*Computer Supported Cooperative Work* **11:** 243–267, 2002.



# Activity Theory and Distributed Cognition: Or What Does CSCW Need to DO with Theories?

CHRISTINE A. HALVERSON
*IBM Research, 650 Harry Rd., San Jose, CA 95120, U.S.A. (E-mail: krys@us.ibm.com)*

**Abstract.** This essay compares activity theory (AT) with distributed cognition theory (DCOG), asking what each can do for CSCW. It approaches this task by proposing that theories – when viewed as conceptual tools for making sense of a domain – have four important attributes: descriptive power; rhetorical power; inferential power; and application power. It observes that AT and DCOG are not so different: both emphasize cognition; both include the social and cultural context of cognition; both share a commitment to ethnographically collected data. Starting with a description of the distributed cognition approach, it uses an example of a DCOG analysis to ground a discussion of the strengths and weaknesses of AT and DCOG as an approach to issues in CSCW. Finally, the essay considers what theoretical work is being done by the attributes of the respective theories, and whether AT, DCOG, or any theory developed outside the context of group work, will work for CSCW.

**Key words:** activity theory, analysis, distributed cognition, methodology, theory

## 1. Introduction

Activity theory (AT) is one of the many theories and approaches being used in CSCW, and the range of papers in this issue illustrates its popularity. Indirectly, they also indicate what researchers are looking for from a theory. They describe 'native' cooperative phenomena (Nardi, Whittaker and Schwarz; Zager) as well as computer support for (cooperative) work (Miettinen and Hasu; Spasser). Some use the theory for meta-level analyses based on theoretical precepts. Barthelmess and Anderson use AT to compare Process Centered Software Development Environments (PCSDE) while Korpela, Mursu and Soriyan examine AT driven Information Design itself. Several papers push AT by developing models to extend the theory (Clases and Wehner), or defining new phenomena (Nardi et al. and Zager). Finally, some take a practical bent by addressing design (Fjeld, Lauche, Bichsel, Voorhorst, Krueger and Rauterberg; Miettinen and Hasu). The fact that papers are mentioned more than once, and in different categories, is not surprising. The range of uses illustrates a broader issue in CSCW. We appropriate theories and methodologies from other fields. Can we continue to satisfy our analytical needs in this way?

In this essay I compare activity theory (AT) to distributed cognition theory (DCog).[1] First, I frame this essay by looking at what we expect from theories. I define four attributes important in theories: *descriptive, rhetorical, inferential,* and

*application* power. I explore the relative strengths and weaknesses of both theories with reference to these attributes, and with respect to what each theory does for CSCW. I touch on whether and how AT and DCog help with design. Finally, I discuss what theoretical work is being done by the attributes named. I explore whether any theory derived outside of the context of group work, whether AT or DCog or something else, will work for CSCW.

## 2. Why theory?

CSCW, like HCI, has adopted a number of theoretical constructs. The approaches used include theories, conceptual frameworks, and descriptive methods, as well as a variety of hybrid forms (Shapiro, 1994). Just a partial list from A to S includes: activity theory (Engeström, 1987; Kuutti, 1996; Nardi, 1996b; Bardram, 1997; Engeström et al., 1999), conversation analysis (Sacks et al., 1978; Frohlich and Luff, 1989; Sacks, 1992; Katzenberg and McDermott, 1994), coordination theory (Schmidt and Simone, 1996; Carstensen and Nielsen, 2000), distributed cognition theory (Rogers and Ellis, 1994; Hutchins, 1995b; Ackerman and Halverson, 1998), ethnomethodology (Garfinkel, 1967; Button, 1991; as well as numerous papers in CSCW such as Bentley et al., 1994; Rouncefield et al., 1994; Heath and Luff, 1996), grounded theory (Glaser and Strauss, 1967; Strauss and Corbin, 1998), situated action (Suchman, 1987; Schiff et al., 1997) and social/symbolic interactionism (Blumer, 1986; Fitzpatrick et al., 1996).

Most of these have been used to study and describe CSCW settings and systems, but few explicitly approach the design of those systems. As Button and Dourish (1996) point out in the case of ethnomethodology, closing the gap between critique and design is quite a challenge. CSCW often turns to other methods to support the design process including: contextual enquiry (Beyer and Holtzblatt, 1998), participatory design (Greenbaum and Kyng, 1991), and user centered design (Norman and Draper, 1986). We draw on still another set of theories when we address the underlying computer system(s)' architecture. What are we doing with these theories, approaches, and methods? Some of us want to evaluate the truth of the world, often through logical manipulations of theoretical constructs, intuition, and thought experiments. Others want to confirm their theoretical musings by empirical reference. Both cases require testable hypotheses to be validated or falsified. In contrast, ethnomethodology rejects theory, reacting to problems operationalizing sociology's theoretical constructs with reference to the observed world (Button, 1991, p. 3). In practice, many of us adopt the view expressed in Barthelmess and Anderson (2002).

> The value of any theory is not 'whether the theory or framework provides an objective representation of reality' (Bardram, 1998), but rather how well a theory can shape an object of study, highlighting relevant issues. In other words, a classification scheme is only useful to the point that it provides relevant

> insights about the objects it is applied to. (Barthelmess and Anderson, 2002, this issue)

From this point of view, theories are more like a pair of dark glasses. We put them on and the world is tinted. The change brings some objects into sharper contrast, while others fade into obscurity. However, by adopting theories from other fields we may be bringing theoretical objects into focus that are not appropriate for CSCW. For example, activity theory and distributed cognition theory are both first and foremost, theories about cognition. What they can say about group interaction is based on what they say about cognition. That may be OK, depending on how we use the theory. But how do we evaluate their usefulness for CSCW?

From a pragmatic view of theory we can identify four attributes we want. First, we require *descriptive* power. Theory in CSCW should provide a conceptual framework that helps us make sense of and describe the world. This includes describing a work setting as well as critiquing an implementation of technology in that setting. Second, we need *rhetorical* power. Theory should help us talk about the world by naming important aspects of the conceptual structure and how it maps to the real world. This is both how we describe things to ourselves and how we communicate about it to others. Further, it should help us persuade others that our view is correct.

The third attribute is *inferential* power. Without engaging in arguments about whether theories are true, or only falsifiable (Popper, 1992), we do want a theory to help us make inferences. In some cases those inferences may be about phenomena that we have not yet understood sufficiently to know where or how to look. We may hope that inferences will lead to insights for design. Or we may want to predict the consequences of introducing change into a particular setting. An important fourth attribute has to do with *application*: how we can apply the theory to the real world for essentially pragmatic reasons. Mostly this translates to our need to inform and guide system design. We need to describe and understand the world at the right level of analysis in order to bridge the gap from description to design.

Understanding what we want to do with a theory is very important. Just having all these attributes is not enough. A theory of particle physics is not likely to map onto CSCW. Additionally, a particular strength in one attribute or another will make a theory more congenial for a particular task. For example, a theory in physics that focuses on the qualities of the appropriate elements in terms of wave-like properties may have difficulty describing their interaction or relationships as quanta, or vice versa. There are two lessons here. First, we need to be aware of what a theory might be predisposed to do – based on the nature of its attributes. Second, we need to be equally aware of what we *want* the theory to do. This second lesson has to do with the scope of the theory. Design of collaborative systems is only one possible use of theory. Another is supporting discourse in a community, while a third is providing the apparatus that allows comparison across empirical observations. All of these are important for a field.

A special issue organized around a theory, like this one, implicitly argues that the theory (AT) can provide the CSCW community theoretical leverage. The

diversity of papers offered here attest to that. However, because what we have in CSCW is a grab bag of theories we need to ask three questions. Can one theory do everything for us? Does AT aid design, support discourse about CSCW, and help us compare across field settings? If not, what value does each theory, or approach, provide?

To explore these questions I compare activity theory with distributed cognition theory (DCog) another cognitively based theory. I use DCog for a number of reasons, but primarily because it is the theory that I use in my own analyses. I find it shows different strengths and weaknesses from AT. Furthermore, where those strengths and weaknesses fall sheds light on what we want from theories in CSCW.

## 3. A tale of two theories

### 3.1. BACKGROUND

There has been a steadily increasing interest in AT during the 1990s in both the HCI and CSCW communities (for example Engeström, 1987; Bødker, 1991; Kuutti, 1991; Raeithel and Velichkovsky, 1995; Nardi, 1996) with a much narrower dissemination of DCog in the same period for example (Rogers and Ellis, 1994; Hutchins, 1995b; Ackerman and Halverson, 1998; Hollan et al., in press).[2] In many ways these theories are closely tied because they share a common intellectual heritage – the emphasis on the cognitive.[3] They are also in contrast, since Western-European and Russian pursuits of cognitive science diverged in the beginning of the 20th century.[4] With a common heritage we might ask whether and how they diverge along the attributes of descriptive, rhetorical, inferential, and application power.

Both diverge from other cognitive theories by incorporating the social and cultural context of cognition. In practice, they do this in different ways. Each theory's approach has much to do with its historical development. As a cognitive scientist, I'm interested in the divergence of their approaches. For me, the many phenomena of human society and activity are the result of human cognition. Much of their power arises from how cognition instantiates itself in the material world. As a practitioner of DCog analyses, and not unlike AT practitioners, I see the world of artifacts, personal history, culture, social, and organizational structure through a filter that labels them as the residua of collaborative cognition, analyzed along numerous time scales.

As a CSCW researcher, however, I am more concerned with how I can use a theory to understand a specific domain, reach insights about collaborative work in general, or design for a particular problem. Each of these puts different demands on the theory – the first on *descriptive* power, the second on *rhetorical* and *inferential* power, and the third on the practical *application* of the inferences.

In many ways I see the differences between AT and DCog as being superficial, at least as they apply to CSCW. Before the arguments begin, let me clarify. A

large part of the power and usefulness of both theories, as with ethnomethodology, is their commitment to ethnographically collected data. That is, practitioners go to where the action is, observe how things really work, and are confronted with how (well or poorly) reality maps onto theoretical constructs.[5] This integration of ethnographic practice with theoretical constructs makes learning and using both theories more difficult. Of course, this begs the questions of how much of each theory's success is due to the ethnography as opposed to theoretical traction. While I believe theories do provide additional leverage, both to the ethnographic practice and to the analysis, I set aside this issue here because both theories share this criticism.

As I read the papers in this issue I began to see several reasons why the AT perspective has become appealing in CSCW. As can be seen in these papers, it is applicable to a range of domains and levels of analysis, and it has descriptive power. Despite early calls that it was too difficult to learn (as reported in Nardi, 1996b) the range of practitioners here – academics, members of large and small companies, as well as researchers – attest to its growing converts.

When I compare AT with DCog several things stand out:

1. *AT has named its theoretical constructs well*. Even though some names may conflict with common use of the terms, naming is very powerful – both for communicative as well as descriptive reasons.

   In contrast, *few theoretical constructs are explicitly named in DCog*. Those that are discussed, either in *Cognition in the Wild* (Hutchins, 1995b) or elsewhere, are not presented in a way that gives them same the rhetorical force of naming as seen in AT. This is important because names are often what you manipulate in a theory. Being able to manipulate data along with the names in AT provides an additional rhetorical advantage.
2. *In AT, the perspective of the individual is at the center of everything.* AT focuses on the cognitive process of an individual *situated* in a social, cultural, historical, and artifactual world.

   In contrast, *DCog focuses on the socio-technical system*, which usually (but not necessarily) includes individuals. DCog uses the same theoretical language for both people and artifacts. This common language has led others to critique the theory for assuming people are equated with artifacts in some way that denies their humanity. This is, in fact, not the case.
3. *Dealing with process is built into the structure of how AT is presented.* Activity system diagrams (e.g. Barthelmess and Anderson, p. 16; Clases and Wehner, p. 43; Collins et al., p. 58; Korpela et al., pp. 113 and 114; Spasser, p. 96) keep process in the foreground for both reader and analyst. This is somewhat ironic, since a static diagram represents essentially dynamic relationships between the key components. Nonetheless, their representation provides both descriptive and rhetorical power.

   In DCog, *process(ing) is so central to the analysis that it may be less obvious* to the uninitiated. Unlike AT there is no iconic structure applied to each situa-

tion. Instead, it is built into the process of analysis, and may or may not be represented in the products of that analysis.

To clarify these statements I need to introduce distributed cognition theory and compare it with aspects of activity theory. For illustration I will draw on work investigating call centers and organizational memory (Ackerman and Halverson, 1998, 1999, 2000) and compare primarily with two papers from this issue: Clases and Wehner, and Collins, Shukla, and Redmiles. By illustrating the similarities and differences between AT and DCog within a comparable domain I explore what we need in CSCW from a theory.

### 3.2. CHARACTERISTICS OF DISTRIBUTED COGNITION

In the last century American cognitive science focused on the cognition of the individual extracted from their social and cultural context (Hutchins, 1995b quoting Gardner, 1984). This may seem odd juxtaposed with the approach taken in AT.[6] However, it is only in the last decade[7] that cognition has been more generally acknowledged as distributed rather than by definition the property of an individual mind (Salomon, 1993; Hutchins, 1995b; Clark, 1997). (Researchers differ on how cognition is distributed, but Spasser's (this issue) casual reference to a distributed cognitive system without any specific citation or definition speaks volumes for the current acceptance of this notion.) With this has come the recognition that collections of individuals have cognitive properties that are different from sole individuals, often emergent from their collective behavior.

Several researchers have used the term *distributed* to mark the difference in perspective from more traditional approaches to cognitive science (Norman, 1991; Zhang and Norman, 1991; Salomon, 1993; Hutchins, 1995b) including people writing in the AT tradition (Cole and Engeström, 1993). I focus on distributed cognition theory developed by Hutchins beginning in the mid-80s, published in his book (1995) and a number of articles (Hutchins, 1988, 1990, 1991, 1995a; Hutchins and Hazlehurst, 1990; Hutchins and Klausen, 1992; Flor and Hutchins, 1992; Hutchins and Palen 1993) and which continues to be developed in his lab (Halverson, 1995; Holder, 1999; Hollan et al., in press) and elsewhere (Rogers and Ellis, 1994; Perry, 1997).

Distributed cognition is not some "new" kind of cognition, rather a recognition of the perspective that all of cognition can be fruitfully viewed as occurring in a distributed manner. As a cognitive theory, DCog is focused on the organization and operation of cognitive systems; that is, with the mechanisms that make up cognitive processes, which result in cognitive accomplishments. It recognizes that "a process is not cognitive simply because it happens in a brain, nor is a process non-cognitive simply because it happens in the interactions among many brains" (Hollan et al., in press). This opens up our notions of cognitive processes to a much wider variety of mechanisms than the classic symbol manipulation of the physical symbol system hypothesis (PSS) (Newell and Simon, 1972; Simon, 1990). Hutchins argues

that PSS works better as "*... a model of the operation of a sociocultural system from which the human actor has been removed*" (p. 363, emphasis his), rather than a model of an individual's internal cognitive processes. Distributed cognition theory capitalizes on this view by refocusing attention on the social-cultural system – the cognitive system which functions by bringing representational media into coordination with one another.

> I do believe that the computation observed in the activity of the larger system can be described in the way cognition has been traditionally described – that is, as computation *realized through the creation, transformation, and propagation of representational states* (Hutchins, 1995b, p. 49, emphasis mine).

Hutchins' approach carries with it a commitment to ethnographic data collection and method. The analysis in *Cognition in the Wild*, following Marr (1983), proceeds through multiple levels of analysis which can be described as: (1) a functional definition of the cognitive system; (2) enumeration of representations and processes within that system; and (3) the physical instantiation of representations and the algorithm(s) that control the processes.

The utility of DCog for CSCW, like AT, is its theoretical commitment to examine this broader socio-cultural-technical system, which is necessary for the collaboration between individuals mediated by artifacts. Furthermore, its focus on representational states and the system level cognitive work they do is extremely useful for design. But how do we define that system?

### 3.3. WHAT'S IN A NAME? THE DEFINITION OF A UNIT OF ANALYSIS

A key tenet of distributed cognition is its commitment to a unit of analysis defined in relation to the complex phenomena being observed. As Hutchins (1995b) shows, the information processing in a navigation team varies with the context and circumstances. Solo watch standing involves the interaction of one individual with various artifacts, structured via well-established procedures and routines. In contrast, entering a harbor requires the effort of several people, again in coordination with specialized tools and with each other, but at a much more rapid pace. While the overall behavior exhibited by the system is the same – navigation – the means change. Thus, we see that within the system there are mechanisms that dynamically reconfigure to bring subsystems into coordination in order to accomplish certain functions.

More specifically, for AT the primary unit of analysis is the *activity* (cf. all of the papers in this issue, as well as Kuutti 1996). Thus you have Collins et al. discussing the *Customer Support Activity System* and the *Knowledge Authoring Activity System.* This naming makes the object of inquiry very clear-cut rhetorically. That is, the primary theoretical concept of activity theory is *activity* and which is comprised of *action*. AT also defines activity as the central unit of analysis. This overlaps with the common sense use of activity as something that one does. For

example, look at how Collins et al. (2002, this issue, §2.1) outline their object of inquiry:

> In Hewlett-Packard's culture, this documentation activity is called "knowledge authoring". The term *Knowledge Authoring Activity System* will be used to refer to this activity. Closely linked but not discussed in detail in this analysis, was an explicit *Knowledge Maintenance Activity System.* Finally, both knowledge authoring and maintenance are part of a larger activity of supporting Hewlett-Packard customers, the *Customer Support Activity System.*

It is immediately evident what aspects they are exploring. Equally, we know which activities they have set aside from consideration. Using Engeström's Activity System Model (Engeström et al., 1999) as a conceptual framework they describe the setting of a help desk, situated within the broader organizational context. They enumerate not only the key parts of the activity system – definition of the activity, as well as subject, object, and outcome – but also what governs the relationships between them. By naming these – rules, division of labor, and mediating artifacts – it becomes easier to communicate about the setting and analysis with those who understand the terms.

This highlights what I meant when I said that the theory is good at naming things. AT has rhetorical power, not because it names things-in-the-world, but because it names conceptual and analytical constructions with which any analyst looking at a collaborative system has had to struggle. Naming a category "mediating artifacts" focuses the analyst's attention around those objects used by the subjects of the activity system. Naming helps communicate to others – particularly when they do not understand the particular domain. (To take a trivial example, if a ruler is a mediating artifact then the analyst is signaling me that the ruler is doing some work that is important for me, the reader, to look at more closely.) Conversely, if a reader understands the domain, they can bridge to the theoretical concepts because they are named and organized and mapped onto the domain. This is not unique to AT, but nonetheless it is powerful.

In apparent contrast, DCog does not have a special name for the unit of analysis. It frames the problem in terms of examining the cognition of a system in terms of its function. The functional requirements drive analytical focus, wherein functional operation is decomposed into smaller units of analysis that make sense with respect to the particular function or task within the system. Like the example from Collins et al. (2002) above, we would begin to define the functional system in a straightforward manner. System operation will re-focus us on an event driven segmentation of the tasks (and subtasks). Taking a perspective that does not privilege the individual (yet also does not exclude the individual as the scope of the unit of analysis) may mean that configurations exist of both multiple or solitary components, human agents as well as human produced artifacts, and social and cultural structures. It is the task requirements that dictate which configuration is the one that counts for understanding a particular task.

This may be more obvious if we look at a more concrete example. In the study of the operation of a hotline for personnel questions, we define the unit of analysis variably. Sometimes it is single customers' call bounded on one side by the initial ringing, and on the other side by the ending disconnect of the telephone (Ackerman and Halverson, 1998, 2000). Elsewhere (Ackerman and Halverson, 1999) it is defined more conceptually, based on events that focused on one issue but whose resolution spans hours or days.

Regardless of the scope of the unit of analysis, the process of analysis is the same. In each case, within the unit of analysis, representational states and the processes that act on them are identified. However, the potential of the analysis is determined by the scope of the unit of analysis, and that scope varies. In the simplest case above, a call to verify employment, the unit of analysis comprised two individuals (the customer calling and the call taker) and several artifacts that both mediate the call and that contain the information in question.

The purpose of drawing this distinction around how the unit of analysis is defined is to highlight different strengths of the two theories. In AT the naming of the unit of analysis as *activity* is just one of many theoretic names at different levels of abstraction. The papers in this issue range from detailing phenomena across a broad range of these levels (Collins et al., 2002) to specifying them at only one level (e.g. Barthelmess and Anderson, 2002). I suspect that having the overhead of naming does make it difficult to learn and master the theory. While it requires additional precision on the part of the analyst – almost every paper here defines and clarifies terms and their use – it also provides precision in communication to other AT practitioners. (But compare with Collins et al.'s report of problems communicating to other researchers and managers at their field site because of the confusion between the theoretical *object* and the common sense *objective*. In addition, there is the careful work of Barthelmess and Anderson to detail the difference between the theoretical language of AT and the technical use of the same or similar terms within their domain.) In addition, the power of naming theoretical constructs and defining their relations allows an analyst to manipulate the theory at the same time she manipulates her data. In the terms I used before, this shows descriptive, rhetorical, and inferential power.

Clases and Wehner (2002) do an exquisite job connecting activity theoretic concepts to the issues they see of importance. They reason through the theoretic concepts until it seems that the conclusions come directly from the theory rather than from an analysis of a specific setting. For example, when talking about how artifacts are a symbolic externalization of a specific practice they draw out an essential knowledge management example.

> One of the core ideas of activity theory is that human activity is *mediated* by societal forms as well as operative means. Figure 2 is based on these schemes and visualizes CSCW systems as mediating the joint activity in or between different communities of practice. The figure shows that the joint activity evolving between different actors is mediated – on the level of societal *forms*

> – by informal rules, self-constraints and a certain division of labor that historically evolve in communities of practice. On the other hand, the interaction between actors in computer-supported work places is being structured – on the level of operative means – by the characteristics of the specific CSCW system in use. The CSCW system will provide actor A with *means of production*, i.e. features to generate certain *objects*, which will then be represented for Actor B by the use of the system providing *means of orientation*. The artifacts produced by means of CSCW systems may be looked upon as symbolic externalizations of a specific practice. Therefore, when using a CSCW system, Actor A has to transform his experiences made and *knowledge* gained into a certain document. For Actor B, this externalization of a specific practice in the first case appears as *codified knowledge*, i.e. *information* that *might be useful* in another context. Depending on the way in which the context of generation the information is presented, Actor B will be more or less able to put it into perspective. In other words: Knowledge may not be immediately be 'transferred' but is *transformed* by processes of codification and interpretation.

Of note, most of the terms italicized in the above excerpt are not just for rhetorical emphasis, but also indicate theoretical terms. While a flavor of the knowledge management domain comes through in this excerpt, overall the example reasons using the higher level of theoretical constructs. In contrast, in Ackerman and Halverson (1998, 1999, 2000) we talk about similar phenomena with reference to the domain of inquiry, that is the specifics of the hotline, rather than the theory. Within the domain there is the problem that knowledge must be de-contextualized from its specific situation before it is stored, but in order to be used, it must be re-contextualized to fit with the new situation. Using DCog (and some AT terminology) our analysis deconstructed actions of a particular actor at the very low level of representational states. With other input derived from field observations we then used those analytical insights and rebuilt a narrative of our understanding situated with respect to the domain. The insights we gained are with respect to the domain, and mostly fall outside the theory. One of our conclusions, the notion that we were seeing information acting as a boundary object, is not an insight into the theory of distributed cognition per se. It is also not an insight extracted by manipulation of DCog's theoretical constructs. But neither was it obscured by the theory.

DCog names almost none of its theoretical constructs, except at the very basic level of representational states. An analyst manipulates data to draw conclusions about the world, but this does not equate to manipulating the theory itself. The chain of inferences that build back from the low level of analysis to higher theoretical constructs is almost completely hidden from others. In DCog, descriptions analogous to 'division of labor' or 'mediating artifact' are higher-level constructs that are not named within the theory. The communicative weight is carried by a description of the phenomena and the higher-level implications. This translates to less rhetorical power and makes discourse in the theoretical community more cumbersome. However, the focus at the level of processes, representational states,

and their meaning (representations), exposes system workings at a level that has considerable descriptive power. This makes DCog particularly useful for those who are focused on design. For those who understand the domain, the detailed description at this level makes it possible to see the implications of changes.

Part of DCog's power lies in its flexible unit of analysis. This provides a mechanism to reconfigure the analytical framework in a situation specific manner. In the case of the hotline group, one can imagine that treating it in terms of the activity of taking calls would be fine. At a high level this is true. But flexibility in drawing the boundary of the unit of analysis exposes how a simple call for employment verification is both like, and unlike, a more complicated call regarding insurance payments. In the simplest call, we can see that with increasing automation that the same call may use only one individual, or even none, while the work done in the more complex call is hard to envision without the intervention of a human. Because DCog deals with humans and artifacts as they contribute to the larger socio-technical system, both possibilities can be analyzed. In contrast, because AT centers the activity system around the subject (individual), analysis of an automated subtask is problematic.

### 3.4. THEORETICAL LANGUAGE AND NON-HUMAN AGENTS

This example raises the issue of how DCog and AT handle people and artifacts. Fjeld et al. (2002, referring to Nardi, 1996a) state "... distributed cognition puts people and things at the same level; they are both 'agents' in a system". They go on to say that this means DCog "ignores the faculties of human beings not found within computers, like motive, emotionality and consciousness. It also ignores for computers their non-human traits, name their ability to execute programs in a precise and predictable manner."

While I agree that DCog does not *focus* on some of what goes on inside humans, I disagree that it ignores all that goes on inside both humans and artifacts, including computers. This misconception of the theory is based on how and why DCog 'treats humans and artifacts the same'. Analysis enumerates the representational states, the media on which they are instantiated, and the observed processing of those states.

> The conduct of the activity proceeds by the operation of functional systems that bring representational media into coordination with one another. The representational media may be inside as well as outside the individuals involved. These functional systems propagate representational state across the media. (Hutchins, 1995, pp. 372–373)

The phrasing may be awkward, but it reflects DCog's theoretical commitment to not privilege the individual. Thus humans are not the only agents that bring representational media into coordination. This is possible because the theoretical language of distributed cognition theory itself does not privilege the individual over

other components of the system. One way to view this is indeed that human and non-human can be cognitive agents, and the focus is on the observable aspects of the cognitive processing. This does mean that emotion may be left out of the analysis, insofar as it occurs hidden from view inside an individual's head. (However, insofar as it is manifested externally in the operations of the cognitive system it may be a valid part of the analysis.)

For our analyses (Ackerman and Halverson, 1998, 1999, 2000), being able to span human and non-human cognitive agents, as well as organizational and cultural structures and norms allows us to cover the diverse manifestations of organizational memory. The common breakdown into representational states and processes provides a way to analyze how the observed details achieve the particular function that is the focus of a unit of analysis. This presents artifacts, human actors, and organizational and social structures on an equal theoretical footing. With a description constructed in these terms we can begin to understand how technologies *and* social structures currently fit a system's operation. Once analyzed into its component representational states and processes, the analyst uses that information to reconstruct the functioning of the system. This allows an analysis with respect to the context of use within an organization. By extension one can speculate about how changes in technologies might affect future operations. What does this look like?

In Ackerman and Halverson (2000) we analyze a very simple call – one about employment verification. As is common for many complicated analytical frameworks, in the paper we skip presenting the full details of the analysis process and instead present what is necessary to support the conclusion that the call taker uses not one memory, but many, and we support this with a diagram showing all the memories used (Figure 1). To highlight both the power of the analysis and how people seem to get "left out" I want to walk through a part of the analysis that we left out of that previous paper.

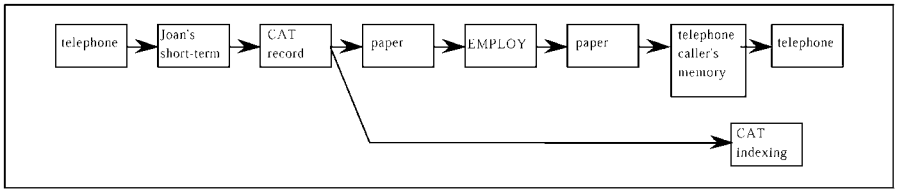


*Figure 1.* Various memories used during the employee verification process.

The setting is a hotline group (here abbreviated HLG) for personnel concerns at large company. HLG takes calls from both inside and outside the company. This particular call is an "employment verification", where a caller (for example, a mortgage lender) contacts HLG to find out if a person is actually an employee. In order to answer this request the agent, Joan, must look up the person in a specific database, the EMPLOY system. Because of technical incompatibilities, the database must be accessed on a terminal separate from the one on her desk. This

*Table I.* First 3 turns of an employment verification call

| | |
|---|---|
| 1 | HR Hotline. This is Joan, thanks for holding . . . |
| 2 | I can do that. |
| 3 | Hold one moment . . . |
| | (Joan closes here old Call Tracking (CAT) record, and starts a new CAT record.) |
| | I just need to get, to get a little more information. M. . .A. . .N. . .D. . .E. . .L? (Joan types name as she is sounding the letters.) Do you have a social security number? (She types the social security number as she listens.) |
| | Okay, hold on, please. |

terminal (with EMPLOY) is shared by all of the agents, and it is located about three meters from Joan's desk. The agent, then, must disconnect her headset from the phone, walk to this central table and look up the person on the EMPLOY system. Furthermore, part of the HLG agent's job is to maintain a record of call requests. To do this they use another computational system, the Call Tracking system (CAT), which is accessed from their desktop system.

The analysis began with the observations – primarily videotape, supported by additional direct observations, semi-structured interviews, social network analysis, and field notes. For the DCog analysis, the unit of analysis, as I discussed above, was clearly circumscribed by the extent of the phone call – because the temporal extent happened to coincide with the functional extent of the verification. Transcribing the call included actions as well as discourse. For privacy reasons we could only record half of the conversation, so we are limited in what we can directly observe. Table I shows the first three turns of this call in the transcript, interleaved with Joan's actions.

Like AT there are many levels at which we can represent this. At the most basic level we detail:

1. Representational states and the media they are instantiated in (or on);
2. The character of the processing (such as creation, propagation, transformation) and a description of its mechanism;
3. Agents which enable the processing, whether human or artifact.

At this stage all agents involved in the processing are enumerated, and only later pruned. Table II shows one detailed representation of the first three boxes from Figure 1, which coincides with most of turn 3 (shown in Table I). Reading down under each spoken fragment we see that the representational state is propagated through a variety of media and agents. A trio of entries details the agent, the representational state and medium, and the kind of processing. So Joan moving a mouse is represented by signifying the agent as Joan. The representational state is her hand position on the medium of the mouse. This is *creating* a physical process. The representational media detailed in the table are coordinated with each other to

*Table II.* Detail of representational states, media, and processes for third turn.

| | Call Tracking (CAT) record | (Joan closes her old and stars a new CAT record.) | I just need to get, to get a little more information. | M. . .A. . .N. . .D. . .E. . .L? (Joan types name as she is sounding the letters.) | Do you have a social security number? |
|---|---|---|---|---|---|
| Agent: | Joan | Joan | Joan | Joan | Joan |
| rep | | hand position/ | I just need to . . ./ | | Do you have a . . ./ |
| state/medium | hand positon/mouse | mouse | voice | M. . .A. . ./voice | vocie |
| processing/kind | create/physical | create/physical | create/verbal | create/physical | create/verbal |
| Agent: | mouse | mouse | telephone | keyboard | telephone |
| rep | | | I just need to . . ./ | | Do you have a . . ./ |
| state/medium | movement/mouse | movement/mouse | voice | M. . .A. . ./mouse | vocie |
| processing/kind | propagate/physical | propagate/physical | propagate/audible | propagate/physical | propagate/audible |
| Agent: | computer | computer | Caller | computer | Caller |
| rep | | | I just need to . . ./ | | Do you have a . . ./ |
| state/medium | moving cursor/CRT | moving cursor/CRT | voice | M. . .A. . ./mouse | vocie |
| | | | propagate/auditory | | propagate/auditory |
| processing/kind | propagate/CRT | propagate/CRT | processing | propagate/CRT | processing |

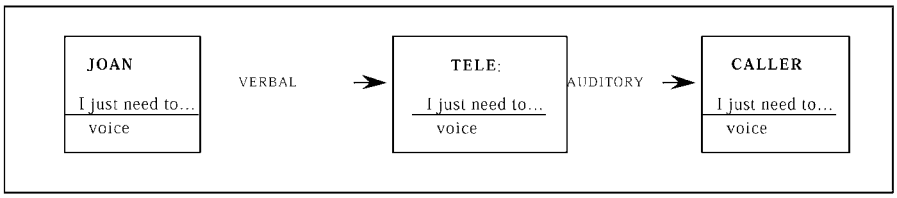


*Figure 2.* The processing from the 3rd column of Table II re-represented as a diagram.

move the representational states through the processing necessary to accomplish the cognitive functioning of the system.

Notice that people and artifacts are treated equally as agents in some cases because they do processing. So Joan's use of the mouse to drive the cursor to close a call tracking record and open another is represented as the propagation of her physical action, transformed through the mouse and the CPU that results in the cursor movement that appears on the CRT (first column Table II). There is internal processing in Joan, and similarly there is implied internal processing that is happening inside the mouse and the computer CPU. In this representation internal processing has been left out. We generally ignore internal processing for two reasons: it is not the focus of the functional system, nor is it observable. But we often know that it is there. We can infer hidden from observable processes. In the case of computers we often have other means to know the internal processing, such as manuals. (Unfortunately we do not have the definitive manual on human processing.)

Figure 2 collapses some of the detail from the third column in this system into a diagrammatic "short hand" that re-represents it. Each box shows the agent, the representational state, and the media it is instantiated on. Joan says, "I just need to get, to get a little more information". In saying this she does some internal processing that creates a representational state of the words carried on a vocal medium (i.e. her voice). This representational state is propagated verbally to the telephone, which then does its own processing, propagating the same representational state to the listener. At this level we presume the same medium. The caller does auditory processing on the same representational state.

One thing this figure points out is the problem of representing representational states and processes sufficiently. Table II is more explicit about what the processes are, while Figure 2 provides a better sense of the movement of the representational state as it gets processed. Figure 3 however gives a better idea of where memory is and foreshadows the result presented in Figure 1. It also gives a better representation of how agents bring representational states into coordination with each other to accomplish processing. Figure 3 uses yet another representation where agents, representational states, processes, and memory are all present. Agents are circular. Triangles represent memory. (The grayed-out triangle 'switch' as part of

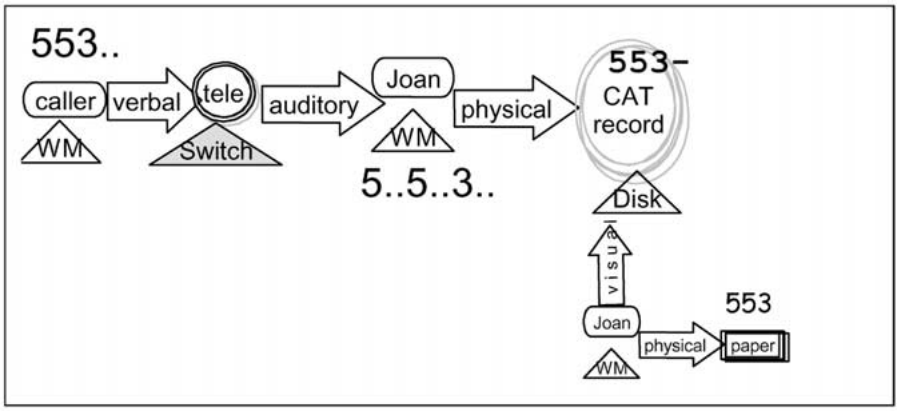


*Figure 3.* Graphical representation of representational states and processes. Agents are circular. Triangles represent memory. (The grayed-out triangle is unused memory.) Arrows represent processing.

the telephone is unused memory.) Arrows represent the character of the processing; in this case more of a memory aid for the analyst to reconstruct what happened.

Thus, the caller says the social security number, which is processed through the telephone and via Joan's auditory processing. Through rehearsal (saying each number) we infer that it is processed in her working memory, while we can observe her coordinating the audible representational state with her physical actions typing the same numbers into the Call Tracking (CAT) record. We again see Joan coordinating representational states as she reads from the CAT record and writes the number on a piece of paper. The paper is a special case because it is both a medium for the representational state and the "memory" of that state.

After this segment Joan writes down the social security number on a piece of paper and carries it over to the EMPLOY database to look up the information. If you compare Figure 3 with Figure 1 you can see that the analytic representation has been further focused to just one aspect of interest – in this case *memory*. Thus in Figure 1 we track the propagation of what the caller says through the mediation of the phone's memory, Joan's and the caller's short term and working memory, the Call tracking record, a piece of paper, and the EMPLOY database.

If this is confusing it is because these diagrams are private for the analyst; quite unlike the diagrams in AT which serve a more public and rhetorical purpose. You may notice that in Figure 1 the telephone is a 'memory' while in Figure 3 it has a component 'switch' memory. I have deliberately excluded some observed representations and processes related to the telephone system, both for clarity and because at this point in the analysis I am not yet concerned with it. If I care about the processing going on in the phone, I can make that its own unit of analysis or expand the detail of this analysis to include it. In fact, the phone turns out to be part of a sophisticated switch and record keeping process that propagates though the

larger system of the hotline group (Ackerman and Halverson, 1999). This is why in Figure 1 it is part of the collection of memories that are the processing substrate for the system. Understanding the telephone's role came from the extended ethnographic work – collecting information about how the phone was used. In turn, examining how the phone contributed to the propagation of those analyses was in fact fruitful.

As the range of presentations in the figures show there is not only one useful representation. Other representations, such as those used in Holder (1999) only indicate on what media the representational states appear at each time segment. This highlights which representational media are brought into coordination at which times. This is particularly useful in the case of an aircraft cockpit where most of the representational state appears on the medium of fixed instruments. Focusing in this way elides some constructs and highlights others. Part of what the ethnographic work contributes is to our understanding where such a representation will provide analytical power.

Because of the simplicity of this example one might imagine design solutions that would eliminate the need for Joan to do this kind of physical coordination of representational media. Electronically linking the CAT record system with the EMPLOY database on the same two computer systems could remove Joan from most of her coordinating role. One might imagine a voice response unit could eliminate her intervention entirely, while providing problems of its own (which would require seeing the results of another analysis). Seeing these alternative solutions is possible I argue, because the diagram is at the level of representational states.

This contrasts with how artifacts are treated as whole objects in AT. For example, Collins et al. adopt Engeström's refinement of Wartofsky's mediating artifacts hierarchy (Engeström et al., 1999) as a framework for both analysis and enumeration based on use. This makes it clear that, for example, conversation was a mediating "what" artifact, which contributes a means of achieving the object of knowledge authoring. ". . . hallway conversations and computer mediated conversations were central mediating artifacts, especially in transforming implicit knowledge to explicit and documented knowledge" (p. 7). From this description we get little information at a level of detail that is useful for altering the design. That kind of information is exposed later when discussing the relationship between two elements of the activity system.

AT is an approach that analyzes the dialectical aspects of the activity system. Thus we learn more about how computer mediated conversation works within this activity system when the authors detail the tensions[8] between the subject and mediating artifact elements. One example they give is about "InstaNote" a broadcast request and response tool. The authors argue that a subject's description of the use of this artifact, in comparison to another artifact the chat window, implicitly provides design requirements. They quote an interviewee who states that InstaNote is more noticeable than the chat room tool because that tool is a little window that

can get hidden. In contrast, InstaNote pops up a big yellow screen that cannot be missed. The implication is that changing notification on the chat room tool for similar obviousness to InstaNote would be good.

While interviews report the salient information, notice that it is a representation (representational state and its interpretation) that is the critical difference here. Direct observation and analysis with DCog, with its direct focus on representations might have suggested the same result. Further, it would also help examine the possibility of such a change, even if the interviewee had neglected to provide this information. Again, for both AT and DCog part of the usefulness of the approach relies on good ethnographic work to have exposed the data for observation. I think DCog makes the implication of the specific representations and the processes that coordinate them more evident. While communicating that is difficult because there are not named categories to put the observations in, the low level nature of analysis exposes the necessary information more directly.

### 3.5. THE ROLE AND EVIDENCE OF PROCESS

Although a DCog analysis is centered on cognitive processing, AT keeps process explicitly in the foreground by diagramming relations between elements within the activity system. In the analysis above we see how the framework of that diagram provides a way to walk through the various elements and their relationships to explain the workings of the system. By using it as a conceptual framework all interactions between subject and mediating artifact are considered together. Similarly, there is an assumed interaction between subject and object, which is mediated by artifacts. Stating it thus is an oversimplification of course, but again this is an example of how naming theoretical constructs helps AT both describe and communicate.

In a DCog analysis the mediating role of an artifact is not assumed between an object and subject, although it may occur that way. The telephone is a mediating artifact, but so is the content of the conversation over the telephone.[9] Breaking down an interaction into its respective representational states and examining the processes helps us understand more clearly what those representational states are doing, but it requires more work to clearly identify and talk about broad classes of processes. The process implicit in the interposition of artifacts mediating between subject and object is not assumed in DCog. The primacy of the *processing* of representational states for the cognitive work to get done and a theoretical language that is common to both individuals and artifacts, places all the processing at low level. Any one artifact may just help propagate a representation along a processing path, or it may play a different role with its own internal processing. After the low-level DCog analysis, and during the process of reconstructing a narrative of the cognitive work, a collection of phenomena might be labeled with the phrase ‘division of labor’ or ‘mediating artifact’. But what motivates the analysis, critique, and any subsequent design choices is not a separation of phenomena into named

categories, but an analysis of the specifics of coordination of representational media.

Unfortunately, this low level focus may mean the reader overlooks the higher-level process. AT's basic structure posits certain kinds of process interrelationships, which are implicit even when the analyst may not make them explicit. DCog obscures those relationships somewhat by focusing on the lower level. Where higher level phenomena have been drawn out of an analysis, such as Hutchins analysis of learning on both an individual and organizational level (Hutchins, 1995, Ch. 6 & 7) constructs are usually not named and propagated through the theory. In addition, the lack of theoretical names makes it harder to bring it up to the higher level, requiring more descriptive work on the part of the analyst to explain the process at a higher level.

## 4. What kind of a theory does CSCW really need

I hope that I have shown that while both AT and DCOG are cognitively based theories they operate very differently. They direct our focus as analysts to different aspects of their respective unit of analyses based on both what they deem as important to analyze (scope of the unit) as well as how they perform the analysis, and how they communicate it. While I have found DCog very useful for analyzing how an organizational memory in a call center works, others have clearly found AT useful to understand other aspects of knowledge management in similar settings.

At the beginning of this essay I proposed four attributes against which to judge the utility of a theory and I raised three questions about how we might view the success of theory for CSCW. I want to revisit those attributes now. Taken together, these four attributes relate to a juxtaposition of evidence that philosophers refer to as warrants and acceptance.

> Warrant is a normative notion; the warrant-status of a proposition is a matter of how good or bad the evidence with respect to that proposition is. Acceptance is a descriptive notion; the acceptance-status of a proposition is a matter of the standing of the claim in the eyes of the scientific community or relevant sub community: rejected as definitely false; regarded as a possible maybe worthy of further investigation; acceptable as definitely true; as established unless and until something unexpected turns up, and so on. Ideally, the acceptance-status of a claim will vary concomitantly with its warrant-status. (Haack, 1998)

What a theory can warrant is not all that is necessary to make it useful. We need some way to compare and situate one setting against another in the natural history sense, and that will provide us with a taxonomy of field settings and their characteristics. From this we might build towards understanding phenomena, which in turn might become a better understanding of group work, if not a theory of it.

Being able to evaluate the warrant-status of theoretical propositions made about group work, coupled with a taxonomy of instances, would help us go beyond

description to prediction. Such a taxonomy would begin to build a cross referenced description like we see in the Nardi et al. (2002) comparison of certain kinds of small groups. Theory helps describe the characteristics that tell us how these groups are the same, as well as how they are distinct, as in Nardi et al.'s intensional networks and Zager's coalitions in this issue. It would be useful to go beyond description of phenomena to prediction. For example, prediction of what might happen to these small groups if a piece of technology was introduced. In the case of the distinction of intensional networks from coalitions, we might want to answer whether the introduction of a certain mobile telephone application would be different in terms of adoption, use, or effect on the group. I believe this would require a theory that encompassed an understanding of group work and technology's role in it.

Instead, what we have now in the CSCW community mostly centers on acceptance-status. Haack's quote alludes to how the acceptance-status of a proposition reflects the how of science because of the process of evaluation. For example, the clarification of the notion of intensional networks in Nardi, et al., in comparison to Zager's coalitions and Engeström's knotworking (Engeström et al., 1999) is a step along the road to the acceptance status of a proposition that might read as "intensional networks as a concept within group work is defined by these characteristics".

Equally, it might be clarified as "intensional networks have this special meaning and is distinct from the theoretical concepts *coalition* and *knot* in these ways" only within Activity Theory. Within a particular theory adoption of a term is evidence of its acceptance status in the community. Thus we see among these papers references to using one or another flavor of activity theory: elements of Engeström, or Bødker's. Adoption and reference to other papers in the community, even when they are not within the same theoretical tradition, also speak to acceptance status – whether it is in agreement or denial (cf. Fjeld et al.'s (2002) comment about distributed cognition). In this we see how theories and frameworks give us ways to describe the world we observe and a common vocabulary for comparison.

What I have illustrated in the comparison of AT and DCog is that each has value for the field of CSCW, but neither will satisfy all our needs. AT is powerful because it names and names well, but this both binds and blinds its practitioners to see things in those terms. Going back to the glasses metaphor, AT brings "anointed" objects of analysis into high relief while backgrounding and obscuring those not called out by the theory. DCog, in contrast, is more flexible. What is anointed by the theory is the observed qualities of the representational states and media, and observing how processes bring those media into coordination. It's more likely to catch the significance of a situation being analyzed because it's more data-driven. DCog is perhaps a more direct route to aid design because it presents data at the right level to impact the design of representations and processes, but we know that this is a hard problem for any approach.

Because of how constructs are named AT is perhaps better at supporting discourse within a community that understands the theory, but both AT and DCog, like ethnomethodology, have to fall back on the "thick" descriptions of their ethnography to explain their findings to others 'not in the know'. While learning AT is difficult because of the complexity of its conceptual structure, DCog is similarly difficult because its power is largely in its application. To quote Rogers (2000) "However, those who hope it will provide them with a methodology to derive system requirements are often disappointed. There is no 'off-the-shelf' method that can be followed, ... because the approach does not lend itself to step-by-step procedures. Whilst it is relatively straightforward to learn about the properties and processes of a distributed system through reading Hutchins and other distributed cognition analyses, it is much more difficult to apply the method to an actual setting" (p. 15).

For the moment it seems we must be satisfied in CSCW with a theoretic grab bag. This places the burden on us as readers to understand each other's theoretical frameworks and as writers to be careful in our presentation so as not to so shorthand the work that it becomes obscure to only those in the know. This special issue, and the dialog it can engender, is a start on that path.

## Acknowledgements

This paper has benefited from the comments of several intrepid readers. My thanks to Mark Ackerman, Tom Erickson, Wendy Kellogg, and Yvonne Rogers who told me what to say and not say. Mostly I listened to them. I am also indebted for previous discussions with Ed Hutchins, Kjeld Schmidt, Dave Randall, John Hughes, the DCog Lab (circa 1990–1995), Pim Techamaunvivit, Matt Holloway, Michael Tschudy, Victor Kaptelinin, and Arne Raeithel. Talking with Arne (through the cloud of cigarette smoke at InterCHI93) was the first time I realized that there might not be so much different between DCog and AT.

## Notes

1. From here on I will use distributed cognition abbreviated as DCog to refer to Hutchins' (Hutchins, 1995b) theory, while written out it will refer to the general phenomena of cognition being distributed.
2. One confusion with the term – distributed cognition – is the use of the name to cover a variety of approaches. The focus here is on Hutchins' use of the term, as distinct from Zhang and Norman's external cognition (Zhang and Norman, 1991) or the more general use found in the chapters in Salomon's book "Distributed Cognitions" (Salomon, 1993). Similarly the articles in this issue frequently identify the *flavor* of AT as Engeström's, Bødker's or Kuutti.
3. Key members of both theories have been debating and educating each other for the last 10–15 years. Yrjö Engeström, Mike Cole, and Ed Hutchins have been involved in several reading groups and have team taught classes at University of California, San Diego. Evidence of cross-pollination can be seen in Hutchins 1986 article about mediation in Mind, Culture, and Activity

(Hutchins, 1986) and Cole and Engeström's (Cole and Engeström, 1993) chapter in Salomon's book Distributed cognitions (Salomon, 1993).

4. In the Introduction to Activity, Consciousness, and Personality Leont'ev (1978) states:

   > It is almost a hundred years since world psychology has been developing under conditions of crisis in its methodology. Having split in this time into humanistic and natural science, descriptive and explanatory, the system of psychological knowledge discloses ever new crevices into which it seems the very subject of psychology disappears. ... Negligence and skepticism in relation to the general theory of the psyche, and the spreading of factologism and scientism characteristic for contemporary American psychology (and not only for it) have become a barrier blocking the road to investigating the principal psychological problems. It is not difficult to see the connection between this development and the disillusionment resulting from unfounded claims of the major Western European and American Trends that they would effect a long-awaited theoretical revolution in psychology.

5. Or not map in the case of ethnomethodology.
6. Especially see Cole and Engeström (1993) where they discuss the early western use of distributed cognition.
7. But compare this with Tolman and Piekkola's (1989) analysis of Dewey's 1896 article on the reflex arc which they argue parallels and anticipates the development of activity theory.
8. Engeström's term is contradictions.
9. While language can clearly be a mediating artifact in AT, it is still mediating between subject and object.

## References

Ackerman, M.S. and C.A. Halverson (1998): *Considering an Organization's Memory.* Proceeding of the Conference on Computer Supported Co.

Ackerman, M.S. and C.A. Halverson (1999): *Organizational Memory: Processes, Boundary Objects, and Trajectories*. *32nd Hawaiian International Conference on Systems Science*. Maui, HI: IEEE.

Ackerman, M.S. and C.A. Halverson (2000): Re-Examining Organizational Memory. *Communications of the ACM*, vol. 43, no. 1, pp. 58–64.

Bardram, J. (1997): Plans as Situated Action: An Activity Theory Approach to Workflow Systems. In John A. Hughes, Wolfgang Prinz, Tom Rodden and Kjeld Schmidt (eds.): *ECSCW 97: Proceedings of the Fifth European Conference on Computer Supported Cooperative Work*. Kluwer Academic Press, pp. 17–32.

Bardram, J. (1998): *Designing for the Dynamics of Cooperative Work Activities*. Conference on Computer-Supported Cooperative Work, Seattle, WA, ACM.

Barthelmess, P. and K.M. Anderson (2002): A View of Software Development Environments Based on Activity Theory. *Computer Supperted Cooperative Work*, vol. 11, nos. 1–2, pp. 13–37.

Bentley, R., T. Rodden, P. Sawyer, I. Sommerville, J. Hughes, et al. (1994): *Ethnographically-informed Systems Design of Air Traffic Control*. Conference on Computer Supported Cooperative Work, Toronto Canada, ACM.

Beyer, H. and K. Holtzblatt (1998): *Contextual Design: Defining Customer-centered Systems*. San Francisco, CA: Morgan Kaufmann.

Blumer, H. (1986): *Symbolic Interactionism: Perspective and Method*. Berkeley: University of California Press.

Bødker, S. (1991): *Through the Interface: A Human activity Approach to User Interface Design.* Hillsdale, NJ: Lawrence Erlbaum.

Button, G. (ed.) (1991): *Ethnomethodology and the Human Sciences*. Cambridge: Cambridge University Press.

Button, G. and P. Dourish (1996): *Technomethodology: Paradoxes and Possibilities*. ACM Conference on Human Factors in Computing Systems CHI '96, Vancouver, BC. Canada, ACM.

Carstensen, P.H. and M. Nielsen (2000): Guiding the Thrust! Analytical Concepts in the Service of Coordination Support Systems. In R. Dieng, A. Giboin, L. Karsenty and G.D. Michelis (eds.), *Designing Cooperative Systems: the Use of Theories and Models. Proceedings of the 5th International Conference on the Design of Cooperative Systems (COOP '2000)*. Amsterdam: IOS Press.

Clark, A. (1997): *Being There: Putting Brain, Body, and World together Again*. Cambridge, MA: The MIT Press.

Clases, C. and T. Wehner (2002): Steps Across the Border – Cooperation, Knowledge Production and Systems Design. *Computer Supported Cooperative Work*, vol. 11, nos. 1–2, pp. 39–54.

Cole, M. and Y. Engeström (1993): A Cultural-historical Approach to Distributed Cognition. In G. Salomon (ed.), *Distributed Cognitions: Psychological and Educational Considerations*. Cambridge: Cambridge University Press, pp. 1–46.

Collins, P., S. Shukla and D. Redmiles (2002): Activity Theory and System Design: A View from the Trenches. *Computer Supported Cooperative Work*, vol. 11, nos. 1–2, pp. 55–80.

Engeström, Y. (1987): *Learning by Expanding: An Activity-theoretical Approach to Developmental Research*. Helsinki, Orienta-Konsultit Oy.

Engeström, Y., R. Engeström and T. Vahaaho (1999): When the Center Does not Hold: the Importance of Knotworking. In S. Chaiklin, M. Hedegaard and U. Jensen (eds.), *Activity Theory and Social Practice: Cultural Historical Approaches*. Aarhus, Denmark: Aarhus University Press.

Engeström, Y., R. Miettinen and R.-L. Punamaki (eds.) (1999): *Perspectives on Activity Theory*. New York: Cambridge University Press.

Fjeld, M., K. Lauche, M. Bichsel, F. Voorhorst, H. Krueger and M. Rauterberg (2002): Physical and Virtual Tools: Activity Theory Applied to the Design of Groupware. *Computer Supported Cooperative Work*, vol. 11, nos. 1–2, pp. 153–180.

Fitzpatrick, G., S. Kaplan and T. Mansfield. (1996): *Physical Spaces, Virtual Places and Social Worlds: A Study of Work in the Virtual.* Proceedings of the Conference on Computer Supported Cooperative Work. Cambridge MA: ACM, pp. 334–343.

Flor, N.V. and E.L. Hutchins (1992): Analyzing Distributed Cognition in Software Teams: a Case Study of Collaborative Programming During Adaptive Software Maintenance. In J. Koenemann-Belliveau, T. Moher and S. Robertson (eds.), *Empirical Studies of Programmers*. Norwood, NJ: Ablex.

Frohlich, D. and P. Luff (1989): Conversational Resources for Situated Action. *Proceedings of the SIGCHI Conference on Wings for the Mind*. Austin, TX: ACM, pp. 253–258.

Gardner, H. (1984): *The Mind's New Science: A History of the Cognitive Revolution*. Basic Books.

Garfinkel, H. (1967): *Studies in Ethnomethodology*. Engelwood Cliffs, NJ: Prentice-Hall.

Glaser, B.G. and A.L. Strauss (1967): *Discovery of Grounded Theory: Strategies for Qualitative Research*. Chicago: Walter De Gruyter.

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

Haack, S. (1998): *Manifesto of a Passionate Moderate*. Chicago: University of Chicago Press.

Halverson, C.A. (1995): *Inside the Cognitive Workplace: New Technology and Air Traffic Control*. Unpublished dissertation, Cognitive Science Department, University of California, San Diego.

Heath, C. and P. Luff (1996): *Documents and Professional Practice: 'Bad' Organisational Reasons for 'Good' Clinical Records*. Conference on Computer Supported Cooperative Work. Boston, MA: ACM.

Holder, B.E. (1999): *Cognition in Flight: Understanding Cockpits as Cognitive Systems*. Dissertation, Cognitive Science, University of California, San Diego.

Hollan, J.D., E.L. Hutchins and D. Kirsh (in press): Distributed Cognition: A New Theoretical Foundation for Numan-computer Interaction Research. *ACM Transactions on Human-Computer Interaction*.

Hutchins, E. (1988): The Technology of Team Navigation. In R.K.J. Galegher and C. Egido (eds.), *Intellectual Teamwork: Social and Technical Bases of Cooperative Work*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Hutchins, E. (1990): Organizing Work by Adaptation. *Organization Science*, vol. 2, no. 1, pp. 14–39.

Hutchins, E. (1991): The Social Organization of Distributed Cognition. In L. Resnick and J. Levine (eds.), *Perspectives on Socially Shared Cognition*. Washington, D.C.: APA Press.

Hutchins, E. (1995a): How a Cockpit Remembers Its Speeds. *Journal of the Cognitive Science Society*, vol. 19.

Hutchins, E. and B. Hazlehurst (1990): Learning in the Cultural Process, Institute of Cognitive Science, University of California, San Diego.

Hutchins, E. and T. Klausen (1992): Distributed Cognition in an Airline Cockpit. In D. Middleton and Y. Engeström (eds.), *Communication and Cognition at Work.* Beverly Hills, CA: Sage Books.

Hutchins, E.L. (1995b): *Cognition in the Wild.* Cambridge, MA: MIT Press.

Hutchins, E.L. and L. Palen (1993): *Constructing Meaning from Space, Gesture, and Talk.* Discourse, tools, and reasoning: situated cognition and technologically supported environments., Lucca, Italy.

Katzenberg, B. and J. McDermott (1994): *Meaning-making in the Creation of Useful Summary Reports.* Proceedings of the Conference on Computer Supported Cooperative Work. Chapel Hill, NC: ACM, pp. 199–206.

Korpela, M., A. Mursu and H.A. Soriyan (2002): Information Systems Development as an Activity. *Computer Supported Cooperative Work*, vol. 11, nos. 1–2, pp. 111–128.

Kuutti, K. (1991): The concept of Activity as a Basic Unit of Analysis for CSCW Research. *Proceedings of the Second European Conference on CSCW*. Amsterdam.

Kuutti, K. (1996): Activity Theory as a Potential Framework for Human-computer Interaction Research. In B.A. Nardi (ed.), *Context and Consciousness: Activity Theory and Human-computer Interaction*. Cambridge, MA: MIT Press, pp. 17–44.

Leont'ev, A.N. (1978): *Activity, Consciousness, and Personality*. Prentice-Hall.

Marr, D. (1983): *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. New York: W. H. Freeman & Co.

Miettinen, R. and M. Hasu (2002): Articulating User Needs in Collaborative Design: Towards an Activity-Theoretical Approach. *Computer Supported Cooperative Work*, vol. 11, nos. 1–2, pp. 129–151.

Nardi, B. (1996a): Studying Context: A Comparison of Activity Theory, Situated Action Models, and Distributed Cognition. In B. Nardi (ed.), *Context and Consciousness: Activity Theory and Human-computer Interaction*. Cambridge, MA: The MIT Press.

Nardi, B.A. (ed.) (1996b): *Context and Consciousness: Activity Theory and Human-computer Interaction*. Cambridge, MA: The MIT Press.

Nardi, B.A., S. Whittaker and H. Schwarz (2002): NetWORKers and their Activity in Intensional Networks. *Computer Supported Cooperative Work*, vol. 11, nos. 1–2, pp. 205–242.

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

Norman, D.A. (1991): Cognitive Artifacts. In J.M. Carroll (ed.), *Designing Interaction. Psychology at the Human-computer Interface.* Cambridge: Cambridge University Press, pp. 17–38.

Norman, D.A. and S.W. Draper (eds.) (1986): *User Centered System Design*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Perry, M. (1997): *Distributed Cognition and Computer Supported Collaborative Design: The Organisation of Work in Construction Engineering*. Department of Information Systems and Computing, Brunel University, UK.

Popper, K.R. (1992 (reprint)): *Logic of Scientific Discovery*. Routledge.

Raeithel, A. and B. Velichkovsky (1995): Joint Attention and Co-construction: New Ways to Foster User-designer Collaboration. In B. Nardi (ed.), *Context and Consciousness: Activity Theory and Human-computer Interaction*. Boston: The MIT Press, pp. 199–233.

Rogers, Y. (2000): Recent Theoretical Developments in HCI: Their Value for Informing System Design.

Rogers, Y. and J. Ellis (1994): Distributed Cognition: an Alternative Framework for Analysing and Explaining Collaborative Working. *Journal of Information Technology*, vol. 9, 119–128.

Rouncefield, M., J. Hughes, T. Rodden and S. Viller (1994): Working with Constant Interruption: CSCW and the Small office. *Conference on Computer Supported Cooperative Work*. Chapel Hill, NC: ACM.

Sacks, H. (1992): *Lectures on Conversation*. Oxford: Blackwell.

Sacks, H., E. Schegloff and G. Jefferson (1978): A Simplest Systematics for the Organization of Turn-taking in Conversation. In J. Schenkein (ed.), *Studies in the Organization of Conversational Interaction*. New York, NY: Academic Press.

Salomon, G. (ed.) (1993): Distributed Cognitions. *Learning in doing: Social, Cognitive, and Computational Perspectives*. Cambridge: Cambridge University Press.

Schiff, L, Van House, N. and M. Butler. (1997): Understanding Complex Information Environments: A Social Analysis of Watershed Planning. *Proceedings of the Conference on Digital Libraries*. Philadelphia, PA: ACM, pp. 161–168.

Schmidt, K. and C. Simone (1996): Coordination Mechanisms: Towards a Conceptual Foundation of CSCW Systems Design. *Computer Supported Cooperative Work, The Journal of Collaborative Computing*, vol. 5, nos. 2–3, pp. 155–200.

Shapiro, D. (1994): *The Limits of Ethnography: Combining Social Sciences for CSCW*. Computer Supported Cooperative Work, Chapel Hill, NC, ACM.

Simon, H.A. (1990): *The Sciences of the Artificial*. Cambridge MA: MIT Press.

Spasser, M.A. (2002): Realist Activity Theory for Digital Library Evaluation: Conceptual Framework and Case Study. *Computer Supported Cooperative Work*, vol. 11, nos. 1–2, pp. 81–110.

Strauss, A.L. and J.M. Corbin (1998): *Basics of Qualitative Research: Techniques and Procedures for Developing Grounded Theory*. Sage Publications.

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

Tolman, C.W. and B. Piekkola (1989): John Dewey and Dialectical Materialism. *Activity Theory*, vol. 1, nos. 3/4, pp. 43–46.

Zager, D. (2002): Collaboration as an Activity. *Computer Supported Cooperative Work*, vol. 11, nos. 1–2, pp. 181–204.

Zhang, J. and A. Norman (1991):Distributed Cognition: The Interaction of Internal and External Representations.