Human Studies 11:305–325 (1988)

Representing practice in cognitive science*

LUCY A. SUCHMAN
System Sciences Laboratory, Xerox Corporation, 3333 Coyote Hill Road, Palo Alto, CA 94304, USA

1. Introduction

Recent social studies of science take as a central concern the relationship between various representational devices and scientific practice (see Tibbett, 1988, and Lynch and Woolgar, 1988, in this special issue.) Representational devices include models, diagrams, formulae, records, traces and a host of other artifacts taken to stand for the structure of an investigated phenomenon. Several premises underlie the study of the relation of such devices to scientific practice. First, that it is through these devices that the regularity, reproducibility and objectivity both of phenomena and of the methods by which they are found are established. Second, that representational devices have a systematic but necessarily contingent and *ad hoc* relation to scientific practices. And third, that representational technologies are central to how scientific work gets done. To date science studies have concentrated on the physical and biological sciences (see for example Collins, 1985; Latour and Woolgar, 1979; Garfinkel et al., 1981; Knorr-Cetina, 1981; Lynch, 1985; Lynch et al., 1983). This paper joins with others (Woolgar, 1985; Collins, 1987) in directing attention to a new arena of scientific practice; namely, cognitive science.

* This paper and the work that it reports have benefited substantially from discussions with my collaborators Randy Trigg and Brigitte Jordan. For developing observations on the use of whiteboards I am indebted to Randy Trigg, John Tang and to members of the Interaction Analysis Group at Xerox PARC; Christina Allen, Stephanie Behrend, Sara Bly, Tom Finholt, George Goodman, Austin Henderson, Brigitte Jordan, Jane Laursen, Susan Newman, Janice Singer, and Debbie Tatar.

In turning to cognitive science as a subject of sociological inquiry we are faced with an outstanding issue concerning the relation of representation to practice. The issue can be formulated, at least initially, as follows: Ethnomethodological studies of the physical and biological sciences eschew any interest in the adequacy of scientific representations as other than a members' concern. The point of such studies is specifically *not* to find ironies in the relation between analysts' constructions of the phenomenon and those of practitioners (Garfinkel, 1967:viii; Woolgar, 1983.) Rather, the analyst's task is to see how it is that practitioners come to whatever understanding of the phenomenon they come to as the identifying accomplishment of their scientific practice. In turning to cognitive science, however, one turns to a science whose phenomenon of interest itself *is* practice. For cognitive science theorizing, the object is mind and its manifestation in rational action. And in designing so-called intelligent computer systems, representations of practice – expert/novice instruction, medical diagnosis, electronic troubleshooting and the like – provide the grounds for achieving rationality in the behavior of the machine.

In this paper I consider two distinct but related conceptions of the notion of "representing practice" with respect to cognitive science, through a discussion of two studies. The first study, recently completed, looks at the ways in which cognitive scientists depict the nature and operation of social practice, as part of their own agenda for the design of intelligent machines. These ways include the representation of practice as logical relations between conditions and actions, and the design of artifacts that embody such representations. The second study, just underway, looks at the representational practices of cognitive scientists, through a detailed analysis of researchers engaged in collaborative design work at a "whiteboard."[1] Together these studies consider representing practice as both the object of cognitive scientists' work and as sociology's subject matter.

2. Artificial intelligence and interactional competence

The term "cognitive science" came into use in the 1970s to refer to a convergence of interest over the preceding 20 years among neurophysiologists, psychologists, linguists, cognitive anthropologists, and later computer scientists, in the possibility of an integrated science of cognition (for an enthusiastic history see Gardner, 1985). The commitment both to cognition and to science was, at least initially, an important part of the story. At the turn of the century, the recognized method for studying human mental life was introspection and, insofar as introspection was not amenable to the emerging canons of scientific method, the study of cognition seemed doomed to be irremediably unscientific. In reaction to that prospect, the behaviorists posited that human action should be investigated in terms of publicly observable, mechanistically describable relations between the organism and its environment. In the name of turning cognitive studies into a science, the study of cognition as the study of something apart from conditioned behavior was effectively abandoned in mainstream psychology.

Cognitive science, in this respect, was a project to bring thought, or meaning, back into the study of human action while preserving the commitment to scientism. Cognitive science reclaims mentalist constructs like beliefs, desires, intentions, planning and problem-solving. Once again human purposes are the basis for cognitive psychology, but this time without the unconstrained speculation of the introspectionists. The study of cognition is to be empiricized not by a strict adherence to behaviorism, but by the use of a new technology; namely, the computer.

The branch of cognitive science most dedicated to the computer is Artificial Intelligence. The sub-field of AI arose as advances in computing technology were tied to developments in neurophysiological and mathematical theories of information. The requirements of computer modeling, of an "information processing psychology," seem both to make theoretical sense and to provide the accountability that will make it possible to pursue a science of otherwise inaccessible mental phenomena. If underlying mental processes can be modelled on the computer so as to produce the right outward behavior, the argument goes, the model can be viewed as having passed at least a sufficiency test of its psychological validity.

A leading idea in cognitive science is that mind is best viewed as neither substance nor as insubstantial, but as an abstractable structure implementable in any number of possible physical substrates. Intelligence, on this view, is only incidentally embodied in the neurophysiology of the human brain. What is essential about intelligence can be abstracted from that particular, albeit highly successful substrate and embodied in an unknown range of alternative forms. The commitment to an abstract, disembodied account of cognition, on the one hand, and to an account of cognition that can be physically embodied in a computer, on the other, has led to a view of intelligence that takes it to be first and foremost mental operations and only secondarily, and as an epiphenomenon, the "execution" of situated actions.

While intelligence is taken by cognitive science, without much question, to be a faculty of individual minds, the measure of success for the AI project is and must be an essentially social one. Evidence for intelligence, after all, is just the observable rationality of the machine's output relative to its input. This sociological basis for machine intelligence is implicit in the so-called Turing Test, by now more an object of cognitive science folklore than a part of working practice. Turing (1950) argued that if a machine could be made to respond to questions in such a way that a person asking the questions could not distinguish between the machine and another human being, the machine would have to be described as intelligent.[2] Turing expressly dismissed as a possible objection to his proposed test that, although the machine might succeed in the game, it could succeed through means that bear no resemblance to human thought. Turing's contention was precisely that success at performing the game, regardless of mechanism, is sufficient evidence for intelligence (1950:435). The Turing test thereby became the canonical form of the argument that if two information-processors, subject to the same input stimuli, produce indistinguishable output behavior, then regardless of the identity of their internal operations one processor is essentially equivalent to the other.

The lines of controversy raised by the Turing test were drawn over a family of programs developed by Joseph Weizenbaum in the 1960s under the name ELIZA, designed to support "natural language conversation" with a computer (Weizenbaum, 1983).

Of the name ELIZA, Weizenbaum writes:

> Its name was chosen to emphasize that it may be incrementally improved by its users, since its language abilities may be continually improved by a "teacher." Like the Eliza of *Pygmalion* fame, it can be made to appear even more civilized, the relation of appearance to reality, however, remaining in the domain of the playwright. (p. 23)

Anecdotal reports of occasions on which people, approaching the teletype to one of the ELIZA programs and believing it to be connected to a colleague, engaged in some amount of "interaction" without detecting the true nature of their respondent led many to assert that Weizenbaum's program had passed a simple form of the Turing test. Weizenbaum himself, however, denied the intelligence of the program on the basis of the underlying mechanism, which he described as "a mere collection of procedures" (p. 23):

> The gross procedure of the program is quite simple; the text [written by the human participant] is read and inspected for the presence of a *keyword*. If such a word is found, the sentence is transformed according to a *rule* associated with the keyword, if not a content-free remark or, under certain conditions, an earlier transformation is retrieved. The text so computed or retrieved is then printed out. (p. 24, original emphasis)

The design of the ELIZA programs exploits the natural inclination of people to make use of the "documentary method of interpretation" (see Garfinkel, 1967:Ch. 3): to take appearances as evidence for, or the document of an ascribed underlying reality while taking the reality so ascribed as a resource for the interpretation of the appearance. In a contrived situation that, though designed independently and not with them in mind, closely parallels both the "Turing test" and encounters with Weizenbaum's ELIZA programs, Garfinkel set out to test the documentatry method in the context of counseling. Students were asked to direct questions concerning their personal problems to someone they knew to be a student counselor, seated in another room. They were restricted to questions that could take yes/no answers, and those answers were given by the counselor on a random basis. For the students, the counselor's answers were motivated by the questions. That is to say, by taking each answer as evidence for what the counselor "had in mind," the students were able to find a deliberate pattern in the exchange that explicated the significance and relevance of each new response as an answer to their question:

> The underlying pattern was elaborated and compounded over the series of exchanges and was accommodated to each present "answer" so as to maintain the "course of advice," to elaborate what had "really been advised" previously, and to motivate the new possibilities as emerging features of the problem. (1967:90)

The ELIZA programs and Garfinkel's counselor experiment demonstrate the generality of the documentary method and the extent to which the meaning of actions is constituted not by actors' intentions but through the interpretive activity of recipients. Users of ELIZA and Garfinkel's students are able to construct out of the mechanical "responses" of the former and the random "responses" of the latter a response to *their* questions. This clearly poses a problem for Turing test criteria of intelligence and for the test of intentionality proposed by Dennett (1978:Ch. 1), who argues that intentional systems are just those whose behavior is conveniently made sense of in intentional terms. ELIZA and the counselor clearly meet that criterion, yet they show at the same time the inadequacy of that measure for intentional interaction. The injunction for the counselor is precisely that he or she *not* interact. The counselor's "responses" are not responses to the student's questions, nor are the interpretations that the student offers subject to any remediation of misunderstanding by the counselor. Or rather, there *is no* notion of misunderstanding, insofar as in the absence of the counselor's point of view any understanding on the part of the student that "works" will do. In human communication in contrast there are two "students," *both* engaged in making sense out of the actions of the other, in making their own actions sensible, in assessing the senses made, and in looking for evidence of misunderstanding. It is just this highly contingent and reciprocal process that we call "interaction."

For behavior to be not only intelligible but intentional, it seems, there must be something about the *actor* that gives her action its senses. As participants in interaction we see our work not as the single-handed construction of meaning but as a kind of reading off from the action of the actor's underlying intent. This common sense view is adopted by cognitive scientists, who take actions to reflect the underlying cognitive mechanism or plans that generate them. The representation of those mechanisms or plans, on this view, is effectively the representation of practice.

3. Plans as determinants of action

The identification of intent with a plan-for-action is explicit in the writing of philosophers of action supportive of artificial intelligence research like Margaret Boden (1973) who writes:

> unless an intention is thought of as an action-plan that can draw upon background knowledge and utilize it in the guidance of behavior one cannot understand how intentions function in real life. (pp. 27–28)

A logical extension of Boden's view, particularly given an interest in rendering it more computable, is the view that plans actually are prescriptions or instructions for action. Traditional sociology similarly posits an instrumentally rational actor whose choice among alternative means to a given end is mediated by norms of behavior that the culture provides – an actor Garfinkel dubbs the "cultural dope":

> By "cultural dope" I refer to the man-in-the-sociologist's-society who produces the stable features of the society by acting in compliance with preestablished and legitimate alternatives of action that the common culture provides. (1967: 68)

Cognitive science embraces this normative view of action in the form of the *planning model.* The model assumes that in acting purposefully actors are constructing and executing plans, condition/action rules, or some other form of representation that controls, and therefore must be prerequisite to, actions-in-the-world. An early and seminal articulation of this view came from Miller, Galanter and Pribram, in *Plans and the Structure of Behavior* (1960):

> Any complete description of behavior should be adequate to serve as a set of instructions, that is, it should have the characteristics of a plan that could guide the action described. When we speak of a plan ... the term will refer to a hierarchy of instructions ... *A plan is any hierarchical process in the organism that can control the order in which a sequence of operations is to be performed.*
>
> A Plan is, for an organism, essentially the same as a program for a computer ... we regard a computer program that simulates certain features of an organism's behavior as a theory about the organismic Plan that generated the behavior.
>
> Moreover, we shall also use the term "Plan" to designate a rough sketch of some course of action ... as well as the completely detailed specification of every detailed operation ... We shall say that a creature is executing a particular Plan when in fact that Plan is controlling the sequence of operations he is carrying out. (p. 17, original emphasis)

With Miller et al., the view that purposeful action is planned assumes the status of a psychological theory compatible with the interest in a mechanistic, computationally tractable account of intelligent action. The identification of intentions with plans, and plans with programs, leads to an identification of representation and action that supports the notion of "designing" intelligent actors. Once representations are taken to control human actions, the possibility of devising formalisms that could specify the actions of "artificial agents" becomes plausible. Actions are described by preconditions, that is, what must be true to enable the action, and postconditions, what must be true after the action has occurred. By improving upon or completing our common sense notions of the structure of action, the structure is now represented not only as an empirically ascertained set of behavioral patterns or a plausible sequence of actions but as an hierarchical plan. The plan reduces, moreover, to a detailed set of instructions that actually serves as the program that controls the action. At this point, the plan as stipulated becomes substitutable for the action, insofar as the action is viewed as derived from the plan. And once this substitution is done, the theory is self-sustaining: the problem of action is *assumed* to be solved by the planning model, and the task that remains is to refine the model.

4. Plans as resources for action

Taken as the determinants of what people do, plans provide both a device by which practice can be represented in cognitive science and a solution to the problem of purposeful action. If we apply an ethnomethodological inversion[3] to the cognitive science view, however, plans take on a different status. Rather than describing the mechanism by which action is generated and a solution to the analysts' problem, plans are common sense constructs produced and used by actors engaged in everyday practice. As such, they are not the solution to the problem of practice but part of the subject matter. While plans provide useful ways of talking and reasoning about action, their relation to the action's production is an open question.

One can see clearly the descriptive or interpretive function of talk about intentions, and its problematic relation to production, in the case of our talk about babies.[4] Nursing babies are very good at finding milk. If you touch a baby on the cheek, it will move its head in the direction of the touch. Similarly, if you put your finger on the baby's lips, it will suck. In some sense we would say, in describing the baby's behavior, that the baby "knows how to get food." Yet to suggest that the baby "has a goal" of finding food in the form of a representation of the actions involved, or performs computations on data structures that include the string "milk" to reach that goal, seems somehow implausible. It is not that all behavior can be reduced to the kind of reflex action of a nursing baby, or that some behavior is not importantly symbolic. The point is that the intentional description, however useful, doesn't distinguish those things.

At the same time, such description is clearly a resource. Our imagined projections and retrospective reconstructions are the principal means by which we catch hold of situated action and reason about it, while situated action itself is essentially transparent to us as actors. In contemplating the descent of a problematic series of rapids in a canoe, for example, one is very likely to sit for a while above the falls and plan one's descent.[5] So one might think something like "I'll get as far over to the left as possible, try to make it between those two large rocks, then backferry hard to the right to make it around that next bunch." A great deal of deliberation, discussion, simulation, and reconstruction may go into such a plan and to the construction of alternate plans as well. But in no case – and this is the crucial point – do such plans control action in any strict sense of the word "control." Whatever their number or the range of their contingency, plans stop short of the actual business of getting you through the falls. When it really comes down to the details of getting the actions done, *in situ*, you rely not on the plan but on whatever embodied skills of handling a canoe, responding to currents and the like are available to you. The purpose of the plan, in other words, is not literally to get you through the rapids, but rather to position you in such a way that you have the best possible conditions under which to use those embodied skills on which, in the final analysis, your success depends.

The planning model takes off from our common sense preoccupation with the anticipation of action and the review of its outcomes and attempts to systematize that reasoning as a model for situated practice itself. These examples, however, suggest an alternative view of the relationship between plans, as representations of conditions and actions, and situated practice. Situated practice comprises moment-by-moment interactions with our environment more and less informed by reference to representations of conditions and of actions, and more and less available to representation themselves. The function of planning is not to provide a specification or control structure for such local interactions, but rather to orient us in a way that will allow us, through the local interactions, to respond to some contingencies of our environment and to avoid others. As Agre and Chapman put it "[m]ost of the work of using a plan is in determining its relevance to the successive concrete situations that occur during the activity it helps to organize" (1987a).[6] Plans specify actions just to the level that specification is useful; they are vague with respect to the details of action precisely at the level at which it makes sense to forego specification and rely on the availability of a contingent and necessarily *ad hoc* response. Plans are not the determinants of action, in sum, but rather are resources to be constructed and consulted by actors before and after the fact.

5. Engineering interaction

Adherents of the planning model in AI view interaction just as an extension of the planning problem from a single individual to two or more individuals acting in concert. In a 1983 paper on recognizing intentions, James Allen puts it this way:

> Let us start with an intuitive description of what we think occurs when one agent A asks a question of another agent B which B then answers. A has some goal; s/he creates a plan (plan construction) that involves asking B a question whose answer will provide some information needed in order to achieve the goal. A then executes this plan, asking B the question. B interprets the question, and attempts to infer A's plan (plan inference). (p. 110)

The problem for interaction, on this view, is to recognize the actions of others as the expression of their underlying plans. The appropriateness of a response turns on that analysis, from which, in turn the hearer then adopts new goals and plans her own utterances to achieve them. On this model, Searle's speech act theory seems to offer some initial guidelines for computational models of communication. Searle's conditions of satisfaction for the successful performance of speech acts are read as the speech act's "preconditions," while its illocutionary force is the desired "effect:"

> Utterances are produced by actions (speech acts) that are executed in order to have some effect on the hearer. This effect typically involves modifying the hearer's beliefs or goals. A speech act, like any other action, may be observed by the hearer and may allow the hearer to infer what the speaker's plan is. (Allen, 1983:108)

Given this view, the design of interactive computer systems affords a kind of natural laboratory in which to see what happens when artifacts embodying the planning model of action encounter people engaged in situated activity. The practical problem with which the designer of an interactive machine must contend is how to ensure that the machine responds appropriately to the user's actions. The design strategy for plan-based systems is essentially to *specify* an appropriate linkage between user actions and machine states. This strategy assumes that the behavior of both user and machine can be represented in advance as a plan that not only projects but determines their local interaction.

A conversation analysis of such encounters, however, reveals that while interaction between people and machines requires essentially the same interpretive work that characterizes interaction between people, fundamentally different resources are available to the "participants" (for a full account see Suchman, 1987). In particular, people make use of a rich array of experience, embodied skill, material evidence, communicative competence and members' knowledge in finding the intelligibility of actions and events, in making their own actions sensible, and in managing the troubles in understanding that inevitably arise. Due to constraints on the machine's access to the situation of the user's inquiry, however, breaches in understanding that for face-to-face interaction would be trivial in terms of detection and repair become "fatal" for human-machine communication (see Jordan and Fuller, 1974). The result is an asymmetry that severely limits the scope of interaction between people and machines.

Because of this asymmetry, engineering human-machine interaction becomes less a matter of simulating human communication than of finding alternatives to interaction's situated properties. Those properties and the subtlety of their operation are nicely illustrated in the following fragment of naturally occurring conversation:

A: Are you going to be here for ten minutes?
B: Go ahead and take your break. Take longer if you want.
A: I'll just be outside on the porch. Call me if you need me.
B: OK. Don't worry.
(Gumperz, 1982:326)

In his analysis of this fragment Gumperz points out that B's response to A's question clearly indicates that B interprets the question as an indirect request that B stay in the office while A takes a break, and by her reply A confirms that interpretation. B's interpretation accords with a categorization of A's question as an indirect speech act (Searle, 1979), and with Grice's discussion of implicature (1975); that is, B assumes that A is cooperating, and that her question must be relevant, therefore B searches her mind for some possible context or interpretive frame that would make sense of the question, and comes up with the break. But, Gumperz points out, this analysis begs the question of how B arrives at the right inference:

> What is it about the situation that leads her to think A is talking about taking a break? A common sociolinguistic procedure in such cases is to attempt to formulate discourse rules such as the following: "If a secretary in an office around break time asks a co-worker a question seeking information about the co-worker's plans for the period usually allotted for breaks, interpret it as a request to take her break." Such rules are difficult to formulate and in any case are neither sufficiently general to cover a wide enough range of situations nor specific enough to predict responses. An alternative approach is to consider the pragmatics of questioning and to argue that questioning is semantically related to requesting, and that there are a number of contexts in which questions can be interpreted as requests. While such semantic processes clearly channel conversational inference, there is nothing in this type of explanation that refers to taking a break. (1982:326)

The problem that Gumperz identifies here applies equally to attempts to account for inferences such as B's by arguing that she "recognizes" A's plan to take a break. Clearly she does: the outstanding question is how. While we can always construct a *post hoc* account that explains interpretation in terms of knowledge of typical situations and motives, it remains the case that neither typifications of intent nor general rules for its expression are sufficient to account for the mutual intelligibility of our situated action. In the final analysis, attempts to represent intentions and rules for their recognition seem to beg the question of situated interpretation, rather than answering it.

6. Cognitive science's situated practice

The decontextualized models of action embraced by the majority of cognitive science researchers stand in contrast to the situated structuring of their own scientific practice.[7] Our current research examines how various "inscription devices" (Latour and Woolgar, 1979) or technologies for representation are used by cognitive scientists and systems designers engaged in the collaborative invention of new computational artifacts. A common technology for representation in our laboratory is the "whiteboard." We begin with the observation, due to Livingston (1978), that the inscriptions on a whiteboard – lists, sketches, lines of code, lines of text and the like – are produced through activities that are not themselves reconstructable from these "docile records" (Garfinkel and Burns, 1979). Methodologically, this means that the core of our data must be audiovisual recordings of the moment-by-moment interactions through which the inscriptions are produced. Made observable, the organization of activities that produce marks on the whiteboard and give them their significance, and the function of marks in the structure of the activity, become our research problem.

Our starting assumption is that the use of the whiteboard both supports and is organized by the structure of face-to-face interaction. On that assumption, our analysis is aimed at uncovering the relationship between (i) the organization of face-to-face interaction, (ii) the collaborative production of the work at hand and (iii) the use of the whiteboard as an interactional and representational resource. From the video corpus we aim to identify systematic practices of whiteboard use, with a focus on just how those practices and the inscriptions they produce constitute resources for particular occasions of technical work. Some initial conjectures are the following:[8]

6.1 The whiteboard is a medium for the construction of concrete conceptual objects

Inscriptions on the whiteboard are conceptual in that they stand for phenomena that are figurative, hypothetical, imagined, proposed or otherwise not immediately present, but they are also concrete – visible, tangible marks that can be pointed to, modified, erased and reproduced. Over the work's course topics of talk are visibly constituted on the board, becoming items to be considered, revised, adopted and reconsidered. Technical objects once represented can be "run," subject to various scenarios, examined for their structure and so on. Conceptual objects rendered concrete, in sum, become available for development and change.

6.2 The whiteboard structures mutual orientation to a shared interactional space

Through their orientation in seating arrangements, body positions, gesture and talk, collaborators turn the whiteboard and its marks into objects in a shared space. We see designers, on first sitting down to work, "referring" in their talk and gestures to a whiteboard on which nothing has yet been written. Mutual engagement is demonstrated (or not) by attention either to the other(s) or to the shared space of the board. Bodily movements of, for example, standing at the board with marker raised or stepping back with folded arms display the status of objects as incomplete, problematical, satisfactory and the like.

6.3 Talk and writing are systematically organized

Skilled work at the whiteboard effectively exploits the "simplest systematics for the organization of turn-taking for conversation" (see Sacks et al., 1974) in the sequential organization of turns at talk and writing. The board provides a second interactional floor, co-extensive and sequentially interleaved with that of talk. So, for example, the board may be used in taking and holding the floor, or in maintaining some writing activity while passing up a turn at talk. Writing done during another's talk (may (a) document the talk and thereby display the writer's understanding, (b) continue the writer's previous turn or (c) project the writer's next turn, providing an object to be introduced in subsequent talk.

6.4 The spatial arrangement of marks on the whiteboard reflects both a conceptual ordering between items and the sequential order of their production

The use of the whiteboard to represent logical relations is a practical, embodied accomplishment. Each next entry onto the board must be organized with reference to the opportunities and limitations provided by previous entries given the physical confines of the available space. At the same time, the necessary juxtaposition of items is a resource for representing meaningful relations among them. The significance of spatial organization among items is to some extent conventionally established (e.g., the list), in other ways dependent on the contingencies of the particular items' production.

6.5 Whiteboards may be delineated into owned territories, or inhabited jointly: Similarly with particular items

Use of the whiteboard varies between more and less exclusive activity by a "scribe" to joint use, and the use of space varies from territoriality (often just on the basis of proximity) to shared access. Territory or items entered by one participant may become joint as others add to or modify them.

6.6 Items entered on the whiteboard may or may not become records of the event

Writing done on the whiteboard may be communicative without being documentary. An extreme case is the "ghost" entry – a gesture at the board that never actually becomes a mark but can nonetheless be referred back to in subsequent talk (Garfinkel and Burns, 1979). Less extreme forms are various cryptic lines, circles and the like that direct attention and accompany talk but are not themselves decipherable. Items can be and often are erased, indicating their status specifically as *not* part of the record, and the status of the talk that produced the item as an aside or digression. Alternatively, an item constructed as illustration may effectively become a document of the talk.

6.7 The whiteboard is a setting for the production and resolution of design dilemmas

Like any practical activity, research and design work encounters both routine and remarkable troubles, the latter becoming objects for reflection and resolution. But in design the dilemmas are not only expected but actively looked for. As a way of proceeding, the designer's task is to make trouble for herself in the form of unsolved problems and unanswered questions. Represented on the board, those problems and questions provide the setting for subsequent actions. Work at the whiteboard thus involves the resolution a series of dilemmas of its own making.

6.8 The whiteboard is embedded in a network of activities

While the whiteboard comprises an unfolding setting for the work at hand, the items on the board also index an horizon of past and future activities. The outcomes of previous actions are reproduced as the basis for what to do now, while what gets done now makes reference to work to be done later. Nonetheless within this network of their own and others' ongoing activities, scientists manage somehow to bound their activities in ways that bring closure each time for this time and place.

7. Conclusion

The situated practice of work at the whiteboard underscores a phenomenon observed elsewhere in social studies of science (Collins, 1985; Knorr-Cetina, 1981, Garfinkel, et al., 1981; Lynch et al., 1983). While scientific reasoning consists in negotiating practical contingencies of shop talk and its technologies, those practices are notably absent from the scientific outcomes and artifacts produced. This absence is not offered by sociologists of science as an irony, but rather as an observation with profound implications for how we understand the status of representations in science and elsewhere: *viz.* we must understand them in relation to, as the product of and resource for, situated practice. Just as instructions presuppose the work of "carrying them out," so representational devices assume the local practice of their production and use. Such situated practice is the taken-for-granted foundation of scientific reasoning.

While the rational artifacts of cognitive scientists' work are programs that run, cognitive scientists' own rationality is an achievement of practices that are only *post hoc* reducible to either general or specific representation. Canonical descriptions do not and cannot capture "the innumerable and singular situations of day to day inquiry" (Lynch et al., 1983:209). The consequence is a disparity between the embodied, contingent rationality of scientists' situated inquiries and the abstract, parameterized constructs of rational behavior represented in computer programs understood to be intelligent. To the extent that cognitive science defines the terms of rational action the disparity is not only theoretically interesting, but has political implications as well. In particular, science studies recommend indifference toward the relation of representation to phenomenon, in favor of a focus on the practices by which representations of phenomena are produced and reproduced. In the case of cognitive science, however, the phenomena are just those things on which our studies take a stand; namely, the organization of practice. In turning to the work of cognitive scientists, therefore, we have a vested interest – not only in the products of cognitive scientists' theorizing but in the adequate rendering of their and others' situated practice.

Notes

1. The study of whiteboard practices is part of a larger project with Randy Trigg to investigate how computer-based technologies might support scientific research practices. "Whiteboards" are just like blackboards but are white, and are written on with colored markers.
2. As Michael Lynch puts it: "Given how easy it is to constitute a docile subject as intentional, it raises the question of how machine intelligence can possibly be extracted from such interactional work" (personal communication).
3. Garfinkel's (1967) original inversion, on which ethnomethodology is founded, has to do with Durkheimian proposals regarding the nature of social facts:

   > Thereby, in contrast to certain versions of Durkheim that teach that the objective reality of social facts is sociology's fundamental principle, the lesson is taken instead, and used as a study policy, that the objective reality of social facts as an ongoing accomplishment of the concerted activities of daily life, with the ordinary, artful ways of that accomplishment being by members known, used and taken for granted, is, for members doing sociology, a fundamental phenomenon. Because, and in the ways it is practical sociology's fundamental phenomenon, it is the prevailing topic for ethnomethodological study. (p. vii)

4. I owe this example to a talk by Terry Winograd. For a wide-ranging critique of the "rationalistic tradition" of cognitive science and alternate proposals for computer design, see Winograd and Flores (1986).
5. I am indebted for this example, and many clarifying discussions of planning, to Randy Trigg.
6. For a recent attempt to develop a computational account of "abstract reasoning as emergent from concrete activity," see Chapman and Agre (1986); and Agre and Chapman (1987b).
7. For an eloquent treatise on the situated structuring of activity, see Lave (in press).
8. For a detailed treatment and the evidence for these conjectures, see Suchman and Trigg (forthcoming).

References

Agre, P., and Chapman, D. (1987a). *What are plans for?* Paper presented for the panel on Representing Plans and Goals, DARPA Planning Workshop, Santa Cruz, CA., MIT Artificial Intelligence Laboratory, Cambridge, MA.

Agre, P., and Chapman, D. (1987b). *Pengi: An implementation of a theory of activity.* Proceedings of the American Association for Artificial Intelligence, Seattle, WA.

Allen, J. (1983). Recognizing intentions from natural language utterances. In M. Brady and R. Berwick (Eds.), *Computational models of discourse*. Cambridge, MA: MIT Press.

Boden, M. (1973). The structure of intentions. *Journal of Theory of Social Behavior* 3:23–46.

Chapman, D., and Agre, P. (1986). Abstract reasoning as emergent from concrete activity. In M. Georgeoff and A. Lansky (Eds.), *Reasoning about actions and plans: Proceedings of the 1986 workshop*. Los Altos, CA: Morgan Kaufmann.

Collins, H. (1985). *Changing order: Replication and induction in scientific practice*. London: Sage.

Collins, H. (1987). Expert systems and the science of knowledge. In W. Bijker, T. Hughes and T. Pinch (Eds.), *The social construction of technological systems*. Cambridge: MA: MIT Press.

Dennett, D. (1978). Brainstorms. Cambridge, MA: MIT Press.

Gardner, H. (1985). *The mind's new science*. New York: Basic Books.

Garfinkel, H. (1967). *Studies in ethnomethodology*. Englewood Cliffs, NJ: Prentice Hall.

Garfinkel, H., and Burns, S. (1979). Lecturing's work of talking introductory sociology, Department of Sociology, UCLA. To appear in *Ethnomethodological studies of work*, Vol. II. London: Routledge and Kegan Paul.

Garfinkel, H., Lynch, M., and Livingston, E. (1981). The work of a discovering science construed with materials from the optically discovered pulsar. *Philosophy of the Social Sciences*, 11(2):131–58.

Grice, H.P. (1975). Logic and conversation. In P. Cole and J. Morgan (Eds.), *Syntax and semantics*, Vol. 3: *Speech Acts*. New York: Academic Press.

Gumperz, J. (1982). The linguistic bases of communicative competence. In D. Tannen (Ed), *Georgetown University roundtable on language and linguistics: Analyzing discourse: text and talk*. Washington, DC: Georgetown University Press.

Jordan, B., and Fuller, N. (1974). On the non-fatal nature of trouble: Sense-making and trouble-managing in *lingua franca* talk. *Semiotica* 13:1–31.

Knorr-Cetina, K. (1981). *The manufacture of knowledge*. Oxford: Pergamon Press.

Latour, B., and Woolgar, S. (1979). *Laboratory life: The social construction of scientific facts*. London and Beverly Hills: Sage.

Lave, J. (1988). *Cognition in practice*. Cambridge, UK: Cambridge University Press.

Livingston, E. (1978). *Mathematicians' work*. Paper presented in the session on Ethnomethodology: Studies of Work, Ninth World Congress of Sociology, Uppsala, Sweden. To appear in Garfinkel, H., *Ethnomethodological studies of work in the discovering sciences*, Vol. II. London: Routledge and Kegan Paul.

Lynch, M. (1981). *Art and artifact in laboratory science*. London: Routledge and Kegan Paul.

Lynch, M., Livingston, E., and Garfinkel, H. (1983). Temporal order in laboratory work. In K. Knorr-Cetina and M. Mulkay (Eds.), *Science observed: Perspectives on the social study of science*. London and Beverly Hills: Sage.

Lynch, M., and Woolgar, S. (1988). Introduction: Sociological Orientations to representational practice in science. *Human Studies* 11(2–3):99–116.

Miller, G., Galanter, E., and Pribram, K. (1960). *Plans and the structure of behavior*. New York: Holt, Rinehard and Winston.

Sacks, H., Schegloff, E., and Jefferson, G. (1974). A simplest systematics for the organization of turn-taking in conversation. *Language* 50(4): 696–735.

Searle, J. (1979). *Speech acts: An essay in the philosophy of language*. Cambridge, UK: Cambridge University Press.

Suchmann, L. (1987). *Plans and situated actions: The problem of human-machine communication*. Cambridge, UK: Cambridge University Press.

Suchman, L., and Trigg, R. (in press). Constructing shared conceptual objects: A study of whiteboard practice. In J. Lave and S. Chaiklin (Eds.), *Situation, occasion, and context in activity*. Cambridge, UK: Cambridge University Press.

Tibbetts, P. (1988). Representation and the realist-constructionist controversy. *Human Studies* 11(2–3):117–132.

Turing, A.M. (1950). Computing machinery and intelligence. *Mind* 59 (236):433–61.

Weizenbaum, J. (1983). ELIZA: A computer program for the study of natural language communication between man and machine. *Communications of the ACM, 25th Anniversary Issue* 26(1):23–3 (reprinted from *Communications of the ACM* 29(1):36–45, January 1966).

Winograd, T., and Flores, F. (1986). *Understanding computers and cognition: A new foundation for design*. Norwood, NJ: Ablex.

Woolgar, S. (1983). Irony in the social study of science. In K. Knorr-Cetina and M. Mulkay (Eds.), *Science observed: Perspectives on the social study of science*. London and Beverly Hills: Sage.

Woolgar, S. (1985). Why not a sociology of machines? The case of sociology and artificial intelligence. *Sociology* 19(4):557–572.