

Organizational Memory as Objects, Processes, and Trajectories: An Examination of Organizational Memory in Use

Mark S. Ackerman¹

School of Information and Department of Electrical Engineering and Computer Science
University of Michigan
Ann Arbor, MI 48104 USA
ackerm@umich.edu
<http://www.eecs.umich.edu/~ackerm>

Christine Halverson¹

Social Computing Group
IBM Research
650 Harry Rd, D2-252
San Jose, CA 95120 USA
+1 408-927-3620
krys@us.ibm.com

0. Keywords

Organizational memory, knowledge management, distributed cognition, memory reuse, information reuse, contextualization, trajectories of information, boundary objects, corporate memory, collective memory.

1. Abstract

For proper knowledge management, organizations must consider how knowledge is kept and reused. The term *organizational memory* is due for an overhaul. Memory appears to be everywhere in organizations; yet, the term has been limited to only a few uses. Based on an ethnographic study of a telephone hotline group, this paper presents a micro-level, distributed cognition analysis of two hotline calls, the work activity surrounding the calls, and the memory used in the work activity. Drawing on the work of Star, Hutchins, and Strauss, the paper focuses on issues of applying past information for current use. Our work extends Strauss' and Hutchins' trajectories to get at the understanding of potential future use by participants and its role in current information storage. We also note the simultaneously shared provenance and governance of multiple memories—human and technical. This analysis and the theoretical framework we construct should be to be useful in further efforts in describing and analyzing organizational memory within the context of knowledge management efforts.

2. Introduction

Central to knowledge management efforts is reuse -- of explicit information, tacit knowledge, or lessons learned -- within an organization. Intuitively, the reuse of previously stored information, or *organizational*

memory, is critical to the success of modern organizations. Yet much remains unanswered about organizational memory. We know relatively little about, for example:

- How do organizations, as collectivities of people, remember and forget?
- Where do knowledge-based organizations store information to be reused? What does it even mean for an organization to "store" information? How are people part of that "stored" information?
- How are memories accrued, located, and then used?
- How is information abandoned?

These are all areas of significant practical and theoretical merit.

If CSCW as a research area wishes to construct organizational memory systems as part of knowledge management in the context of organizational life, developing the ability to theorize at an appropriate level (Strauss 1991) will be required to guide efforts. Theoretical conceptions of organizational memory, however, are due for an overhaul. The term organizational memory has been unnecessarily restricted to only a few uses centered around particular technologies; yet, memory appears to be everywhere in organizations.

Accordingly, we wish to step back in this paper and reexamine organizational memory to find suitable underlying theoretical constructions.² Our intent is not (yet) to create technological organizational memory systems. Instead our goal is to examine and to understand where memory exists within an organizational setting

The major portion of this paper consists of a detailed analysis of two hotline calls, progressively describing each call and the work activity surrounding the call. We found our field site, a telephone helpline for personnel issues, a particularly useful domain for studying organizational memory; largely because the repetition of questions and answers facilitated our observation and recognition of particular patterns. A descriptive examination, rooted in an organizational field study, that allows both a micro-scale analysis along with a theoretical development is most likely to be useful to later system construction. Accordingly, we necessarily have restricted our examination to a very small scale of operation because of the detail required. We approach this examination from distributed cognition theory (Halverson 1995; Hutchins 1995; Rogers 1992), described below, because its theoretical language spans the diverse manifestations of organizational memory – from private to public and small scale to large.

The paper begins with a brief synopsis of the organizational memory literature and its need for empirically-based analyses of organizational memory. The next sections provide a brief introduction to distributed cognition theory. We follow this with a description of our field site and ethnographic data collection. As mentioned, the majority of the paper is an explication of two hotline calls along with the construction of an appropriate set of theoretical concepts. As with most ethnographically-based studies, the explication of the data and the findings are interwoven in the paper. The paper concludes with general implications for organizational memory research and a look at future research directions.

3. Organizational Memory

Organizational memory, while it serves as an entry point for considerations of knowledge management has many varying, and occasionally competing, definitions in the literature. As will be seen below, few of the literature's definitions and conjectures, to date, rest on empirical examinations of organizational memory within a context of use.

There have been repeated calls for empirical studies.³ Walsh and Ungson (1991), for example,³ note that:

Despite the general use of the term organizational memory, it is not clear that we have understood the concept or its implications for the management of organizations. To date, a

myriad of unexamined conjectures has defined a concept that has even served as a basis for prescriptive management advice. (pp. 84-85)

Studies of organizational memory fall into three groups. The first are theoretical studies, providing a general description of memory at the organizational level. Generally, papers theorizing about organizational memory also theorize at a very grand scale, not relying on empirical data. They are often founded on a March and Simon model of organizations as information processors using search (March and Simon 1958). For example, Stein and Zwass (1995) argued that the organization uses its mnemonic functions to attain its goals; no detailed examination of how this is accomplished was provided. Huber (1990) argues that organizational learning and memory support would be useful, but the paper does not distinguish clearly what constitutes organizational memory. Smith (1994) uses a similar model. At a similar level, Sandoe and Olfman (1992) examined how remembering and forgetting might be useful organizational mechanisms, and Jennex and Olfman (2002) investigated the efficacy of organizational memory systems.

The second set of organizational memory studies has examined the use of particular computer systems designed to augment an organization's memory. Most such studies have largely focused on the technology systems designed to replace human and paper-based memory systems. For example, Ackerman studied the Answer Garden system in a series of technical studies (Ackerman and Malone 1990; Ackerman 1993) and field studies (Ackerman 1996). Answer Garden attempted to "grow" an organizational memory by providing a way to augment an information database and help facility. As users asked questions about their work or other tasks, the experts who answered them could add the questions and their answers to the information database, thus growing the organizational memory in a manner useful to the organizational members. A redesign of Answer Garden, based on additional case studies of information seeking (Ackerman and Palen 1996), was reported in (Ackerman and McDonald 1996) and (Ackerman and Starr 1996).

Similarly, Conklin examined the use of his Corporate Memory design rationale system in (Conklin and Begeman 1988) and (Conklin 1992). Morrison and Weiser in a series of papers (Morrison 1993; Morrison and Weiser 1996) examined the use of a project memory system called TeamMemory, and Mandviwalla, Clark, and Sandoe (Mandviwalla et al. 1995) studied collaborative writing as a form of organizational memory. Other systems, such as Grassroots (Kamiya et al. 1996), have also been built and are being studied. Commercially, Lotus Notes has been a successful technology for constructing organizational knowledge, and Intranet uses of the World Wide Web are now heavily deployed. Recently, Robinson et al. (2000) examined the use of a work "diary" in a factory setting (a paper mill). They uncovered the use of asynchronous "remark-aloud" behavior, similar to that on a trading floor or communication facility, where one worker remarks on a problem or issue for later use.

These studies of specific systems have been limited. They often rely on narrow definitions of organizational memory or organizational tasks. More importantly, examining particular systems, often prototypes, brings a concomitant difficulty in theoretical generalization.

The third set of organizational memory studies examines the use of organizational memory through fine-grained, field-based empirical studies. Although these would be the most useful for both designing systems and for understanding the nuances of practice, relatively few such studies exist. In one study, Hughes and King (1992) examined the use of paperwork in an accounting office. They found that the organizational members used paperwork to account for the work accomplished so far. Other studies of paper artifacts have found similar uses for paper. Sacks (1994), in an organizational ethnography, examined organizational learning among software engineers; he found that his study participants learned from other engineers, not the code itself. Orlikowski (1992) noted the importance of incentives for information reuse in her Lotus Notes field site. Many studies (e.g., (Cicourel 1990; Bannon and Bødker 1997); Constant et al. 1994; Suchman and Wynn 1984) echo these findings; they suggest that use of information is extremely nuanced, exception handling is the norm, and social arrangements are critical.

Despite these foundational studies, there has been little follow-on to the carefully detailed, field-based, empirically-grounded analyses upon which to base either theory or system design. In this paper we focus on theoretical development based on such an analysis. We agree with Bannon and Kuutti (1996) that despite the conceptual problems, there is something arresting about the idea that an organization can remember. They state:

...that such a concept is appealed to across a wide range of studies, even if its definition is disputed, is testimony to the fact that even if people cannot agree on what exactly the term means, there must be some set of issues that can be subsumed under its umbrella that people feel are important and worth discussing. (pp. 156-157)

The relative dearth of empirical examinations specifically about organizational memory is unfortunate. The need for systematic work to examine organizational memory is even more pressing than for many other organizational concepts. Organizational memory as a concept lends itself to a number of theoretical problems. Relying on an mistaken characterization of how individual memory works propagates the disservice when applied to groups and organizations. Organizations are hardly a single, unified entity, as the metaphor implies, nor is their memory. Furthermore, organizational memory, as a collective function, must also be socially constructed, maintained, and driven.

Organizational memory as a theoretical concept, then, should be grounded in studies within organizational field settings; that is, within a context of everyday use. With such a basis, system construction can be instituted upon these empirically-determined insights – instead of just building systems blindly. Indeed, others (Bannon and Kuutti 1996; Hughes et al. 1995) have strongly argued that more empirical field-based studies are required to understand organizational memory.

In the next section we introduce distributed cognition theory, explain why we think it is useful, and outline how to go about an analysis that includes tasks, individuals, technologies, and social arrangements.

4. Distributed Cognition Theory and Analysis

Organizational memory is a messy problem. Instead of clear-cut mechanisms making it easy to map the real world to some theoretical expectations we see contradictions or border cases in any carefully detailed field study. To find a deeply grounded theoretical coherence, we needed come to terms with all the convergent and divergent observations. When we first began looking at this data, we were struck by similarities with Halverson's work with air traffic controllers (Halverson 1995) and Hutchins' work with large ship navigation (Hutchins 1995). These similarities, coupled with the obvious cognitive aspects of talking about memory, led us to look at Hutchins' distributed cognition theory for our analysis.

Distributed cognition theory (Hollan et al. 2000; Hutchins 1995), described below, has a theoretical language that can span the diverse manifestations of organizational memory – from private to public and small scale to large. Both based in, and reacting to, cognitive science's characterization of cognition as computation, and anthropology's emphases on culture, distributed cognition theory serves as a useful conceptual framework for the analysis of human organizational systems. To explain why we find this theory so useful for examining organizational memory we will first outline the basic theoretical constructs and then discuss what this means for studying organizational memory. Moving from the theoretical to the practical, we introduce the methods, our setting, and outline the scope of the analysis.

4.1. Cognition Is Distributed

In the last two decades there has been a growing recognition that removing human behavior from its surrounding environment carries with it as many dangers as advantages. This has been true in a number of fields, including CSCW (Suchman 1987). Since the mid 1990s there has been a parallel recognition in

cognitive science that cognition is more fruitfully viewed as distributed, rather than assuming that it is the property of only an individual mind (Clark 1997; Hutchins 1995; Salomon 1993).

Several researchers have used the term *distributed* to mark this difference in perspective from more traditional approaches to cognitive science (Cole and Engeström 1993; Hutchins 1991; Hutchins and Klausen 1996; Norman 1991; Salomon 1993). How cognition is distributed varies by the circumstances and the researcher. At the same time there is increasing recognition that societies and organizations have cognitive properties that are different from those of individuals.

We focus on the theory of distributed cognition developed by Hutchins beginning in the mid-80s, published in his book (1995) and continuing today (Hollan et al. 2000). As a term, *distributed cognition* refers to the perspective that cognition occurs in a distributed manner. *Distributed cognition theory* is concerned with the organization and operation of cognitive systems; that is, with the mechanisms that make up cognitive processes, which result in cognitive accomplishments. Furthermore, it re-situates cognition in the socio-cultural context.

As Hutchins points out (1995), there were many costs when cognitive science set aside history, culture, context and emotion from the study of cognition. One mistakes the properties of the system for those of the individual, putting too much focus on the boundary between the individual and everything else. In turn “when one commits to the notion that all intelligence is inside the inside/outside boundary, one is forced to cram inside everything that is required to produce the observed behaviors” (p355). Instead distributed cognition takes the position 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. 2000), p175) This opens up our notions of cognitive processes to a much wider variety of mechanisms than classic symbol manipulation (Newell and Simon 1972; Simon 1990). This is one of the two key theoretical commitments that define distributed cognition.

A second and related tenant of distributed cognition is its commitment to a unit of analysis defined in relation to the complex phenomena being observed. In *Cognition in the Wild* (Hutchins 1995), Hutchins shows the information processing in a navigation team varies with the context and circumstances. Solo watch standing involves one individual in structured interaction with various artifacts via well-established procedures and routines. In contrast, entering a harbor is a high-tempo activity that requires the effort of several people coordinating with specialized tools and each other. While the overall functional behavior exhibited by the system is the same, that is safe navigation of the ship, the means change.

The power of these two tenets, taken together, is extremely useful for studying how organizational memory really happens in organizations. Distributed cognition gives us a way to approach the processes concerned with information use, while still considering that the specifics of how information is represented, and where it is represented, might affect the outcomes of both storing and retrieving information. There are three basic interactions that involve the distribution of cognitive processes and affect the phenomena of organizational memory. These are the distributions of cognitive process (a) between internal and external representations, (b) through time, and (c) across members of a group. In the sections below we discuss each of these.

4.1.1 *Coordination of Cognitive Processing Between Internal and External Representational States*

When the unit of analysis is widened to include other observable cognitive agents we become more aware of how representations are coordinated among agents and across boundaries. Cognitive processing necessarily requires processing that is both internal and external to individuals, as well as some artifacts.

For example, Hutchins deeply analyzes the socio-cultural-technical system of navigation aboard a large military vessel. Building on Simon’s conception that problem solving is re-representing the problem until the solution becomes transparent, Hutchins lays out how the cognitive computation of the system solving the “problem” of navigation can be detailed. As he states: “The fix cycle⁵ is accomplished by the *propagation of*

representational state across a series of *representational media*. The representations of the position of the ship take different forms in the different media as they make their way from the sighting telescopes of the alidades to the chart.” (Hutchins 1995), p. 117) That propagation requires the careful coordination of information among individuals and artifacts so that the representational state being propagated is meaningful.

4.1.2 *Temporally Distributed Cognitive Processes*

Culture is a process that accumulates partial solutions to frequently encountered problems (Hollan et al. 2000), p. 178).

Cognition is distributed through time in such a way that the products of earlier events transform the nature of later events. Culturally defined routines, such as a routine about the “right way” to do something, are examples of this. Artifacts can also embed a history of their use. For example, a ruler is marked with lines that have meaning that affect their current use. Saying that a piece of paper is 8 ½ by 11 inches assumes that the measurement, as well as the number system is meaningful to both the speaker and the hearer. That meaning is based on their previous use and the history of how those markers came to have meaning with respect to socially agreed measurements. The artifact of the ruler then serves as a way to save and carry forward a solution to the problem of measurement. In some cases the solution is more specific, for example, in the specialized artifacts for navigation, such as the Mercator projection (Hutchins 1995), which helps solve certain kinds of navigation problems.

Both the ruler and the map show one way that artifacts, which save and carry forward a partial solution, can affect the process of an event far removed in time. To be successful however, such artifacts are knit together with the mechanisms of current social practice. In the case of a well-defined problem like navigation, where the parameters that are important have been well understood for over a century, the social practice has been routinized and carried forward with the artifact. Thus, contextualizing the artifact (which is necessary for its use) has been carried forward with it in time.

One possible inference is that this embedding of solutions in artifacts and their accompanying practice is only possible where systems are usually solving repetitive well-defined problems such as ship navigation (Hutchins 1988; Hutchins 1995), and aviation (Hutchins 1995; Hutchins and Klausen 1996). Such problems are often complicated but well structured. In addition, the repetition means that the high cost of proceduralizing and embedding the solution in an artifact is well worth the effort. However, these kinds of solutions also happen on a local scale, as Halverson observed in air traffic control (Halverson and Ackerman 2003; Halverson 1995). One easily recognizable example is so-called cheat sheets, such as the phone list of a work group or notes on how to back up files to a shared file system. In these cases the problem to be solved is repetitive and well-defined within a very localized setting (perhaps only one individual) and over a comparatively short time span.

4.1.3 *Socially Distributed Cognition*

The application of these abilities must be “organized” in the sense that the work done by each component ability must be coordinated with that done by others. (Hutchins 1995), p. 154)

Cognitive processing can be distributed across members of a group as well. If artifacts act as repositories, encoding partial solutions to problems, then how might the cognitive properties of a group differ from those of an individual? Hutchins points out that since Durkheim at the turn of the last century there have been largely programmatic assertions that groups might have cognitive properties differing from those of an individual. Roberts (1964) began to demonstrate how such properties might be manifested at the individual and group levels. Roberts suggested, as Hutchins notes, “that a cultural group can be seen as a kind of widely distributed memory. Such a memory is clearly more robust than the memory of any individual and undoubtedly has a much greater capacity.” (Hutchins 1995), p. 177) Roberts studied four American Indian tribes and compared their ability at information retrieval at the tribal level, linking the observed properties to particular features of

group organization. Hutchins builds on this by discussing how both communication and learning processes differ between individuals and groups involved with the navigation of a ship.

Hutchins argues that different social organizations help a system solve different sorts of problems. Furthermore, the organization, for example of the team and artifacts used to navigate, can change in response to the changing needs of the system. Thus, while existing social structure may proscribe the path of certain processes it does not restrict potential paths completely. The “cognitive accomplishments of any two groups might depend entirely on differences in the social organization of distributed cognition and not at all on differences in the cognitive properties of individuals in the two groups.” Hutchins explores this further by investigating the difference between learning at an individual and organizational level.

Our observations resonated with many of Hutchins’ findings. As we discuss below, organizational memory is necessarily distributed through time and across not only group members, but groups. In addition, a distributed cognition analysis exposes the coordination necessary between representational states that we found critical for day-to-day operations.

4.2. Methods for Data Collection and Analysis

It is one thing to talk about observing the coordination of representational media, but what does that mean in terms of data collection and analysis? In this section, we discuss how the theoretical stance affects both methods and analysis.

4.2.1 Analysis

To begin, distributed cognition theory frames the problem in terms of examining a cognitive system in terms of its functions. The first task in a distributed cognition analysis is to identify how a functional system works, good and bad (Rogers 1992). There are low-level functional systems often embedded in larger functional systems that contribute to the overall activity being observed. Functional operation is decomposed into smaller units of analysis that make sense with respect to the particular system. While appearing in some cases to be straightforward task decomposition, it is more essentially an event driven segmentation.

As with other cognitive theories, distributed cognition identifies the observed informational inputs entering, as well as those outputs leaving, a system. The functional definition of the unit of analysis helps with this identification. Stepping down in grain size, inside the system the focus is about how information is represented, and how these representations are transformed, combined, and propagated through that system in order to produce the system’s observable behavior (Simon 1990). It is the detailing of representational states and processes that helps the analyst to understand much of the system processing as it involves transitions between humans and artifacts. One records the representational state, the material media on which it is instantiated, as well as the processes that transform it. Video recording is particularly helpful here because of the ability to review the record repetitively and glean details that are observable but difficult to catch in just “one-pass”.

Thinking of organizations as cognitive systems is not new, of course. What is new is the examination of the role of the material media in which representations are embodied, and in the physical processes that propagate representations across media. Applying the cognitive science approach to a larger unit of analysis requires attention to the details of these processes as they are enacted in the activities of real persons interacting with real material media. (Hutchins 1995), p 266)

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. We believe that looking at the phenomena of organizational memory is well supported by taking this essentially cognitive view of a system, and in our interpretation, giving it a certain social twist. However, before discussing how distributed cognition can be applied to a specific organizational memory, we must first present the details of the field setting.

5. Setting and data collection

This study is based on field observations of a telephone hotline group (called HLG here) at a well-established company, CyberCorp, headquartered in Silicon Valley. HLG answers human resource questions for CyberCorp, primarily about benefits and personnel policies for the company's thousands of employees. In general, telephone hotlines are of interest in the study of organizational memory, largely because they are so information intensive. HLG agents have to start forming their answer within 45 to 60 seconds while simultaneously listening to the caller's elaborations and information. Many answers came directly from the hotline member's memory; hotline questions tend to be repetitive. There is also a great need for additional information sources: Facts must be double-checked, new questions arise, and answers become obsolete with new conditions.

The field study took place over a period of 18 months. A variety of data collection methods were used, including direct observation, video, semi-structured interviews, and social network analyses. Here we describe and analyze the responses to two calls captured on video. Both calls involve Joan⁶. Joan was an experienced agent, having been at CyberCorp for five years and at HLG for one year.

Because of privacy reasons, only one side of the conversations was taped. Joan, however, described each call fully to the camera. (Moreover, we are concerned here with Joan's actions, rather than the intent of the caller.) Throughout, Joan appeared to be natural and relaxed.

Joan and the other agents work in cubicles that are open to a central corridor (Figure 1). Like many control room settings (e.g., (Seifert and Hutchins 1992), the cubicles are close enough to easily hear the activities of other agents. Her chair faces away from the corridor between the cubicles. This arrangement is important for accomplishing her work.




Figure 1: Joan's workplace. The left-hand portion shows how the cubicles are related to each other, as well as the shared terminals. The right is a larger picture of Joan's office.

Each agent has two monitors, where she uses a number of software packages. The telephone to the right of her monitors is another computational system that plays a role in the work. These computational systems and software will be discussed at length below.

5.1. Applying distributed cognition to HLG

Like other functional human systems, HLG has the property that many of the system properties are directly observable. We can bound the portion of the system to be analyzed, based on the observed function and initially its temporal limits. Within this unit, we expose information about the task, its resources, and organization.

In the case of HLG, the cognitive system's purpose is to answer or solve the caller's problem. The caller can be seen as the input to this system, and her question is a representational state traveling via the material media of the phone to one of the HLG agents. For each call the unit of consideration may vary. We might initially *bound* the system to include the caller, Joan, the telephone, and other materials (or *resources*) available to Joan. While her use of the telephone is an indicator of where the process is occurring, it does not completely define the limits of the task. In general, physical, resource, and temporal limits bound the functional system observed at any point in time. In the analyses that follow some calls are bounded temporally from when the receiver is picked up to when it is set down, a matter of a few minutes. In other cases, the temporal extent is considerably longer but the extent is framed by the focus on the solution of a single problem.

In the case of HLG, a distributed cognition analysis highlights how the work of the system is organized and reorganized to meet changing needs. A plethora of information is available in HLG. In some cases the required information is easily accessible, while in others it must be located or uncovered. In any case, information must be monitored, managed, and communicated in order to do the job. The details of this information use form the basis for our description and analysis.

In the remainder of the paper we will present the results of the analysis. We present the analysis of two HLG telephone calls. (We observed at least 300 calls, taped approximately 60, and chose 10 for analysis in a manner similar to (Hutchins and Palen 1993) and (Halverson 1995).) Both of the calls presented here are necessarily simple. Space limitations prevent us from presenting longer calls, but more importantly, even simple calls turn out to be surprisingly complex and rich when fully analyzed. Indeed, we selected these specific calls to best show the organization and function of organizational memory, as these simple calls are easier to explain but still involve the same underlying issues.

6. Correcting a Database Entry

Below is a transcript from videotape of an HLG call. A walkthrough of this call will begin to unpack the working of organizational memory. We reproduce and discuss the beginning of this call at length; later, we will summarize sections of the call to be more concise. Because of privacy reasons, only one side of the exchange was taped.

6.1. Answering the Phone

The first four turns of the call are routine, but even so they demonstrate critical aspects of the organization's memory. The call begins with a routine set of conversational markers, Joan's standard opening (Drew and Heritage 1992). Turn 1 consists of Joan greeting the caller, and acknowledging that the caller was forced to hold for some period of time. While the telephone system automatically routed a call to her as soon as she finished the previous call, the caller was forced to wait for an available agent. Turn 2 acknowledges the caller's request. In this case an employee has called because a benefits provider does not show her as receiving their benefit. This benefits provider, Eye and Vision Associates (EVA), will not process the employee's claim (or rather, allow the employee to purchase eyeglasses at a discount) until this is corrected.

1	HR Helpline. This is Joan. Thanks for holding. (Joan nods slowly to herself.)
---	---

2	You called EVA directly? I'll find out what our system shows for you. Oh I see. Oh, so right, right...yeah... (Encouragingly) Let me do this. (In an aside tone) The system is being very slow, so it'll take a second to get a new record up for you.
---	--

The employee says that she called EVA directly and that EVA did not show her as having coverage. Joan must then check whether the employee, Michelle, actually showed coverage within the CARL database: This database, built from payroll data, shows the employee benefits for each employee in CyberCorp.

At the end of turn 2, Joan starts a new call-tracking record in the CAT (Call Tracking) system for the new call, closing out the old one that she had not quite finished. Turn 3 is more complex, as Joan does several things simultaneously. She asks for the relevant information, namely the caller's employee number. Joan has now opened a new CAT record for the call and checked the employee phone list for the caller. She double-checks the employee's social security number for later use with the CARL database. As she is confirming the employee information with the caller, Michelle, she is also typing the information into the call-tracking record. She also writes Michelle's social security number on a piece of scrap paper, because she must go to the CARL terminal to look up the employee data. When she is finished obtaining and double-checking this information, she asks Michelle to hold, knowing that the next step will take a minute or two.

3	I'm sorry. Oh what a shame. Well you know, these funny things. (More officially) What is your employee number?
4	Okay, what I'm going to do is, I'm going to check the, uh, the CARL. Is this Michelle? Let me check your social security number is the one I have, 1-8-2-4-zero-7-3-5-5. Okay, what I'm going to do is to check the system and find out what your coverage is, and if it's not showing...coverage on CARL, uh, then we'll, hmm, we'll have to do something else, okay? Hold just a second, Michelle

At the end of turn 4, Joan goes to the CARL terminal, which is physically in another location behind her cubicle. (There are separate terminals for two different employee databases, and part of Joan's routine is selecting the correct database.) She types, looking at the piece of scrap paper, and pulls up the appropriate record. This takes 23 seconds. After obtaining the data, she visibly pauses and stares abstractedly at the ceiling.

6.2. Many Small Memories

We will proceed with the rest of the call below, but even the first four turns show many cognitive, social, and institutional arrangements in the organization's memory. In her execution Joan uses not one monolithic memory, as many technocentric models would have it. Instead, she uses many small memories.

To reframe, the call was triggered by the *telephone system's short-term memory* of the group's activity. The system state shows that Joan's station is free, and the call falls to her. Hearing the caller's information, Joan uses her *own memory* as she then types that information into the *CAT record*. (For convenience, other uses of Joan's memory are omitted here.)

Joan then takes the information in CAT and reproduces it onto a *piece of paper*, creating a mobile form of memory. She appears not trust to her own memory, but resorts to something that can help her reconstruct the circumstances of the call. She then types the information from the paper into the CARL system. CARL is a typical type of *organizational memory*, a corporate database with employee records. Joan places the CARL output onto the paper again.

Within just these three turns, Joan has used three separate software systems (CAT, the telephone system, and CARL), her own memory, and scratch paper, all of which maintain representational state for Joan, the group, or both. Figure 2 below graphically displays the order of processing of those memories; it consists largely of transferring information from memory to memory in order to answer the query. Within what is so




Figure 2. The trajectory of representational states through memories in the first 4 turns of Joan's call.

far a very structured process, Joan's cognitive work consists largely in knowing which memories to trigger. Little processing has been done on the actual information.

Although in Figure 2 we present the trajectory that representational states take through various memories as an individual process, there are actually multiple group and organizational processes occurring. In distributed cognition theory, expanding the boundaries of the analysis is required as we enlarge the relevant task in order to understand a memory's use in its full organizational context. This will become clearer at the end of the call, when Joan indexes the call-tracking record and changes the telephone state, but at this point, Joan is already embedded within these processes. For example, Joan uses CAT, the call-tracking system, seemingly as a short-term memory aid. Its major use, however, is to provide other agents with the ability to reconstruct the history of a caller's problem. In addition to maintaining a history of calls for the group, the CAT program also creates transformed, longer-term memory in the form of statistics, based on the indexing done by the agent (later) during her wrap-up period. These statistics are used by management to govern the group's future behavior, as has been typical in organizations since the late nineteenth century (Yates 1989; Yates 1990). The telephone system, in addition to being a primitive form of group memory that coordinates and paces the HLG agents' activities, also creates summary statistics for the group and the organization as a whole.

To recap, even within these four turns, Joan uses five discrete memories. Sometimes the memory used is individual and private; sometimes it is group and public. But all of these memories must be used together seamlessly (or nearly so) to create an organizational product (the product being not only the solution to the call but all of the institutional arrangements surrounding it). The density and connectedness of memories used as resources in this environment are remarkable.

The call, of course, continues past Joan's finding the employee's record. Next Joan must determine what to do with the facts she has uncovered.

6.3. Handling an Exception

As mentioned, Joan pulled up Michelle's record within the CARL database. After this, Joan visibly paused and stared at the ceiling. As will be seen below, Michelle should have EVA coverage according to her CARL record; yet, EVA shows no such coverage. Presumably, Joan is trying to figure out what to do about this unusual discrepancy.

After five seconds, she begins a conversation with one of the senior telephone agents who had been wrapping up her call. In turns 5 through 15, Joan asks the senior agent, Nichole, how to proceed.

Joan appears in turns 5 through 15 to be trying to understand whether to escalate the problem to the benefits group. Organizationally, the HRG is dependent on other groups to handle more complicated or complex situations; these are called *escalations*. The senior agent confirms that she should escalate, and tells her what information is necessary to properly create the escalation. (In the following, the angle brackets <> indicate a section of the tape that was indistinct. The slashes \ / and / \ indicate overlapping conversational areas on the tape; a = sign indicates that there was no appreciable pause between the two words.)

5	Joan	Nikki, can I get your...
6	Nichole	Uh, huh. (Nichole nods her head and walks towards Joan.)
7	Joan	If this employee, uh, called EVA and they say they show no coverage for her, but she's showing, showing active on CARL. She's has confirmed no problems, but what should I do as far as the referral?
8	Nichole	Call, call them [EVA] yourself and find out what they want.
9	Joan	Okay.
10	Nichole	Then escalate it to the <Benefits>. EVA is asking for this information on this <indistinguishable>. (Official tone, as though acting out a scenario) We show it as having such-and-such. That way /we\
11	Joan	\Okay./ Okay, /I\
12	Nichole	\can/ tell <them>... what EVA is looking for.
13	Joan	I can tell the employee it's showing there's not, shouldn't be a problem, so=
14	Nichole	=we're trying to fix it.
15	Joan	I'm trying to get a <feel> for it. Thanks, Nik.

Nichole tells Joan, then, to escalate the problem to the Benefits group and to obtain from EVA what they need to resolve the problem. Joan walks away from Nichole at the end of this interaction, and returns to her seat.

6.4. Organizing the Work

We showed earlier how Joan's processing was dependent on many small memories, including her own. However, solving a problem may not be dependent solely on an individual's cognition and the artifacts (memory or otherwise) within the immediate environment. These turns show how Joan's work is socially organized as well as the role of memory in that organizing.

Unlike ship navigation (Hutchins 1988, 1995) or air traffic control (Halverson 1992, 1994, 1995) with their standard operating procedures, HLG has fewer pre-specified routines. Yet, as Pentland (Pentland 1991)

observed in his study of a software hotline and as Katzenberg et al. (1996) observed in hospital situations, the HLG group has developed a set of informal routines that can be combined flexibly to solve a large range of problems. Indeed, the HLG manager repeatedly mentioned during the study that he was trying to balance flexible diagnosis and service with transaction efficiency. There were numerous minor task reallocations during the study period as the group attempted to juggle the two demands and build their repertoire of small routines.

In this case, Joan has a routine to uncover a discrepancy between what an employee wishes to have for benefits and what he or she currently has – she looks it up. Joan can then explain that discrepancy as well as potential solutions to the employee. The discrepancy, however, between the internal CyberCorp databases and the benefit provider's database is not routine. Joan later demarked the situation as neither usual nor unusual, and as such, one without a given routine.

Instead, Joan relies on Nichole. Nichole was considered the “expert” on more obscure situations; HLG agents would often pose difficult questions to Nichole. Like Hutchins' Navy navigation crew, the HLG is organized such that the more senior agents have served in all simpler positions. (The HLG roles are less differentiated than are those for shipboard navigation, but the coverage of experts' knowledge over novices' tasks is the same.) In addition, Joan liked asking questions of other agents as a way of obtaining needed information, so this interaction suited her information seeking style. The intent of many organizational memory systems would be to replace Nichole but without losing the supposedly reusable information. Instead, many exceptions may be better serviced by task experts; as we see here even with this short exchange.

The production is divided not only among group members but among groups as well. Relaxing the boundaries of the task shows Joan's work is also connected organizationally across other group's processes. There are two larger tasks that are invoked in these turns.

First, Joan and HLG are dependent on another organizational group to maintain the database used to verify an employee's benefits. Considering the call more broadly demonstrates the important set of organizing arrangements required to imbue and inscribe the memory with authenticity and veracity (Anderson and Sharrock 1993). The HLG agent relies on the payroll group for the correctness of the information, when they create and maintain an employee record, which serves as a boundary object (Star 1989). Joan knows none of the details of the record's creation or maintenance; almost all of the context has been lost. She does not know whether there are problems with the employee's employment or whether there are extenuating circumstances. Indeed, it is assumed that the CARL database is more authoritatively correct than the employee database (maintained by the accounting department), since CARL is more authoritative for deducted benefits.

In addition, escalations are problematic organizationally because there is a fine line between taking action inappropriately or incorrectly and between wasting the other group's time. In fact, there was a fair amount of tension between the Benefits group and HLG. HLG felt that Benefits looked down on them for not knowing the Human Resources (HR) subject area sufficiently, and HLG agents wished to be regarded as professionals. Indeed, HLG perceived themselves as the future of the HR profession, as did the CyberCorp management. But this was a future unwanted by most of Benefits, since hotline work was perceived by the Benefits professionals as leading to HR deskilling in CyberCorp. Therefore, inappropriate escalations sometimes furthered political tensions. Since escalations were always to some extent problematic, the procedures around them often shifted during the study, thus the need for the question to Nichole.

6.5. Playing a Hunch

In the next segment of the call, Joan returns to her chair, reconnects her telephone headset, and begins to speak with Michelle again. Procedurally, all Joan must do is to tell Michelle that she will call EVA and escalate the problem. However, this is derailed by a side discussion over Michelle's incorrect telephone number.

16	Michelle, when you spoke with the folks at EVA, what, what is it that they said to you?
17	Okay, because you do in mine. ...So I don't know what the situation is.
	[Segment of the call deleted; Joan is dealing with Michelle's phone number.]

In the omitted portion of the transcript, Michelle reports that she cannot seem to correct her telephone number in the employee phone list, even though she has tried many times. She discusses this at some length with Joan, asking her to change it. At first analysis, the incorrect telephone number appeared to be extraneous to this call, but it is an interesting side conversation. The employee, having found someone to correct her employee records, now asks Joan to correct her telephone number in all CyberCorp databases. We believe that to Michelle, it appears that she has found the person that is in charge of correcting the appropriate memory, whether it is her telephone number or her benefits status. Joan cannot, because the responsibility for changing something so seemingly mundane as a telephone number lies with another group. In the actual production, there are distinct organizational boundaries to the memories. To the organizational member, however, these boundaries are arbitrary and frustrating.

Joan continues with her explanation of the escalating process (as it needs to be seen by the employee). However, this is derailed several times by a side discussion about obtaining an electronic form for EVA. This side discussion triggers off a recollection and hunch in Joan.

24	Okay, what we'll do is, I'm going to call EVA (voice rising in question tone) directly and find out what the situation is, and, uh, it's possible they might need some other information that the benefits department can provide for them, but either way, you'll get a call back from somebody by the end of the day tomorrow and let you know what the status is, uh, go ahead...
25	Oh. (surprised) I can request a form for you, but the, the thing is if I request the form and for some reason they don't have coverage for you they're going to reject it. So what I'll do is I'll do two things. Uh, EVA, did you, did you call the 800 number? Did you call, did the hotline direct you to that number?
26	Right
27	Oh, I see. Well, you know what? I think I remember this happening once before, Michelle, and since my system showed that, that the employee that called was covered, they said, okay, I'll put that in the system and we'll go ahead and send the form out ... so who knows? Did...
	[Segment of the call deleted; Joan and Michelle discuss Michelle's address for the EVA form.]

Finally, Joan ends the call with the employee. Interestingly, Joan has formulated a plan of action that attempts to solve the problem, rather than following Nichole's advice literally. In this, she uses Nichole's statements as advice to formulate her own activity, rather than a requirement.

At the end of her conversation with the employee, Joan turns to the camera and explains the call. In the time that it takes Joan to explain the situation, she pulls up EVA's telephone number (using a computerized rolodex program) and dials the call. In fact, she appears to end her explanation when the call is connected. In the final section of the call, Joan convinces the EVA agent that Michelle is a CyberCorp employee and that she should have EVA benefits. (This has been compressed for space reasons.)

	You don't [show her on your system]? No. Is it, is it possible for me to, to uh, verify she's on our (rising, questioning) system? And you can go ahead and send the form to her? Excellent. I have an address (rising, questioning), if that would help you?
--	---

After finishing her call with the EVA agent, Joan wraps up by indexing the call in the CAT system and then signaling to the phone system that she is ready for a new call.

At the end of the call, one might surmise that Joan now has a new routine, calling EVA directly and getting an employee her benefits. She could assume that this kind of transaction is routine, since the EVA agent treats it as relatively straightforward. However, as Joan says: "I think I remember this happening once before." There is no reason to suppose that this will not be recreated again from a very hazy human memory.

We next turn to a discussion of the call as a whole and its implications for organizational memory.

6.6. Context vs. Contextualizing

The term "context" is often used as a catchall to denote the wealth of organizational and cultural knowledge one effortlessly brings to bear in a given situation. Joan shows in this call that she has knowledge of how her world works and that this knowledge is necessary for getting her work done. Within the call, she displays an understanding of her limitations, the range of knowledge within HLG, Nichole's role, and the complex relationship of HLG to Benefits and to EVA.

If one were primarily interested in implementing a technocentric view of organizational memory or knowledge management, with their immediate emphasis on capture and reuse, one might assume that all Joan needs is more information, whether in a computer system or on paper. Instead, the story is more complex. We see Joan moving seamlessly through this real setting with its many process states, using the memories and other artifacts that she judges to be critical to finishing her task. To understand her use of memory resources in accomplishing this, however, it will be important to tease apart the concurrent use of contextualization, decontextualization, and recontextualization in obtaining a solution to the situation. We do this more fully in the second call.

6.7. The Missing Bark

The need for context, and Joan's contextualization of the environment, is revealed in the order and access of various resources. As with Sherlock Holmes' famous hound (which remained silent while the supposed crime was being committed), the interesting thing about this call is all of the things that did not happen. Joan did not stumble or falter, searching for her next step. She needed help in the face of a breakdown, but she carefully and quickly determined a resolution. In her selections, she showed an enormous understanding of her environment and the role of artifacts and people within it.

On the other hand, Joan did not use paper or on-line documents to formulate her plan of action for this situation. The key was not that detail was missing from the on-line sources. If this were so, the solution would be indeed to add even more information on-line. To the contrary, the key to Joan's lack of use is the overwhelming amount of detail available. Since an HLG agent must have an answer within 45-60 seconds, it is impossible to sift through megabytes of information. A search with terms "EVA" or "enrollment" might turn up dozens of documents. Even if the retrieval is ranked (i.e., the search engine evaluates the likelihood of fit), the agent does not have the time to consider the retrieved materials.

The information use that does occur is situated (Suchman 1987). Joan does not view all sources equally. She went to Nichole, rather than to another agent. She avoided one database for another. Despite Nichole's advice, she relied on her own memory. Like the other agents, Joan considers some sources as authoritative.

Sources have different costs, psychological or financial, associated with their access. She selects and chooses resources in her environment, according to her immediate understanding of the situation.

Importantly, the caller, Michelle, does not need to know the space within Joan is operating. All she needs is her problem solved.

6.8. Memories as Boundary Objects

In a distributed cognition view, an organization's memory consists of many states, instantiated by people and artifacts, all within a single system allowing the participants to get their work done. As Joan works to solve Michelle's problem, we see the process is apparently losing contextual information in many places – from payroll to HLG, from Nichole to Joan, from HLG to Benefits, from Joan to Michelle (in explaining the process), from Joan to the EVA representative, and finally from Joan to the CAT record.

Some of these differentiations correspond to boundaries with respect to the propagation of representational state and its re-interpretation. These boundaries often correspond with organizational divisions. Three of the four individuals involved in the call are all employees of CyberCorp. However, their different roles imply different meanings for the same representations. For the caller Michelle, Joan is the expert at HLG who will solve her problem. Pushing the notion of social boundaries to its extreme, even though Joan and Nichole are in the same department, Nichole's acknowledged expertise creates a boundary, and thus different contextualizations between her and Joan. A comparison of these boundaries with the propagation of representations through the processing of the call illustrates how the cognition involved (i.e., the use of memory) is constrained by social arrangements.

As representational state propagates across individuals, inter-organizational, and intra-organizational boundaries, it must necessarily lose some of its context. As Star (1989) points out, boundary objects in an organization work because they necessarily contain sufficient detail to be understandable by both parties, but at the same time, neither party must understand the full context of use by the other. As boundary objects, artifacts and other potential resources are given to other people or in this case, stored for later use by others. This requires the information to be decontextualized. Otherwise, the secondary users will drown in unnecessary, unhelpful, or conflicting data. (One may also wish to hide or obscure some institutional and social arrangements from the other group.) As an example, the payroll records are necessarily standardized and stripped of much informational content before they can be given to groups such as HLG. Decontextualization, and perhaps commodification, must be expected by those who will try to use the memory (Schmidt and Bannon 1992); Ackerman 1993).

To reuse a memory, the user must then recontextualize that information (Schmidt and Bannon 1992; Ackerman 1993; Oakshott 1983; Carr 1961). The information, if not supplied by the same individual, must be reunderstood for the user's current purposes. Elsewhere we have written about the recontextualization problem in organizational memory systems (Ackerman 1996), but the call shows that a difficulty, or even outright inability, in recontextualization, would make the memory useless or nearly so. For example, reused information must be imbued with attributes of veracity, authenticity, and even status (Anderson and Sharrock 1993); it must be provided with everyday organizational context.

Joan understands that the CARL database record is authoritative; she knows enough of its meaning within HLG's context. Joan also knows she can ignore Nichole's advice, based on her understanding of HLG, Benefits, and EVA. Rather than adopt Nichole's advice – find out what EVA wants and then escalate – Joan proceeds to convince EVA herself that Michelle is an employee.

6.9. The Complexity of Reuse

Obviously reuse is possible. Joan's call showed many memories reused by Joan, other members of the group, and the organization as a whole. We have also shown that these memories can be rearranged and reused in different ways, depending on the production that is required.

The above sections, however, cautioned that the informational requirements for reuse are not trivial. For reuse, decontextualized information must be recontextualized sufficiently by the secondary user, according to his or her purposes. This is most likely to be done within a familiar process with little time lapse (Ackerman 1996). In distributed cognition terms, the cost of reuse is most likely to be borne where the traffic or the processing utility is highest – within heavily used or high payoff processes.

The problem with reuse, then, is the coexisting requirement for contextualization, decontextualization and recontextualization. To use information as a memory, one must remove the detail that provides context, making the information into a boundary object. However, at the same time one must consider how others will use it later as a resource in their processes; otherwise, subsequent users of the memory will not be able to properly recontextualize it

On the other side, based on their understanding of their own situations, users of a memory must determine what contextual aspects of that memory are important and whether they have changed. This determination may be nearly impossible, given people's limited view of an organization as well as the decontextualization of the information. Users must then decide how their differing context should effect their use of the memory, and finally they must absorb the memory into their current situation. The simpler the memory, the more likely this all can occur.

6.10. Memories as Processes

To summarize, we described an employee verification, one of the simplest procedures that HLG performs. Yet, this simple procedure exposed a number of interesting aspects of organizational memory. The procedure involved nine different *memories*, and the human agent involved either translated among representational states or reconstructed memory states. We also noted that this process was *simultaneously embedded* within several short-term and long-term memory processes (such as the formation of suitable group statistics and a call-history).

Perhaps most importantly, memories had *mixed provenance*. They were sometimes the provenance of the individual (e.g., Joan's scratch notes) or the group (e.g., the call-handling procedure embedded in the telephone system). But, often enough, the memory that served Joan as individual memory also had a definition as a group and even an organizational memory. The call-tracking record, for example, was used by Joan to aid her short-term memory, by the group to reconstruct a call-history, and by the organization to monitor the activities of the group.

Table 1. Nine distinct memories used in the first call (types of memories).

Technical systems	Databases	Human	Other
CARL	Employee phone list	Joan's memory	Paper
Phone system	Computerized Rolodex	Nichole's memory	
Call Tracking System	Benefit providers database		

We also noted the importance of *boundary objects* and *recontextualization* on the part of the user in this analysis. In this verification call, the boundary object was standardized and recontextualization was straightforward. Both Joan and the caller do not expect there to be exceptions or breakdowns in the process. In the next situation, however, Joan must deal with an exceptional call, one that clearly involves breakdowns. This next situation more clearly delineates what we see as the critical issues in recontextualization, namely trajectories and their projected consequences.

7. Re-reading the record

The second situation (below) involves re-reading a previously created call record. The situation spans a number of hours, and we present only portions of the transcript here.

The “COBRA man”, as Joan labels him in the call, had called earlier in the day, checking his coverage. In the U. S. companies are required to provide medical insurance for laid-off employees for one year. The law is called COBRA, and at CyberCorp the medical insurance was called “COBRA benefits” or simply “COBRA”. For the ex-employee to obtain his COBRA benefits, he had to periodically send in checks to CyberCorp. The ex-employee called to determine whether his current check had been received, because he had already had a problem with an earlier payment. The CyberCorp record, a spreadsheet, did not show his current payment. The ex-employee then pointed out to Joan that this had happened before, and at the time, he had been found to be fully paid. He therefore asked whether this could have happened again. Joan said she would investigate.

As part of her investigation, Joan looked at the earlier CAT records for this ex-employee. As mentioned, the CAT system allows an agent to pull up previous call records for a caller. Generally, this was used as a group memory of previous calls on a problem. Occasionally, the CAT records were used across problems to consider the history of a caller, especially one with an organizationally problematic situation. For example, one such call involved a caller who was trying to weasel an extension to his short-term disability. In the case of the COBRA man, Joan tried to use the previous CAT records to shed some light on his reliability and his current problem. Unfortunately, the earlier CAT records were not sufficiently complete – the ex-employee was providing details that were not in the record.

In the following exchange, Joan double-checks with Lisa about one of the calls, one that occurred three months earlier. Because of the conversation’s length, we include only excerpts here.

Joan begins by restarting the COBRA conversation as Lisa is about to leave for the day. Joan briefly describes the situation and what she can reconstruct of the previous call. Lisa initially says rather sharply that she remembers nothing of the call. Nonetheless, in turn 4 she offers to look at the call record, and by turn 8, she is engaged in the problem.

7	Joan	I've got the calls. There are like four COBRA calls, so I've got the, uh, ..., uh the serial numbers.
8	Lisa	Give me mine. (helpfully) It might ring a bell.

While Lisa reads her record, Joan provides a description of the caller’s state-of-mind. Lisa’s actions cannot be determined from the video, but during this period, there are conversational cues that she is reading her CAT record. She does not provide any cues that she remembers very much, if anything. Joan continues to describe the ex-employee’s comments:

21	Joan	And, you know, the feeling that nothing is happening for him. So he says, so, what if I get hit by a truck when I go outside and I don't have coverage? What's going to happen to me?
----	------	---

The description in turn 21 seems to trigger Lisa's recollection. As the recollection continues in turns 26 through 38, Joan and Lisa show many conversational signs, such as overlapping speech, that they are now acting in concert. In turns 26 through 34, Lisa provides the memory that a person from the Benefits group came down to Lisa's cubicle with the ex-employee's file. Together they determined that there had been an error: The payment had been logged incorrectly. (There is a hint in turn 28's intonation that the Benefits person might have brought the file down with the intention of proving the ex-employee wrong, but had been proven mistaken herself by Lisa.)

26	Lisa	\I actually do remember that./ I remember she had his file. She brought it down=
27	Joan	(agreeing) =mm-hm
28	Lisa	and she went (quickly, mock tone) oops (normal voice), it got logged wrong. He is paid.
29	Joan	Okay.
30	Lisa	Yeah it was our error...
31	Joan	Yeah.
32	Lisa	Not ours, but it's their error in logging it. They didn't log it in...
33	Joan	Well... (Joan starts to put away a cup on her desk)
34	Lisa	When they received it.
35	Joan	I'm going to go ahead and...and do the escalation. I didn't think that you'd have any impact to the call. That's fine. I was just wondering because he's said=
36	Lisa	=I know that CyberCorp had made an error there
37	Joan	/Um, okay\
38	Lisa	\At the ti/me

Joan declares in turn 35 that she will go ahead and send the problem on to Benefits as an "escalation". As before, inappropriate escalations sometimes furthered political tensions. Since escalations were always to some extent problematic, Joan and HLG felt the necessity to confirm the details of a problem before sending it on to Benefits.

Through this exchange, Joan recovered critical details of the previous call. The ex-employee was providing an accurate account. By checking into the record, Joan determined that he was probably telling the truth about the current situation. More importantly, Joan recovered key contextual information about Benefit's analysis of the situation. In order to accurately weigh the evidence to be presented in the escalation, the CAT record, being incomplete, needed extraordinary recontextualization.

Of course, if the record had been available only as a decontextualized object, it may have been possible for the escalation to succeed. However, one can easily imagine that extra work (if not the repeat of the diagnostic effort) would have been required by the organization. Nonetheless, one might wonder why so little was written in the CAT record.

7.1. Trajectories and the memory process

...consider the cognitive properties of the team as a whole. ...It is not the case that two or more heads are always better than one. (Hutchins 1995, p. xvi)

Lisa chose to write up a terse CAT record on the previous call, necessitating Joan's extraordinary recontextualization. Agents wrote very short, terse CAT records when they felt that no one would later need the record. For example, the CAT record where a caller verifies that a person is a CyberCorp employee (e.g., for a mortgage application) might be a single line, if the agent even wrote a record. Employee verifications were not revisited.

In Lisa's call, she assumed that she should hurry through the write-up. HLG gets rewarded for fast wrap-ups, and avoiding copious notes for all calls is required. The ex-employee's payment had been found, the logging corrected, and the problem rectified. The *trajectory* of future use appeared to be obvious.

Trajectory (Hutchins 1995; Strauss 1993) describes the path of an event; in this case, we mean it to be the likely trajectory as anticipated. Strauss (1993) explains trajectory by using the example from medical diagnosis: The treatment of a fever is dependent on its anticipated trajectory. Treatment of a simple fever for a healthy 25-year-old male is quite different than that for an HIV-positive male. The anticipated trajectory of a situation is often critical to organizational processes and is usually based on trajectories of past experience – what Hutchins calls developmental trajectory. Past experience in the development of the practice, the practitioners and the conduct of the activity affect the present of the activity.

The incentives for keeping memory follow the developmental trajectory, the assumed trajectory, and its *projected consequences*. In this case, Lisa assumed that the call would never be referenced again; she had little incentive to write a complete call record. Joan had to deal with the unanticipated (and perhaps unanticipatable) consequences of Lisa's projecting the trajectory incorrectly.

One can perhaps see this even more clearly in Joan's escalation of the problem. She must weigh the event's trajectory and projected consequences to create a correctly formed escalation. As mentioned, she must now escalate the problem to the Benefits group; she must therefore create a boundary object. Slightly over 20 minutes after her conversation with Lisa, Joan begins to rewrite the CAT record for the current call; this will also be sent to Benefits as the escalation record. The trajectory is now that the CAT record will be reused as a boundary object, so that Benefits can understand the problem.

In a lengthy aside to the observer as she prepared the escalation, Joan pointed out the facets of "appropriate" escalation for this case. Clearly part of what Joan must do is to lay out the "facts" of the case. Initially, this is what she traces through as she describes the process of creating an escalation. She has verified that the COBRA man is providing accurate details, and she states this is important.

She [Lisa] did give me good information because, um, he, the employee had told me that he had called before and spoken with Lisa and, um. (pause) The Benefits specialist came down, and said well, I know, he's covered. [...] So that is, uh, something I can put in this call...

She goes on to say that she needs to suggest to Benefits that the ex-employee may be correct.

Lisa's given me information I can share with Sally [a Benefits person] saying, I've spoken to Lisa and she said yeah, it was logged wrong, and you know, you have his file. Is this possible ... that this has happened again, or what is the situation.

So far, she is merely detailing the “facts” as she has uncovered them. However, Joan then continues her comment by explaining why she is spending so much time creating the appropriate escalation record. A considerable amount of time is spent removing extraneous detail from the CAT record. She tells the camera that she must pick her facts with care, removing extraneous detail (i.e., decontextualizing the actual record) and double-checking any detail that remains.

A lot of this information I just take as I'm writing, as I'm, um, taking the call, and it's not just pertinent information to share with the Benefits folk, but that's, you know, having to go back and just kinda clean up the call before we send it over...as an escalation.

7.2. Objects in Processes and Processes in Objects

Joan edits and re-edits the CAT record for the escalation. While this is partly a result of Joan's inexperience with COBRA benefits, the following suggests that she is also rewriting the record to cover any uncertainty or incompetency as well:

Part of the amount of time it takes is really making sure that the information we're sending over is really accurate and ... when there are um, areas, like, I don't live and breathe COBRA information, so what might be important, what I might think is important information, might be nothing to them.

As Joan traces through what makes an appropriate escalation record, her explanation begins to deviate from merely detailing the facts. The escalation record will go to Benefits, as discussed above, a group with which HLG has had problematic relations. Indeed, it may become part of the evidence in any later status conflicts. This escalation record, then, serves as part of HLG's organizational “face”.

Joan points out to the camera that a proper escalation must be complete and accurate - in a word, what Benefits requires and nothing more. This CAT record, and Joan's anticipated consequence, is that the record will be thoroughly scrutinized by Joan's supervisor and by some part of Benefits. Both will judge the record's quality. (Joan was seen as a good employee by her supervisors, although she was unsure why she had not been promoted to supervisory status herself. We have already described Benefits' interactions with the HLG.)

One of the supervisors, um, will check and make sure it's accurate information and the call notes are complete. And she'll say did you check, you know, the system, or did you check this out.

Presumably this judgment will be part of future interactions between Joan and her supervisor as well as HLG and Benefits. In a sense, then, a boundary object is only partially an “object”. It is also an event in many different organizational processes. Later boundary objects (and possibly previous objects) will be understood in light of this particular boundary object and its implications. An organizational member's assumptions about anticipated consequences and potential trajectories will include not only the current situation, but future situations as well.

In summary, this situation shows that raw CAT records are hardly useful. They can jog an individual's memory, but we saw that Lisa needed additional cues to actually recover her memory. They can serve as a group memory, but Joan could not adequately recontextualize Lisa's record. Finally, the CAT records can serve as boundary objects between groups, but this can take considerable work.

Obviously, not all attempts to reuse CAT records were failures. Most reuse appeared effortless and seamless. This example, however, highlights the everyday work that must be done to make reuse appear

effortless in its success. Much of that work involves being able to properly recontextualize a record, given that the person who created it properly understood its later reuse. In other words, later recontextualization succeeds or suffers from earlier assumptions about the record's trajectory.

8. Conclusions

This study has many limitations. As with many ethnographies, the representativeness of HLG and its activities could be argued. We have been able to present only a few simple cases here. Yet, within these limits, we have tried to show that:

- There is no such thing as *an* organizational memory per se, as the organizational memory metaphor attempts to invoke. We have tried to detail how a supra-individual memory works in its use of multiple people and many artifacts. Our analysis also demonstrated the utility of distributed cognition theory in understanding such a supra-individual memory.

These simple calls and their distributed cognition analysis exposed a number of interesting aspects of this organization's memory at a micro-level. We also showed how:

- The employment verification procedure involved nine different memory states and the human agent involved either translated among representational states or reconstructed memory states.
 - Even this simple procedure was a complex case of distributed memory. *Memories* were complexly distributed, interwoven, and occasionally overlaid. They had a *mixed provenance*. Sometimes they belonged to the individual (e.g., Joan's scratch notes) or the group (e.g., the call-handling procedure embedded in the telephone system). But, often enough, the memory that served as individual memory also had a definition as a group and even an organizational memory.
 - While knowledge management largely restricts itself to viewing organizational memory as repositories of experience "objects" that are magically reusable, we have tried to show that it is more fruitful to consider organizational memory as both object and process. Memory is both an artifact that holds its state and an artifact that is *simultaneously embedded* in many organizational and individual processes.
 - The container metaphor is easier to consider computationally, but it is extremely limited organizationally. The distributed cognition view of a network of artifacts and people, of memory and of processing, bound by social arrangements, provides a deeper and ultimately more usable understanding of organizational life. It describes how memory as representational states can be both separated from organizational actors, and is at the same time, necessarily bound to their actions and understandings.
 - The second call highlighted the issues of *decontextualization* and *recontextualization* that are required to effectively turn a memory "object" into a memory process. As memory crosses between groups or even across time, it becomes a *boundary object*, attempting to serve the needs of both creator and reader but lacking the full context of either. To properly serve the reader or re-user of the memory, the creator must properly project the consequences of the memory's later use, or *trajectory*. This can be a difficult matter, although people do it everyday in their work.

This paper has highlighted many of the issues and problems in creating memories that are used and are found usable by groups and organizations. We have also proposed a number of theoretical concepts (Strauss

1991) – distributed memories, simultaneous embedment, mixed provenance, boundary objects, recontextualization, and trajectory – that enabled us to analyze the use of memory in these calls.

However, we have seen, even in these examples, that memories do get created and reused. More work will be required to examine the details of how work activities evolve memories. Our current work examines how memories evolve, becoming resources in similar complex environments. We are also examining how much memories are maintained and reused over time.

Acknowledgements

This paper has benefited greatly from conversations with many people, including Kjeld Schmidt, Ed Hutchins, John King, Kate Ehrlich, Tom Gruber, Tom Malone, John Quimby, Charlie Osborne, Gerhard Fischer, John Hughes, Kari Kutti, Leysia Palen, Yvonne Rogers, Liam Bannon, and Dave Randall. We would like to especially thank Lorne Olfman, who encouraged us to write up our thoughts on boundary objects, trajectories, and memory processes. We would also like to thank Jonathan Grudin, who posed the initial question of when history was useful, and Wayne Lutters, who has furthered our thinking about boundary objects with his valuable insights. We also thank the participants at the HICSS'95 organizational memory tutorial, 1996 HCI Consortium workshop, ECSCW95 organizational memory workshop, and ECSCW97 Fringe Event workshop.

This project has been funded, in part, by grants from National Science Foundation (IRI-9702904 and IRI-0124878), QSI, and the University of California MICRO program.

We also thank Wendy Kellogg and Tom Erickson for their support at IBM.

References

- Ackerman, Mark S. (1993): Answer Garden: A Tool for Growing Organizational Memory. Massachusetts Institute of Technology, Ph.D. Thesis.
- Ackerman, Mark S. (1996): Definitional and Contextual Issues in Organizational and Group Memories. *Information Technology and People*, vol. 9, no. 1, pp. 10-24.
- Ackerman, Mark S. and Halverson, Christine. (2000): Re-Examining Organizational Memory. *Communications of the ACM*, vol. 43, no. 1, pp. 58-63.
- Ackerman, Mark S. and Malone, Thomas W. (1990): Answer Garden: A Tool for Growing Organizational Memory. In *Proceedings of ACM Conference on Office Information Systems*, Cambridge, MA. New York: ACM Press, pp. 31-39.
- Ackerman, Mark S. and McDonald, David W. (1996): Answer Garden 2: Merging Organizational Memory with Collective Help. In *Proceedings of ACM Conference on Computer-Supported Cooperative Work (CSCW'96)*, Cambridge, MA. New York: ACM Press, pp. 97-105.

Ackerman, Mark S. and Palen, Leysia. (1996): The Zephyr Help Instance: Promoting Ongoing Activity in a CSCW System. In Proceedings of ACM Conference on Human Factors in Computing Systems (CHI'96), Vancouver. New York: ACM Press, pp. 268-275.

Ackerman, Mark S. and Starr, Brian. (1996): Social Activity Indicators for Groupware. IEEE Computer, vol. 29, no. 6, pp. 37-44.

Anderson, R. and Sharrock, W. (1993): Can Organisations Afford Knowledge? Computer Supported Cooperative Work, vol. 1, no. 3, pp. 123-142.

Bannon, Liam and Bødker, Susan. (1997): Constructing Common Information Spaces. In Proceedings of European Computer Supported Cooperative Work conference. New York: Kluwer, pp. 81-96.

Bannon, Liam J. and Kuutti, Kari. (1996): Shifting Perspectives on Organizational Memory: From Storage to Active Remembering. In Proceedings of Hawaii International Conference on System Sciences (HICSS-29). Piscataway, NJ: IEEE Press, pp. 156-167.

Carr, Edward H. (1961): What Is History? New York: Vintage Books.

Cicourel, Aaron. (1990): The Integration of Distributed Knowledge in Collaborative Medical Diagnosis. In J. Galegher and R. Kraut (eds.): Intellectual Teamwork: Social and Technological Foundations of Cooperative Work. Hillsdale, NJ: Lawrence Erlbaum, pp. 221-242.

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

Cole, Michael and Engeström, Yrjo. (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.

Conklin, Jeff. (1992): Corporate Memory. In Proceedings of Groupware '92, San Jose, CA. San Mateo: Morgan-Kaufmann, pp. 131-137.

Conklin, Jeff and Begeman, Michael L. (1988): gIBIS: A Hypertext Tool for Exploratory Policy Discussion. In Proceedings of ACM Conference on Computer Supported Cooperative Work, Portland, OR. New York: ACM Press, pp. 140-152.

Constant, David, Kiesler, Sara and Sproull, Lee. (1994): What's Mine Is Ours, or Is It? A Study of Attitudes about Information Sharing. Information Systems Research, vol. 5, no. 4, pp. 400-421.

Drew, Paul and Heritage, John. (1992): *Talk at Work: Interaction in Institutional Settings*. New York: Cambridge University Press.

Halverson, Christine A. (1992): *Analyzing a Cognitively Distributed System: A Terminal Radar Approach Control*. La Holla: Cognitive Science, University of California, San Diego, Masters Thesis.

Halverson, Christine A. (1994): *Distributed Cognition as a Theoretical Framework for HCI: Don't Throw the Baby Out with the Bathwater - The Importance of the Cursor in Air Traffic Control*. Technical Report 94-03.

Halverson, Christine A. (1995): *Inside the Cognitive Workplace: New Technology and Air Traffic Control*. La Jolla: Cognitive Science, University of California, San Diego, Ph.D. thesis..

Halverson, Christine A. and Ackerman, Mark S. (2003): *Yeah, the Rush ain't here yet -- Take a break*. In *Proceedings of IEEE Hawaii International Conference of System Sciences (HICSS'03)*. Piscataway, NJ: IEEE Press.

Hollan, James D., Hutchins, Edwin L. and Kirsh., David. (2000): *Distributed cognition: toward a new theoretical foundation for human-computer interaction research*. *ACM Transactions on Computer-Human Interaction*, vol. 7, no. 2, pp. 174-196.

Huber, George P. (1990): *A Theory of the Effects of Advanced Information Technologies on Organizational Design, Intelligence, and Decision Making*. *Academy of Management Review*, vol. 15, no. 1, pp. 47-71.

Hughes, John and King, V. (1992): *Paperwork*. University of Lancaster, Esprit COMIC Project, COMIC-LANCS-4-1.

Hughes, John, O'Brien, Jon and Rouncefield, Mark. (1995): *Organisational Memory: Or How Can We Sack Mavis But Keep Her Brain*. Lancaster: European CoTech Action Project (Common Organisational Spaces and Organizational Memory), Lancaster University, working paper COST-14.

Hutchins, Edwin. (1988): *The Technology of Team Navigation*. In R. K. J. Galegher, and C. Egidio (ed): *Intellectual Teamwork: Social and Technical Bases of Cooperative Work*. Hillsdale, NJ.: Lawrence Erlbaum Associates, pp. 191-220.

Hutchins, Edwin. (1991): *The Social Organization of Distributed Cognition*. In L. Resnick and J. Levine (ed): *Perspectives on Socially Shared Cognition*. Washington, D.C.: APA Press, pp. 283-307.

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

Hutchins, Edwin. (1995): How a Cockpit Remembers Its Speeds. *Cognitive Science*, vol. 19, pp. 265-288.

Hutchins, Edwin and Klausen, T. (1996): Distributed Cognition in an Airline Cockpit. In D. Middleton and Y. Engestrom (ed): *Communication and Cognition at Work*. Beverly Hills, CA.: Sage Books, pp. 15-34.

Hutchins, Edwin and Palen, Leysia. (1993): Constructing Meaning from Space, Gesture and Talk. In *Proceedings of NATO Workshop on Discourse, Tools and Reasoning: Situated Cognition and Technologically Supported Environments*, pp. 23-29.

Jennex, Murray E. and Olfman, Lorne. (2002): Organizational Memory/Knowledge Effects on Productivity, a Longitudinal Study. In *Proceedings of Hawaii International Conference on System Sciences (HICSS)*. Piscataway, NJ: IEEE Press, pp. 1029-1038.

Kamiya, Kenichi, Roscheisen, Martin and Winograd, Terry. (1996): Grassroots: A System Providing a Uniform Framework for Communicating, Structuring, Sharing Information, and Organizing People. *Computer Networks and ISDN Systems*, vol. 28, no. 7-11, pp. 1157-1174.

Katzenberg, Barbara, Pickard, Fred and McDermott, John. (1996): Computer Support for Clinical Practice: Embedding and Evolving Protocols of Care. In *Proceedings of the Conference on Computer-Supported Cooperative Work*, Cambridge, MA. New York: ACM Press, pp. 364-369.

Mandviwalla, Munir, Clark, Stanley D., II and Sandoe, Kent. (1995): Collaborative Writing as a Process of Formalizing Group Memory. In *Proceedings of 28th Annual Hawaii International Conference on System Sciences*. Piscataway, NJ: IEEE Press, pp. 342-350.

March, James G. and Simon, Herbert A. (1958): *Organizations*. New York: Wiley.

Morrison, J. (1993): Team Memory: Information Management for Business Teams. In *Proceedings of 26th Hawaii International Conference on System Sciences*. Piscataway, NJ: IEEE Press, pp. 122-131.

Morrison, Joline and Weiser, Mark. (1996): A Research Framework for Empirical Studies in Organizational Memory. In *Proceedings of 29th Annual Hawaii International Conference on System Sciences*. Piscataway, NJ: IEEE Press, pp. 178-187.

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

Norman, Donald A. (1991): Cognitive Artifacts. In J. M. Carroll (ed): *Designing interaction. Psychology at the human-computer interface*. Cambridge: Cambridge University Press, pp. 17-38.

Oakeshott, Michael. (1983): *On History and Other Essays*. Totowa, NJ: Barnes and Noble Books.

Orlikowski, Wanda J. (1992): Learning from Notes: Organizational Issues in Groupware Implementation. In Proceedings of Computer Supported Cooperative Work, Toronto. New York: ACM Press, pp. 362-369.

Pentland, Brian T. (1991): Making The Right Moves: Toward a Social Grammar of Software Support Hot Lines. Cambridge, MA: Massachusetts Institute of Technology, Ph.D. Thesis.

Roberts, John M. (1964): The self-management of cultures. In W. Goodenough (ed.): Explorations in Cultural Anthropology: Essays in Honor of George Peter Murdock. New York: McGraw-Hill, pp. 433-454.

Robinson, M., Kovalainen, M. and Auramaki, E. (2000): Diary as Dialogue in Papermill Process Control. Communications of the ACM, vol. 43, no. 1, pp. 65+.

Rogers, Yvonne. (1992): Ghosts in the Network: Distributed Trouble-shooting in a Shared Working Environment. In Proceedings of ACM Conference on Computer Supported Cooperative Work, Toronto. New York: ACM Press, pp. 346-355.

Sacks, Marc. (1994): On-the-Job Learning in the Software Industry: Corporate Culture and the Acquisition of Knowledge. Westport, CT: Quorum Books.

Salomon, Gavriel. (1993): Distributed Cognitions. Learning in Doing: Social, Cognitive, and Computational Perspectives. In R. Pea and J. S. Brown (eds.), Learning in Doing: Social, Cognitive, and Computational Perspectives. Cambridge: Cambridge University Press.

Sandoe, Kent and Olfman, Lorne. (1992): Anticipating the Mnemonic Shift: Organizational Remembering and Forgetting in 2001. In Proceedings of International Conference on Information Systems (ICIS), Dallas. Atlanta: Association for Information Systems, pp. 127-137.

Schmidt, Kjeld and Bannon, Liam. (1992): Taking CSCW Seriously: Supporting Articulation Work. Journal of Computer Supported Cooperative Work, vol. 1, no. 1-2, pp. 7-40.

Seifert, C. M. and Hutchins, E. L. (1992): Error as opportunity: learning in a cooperative task. Human-Computer Interaction vol. 7, pp. 409-435.

Simon, H. A. (1990): The sciences of the artificial. Cambridge, MA: The MIT Press.

Smith, John B. (1994): Collective Intelligence in Computer-based Collaboration. Hillsdale, NJ: Lawrence Erlbaum Associates.

Star, Susan Leigh. (1989): The Structure of Ill-Structured Solutions: Boundary Objects and Heterogeneous Distributed Problem Solving. In L. Gasser and M. Huhns (ed): Distributed Artificial Intelligence. San Mateo: Morgan Kaufmann, pp. 37-54.

Stein, Eric W. and Zwass, Vladimir. (1995): Actualizing Organizational Memory with Information Systems. *Information Systems Research*, vol. 6, no. 2, pp. 85-117.

Strauss, Anselm. (1991): *Creating Sociological Awareness: Collective Images and Symbolic Representations*. New Brunswick: Transaction.

Strauss, Anselm L. (1993): *Continual Permutations of Action*. New York: Aldine de Gruyter.

Suchman, Lucy and Wynn, Eleanor. (1984): Procedures and Problems in the Office. *Office: Technology and People*, vol. 2, pp. 133-154.

Suchman, Lucy A. (1987): *Plans and Situated Actions: The Problem of Human-Computer Communication*. New York: Cambridge University Press.

Walsh, James P. and Ungson, Gerardo Rivera. (1991): Organizational Memory. *The Academy of Management Review*, vol. 16, no. 1, pp. 57-91.

Yates, JoAnne. (1989): *Control through Communication: The Rise of System in American Management*. Baltimore: John Hopkins Press.

Yates, JoAnne. (1990): For the Record: The Embodiment of Organizational Memory, 1850-1920. *Business and Economic History*, vol. 19, pp. 172-182.

Endnotes

¹ Authors in alphabetical order.

² An earlier version of this theoretical description was published in *Communications of the ACM* (Ackerman and Halverson 2000). The earlier version was substantially less complete (less than one-third this length) and designed for a popular audience.

³ We note the difference between empirical and empiricist studies. We are calling for empirically-grounded, interpretivist studies to define organizational memory.

⁵“The basic procedures of navigation are accomplished by a cycle of activity, called the fix cycle, in which representations of the spatial relationship of the ship to known landmarks are created, transformed, and combined in such a way that the solution to the problem of position fixing is transparent.” (Hutchins 1995), p 117)

⁶All participants and their individual attributes have been disguised for publication.