as it is moved between and used in diverse settings. The cumulative, improvisatory and serendipitous process by which this representation developed are typical, illustrating how concepts develop in conversations that have different purposes and occur in different places over time. Representation "reuse" is therefore a process of adding elements and "reseeing" other activities, so the evolving model subsumes different sources and purposes—precisely what Fig. 5 aims to show.

In all of these diagram modifications and reconceptualizations, I was aware that I was improving prevalent views about the concept of "community of practice," multidisciplinary research groups, and how knowledge acquisition proceeds. That is, the driving force was not just a generalization which subsumed more and more territory, but a sense of being a deeper, more coherent descriptor of complex interactions occurring in any given domain. Each step is sensed as being constructive, in that something more generally useful is being developed. This theme is captured by the title of Engeström's (1987) dissertation, *Learning by Expanding.*

8 Conclusion

In *promoting* organizational learning, we are designing interactions—between people, machines, and their environment. We are developing a socio-technical system of social, physical, and information processing environments. Just as the architect, the writer, and the programmer cannot strictly predict or plan what they will produce, the organizational theorist and manager cannot strictly predict what workplace designs will evolve. Management and all participants must accept the fact that every intervention, every restructuring, every new tool for improved performance, is a naturally occurring experiment. The ongoing process of perceiving how we are doing, finding patterns, and justifying developments through new theories, is a *process of inquiry and discovery*. The manager, like the architect, the writer, and the programmer, must be open to serendipity and unashamed by unexpected interactions. The developing organization is like the emerging building design, the journal article, or the complex software system—there will always be surprises. What happens in practice is always one step ahead of theory.

Acknowledgments

I am indebted to my colleagues at the Institute for Research on Learning who have introduced and discussed many of these ideas: Kären Wieckert, Etienne Wenger, Susan Stucky, Jeremy Roschelle, Brigitte Jordan, Charlotte Linde, Ted Kahn, Peter Henschel, and Penny Eckert. Funding for this research is partially provided by grants from the Xerox Foundation and the NSF Knowledge Systems and Cognitive Models program, and project contracts with NYNEX and the Xerox Corporation.

References

Allen, C., Pea, R., de Vet, J., de Vogel, J. 1991 *Picasso Project Final Report.* IRL Technical Report, November.

Bamberger, J. (1991). *The mind behind the musical ear.* Cambridge, MA: Harvard University Press.

Bamberger, J. and Schön, D.A. (1983). Learning as reflective conversation with materials: Notes from work in progress. *Art Education*, March, pp. 68-73.

Bannon, L. 1991. From human factors to human actors, in J. Greenbaum and M. Kyng (editors), *Design at Work: Cooperative Design of Computer Systems*. Hillsdale, NJ: Lawrence Erlbaum. pp. 25-44.

Barr, A. Bennett, J. And Clancey, W.J. (1979) Transfer of expertise: A theme for AI research. Working Paper HP-79-11, Stanford University Department of Computer Science.

Bartlett, F. C. [1932] 1977. *Remembering--A Study in Experimental and Social Psychology*. Cambridge: Cambridge University Press.

Bateson, G. 1972. *Steps to an Ecology of Mind.* New York: Ballentine Books.

Bickhard, M. H. and Terveen, L. (in preparation). *The Impasse of Artificial Intelligence and Cognitive Science.*

Brown, J.S. 1991. Research that reinvents the corporation. *Harvard Business Review*, January-February, 102-111.

Clancey, W.J. 1989. Viewing knowledge bases as qualitative models. *IEEE Expert*, (Summer 1989):9-23.

Clancey, W.J. 1991a. Situated cognition: Stepping out of representational flatland. *AI Communications*, **4**(2/3):107-112.

Clancey, W.J. 1991b. Review of Rosenfield's "The Invention of Memory," *Artificial Intelligence*, **50**(2):241-284, 1991.

Clancey, W.J. 1992a. Overview of the Institute for Research on Learning. In proceedings of *CHI, 1992* (New York: ACM). Monterey, CA. pp. 571-572.

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. (in press a). Situated action: A neuropsychological interpretation: Response to Vera and Simon. To appear in *Cognitive Science*.

Clancey, W.J. (in press b). Notes on "Epistemology of a rule-Based expert system." To appear in *Artificial Intelligence*.

Clancey, W.J. (in press c). Guidon-Manage revisited: A socio-technical systems approach. Submitted to the *Journal of AI and Education*.

Dewey, J. [1896] (1981). The reflex arc concept in psychology. *Psychological Review*, III:357-70, July. Reprinted in J.J. McDermott (ed), *The Philosophy of John Dewey*, Chicago: University of Chicago Press, pp. 136-148.

Dewey, J. [1938] (1981). The criteria of experience. In *Experience and Education*, New York: Macmillan Company, pp. 23-52. Reprinted in J.J. McDermott (ed), *The Philosophy of John Dewey*, Chicago: University of Chicago Press, pp. 511-523.

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

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

Ehn, P. 1988. *Work-Oriented Design of Computer Artifacts*, Stockholm: Arbeslivscentrum.

Engeström, Y. 1987. *Learning by Expanding.* Orienta-Konsultit, Helsinki. Unpublished dissertation.

Gasser, L. 1991. Social conceptions of knowledge and action, *Artificial Intelligence*, **47**(1-3)107-138., January.

Greenbaum, J. and Kyng, M. 1991. *Design at Work: Cooperative Design of Computer Systems*. Hillsdale, NJ: Lawrence Erlbaum.

Hammer, M. Reengineering work: Don't automate, obliterate. 1990. *Harvard Business Review*, 90(4): 104-112.

Hirschorn, L. 1984. *Beyond Mechanization: Work and Technology in the Postindustrial Age.* Cambridge, MA: The MIT Press.

Jordan, B. 1992. New research methods for looking at productivity in knowledge-intensive organizations. In H Van Dyke Parunak (editor) *Productivity in Knowledge-Intensive Organizations: Integrating the Physical, Social, and Informational Environments*. Working papers of the Grand Rapids Workshop, April 8-9, 1992. Industrial Technology Institute Technical Report 92-01. Ann Arbor, Michigan. pp. 194-216.

Kling, R. 1991. Cooperation, coordination and control in computer-supported work. *Communications of the ACM*, **34**(12)83-88.

Kukla, C.D., Clemens, E.A., Morse, R.S., and Cash, D. (in press). An approach to designing effective manufacturing systems. To appear in *Technology and the Future of Work.*

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.

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

Roschelle, J, and Clancey, W. J. (in press). Learning as Social and Neural. To appear in *Educational Psychologist.*

Rosenfield, I. (1988). *The Invention of Memory: A New View of the Brain.* New York: Basic Books.

Ryle, G. (1949). *The Concept of Mind.* New York: Barnes & Noble, Inc.

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

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

Scott, A. C., Clancey, W.J., Davis, R., and Shortliffe, E.H. (1984) Methods for generating explanations. In Bruce G. Buchanan and Edward H. Shortliffe (editors), *Rule-Based Expert Systems: The MYCIN Experiments of the Heuristic Programming Project.* Reading, MA: Addison-Wesley, pp. 338-362.

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.

Slezak, P. 1992. Situated cognition: Minds in machines or friendly photocopiers? Presented at the *McDonnell Foundation Conference on The Science of Cognition,* Sante Fe. Unpublished manuscript.

Smith, B.C. (1991). The owl and the electric encyclopedia. *Artificial Intelligence* **47**(1-3):251-288.

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

Tyler, S. 1978. The Said and the Unsaid: Mind, Meaning, and Culture. New York: Academic Press.

Wallace, J.W. and Shortliffe, E.H. (1984) Customizing explanations using causal knowledge. In Bruce G. Buchanan and Edward H. Shortliffe (editors), *Rule-Based Expert Systems: The MYCIN Experiments of the Heuristic Programming Project.* Reading, MA: Addison-Wesley, pp. 371-388.

Wenger, E. (in preparation) *Toward a Theory of Cultural Transparency*. Cambridge: Cambridge 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 (eds), *Design at Work: Cooperative design of computer systems.* Hillsdale, NJ: Lawrence Erlbaum Associates, pp. 45-64.

Zuboff, S. 1988. In the Age of the Smart Machine: The future of work and power. New York: Basic Books.