

Sharing Knowledge and Expertise: The CSCW View of Knowledge Management

Mark S. Ackerman, Juri Dachtera, Volkmar Pipek, Volker Wulf

Abstract

Knowledge Management (KM) is a diffuse and controversial term, which has been used by a large number of research disciplines. CSCW, over the last 20 years, has taken a critical stance towards most of these approaches, and instead, CSCW shifted the focus towards a practice-based perspective.

This paper surveys CSCW researchers' viewpoints on what has become called 'knowledge sharing' and 'expertise sharing'. These are based in an understanding of the social contexts of knowledge work and practices, as well as in an emphasis on communication among knowledgeable humans. The paper provides a summary and overview of the two strands of knowledge and expertise sharing in CSCW, which, from an analytical standpoint, roughly represent 'generations' of research: an 'object-centric' and a 'people-centric' view. We also survey the challenges and opportunities ahead.

1. Introduction

One of the major research streams within CSCW has been the role of information in organizational settings. A considerable number of papers have explored the situated role of information within various kinds of collectivities, and many of these are among the great intellectual accomplishments of the CSCW community.

In this paper, we review the CSCW literature on a particular aspect of information - how organizations, or other collectivities, can harness and use the knowledge and expertise that their members have to solve problems or get work done. CSCW has a rich history with this problem.

CSCW has focused on the social practices involved in knowledge sharing, as well as the systems that could support knowledge sharing. While this thematic area is also treated by other research communities, their research epistemologies are quite distinct. We will position them broadly. In Information Systems (IS), typically a business school discipline, this research area is called Knowledge Management. In its mainstream it is largely normatively or prescriptively oriented, and it lacks a technology-design orientation as well as a micro-scale, empirical perspective on the practices of knowledge sharing. In Learning Science, Computer Supported Collaborative Learning (CSCL) investigates the role computers can play in facilitating learning and in improving didactic methodologies. While there has been a recent shift in emphasis and some overlapping work (e.g., Goggins et al. 2013), in its mainstream the CSCL community is oriented towards educational institutions, mainly schools and universities. In Computer Science, knowledge-related topics are typically treated by researchers with an Artificial Intelligence or (Large) Database background. While these communities have been successful in developing interesting algorithms to create and manipulate complex data sets, they generally treated these problems in a strictly technical manner. Empirical explorations into the practice of knowledge-intensive work have been typically lacking in this discourse. Design oriented work in the CSCW community is based on some of the algorithms created in Computer Science,

but the merit of the CSCW community is to have explored the relationship between innovative computational artifacts and knowledge work – from a micro-perspective.

To characterize the specific stance of our community, we will use the terms *knowledge sharing* and *expertise sharing* in this paper. Both terms connote CSCW's spin on the problem - that knowledge is situated in people and in location, and that the social is an essential part of using any knowledge (Ackerman et al. 2002b). Knowledge Management, in mainstream IS terms, often connotes ignoring the nuance and complexity of real organizational and social life, not focusing on the fact that work and knowledge are closely intertwined. As Normark and Randall (2005) also state, far more useful systems can be developed if they are grounded in an analysis of work practices and do not ignore the social aspects of knowledge sharing¹.

Within the CSCW discourse we distinguish here between knowledge and expertise sharing. The term *knowledge sharing* takes a perspective in which externalization of knowledge in the form of computational or information technology artifacts or repositories play an important role. We use the term *expertise sharing* when the capability to get the work done or to solve a problem is instead based on discussions among knowledgeable actors and less significantly supported by a priori externalizations (Pipek et al. 2012; Ackerman et al. 2002b). These perspectives also reflect a certain temporal sequencibility echoing the (technical) developments in the field of Knowledge Management in general. Therefore, we speak of 'generations'² in the following.

The first generation we examine, repository models of information and knowledge, was rooted in documents or computer records. Within CSCW, this had the rubric of organizational memory, although there were many earlier efforts to examine document and information flows. This body of work was concerned with information as an externalized artifact or object, although the information was understood to be within a social context. This work produced many insights into the social practices around knowledge sharing.

The shortcomings of repository models were quickly recognized and empirically investigated both in CSCW and knowledge management practice, and a second generation began to examine tying communication among people into knowledge work. In CSCW, this body of work about expertise sharing bifurcated: There has been research on how to find others in an organization or organizations with the requisite expertise or skills and on online communities (initially restricted in scope and size) that allowed people to share their knowledge (Ackerman and Halverson 2004b). Expertise finders, recommenders, and collaborative help systems were some of the CSCW technical efforts in this direction. As well, there have been many studies of how people helped one another within organizations.

As the span of networks and computational systems expanded, the scope of study expanded to Internet-scale systems and communities. Recent work has also expanded from knowledge and expertise sharing within organizations to include other forms of collectivities, including hobbyist groups (e.g., Torrey et al. 2007), open-source communities (e.g., Singh and Twidale 2008, Dabbish et al. 2012), and social movements (e.g., Saeed et al. 2010). There have also been examinations of information and knowledge use within collectivities of increasing complexity, for example, hospitals (e.g., Zhou et al. 2011), inter-organizational networks (e.g., Pipek and Wulf 2003), and NGOs (e.g., Saeed et al. 2010).

We will cover each of these generations in turn. We will largely follow these generations, and mostly follow chronological order, while recognizing that our structure is, to some extent, merely an expository device.

2. The First Generation: The Repository Model

Discussions of documents, information, and repositories were certainly not new. Libraries have been around for millennia, and early information retrieval systems included document retrieval systems and scientific dissemination of information (SDI) systems, which routed new scientific papers.

However, with the advent of early networked and distributed systems in the late 1980s and early 1990s, interest rapidly grew in organizations to 'harness what they knew'. As one early popular press article gushed:

Such collective knowledge is hard to identify and harder still to deploy effectively. But once you find it and exploit it, you win. (Stewart 1991, p. 44)

CSCW researchers quickly moved in to examine the realities of knowledge management: knowledge sharing in organizations. CSCW, in particular, grew interested in understanding how to use, maintain, and reuse both formal and informal information, perhaps with an emphasis on the latter. Studies included both technical and social-analytic, and often both in sequence. We will start with the former.

2.1 Technical research

The first generation in knowledge management used a repository model; organizations were to build vast repositories of what they knew. (For an early popular account, see Stewart 1991). These repositories were to include digitized manuals, standard procedures, process maps, and so on, but they would also include all sorts of other information including best practices, narratives, e-mail, and other forms of informal information. We will discuss the repository model in the following paragraphs; however, it should be noted that CSCW never really accepted that this model would work in practice (e.g., Bannon and Kuutti 1996). Numerous CSCW studies examined how people actually used information, especially informal information, and even the earliest CSCW systems examined ways of including people in information seeking and use. Nonetheless, efforts in this generation of work did focus on information artifacts that were, to some extent, separable from their authors and users.

CSCW produced a number of repository-model systems. Many of these systems considered technical issues that were worked out enough to produce commercially viable systems; most of the research systems are now obsolete. These systems have to be understood within their historical context.

Early CSCW efforts included attempts to create repositories of design and decision rationales, argumentation, and informal information. These investigations led to systems that primarily supported the sharing of information. Tapestry's (Goldberg et al. 1992) central focus was on sharing through tagging and filtering. Its users could collaboratively tag documents and informal information, thus allowing a continuously filtered stream of material going into a repository to be rerouted back to the users. OVAL (Malone et al. 1995) focused instead on the mechanics of sharing. It provided four primitives for handling

information and sharing: information objects, users' views (or organizations) of their information, links between objects, and semi-autonomous agents that structured information for the user. All of these could be shared. As well, commercial interest manifested itself in systems such as Lotus Notes during this period. Xerox deployed a number of systems, and Orr's important study (Orr 1986; Orr 1996) demonstrated that organizational members, in this case Xerox repair people, could learn from and help one another become more effective, including through technical systems.

In the mid-1990s, a number of CSCW papers gathered under the rubric of organizational memory, which was to examine how an organization could retain knowledge from its past (Walsh and Ungson 1991). Organizational memory later expanded to include other forms of collective memory (e.g., community memories and group memories).

Answer Garden was an early system in this area. Answer Garden (Ackerman 1998; Ackerman and Malone 1990) argued for a repository of informal information, specifically answers to technical questions. However, if users could not find the answer to their question, they could ask their question through the system. Answer Garden would route the question to an appropriate expert, including people as first class information sources. The system also had properties that would be found repeatedly useful. It had matched motivations, with experts getting rid of repeated questions and with users getting answers to their questions. An information repository in the Answer Garden system also 'grew' where people needed information. Yet, some of the standard issues with repository systems were apparent in its field study (Ackerman 1993; Ackerman 1994a), particularly difficulty with motivating users to author and organize the material and to maintain the information and its navigation. Answer Garden 2 (Ackerman and McDonald 1996) (AG2), with a more nuanced view of expertise, tackled the problem of context. It tried to keep answer-providing local as much as possible, so context could be preserved. It assumed expertise was variable, and everyone could provide some help. AG2 also assumed that people nearby an asker would know more about local context and might be better at explaining than might experts. However, if no one nearby could provide an answer, AG2 incrementally escalated questions to larger and larger groups of people (or to paid consultants) through a range of computer-mediated communication mechanisms. Pipek and Wulf (2003) summarized many of the issues in Answer Garden and behind the intuitions that led to AG2, they also pointed out the difficulties of reuse and the organization of the information in repositories over time, especially when context changed. Organizational memory, as a concept, was re-examined in Ackerman and Halverson (1998; 2004a): The study showed that no organizational memory per se existed; the perfect repository was a myth. Instead, people in organizations used a number of memory artifacts, they used the artifacts together as resources, and each memory artifact had different properties.

Terveen et al. (1995) investigated many of the same issues for a design rationale repository for large-scale software engineering. Design Assistant was designed to create 'a living design memory'. The system provided advice in much the same manner as an expert system, using many kinds of information, including informal histories, email, knowledge-bases, and automatic software messages. The study found that such a memory could be constructed and used, but the researchers also found they needed to embed both the system and the information in both practice and in the organizational context. The study included, notably, two years of use. Other organizational and group memory system studies of the period included GIMMe (Lindstaedt 1996), Team Memory (Morrison 1993), and TeamInfo (Berlin et al. 1993).

Other CSCW systems of the same period examined alternative sharing and reuse models for knowledge sharing. Grassroots' (Kamiya et al. 1996) users, presumed to be in a small group with a resulting high level of trust, could select and push information into other people's reading lists. Grassroots supported a hierarchical model of knowledge sharing, and also made sharing be an active process. Boland et al. (1992) used shared cognitive maps, linked to users' assumptions, to share perspectives about their underlying assumptions for decisions and actions. In this case, users at the same organizational level (managers) could enact a common sense-making of their situations.

CSCW technical research on repositories and information artifacts has since diminished. In general, work on repository systems became commercialized, and as well, CSCW researchers began focusing on communications as a more viable and central concern. This shifting of emphasis can be seen in the progression of Answer Garden, with its central concern of building an information repository, to Answer Garden 2, with its augmentation of the information repository with many forms of communication, to later studies of expertise sharing and hybrid systems, to be discussed below.

There have been notable exceptions, however. The Placeless project (Dourish 2003), along with its various facets (Dourish et al. 2000; Dourish et al. 1999), examined the potential roles of properties, or name-value pairs, that could be attached to documents or other information artifacts. The studies found that properties permitted better retrieval and organization than did hierarchical repositories. This provided for both mutual intelligibility, or being able to quickly comprehend how others viewed the information, as well as appropriation and customization. Computationally-based properties, called active properties in Placeless, blurred the line between application and infrastructure, since the infrastructure could serve as components of applications; documents could be maintained and organized over time. These studies repeated the findings of the earlier OVAL work (Malone et al. 1995), but with substantially greater detail and nuance. The problem of (re-)organizing the information was also tackled by Shipman and Marshall (1999), Shipman and McCall (1999; 1994), and Shipman et al. (Shipman et al. 2001), all of which investigated how to incrementally add formalization. In their systems, information could start as informal and unstructured, and slowly be formalized over time. This not only provided more structure to informal information, it also allowed greater computational tractability in retrieval and reuse; the finding is a complement to the Placeless findings. Finally, distillation of informal information into more concise and well-organized forms was examined in Nam and Ackerman (2007). Using Shipman and Marshall's principle, Arkose allowed users to incrementally and collaboratively organize and summarize collections of informal information such as brainstorming sessions and e-community posts.

Recently, technical work into information artifacts has picked up again. Researchers have resurfaced several old research streams, largely by using machine learning and crowdsourcing mechanisms. For example, by using Opinion Space (Faridani et al. 2010), Internet users can more easily find and peruse comments that were found insightful by other users. KnowDis? (Laqua et al. 2011) provided additional, appropriate documents when users read email.

Next we examine the social-analytical studies that argued against and resulted from the repository model. While the technical studies detailed new design spaces, arguably, these studies could be regarded as a major contribution to knowledge sharing. Indeed, even many

of the technical studies were socio-technical, exploring both the technical and social issues simultaneously.

2.2 Social-analytic studies of the repository model

Problems with the repository model became quickly apparent after systems were introduced and field-tested (Bannon and Kuutti 1996, Ackerman 1994b, Brown and Duiguid 1991). Reducing the richness of collective memory to specific information artifacts was utopian. As Bannon and Kuutti (1996), after raising a number of problems with the repository model, so eloquently argued:

...information does not simply exist 'out there', but is produced by specific people in specific contexts for specific purposes. While this does not imply that it is bound solely to that whole context, it does mean that one cannot in any straightforward way extract and abstract from this web of signification items of 'information' which can be stored in some central resource for later use.... (p. 163).

Uncovering and investigating the social issues in knowledge sharing fit well with standard CSCW concerns. CSCW in general has assumed that understanding situated use was critical to producing useful, and usable, systems (Suchman 1987; Suchman and Wynn 1984) and that usability and usefulness are social and collaborative in nature. Exceptions in organizational activities, instead of being assumed to be deviations from correct procedures, were held to be 'normal' in organizational life (Suchman 1983) and to be examined for what they said about organizational activity, including information handling (Randall et al. 2007; Schmidt 1999).

In the following sections, we survey this literature in terms of what the studies say to the general social context of information use and then the specific concerns of boundary objects, common information spaces, and assemblies (collections). These papers studied and analyzed the concerns about information-as-artifact and the repository model, but many of them also apply to expertise sharing where communication is at the core.

Three points need to be made before walking through this literature. First, it must be noted that these concerns were not begun by CSCW nor were they exclusively examined by CSCW researchers; however, CSCW as a field expanded and deepened the earlier studies. Nor are the concerns limited to the first generation systems. Many generalize to any knowledge sharing.

Second, some results come from socio-technical studies. In today's terms, the technical work also served as exploratory probes into work and practice (Wulf et al. 2011). For example, an early socio-technical investigation was Dourish et al. (1993). This study deployed two calendar systems and showed different use resulting from the differing sources of authoring and authority. There was a sharp relationship between who entered the data and the politics of its use: The secretaries entered official, authoritative event information into one, and everyone could enter unauthoritative, but timely, event information into the other. That is, the paper not only included system implementation, it also worked through a social-analytic discussion of system use and social implications. Other socio-technical studies include Ackerman 1994a, Dourish 2003, and Hinrichs et al. 2005.

Finally, most of these studies are ethnographically-based and/or interpretivist studies. These ethnographically-based findings should not be viewed in empiricist terms. Viewing them as resulting in *sensitizing* concepts (Blumer 1986; Strauss 1993) is more useful, especially in specific situations. Sensitizing concepts do not have closure; that is, they are often present and should be considered in any research study or requirements analysis, but they are not always present

Social context. A number of ethnographically-based studies have tackled the thorny issues involved in the *social* creation, use, and reuse of information. There are many of these studies, and it is a rare CSCW ethnographically-based study that does not touch at all on issues of information and knowledge sharing. We cannot hope to cover all of the relevant studies; we include a necessarily brief subset of them. The studies presented here have investigated issues of motivation, context in reuse, assessments of reliability and authoritativeness, organizational politics, maintenance, and reification. Each issue will be discussed in turn.

First, obtaining information for repository systems was found, in research studies and in practice, to be a difficult challenge. Motivating people to add knowledge to a repository was found to be key by Orlikowski (1992), and in her site, largely required dealing with the difficult organizational issues of reward structures and changing mental models of use. The organization Orlikowski studied, a consulting company, had trouble motivating the authoring of material. Although partners wanted to garner the firm's knowledge, the lower-level consultants who were supposed to contribute were not rewarded for their time or effort. The consulting firm later moved to rewarding its authors, but her findings hold up well.

In general, the difficulties in aligning reward and incentive structures with desired contribution, especially among differing groups, was a constant problem with repository projects. The initial vision in these projects was capturing an organization's knowledge rather than sharing it. That vision was often utopian, and it ignored misaligned people and groups that had differing career anchors (Schein 1978), reward and incentive structures, and mental models of how to get their work done. The issue of motivation remains very salient. For example, Hinds and Pfeffer (2003) identified motivational barriers mainly located in the competitive structure within companies. In their study they found that, if knowledge sharing is not rewarded, employees have no incentive to engage in it. Other scholars share this view (Alvesson and Kärreman 2001; Schmidt 2012).

Second, a substantial number of studies have examined the social context of information reuse. The processes of decontextualization and recontextualization loomed over the repository model, and have been heavily examined (Ackerman and Halverson 2004a; Schmidt and Bannon 1992). (This issue was also called embeddedness/disembeddedness in Brown and Duguid 2000.) In short, information as it occurs in a social setting has to be decontextualized by an author or editor to remove extraneous details. The information has to be recontextualized by the re-user in order to understand whether and how to reuse the information object. For example, one might not need to include the specific location of a printer in an office, but details about the current network configuration might be important for reuse. This is difficult to achieve, and even harder to achieve for complex problems.

Third, several key papers point to how this recontextualization is achieved as a situated, social action. Part of what is required for recontextualization is to understand the expertise

of the author (Schmidt 1999). Information is also assessed by potential re-users in many ways to assess its quality: Studies have examined trustworthiness (Van House et al. 1998); reliability, authoritativeness, and quality (Brown and Duguid 2000; Van House et al. 1998); understandability (Lutters and Ackerman 2002; Lutters and Ackerman 2007); the provisional or final nature of the information (Hardstone et al. 2004); and, the obsolescence and completeness of the material (Pipek and Wulf 2003). Potential re-users in organizations also situate the information in its original institutional context. People notice whether the material was officially sanctioned or vetted (Harper 1998; Harper et al. 2001; Van House et al. 1998). Furthermore, Heath and Luff (1996) pointed out that information is shared within mutual practices. In their study, entries in a patient record can connote relative expertise and understanding of a patient's condition, thereby requiring a shared understanding of how medical entries are used for proper recontextualization.

Fourth, as authoring and re-use may be situated in organizational contexts, that information necessarily carries with it organizational politics. A long line of work has discussed how knowledge sharing has politics (Harper 1998; Harper and Sellen 1995; Kling 1991; Mambrey and Robinson 1997; Markus 1983). This can be seen in several of the studies already discussed. Dourish et al. (1993), as mentioned, noted that organizational members understood to which calendar entries to pay attention by interpreting the organizational role of the person entering the event. As well, who can author, maintain, or reuse information may be mandated or controlled within an organization (Ackerman and Halverson 1998; 2004a). In Ackerman and Halverson's hotline study, there were many information sources, and for each, only members of specific groups were allowed to change the data. However, Wulf (1997) showed that the division of labor with regard to categorizing documents can shift when introducing an IT-based repository into a hierarchical organization. Halverson et al.'s (2004) study of a help desk for technical consultants added organizational policies, sometimes embedded in infrastructures, to the set of concerns.

Information's meta-data also has politics. As Suchman and others (Suchman 1994 and commentaries in the same issue) have pointed out, categories in general are linked to political considerations; this is also true in the practices around categorizing documents at the time of their storage. However, Wulf's (1997) study indicates that certain actors are not always aware of their interests when deciding on categorization schemes. This seems to be specifically the case when they are still unfamiliar with the long-term use of shared repositories.

Fifth, information objects, as well as classification schemes, in a repository need to be ongoingly maintained to keep them both up-to-date and complete. Hinrichs et al. (2005), Stevens and Wulf (2002), and Pipek and Wulf (2003) conducted long term studies into large repositories of technical drawings in engineering domains. Their investigations included a steel mill's central drawing archive, which included some that were 100 years old, and a sewage company that just had been privatized. It turned out that the proper functioning of these archives was severely challenged over time by changing storage media and classification schemes, as well as organizational and economic changes. Hinrichs et al. (2005) also argued that appropriate tool support can help in maintaining classification schemes within archives by assigning meta-data efficiently. Their toolkit was integrated into the sewage company's document management system and allowed assigning a specific attribute value to a predefined set of existing documents as well as a set of predefined attribute values to newly created documents.

Other studies have raised many of the same issues for long-term reuse. Lutters and Ackerman (2002; 2007) examined repair records in an aircraft technical archive and also found that older records tended not to be considered useful. The older records could be difficult to recontextualize, and prospective users could then not evaluate the repair record's authoritativeness and reliability: For example, the authors might no longer be known (i.e., they have retired or moved on), procedures might have changed, and some information required for understanding the record later might not have been written down for social reasons (e.g., how competent the author was).

Even retrieval has been found to be difficult over time. While meta-data and proper classification would help this, people find adding meta-data to be laborious (Hinrichs et al. 2005) or tend to use the category with the lowest cognitive effort, what Anderson et al. (2007) called 'residual categories'. This makes the reconstruction of classification schemes a knowledge-intensive practice (Pipek and Wulf 2003; Pipek et al. 2003) and complicates later retrieval.

Finally, repository systems promote an objectified view of knowledge. Ackerman and Halverson (1998) found it more useful to analyze information as a duality of process and object: An information object demarks only a snapshot of various organizational or collective processes: An information object is a 'punctuated crystallization' of a knowledge process, using Shapiro's (1994) terminology.

Despite the many uncovered issues, these studies did point to important abstractions in designing, implementing, and adopting knowledge sharing systems. Three of these are discussed in detail below, both for both their importance and because they remain ongoing research themes in CSCW. We include some research outside of CSCW here, when it is important for continuity.

Boundary objects and other knowledge sharing objects. Star and Griesemer's (1989) seminal paper defines boundary objects as:

...objects which both inhabit several intersecting social worlds and satisfy the informational requirements of each of them. Boundary objects are objects which are plastic enough to adapt to local needs and the constraints of the several parties employing them, yet robust enough to maintain a common identity across sites. (p. 393)

The paper pointed out boundary objects in science, including databases, standardizations, and maps. The concept was expanded in Bowker and Star (1999), which considered the development of classification systems.

Boundary objects, therefore, are critical to knowledge sharing. Because of their plasticity of meaning boundary objects serve as translation mechanisms for ideas, viewpoints, and values across otherwise difficult to traverse social boundaries. Boundary objects are bridges between different communities of practice (Wenger 1998) or social worlds (Strauss 1993).

Boundary objects, for CSCW's purposes, are information objects that have meaning on both sides of an intra-organizational or inter-organizational boundary. However, in most cases, some meaning is lost - and necessarily so - as the information crosses the boundary. As an example, in an organizational setting (Ackerman and Halverson 2004a), management

reviews of an employee may be very critical to the human resources department or the line department, but are not necessary to the payroll department. This is actually quite functional—we would drown in information if all context and detail were provided always. Organizational routines across the boundary can assume summaries of the context and detail.

A number of CSCW researchers have examined how boundary objects are used in work practices. Berg and Bowker (1997) detailed how patient records in hospitals act as boundary objects. A patient record, in their view, ‘produces’ the patient for physicians, technicians, and nursing staff by serving as surrogate representations of the patient in the medical record. The information artifact used is inherently not a complete representation of the patient. Mambrey and Robinson (1997) examined information workflows in a German ministry for the POLITeam project. The study found that additional data, inscriptions in their case, could be used to understand the relative meanings of information artifacts being circulated among different groups. The study also noted that boundary objects could be compound: Folders included papers and documents. Other use studies include how ‘diary’ entries mediate between shift changes at a paper mill (Auramäki et al. 1996; Kovalainen et al. 1998), asynchronously providing summaries of important events. As mentioned, Ackerman and Halverson (2004a) reported on a personnel hotline, detailing the information flows within telephone calls and the construction of the answers. In all of these, the recontextualization of the boundary object was nuanced and difficult, but often achieved, allowing for the reuse of information in organizations.

Other studies have focused on how boundary objects play in the micro-negotiations within developing shared understanding. For example, Yamauchi et al. (2000) addressed information mechanisms used by dispersed volunteer open-source programmers, and Bergman et al. (2004) examined the evolution of proto-architectures and project plans at NASA. The emergence of a new artifact serving as a boundary object was observed in Halverson and Ackerman (2003), which described how a ‘cheat sheet,’ used by both airlines and air traffic control to guide their daily activities, came to be. The study also described how differing representations evolved.

Perhaps due to the tool development focus in HCI and CSCW, certain characteristics of boundary objects have been promoted to the foreground in the literature. Star’s original definition includes such diverse notions of boundary objects as coincident boundaries (e.g., the shared understanding of the geopolitical boundary of the state of California) and ideal types. CSCW has focused on the more tractable information processing objects such as standardized forms and information in repositories. These are easier to implement and of more commercial interest. In so focusing, there has often been an over-emphasis on boundary objects as material artifacts, which can limit the analytical power that boundary objects bring to understanding negotiation and mediation in routine work.

Nonetheless, some important studies have considered how to support both process and information objects. Bandini et al. (2003) described a ‘T-Matrix’ artifact, technically supported, that covered all relevant aspects of producing a tire; it supported capturing past experiences as well as innovation through the emergence of knowledge items concerning tire performance, blend features, and recipe interventions. In later studies (Cabitza et al. 2008; Cabitza et al. 2009), additional artifacts at the intersection of coordinating work practices and fostering knowledge sharing were constructed and analyzed. Their ‘clinical pathways’ helped coordinate a patient’s treatment plan and also encoded best practices of

treating patients with specific health problems. Cabitza and Simone (2012) argue for a perspective where artifacts carry two levels of information: awareness promoting information (API, information on the current state of a collaborative activity) and knowledge-evoking information (KEI, information triggering previously acquired knowledge or triggering and supporting learning as well as innovation). They suggest 'affording mechanisms' (extending previous definitions of 'coordination mechanisms' in Schmidt and Simone 1996) as a formal construct to describe and manage both kinds of information in an integrated way.

As well, work in CSCW has begun to extend the concept of boundary object to include new types of information objects. Lutters and Ackerman (2007) and Lee (2007) suggested classes of more free-form boundary negotiating artifacts, less structured and standardized than boundary objects. Lutters and Ackerman, in their study of aircraft repair work, pointed out that the 'boundary object' concept was too limited, and their study extended the concept to include the existence of higher-level 'event streams' and 'meta-negotiations'. In their site, boundary objects (repair orders) were interpreted in light of past and future interactions between the parties (airframe manufacturer and airline), based in previously established levels of trust and expectation, as well as being guided by current context. Lee, in a study of museum design, pointed to the existence of other information objects that did not have quite the same properties as boundary objects, especially their standardization. Lee termed these objects 'boundary negotiating' objects, as they often helped participants discuss, negotiate, and sometimes settle on boundaries. This work was continued in Bietz and Lee (2009), which examined the use of metagenomic databases as 'sites for negotiation of community boundaries (p. 11)'. This finding echoed O'Day et al. (2001) and Boujut and Blanco (2003), both of which also found information objects that served as a mechanism for discussions. Bietz and Lee also pointed at the prevalence of unfinished or incomplete information objects, especially those designed as boundary objects. Bossen's (2002) work on medication documentation examined the 'boundary spanning' role as mediation among the different social worlds of surgeons, physicians, technicians, nurses and administrators. Reddy et al. (2001) and Halverson and Ackerman (2003, 2007) examined how the same data could be represented in different ways, essentially creating alternative boundary objects.

Cabitza et al. (2013) also expanded the investigation of information objects, again bringing 'boundary objects' into a larger, more abstracted viewpoint. Summarizing and extending earlier studies, Cabitza et al. describe what they term a 'knowledge artifact', which is:

a physical, i.e., material but not necessarily tangible, inscribed artifact that is collaboratively created, maintained and used to support knowledge-oriented social processes (among which knowledge creation and exploitation, collaborative problem solving and decision making) within or across communities of practice....
(p. 35)

Since knowledge artifacts emerge from practice and inform practice they are inherently local; they provide a 'bounded openness' that allows participants to establish a shared meaning on the one hand, while remaining open for modifications on the other. The knowledge artifacts provide shared representations not in a sense that they contain knowledge, but in the sense that they stimulate socialization and internalization of knowledge. They are always underspecified, but at an appropriate level of underspecification that can be described in terms of granularity (level of detail provided),

approximation (degree of interpretable closeness) and rigidity (degree to which it allows useful interpretations in a variety of contexts and situations). Thereby knowledge artifacts, although not as formalized as some boundary objects, can serve in the capacity of boundary objects, but are not limited to doing so.

Important issues remain in considering boundary objects and other information objects within a repository view. Boundary objects clearly do not fully describe the complexity and variety of information objects in use. Although there have been substantial strides, we have seen only the beginning of detailed fieldwork examining the complexity and variety of modern information objects. The types of information objects, *as used in practice*, remain to be fully inventoried. (We remind the reader that, in the standard CSCW view, an information 'object' cannot be separated from its social context and the practices that produce and use that 'object'. Therefore, object and process must be considered together.) Additionally, important details are still lacking in our understanding of how to handle exceptional cases and to support the nuances of trust, context, and authoritativeness.

We are also at the beginning of considering how to design these information objects. The work of Cabitza, Simone, and colleagues, in their series of studies, shows how to consider the social issues in information use and then investigate those issues by designing and testing new information objects with similar properties. This work, and to a lesser extent that of Lutters and Ackerman, points to new types of information artifacts that add to boundary object capabilities. These objects might continuously change, incorporating traces of coordination protocols, process representations, or other forms of negotiated practice knowledge. Finally, the complexities of digital objects are starting to stretch the definition of 'an' object, especially in electronic health records (EHRs) (Bossen et al. 2012; Zhou et al. 2011), which are assemblies of many records for any given patient.

Common information spaces. The second abstraction we wish to consider is that of 'common information spaces' (CIS) (Bannon and Bødker 1997; Schmidt and Bannon 1992). Schmidt and Bannon (1992) defined CIS as:

...the focus is on how people in a distributed setting can work cooperatively in a common information space - i.e. by maintaining a central archive of organizational information with some level of 'shared' agreement as to the meaning of this information (locally constructed), despite the marked differences concerning the origins and context of these information items. The space is constituted and maintained by different actors employing different conceptualizations and multiple decision making strategies, supported by technology. (p. 22)

Bannon and Bødker (1997) continued this analysis by noting that the extent of sharing and the 'the mechanisms used to support "holding in common" the information' (p. 83) can vary. CISes are open and malleable by nature, but also require 'some form of closure, to allow for forms of translation and portability between communities (p. 86)'. Furthermore, participants must be able to interpret the information so as to collaborate, and they must do this actively. Bannon and Bødker (1997) summarizes the differences between CISes and other views of knowledge management repositories:

A common information space is not just a repository of information constituted once and for all, which raises interesting concerns for design.... (p. 94)

Bossen (2002), based on work in Reddy et al. (2001) and Reddy and Dourish (2002) as well as his own fieldwork, attempted to address why CISEs might differ so much (as was pointed out in Bannon and Bødker). Bossen proposed 7 dimensions of analysis for CISEs: the degree of distribution of work, the 'webs of significance' and their number, the level of required articulation work, the means of communication, the web of artifacts, and the need for precision and promptness of interpretation. All of these dimensions could vary, and would account for the large differences in CISEs. Fields et al. (2005) pointed out that these dimensions could be summarized as the degrees of distribution, the kinds of articulation work and artifacts required, the means of communication, and the differences in frames of reference among participants.

Fields et al. (2005) also began an analysis, continued in Rolland et al. (2006) and Munkvold and Ellingsen (2007), that pointed to a conceptual problem with CIS: Any given situation has 'zones of overlapping meaning' (Fields et al. 2005, p. 129). In fact, there are overlapping zones of partially overlapping meaning; boundaries are, if anything, difficult to draw. Munkvold and Ellingsen (2007), in their analysis of nursing plans, extended the issue of 'zones' to include temporality. For Munkvold and Ellingsen, CISEs exist often only temporarily, as separate groups come together to work in a coordinated fashion for a time and then go back into their parallel work.

This problem in the concept of CIS can be found in many similar concepts that try to draw boundaries around situated activity. The Chicago School's concept of 'social world' (Strauss 1993) suffers from the same problem. Nonetheless, the concept of CIS has been found to be useful when considering information artifacts and their use. CIS draws attention to the ongoing work involved in using and reusing information in social settings. It also calls attention to the necessity of the dialectical tension between the regularization and routinization required by boundary objects (or, often, similar information objects) and the openness and implicit assumptions in organizational or social activities (Cabitza et al. 2013). This cannot be resolved for knowledge sharing, especially in its repository form. Since this tension cannot be resolved in the abstract, managing it is a necessary requirement for information environments.

Assemblies. The final abstraction discussed by CSCW researchers, useful in the consideration of information artifacts and repositories, is that of assemblies. (We use 'assembly' to denote an organized collection of information objects. 'Assemblages,' deeply rooted in the Science, Technology, and Society literature would include the surrounding practices and culture around an object or collection.) Implicit in the concept of CIS and in boundary objects is the question of assemblies. A line of work that deserves to be expanded further concerns how assemblies of information objects work together. It is clear that information objects are used together, but how do they work together? How are assemblies of information objects bound together in social processes? No single information artifact or social practice works in isolation; they are always in some socially situated context.

Schmidt and Wagner (2003; 2004) began this line of work, calling for an investigation into what they called 'ordering systems'. An ordering system is 'a complex of interrelated coordinative practices and artifacts (Schmidt and Wagner 2003, p. 274)'. Based on a study of an architect's work, they were able to identify 'clusters of coordinative artifacts, each of which supports a coupled set of practices (p. 280)', including workflows, classification, validation, and other procedures. In Schmidt and Wagner's work is a call to consider

ordering systems, especially as they develop, collide, and adapt. The key for them is the acknowledgement of 'clusters' and their 'coupled set' of practices.

This call was echoed in the later Schmidt et al. (2007), which analyzed work practices in two cancer clinics in Vienna. They show that the work is essentially the same between both clinics, but it is done in very different ways using different information artifacts. This paper again calls for more studies analyzing 'the "higher-order" practices of endless combination and recombination of artifacts, formats, notations, etc. that are found across such sites (p. 1)'.

While assemblies were not the explicit target for Ackerman and Halverson or for Lee, the question is clearly present in both sets of work. As mentioned above, Ackerman and Halverson (2004a) examined the use of many different information artifacts together in their distributed cognition analysis of organizational memory. In their field study of a hotline, they uncovered many small memories, wrapped in many social practices, used in combination. These memories, which were all information artifacts, and the practices were used as resources in flexible arrangements by the hotline's members. That is, specific assemblies were situated as arrangements of many possible artifacts and practices (as in Katzenberg et al. 1996; Pentland 1992). And while Lee's (2007) boundary negotiating objects had distinct roles in participants' practices (perhaps the result of her process of theoretical categorization), the roles were complementary and had to work together in order to get the participants' museum exhibit built. Again, combinations were critical. The concept is also considered in other research, including Munkvold and Ellingsen (2007) and Reddy and Dourish (2002).

We now turn to the next generation, expertise sharing. As one can see, CSCW presaged this move in Knowledge Management, with its emphasis on the social situatedness of knowledge production and use, as well as its interest in tying in computer-mediated communication.

3. The Second Generation: Sharing Expertise

A second generation of CSCW studies emphasized interpersonal communications of knowledgeable actors over externalizations in (IT) artifacts. The shift away from the repository model ascribed a more crucial role to the practices of individuals engaging in knowledge- or expertise-sharing. Emphasis was on finding an appropriate person. In this second generation, sharing tacit knowledge, including that contextual knowledge that might be required to understand information, became critical. The distinction between 'explicit' and 'tacit' knowledge originated from Polanyi's work (1967) and became widespread with Nonaka and Takeuchi (1995). It is worth noticing that Nonaka and Takeuchi's interpretation of Polanyi's terminology led to the idea that 'making tacit knowledge explicit' could be one of the design goals to associate with IT tools particularly of the 'first generation'. But Polanyi introduced the term 'tacit ways of knowing' to exactly describe the kinds of knowing that are difficult, if not impossible to verbalize. Tacit knowledge can be learned only through common experiences, and therefore, contact with others, in some form, is required for full use of the information.

Two concepts were critical in stimulating the shift towards the second generation and have strongly influenced this line of work:

Community of Practice. The concept of 'Community of Practice' (CoP) (Wenger 1998) became a useful analytical lens to identify relevant groups of people. A CoP can roughly be

defined as a group that works together in a certain domain and whose members share a common practice. Most important about the concept is its account of knowledge sharing and socialization, which describes a certain way to learn and internalize a characterizing set of practices: Learning within a CoP takes place as 'legitimate peripheral participation' (Lave and Wenger 1991), a process of apprenticeship in which participants in the group move from the periphery to the center, i.e. become increasingly knowledgeable (Fischer 2007; Lave and Wenger 1991). To better differentiate occasions for learning, Brown and Duguid (2000) added the concept of networks of practice (NoP). Members of a NoP do not necessarily work together but work on similar issues in a similar way, e.g., the carpenters of a certain region. The similarities of their work practices provide a common ground to foster better knowledge and expertise sharing. Another concept, which built on the notion of CoP is the 'Community of Interest' (CoI). Contrary to CoPs, members of a CoI are not defined by common practices but by a common interest. While the latter holds them together, the diversity of their backgrounds provides a rich source for creativity and innovation (Fischer 2001). A comprehensive review of different communities and their functioning can be found in Andriessen (2005). Intrinsically based in their view of 'tacit knowledge,' the Knowledge Management community appropriated CoP in an interventionist manner. CoPs were to be cultivated or even created (Wenger et al. 2002), and they became fashionable as 'the killer application for knowledge management practitioners' (Su and Wilensky 2011, p. 10) with supposedly beneficial effects on knowledge exchange within groups. However, this interventionist adoption in Knowledge Management was disputed among its intellectual fathers (see for a discussion again see Su and Wilensky 2011, p. 30f.).

The CSCW community has generally not followed such an interventionist appropriation of the concept (for an exception, see Stevens et al. 2005). The concept of CoP was applied in a different way - to understand expertise sharing, serving as an analytical lens to get a hold on the situatedness of knowledge-related processes. While these practice-based conceptions of learning had a strong influence on the CSCW discourse, it is important to keep in mind that they describe just one dimension of the enabling conditions.

Social capital. Another useful concept for expertise sharing has been 'social capital' (Cohen and Prusak 2001; Huysman and Wulf 2004; Lesser et al. 2000). It originated from Bourdieu (1977) and refers to collective abilities derived from social networks. The core intuition that makes social capital conceptually interesting is that the goodwill that people have toward others is a valuable resource for expertise sharing (Adler and Kwon 2002). While the concept of CoP enables us to understand learning opportunities to be inherently related to shared practices, social capital sees resources for expertise sharing rather in the relational and empathic dimension of social networks. In this sense it offers an additional intuition to understand conditions which foster expertise sharing. Nahapiet and Ghoshal (1998) argued that social capital consists of three dimensions: structural opportunity, cognitive ability, and relational motivation. Technical support for expertise sharing mainly addresses the structural opportunity dimension by providing an infrastructure for knowledge exchange. The other two dimensions are, however, not very much addressed by system designers (Huysman and Wulf 2006, p. 41). When analyzing actors' cognitive abilities to share knowledge, a practice-based perspective may need to be taken into account. However, social capital is not necessarily grounded in common practices. CoPs have a high degree of social capital, if the degree of trust, mutuality, and recognition by peers is high. The three dimensions of social capital can be translated into the questions of 'who' shares knowledge and 'how' it is shared (structural dimension), 'what' is shared (cognitive dimension) and 'why' and 'when' people engage in knowledge sharing (relational

dimension) (ibid, p. 44). By informing technology design with empirical research along these questions, all three dimensions of social capital can be tackled (Huysman and Wulf 2005).

What, then, needs to be supported in sharing expertise among people? As mentioned, Ackerman and Halverson's (1998) study of a call-center workplace arrived at the conclusion that a techno-centric approach, which concentrates solely on making more knowledge artifacts available, would be misleading. Rather, contextual knowledge is required to select the right source, strip off the information from its original context and then re-contextualize it according to the situation at hand. This knowledge is often found in conjunction with knowing or asking other people. Normark and Randall's (2005) case study of an emergency call center also stressed the socially distributed nature of knowledge. They emphasized efforts of 'finding-out' (Fitzpatrick 2003) when analyzing knowledge-sharing practices. Rather than looking for different kinds of sources, however, they distinguished between knowledge 'in the small' and 'in the large' (Ackerman and Mandel 1995; 1999). While the latter referred to what is termed 'explicit knowledge', knowledge 'in the small' is described as 'local knowledge'. Fitzpatrick (2003) notes that it is especially characterized by its emergent nature. Which knowledge becomes important and who holds it only evolves in the concrete situation at hand (Normark and Randall 2005); this would be particularly true for their time- and safety-critical emergency settings. Which knowledge is needed can hardly be planned in advance.

While the two studies differed slightly in their analytic stances towards knowledge distribution in social settings, they both point out that there is 'finding-out' work involved in gathering and making sense of the available information. While Ackerman and Halverson (1998) highlighted the situated effort of de- and re-contextualization, Normark and Randall (2005) highlighted the process-character and contingent nature of what they call 'expertise combining'. Bannon and Bødker (1997) used the expressions 'packaging' and 'unpacking' information with its sense-making context. All these studies concentrate on particular aspects of knowledge exchange in social settings, which Schmidt (2012), who is critical towards the sources of the concept of 'tacit knowledge,' termed 'didactic practices required for a cooperative ensemble to be able to maintain and develop its own competencies' (p. 208).

3.1 People finding and expertise location

If systems are to play a meaningful role in expertise sharing, they should support this 'finding-out' work in seeking appropriate people. To address both the explicit and tacit dimensions of knowledge, and address the limitations of the repository model resulting from a focus on visible and storable representations rather than the interactions producing and modifying them, expertise sharing systems in the second generation sought to facilitate knowledge exchange at the level of individuals and collectives. These tools often tackled technical, organizational, and social issues simultaneously.

People-finding, as mentioned, really began with some of the first CSCW systems; however, some even earlier attempts existed, using information retrieval techniques – Who Knows (Streeter and Lochbaum 1988) from Bellcore was arguably the first such system, and Information Lens (Malone et al. 1987) had a prototype 'anyone server,' which allowed users to send a query that the system would route to some appropriate organizational member. We have already mentioned several others: Answer Garden (Ackerman 1994a) assumed

users could ask questions of experts, and if the answer were not present already, the questions would be routed in a primitive escalation cascade. Answer Garden 2 (AG2) (Ackerman and McDonald 1996) used a more flexible (and rule-based) escalation system to route the question to potentially larger and larger groups. All of these systems assumed that sometimes knowledge could be found only in other people.

More recent systems focusing on 'people finding' in CSCW are based largely on a combination of approaches for determining social relationships through network analysis, topics of interest, and then expertise determination. Social relationships are usually based on some sort of Social Network Analysis (SNA) (McDonald 2003; Terveen and McDonald 2005), and relationships among people are quantified in terms of tie strength, direction, and type of tie. Social networks can include networks for sharing communications, sociability, or expertise. In general, expertise sharing systems vary considerably in how people's profiles are created, including what data is used and how users are involved; how relative expertise is determined; how queries or people are matched to people; and, how recommendations are made.

Profile construction can be contentious. Below, we will discuss privacy issues, but expertise location systems must depend on their searchable profiles, which store information about individuals. Ehrlich (2003), based on focus group results, established three key factors that should be covered by profiles: the person's credentials, likeliness of responding, and accessibility. Relative expertise is also a requirement in profiles. Early efforts found that it was difficult to have people rate themselves or one another (Ackerman et al. 2002a), and therefore later efforts tried to automate the evaluation of expertise levels (e.g., Zhang et al. 2007). By searching a profile database, people extend their awareness of other people's qualities and may consider contacting them. According to Ehrlich (2003), an expertise locator system aims to bring people together who do not know each other, and therefore relies on weak social ties. Profile maintenance can also be an issue, as it was for repositories. Ehrlich et al. (2007) conducted a series of interviews with consultants and found personal networks to be extremely important to these consultants. At the same time, these networks needed continuous care in order to support time-critical tasks.

Following, we describe the various approaches to expertise sharing in CSCW. These lines of research used a combination of interpretivist analyses of organizational needs, design and construction of technical prototypes, and social studies of system use. They are, then, really socio-technical lines of examination, as one often finds in current CSCW.

McDonald and Ackerman (1998) examined expertise sharing in a small-to-medium sized software company, which they called MSC. They found that MSC participants used a variety of strategies for finding the right expertise in the company. McDonald and Ackerman divided the expertise finding seen in MSC into three stages: identification, selection, and escalation. The identification stage concerns finding an appropriate group of people (candidates), selection concerns picking which candidates to approach, and escalation is identification and selection from an ever-expanding list of people if no satisfactory candidate or answer is found. Selection might include whether a candidate was busy, etc. It should be noted that these stages were merely analytically separable; in practice, they are often mixed and combined, since they are situated activities. However, the authors found the analytic distinctions useful in subsequent technical work. Groth and Bowers (2001), in another study of a software company (in Norway rather than the US), saw expertise finding as much more situated, and availability was key. Their conclusion was that an appropriate

system would have to be one promoting awareness of expertise and availability based in 'seeing' others' activities. These two approaches are largely incompatible. McDonald and Ackerman's approach is more scalable, but at the cost of reifying social activities that are, in fact, often situated and combined; Groth and Bowers' approach can be more fideliou to the contingencies of organizational life, but does not scale well.

Based on the McDonald and Ackerman study, McDonald built a system, Expertise Recommender (ER) (McDonald 2000; McDonald and Ackerman 2000) which architecturally combined a number of different, replaceable identification and selection components. For the evaluation prototype, ER used primarily an identification mechanism based in a local software engineering practice within MSC, a practice of documenting who had handled the software component in question. For this organization, this specific practice served as a convenient and more accurate mechanism for identifying appropriate people with the necessary expertise, and local arrangements could be more meaningful since the expertise finding and escalation itself occurred locally. Mockus and Herbsleb (2002) similarly analyzed who had touched the code. ER also used a variety of other matching mechanisms, including keyword vectors, based on local documents, and so on. ER's flexible architecture allowed both general matching mechanisms (e.g., position in the organizational chart) and local matching mechanisms (e.g., who had maintained the code last).

For ER, selection was key; identification provided recommendations, and selection consisted of filters for determining whom to approach. The prototype ER included selection filters based on graphs for organizational structure, sociability, and knowledge domain. That is, McDonald constructed a graph based on two SNAs, one for 'who hangs out with whom' and who knows what topic in the organization (McDonald and Ackerman 2000); these were constructed using pile-sort techniques for SNA. In his lab-based evaluation studies (McDonald 2001; 2003), McDonald found that ER was capable of providing excellent expertise identification, but his selection filters required additional capabilities. Users wanted filtering based on ego-focused networks rather than aggregated views of the organizational social network.

In another series of studies, Reichling and Veith (2005) also preceded their system development with an ethnographic study; they studied a knowledge-intensive network-(and lobbying-) organization, NIA. They found many barriers to expertise sharing originating from NIA's complex organizational structure. Within the organization, multiple sub-units existed fairly independent from one another. Employees reported a low degree of transparency and weak knowledge about the tasks and competences of colleagues from other units. Between the departments, there was a subtle competition. All these factors resulted in social networks being critically important for accessing expertise; outsiders needed to maintain personal contacts with insiders in order to access services (which was the actual task of the organization).

On the one hand, some opportunities to improve the situation by means of information technology were identified: Employees suggested implementing a Yellow-Pages system which could be accessed by means of a search engine for personal attributes ('like Google'), Another idea was a 'virtual notice board', where external service requests and other work could be posted. On the other hand it became clear that changes on the organizational level would be necessary to create the structural preconditions for successful expertise sharing. For this reason, they used the integrated OTD approach (Wulf and Rohde 1995) that takes into account both organizational and technical factors. Based on this study, the design for

the 'Expert Finder' implementation for NIA was component-based (Reichling et al. 2007). It combined 'push' and 'pull' mechanisms. Expertise profiles consisted of three components: The first was the Yellow Pages component, where users could enter general information about themselves and their position in the organization, and the second was a keyword vector, automatically generated from a user's textual documents (users could also manually select textual documents and choose the data sources for profile generation). The third part was a feedback component. Expert recommendations could be judged on a five-point scale, where no negative scores were permitted and giving feedback was not mandatory.

Technically, the system used the 'ExpertFinding Framework': After language recognition, the textual documents were filtered through a language-specific catalogue of stop words, resulting in a list of keywords. This list was then further processed to reduce the number of terms (Reichling et al. 2005). Matching took place via specifically configured algorithms.

Privacy concerns played a major role in the selection of design criteria for the Expert Finder. Due to the legal situation in Germany, the software needed to be built in a way that respected the co-determination rights of the workers. For this reason, it was decided not to include email sources in the construction of the profiles. Moreover, text documents were selected by the user via their folders and analyzed automatically. However, the documents never left the user's computer. The resulting list of key words was uploaded to the server only with the explicit permission of users and could be edited beforehand (Reichling et al. 2007). Reichling and Wulf (2009) evaluated the Expert Finder's design by introducing it to NIA and observing its appropriation for several months. Users took the semi-automatically generated rankings very seriously and carefully evaluating their own positioning for topic areas in which they claimed to be experts. The authors also found that semi-automatically generated rankings could interfere with formal aspects of an organization. For instance, people expected that the expertise profiles could be taken as job descriptions.

IBM researchers have also studied a number of expertise sharing systems; these systems were designed and used within IBM. Among them was ReachOut, a search tool that could be used to find potential candidates for advice or help (Ribak et al. 2002). Another system was 'SmallBlue', a 'people mining'-system, which included 'SmallBlue Find', a search engine, that ranked people according to specific search terms, which should represent knowledge or skills (Lin et al. 2008). Other parts of SmallBlue also allowed searching for other criteria, but we will focus here on the expertise search function. The system drew heavily on private data sources: outgoing email and chat messages were complemented by data from an enterprise directory. The authors were aware that using private data is a delicate issue, and for this reason, they applied limitations to data collection, usage, and distribution. Users needed to actively opt-in, only outgoing mail was used, and confidential mails as well as external names were ignored. SmallBlue Find displayed the usual contact information (room number, affiliation, etc.), and the system also provided two features that the authors called 'Social Distance' and 'Expertise in Context'. The first showed the shortest (or an alternate) social path from the seeker to the knowledgeable person. The second function provided additional information about the knowledgeable person's sharing-related activities (Ehrlich et al. 2007). The rationale behind these functions was to support the searcher's selection with additional information.

Shortly after its first implementation in consultancy work, SmallBlue was evaluated with regard to its adoption and the users' opinions about privacy (Ehrlich et al. 2007). The evaluation was conducted using a combination of semi-structured interviews, log data, and

an online survey. After making the system publicly available within IBM, people began opting in with their data sources, which led to the rapid indexing of other employees as well. After six months, 1600 people had opted in and over 150,000 people were indexed in the system. In the interviews, most people stated that searching for a person is not a daily activity for them, and the search results were judged to be fairly accurate. When asked about privacy matters, no concerns were expressed. The log data revealed that while most users probably used the tool just out of curiosity, there were also some who used it again and again. The 'social distance' tool was judged useful, because when contacting an unknown person, users found it helpful to mention the social path to him.

Shami et al. (2008) also evaluated how a particular person was selected from a result list. Their study suggested that people prefer others with whom they share a social connection to complete strangers. However, those people whom the participants knew directly were not the first choice to be contacted: Participants did not expect them to hold considerably more expertise. More valuable in terms of selection were people who were 2 or 3 degrees away. This is consistent with McDonald (2003), where it was found that users preferred seeing choices they did not know already. For Shami et al., rank order strongly influenced selection; higher ranked results were considered more frequently for further exploration. Yarosh et al. (2012) presented a later lab study of SmallBlue investigating the relationship between the presentation of additional information and frequently occurring types of expertise-search related tasks. That paper presented a list of ten entries that were considered most useful for accomplishing tasks.

Above, we have mostly equated 'people finding' with 'expertise location'. In fact, however, expertise in a certain field is only one criterion by which people can be searched. Guy et al. (2008b) introduced SONAR (Social Networks Architecture), an API, which collected data from Web 2.0 providers and made them available to SONAR clients. SONAR aimed to show who is related to whom and how. Clients could ask questions about these relationships, which might concern expertise (e.g., on a certain topic) or other criteria.

An interesting feature of SONAR is that it distinguished between public and private sources. Two modes of network analysis were available: a sociocentric view that used only public sources, and an egocentric view that also included private sources. The sociocentric view could be invoked by anyone to gain information about the connections between any individuals and groups. In the egocentric view, users could view only their own connections with others. The distinction between different kinds of sources is crucial for privacy reasons. Guy et al. (2008a) examined how far public sources can approximate the quality of private sources. When implemented within IBM, over ten sources from the company's intranet were used, such as blogs, social tagging systems and the organizational chart (public), as well as email and chat transcripts (private) (Guy et al. 2008b). This exemplary implementation calculated a buddy list based on the strength of ties. In order to account for a different significance of sources and providers, users could also manually change the latter's relative weighting.

Faces (Guy et al. 2012) implemented SONAR, and also drew on other company-specific sources: a social networking site and a corporate directory. Faces' web-based frontend displayed faces alongside organizational information (such as department, related people and contact information). While the SONAR implementation reported above calculated the strength of ties generally on the basis of a user profile, detailed information about the user was not required for Faces. Search results could, however, be optimized by giving more

details about the searcher. Guy et al. (2012) also provided details on the calculation used by the backend.

Won and Pipek (2003) tried to tackle the problem of expertise location by borrowing from the CSCW concept of 'peripheral awareness' (e.g., Rodden 1996, Schmidt 2002). They noticed that expert profiles were often not up-to-date, and that various behavior patterns around self-representation of experts and ontological differences hinder finding the 'right' expert in a profile-based approach. Instead, they suggested supporting 'peripheral expertise awareness' based on expertise-indicating activities in an organization. Indications of expertise can be found when a colleague hands in a reimbursement form for a seminar he attended (i.e., he may be an expert in the seminar topic now), or when a number of Wikipedia edits on a topic seem to qualify a colleague as competent in that topic. As single indications are insecure, users could logically combine them into complex expertise indicators that could be shared and rated by colleagues, to indicate the trustworthiness of expertise hypotheses. Users could then ask to be notified whenever trusted indicators were triggered by activities in the organization. Dörner et al. (2007) also showed in a small evaluation that several privacy issues required a sensitive further development of the basic idea, although a rich infrastructure of software sensors, for which the approach would benefit, was not available at the time.

3.2 Open issues in expertise and people finding

The overview of the studies we presented here allows us to summarize a number of open issues for further research in expertise sharing:

Development of personal profiles. All expertise finders are influenced by the particular environment for which they are developed. SONAR, Faces, and SmallBlue (formerly known as IBM Atlas) relied on data that was collected to reconstruct the social network of the actors. The environment within IBM allowed the systems to draw on many sources, for which content is publicly available (at least inside the company). Examples of their sources included blogs, social tagging systems, and organizational charts. In addition, private sources such as email and chat protocols were available and could in principle also be used. On the other hand, the network organization where Expert Finder (Reichling et al. 2007) was implemented had many constraints concerning data; there was little information easily available or accessible. Although a reconstruction of the actors' social networks would have been highly valuable, it could not be conducted, since the analysis of emails was considered too invasive to the employees' privacy. As a result, the idea to derive profiles from self-selected documents of the experts was an interesting compromise for profile generation without users' losing control of their personal information. It is also worth noting, that the very same users who asked for very specific and targeted selection processes when looking for an expert, were still keen on presenting themselves as expert even at the periphery of their knowledge spectrum – thus blurring the results of an expert finding process (Dörner et al. 2007).

Privacy and Control. Privacy protection is an important topic for all systems. While Reichling et al. (2007) allow users to select indicative folders and chose to completely abstain from emails, Ehrlich et al. (2007) took substantial measures to preserve privacy while still exploiting the email's value. The SONAR API uses email only for private analyses – which means that only the user's own emails are used for her own analysis. Privacy was similarly important for McDonald's participants, and aggregated social network views (so as to avoid

privacy violations that might occur in ego-networks) were not always useful (McDonald 2003). We can observe here different attempts to make sure that the user's private data is not disseminated throughout the organization without his control (e.g., Dörner et al. 2007). Dörner et al. 2007 (2007) found that even if no continuous profile is visible, users want to be in control of the information flows that influence their appearance in an organization. Therefore, Reichling et al. (2007) allowed users to observe and even change their profiles.

Accuracy. It is a methodological challenge to evaluate the quality of algorithms that recommend people. There is not any 'objective' truth and people's judgment on other's capability can vary considerably. Ackerman et al. (2002a) calculated one's expertise in organizational tasks, and then asked people to rate one another's expertise levels. The study determined that about 20 organizational participants were able to do as well as 3 'expertise concierges', managers with high transactive memory about organizational projects. McDonald (2001) also examined the same data, and found that participants' assessments were correct approximately 79% of the time (although the expertise concierges were better with 85%). He argued this accuracy in expertise determination was as high as might be expected for systems. Reichling et al. (2005) also compared automatically generated expertise profiles with those generated by members of a network of researchers about themselves. The results indicated that the deviations between the different human assessments were within the same range of variation as those between human and automatically generated ones. A deeper problem for providing 'accurate' recommendations also lies in the issues around privacy and self-representation that lead to conflicting requirements for expertise finding and to an approach where users can codify for themselves their notions of accuracy (Dörner et al. 2007).

3.3 Finding others in a communication space

Besides expertise finding, another possibility for expertise sharing is to create a place where people with expertise can be found (Ackerman and Halverson 2004b). The use of computer-mediated communication (CMC) for the purposes of obtaining help and information has been extensively studied in CSCW, and we cannot attempt to review of all those studies here. (For general monographs and surveys see, e.g., Preece 2000 and Wellman 1999.) There have been numerous studies of providing help on CMCs, including studies of email distribution lists (e.g., Constant et al. 1994; Finholt and Sproull 1990), Usenet forums (e.g., Fiore et al. 2002; Fisher et al. 2006; Turner et al. 2005), bulletin boards (e.g., Hiltz and Turoff 1993), and synchronous chat (e.g., Ackerman and Palen 1996; Bradner et al. 1999). Much of this literature is germane to expertise sharing. For example, Constant et al. (1994) studied the use of weak ties to obtain help in organizational settings, Fiore et al. (2002) examined how users made judgments about Usenet authors, and Ackerman and Palen (1996) investigated how the technical affordances of a chat-based help system were related to its emergent social characteristics.

With the advent of the Web, Question-and-Answer (Q&A) communities sprung up. They were similar to Communities of Interest (CoI) and older bulletin board systems, but were Internet-based, public facing, and often quite large. There have been a number of studies of Q&A communities, and Gazan (2011) provides a review of this literature up to 2010. (Gazan uses "social Q&A" for these communities. "Social Q&A" has also been used more broadly, as in Mamykina et al. 2011, to include Q&A through social network systems. We will use "social Q&A" in its broader sense below. The terms "social Q&A", "QnA", "CQA", and "QA" communities have all been used as well for Q&A communities.)

As Gazan notes, researchers have extensively studied question classification and answer quality as measures, as well as user satisfaction, motivation, and reputation in Q&A communities. For example, Welser et al. (2009) was the first among an extensive line of research into answerers, people essential to the health of any Q&A community.

Gazan's (2011) categories have continued to be research topics since 2010. The most heavily studied have been motivation and incentives. Hsieh et al. (2010) examined the relationship between quality of answer and payment. They found that more money leads to more answers, but not necessarily to better answers. Jeon et al. (2010), based on a re-examination of two studies' data, found similar results. Hsieh et al. (2011) examined the use of charitable contributions as incentives; they found that such contributions spur credibility of answers, although in a nuanced way, for users. Tausczik and Pennebaker (2012) used both self-reports and inferred measures to determine user's motivations on MathOverflow, a site for professional mathematics. They found that altruism and reputation building to be the two important motivations. Dearman and Truong (2010) reversed the question of why people answer and investigated why people do not answer.

Pal et al. (2012) continued the investigation into answerers. Using StackOverflow data, they were able to build models of expert evolution and identify different categories of experts. This is a promising first step in looking at answerers over time and in trying to understand the dynamic characteristics of these sites. In other work, Yang et al. (2010) examined the answerers' activity, including its burstiness, on three platforms, Baidu, Naver, and Yahoo. They found substantial differences, perhaps due to cross-cultural factors among others.

Work on the answers themselves has been more sparse. Pal et al. (2012) opened up a new line of research into how questions age and become obsolete. This is critical for retrieval at a later date.

More recently, researchers have begun to investigate how people use social media to obtain answers, often called 'social Q&A'. Two lines of research exist, echoing the classic division of social and technical analyses. First, there has been work on how people use existing social network systems, largely Facebook and Twitter, to ask questions and answer (e.g., Lampe et al. 2006; Morris et al. 2010a; Morris et al. 2010b). Panovich et al. (2012) investigated the relation of tie strength (i.e., how intimate two people are) to answer quality. Zhang et al. (2010) examined questions and answers in an organizational setting. Nichols and Kang (2012) reported reasonable response rates in merely asking strangers for help; they found that help was more likely when the query requested quick, non-personal answers. It remains to be seen whether Nichols and Kang's approach can scale. White et al. (2011) investigated the relationship between community size and contact rate in a social Q&A system. In their experiment they found a social Q&A system would perform better as community size increased, but users preferred not being bothered too often.

Second, there has been work on innovative new systems. Aardvark, a commercial system, routed questions to appropriate others via email. Horowitz and Kamvar (2012) provided a detailed description of Aardvark's system and its algorithms.

4. Sharing Knowledge and Expertise: Future Directions

4.1 Interconnected practices: expertise infrastructures

It may seem a paradox that while practices are becoming interconnected (as argued above), the 'expert' may become an anonymous or even intangible actor. But the new technological options also propel new types of knowledge-management interactions based on micro-activities: Community, crowdsourcing, and microtasking environments make it very easy and convenient to help others by sharing one's expertise. As was mentioned above, one phenomenon is the growing importance of 'social search', which makes it easier to help knowledge seekers navigate to someone with expertise. It is clear that social search has different characteristics, and different capabilities, from Q&A communities, but more work will be required to fully determine these. As well, social search combines elements of expertise sharing through online places with expertise finding, providing additional possibilities for design.

There is a wide array of new sources of collaboratively constructed information spaces becoming available. Numerous papers have examined how Wikipedia has been constructed (e.g., Nagar 2012 and Forte et al. 2012 to pick only two). Other wikis have also been examined. Hansen et al. (2007) looked at a wiki associated with a technical distribution list. Kittur and Kraut (2010) in an impressive study examined 6811 wikis for the quality and the processes of activities potentially contributing to knowledge construction. While the knowledge sharing of Wikipedia is obvious, many other forms of informal information now exist on the Web. It remains to be seen how valuable these will be over time, but they are often reused and interestingly, their systems often combine some level of repository and communication. For example, Ko and Chilana (2010) looked at Mozilla bug reports, and Gilbert and Karahalios (2010) looked at Amazon reviews. Dabbish et al. (2012) in their study of Github, a code repository, found that code projects were regularly used by participants to stay current and follow new technical knowledge. Even community-constructed mathematical proofs (Cranshaw and Kittur 2011) have been examined.

Among the social media used for informal communication, microblogging platforms like Twitter with its openness and visibility have spawned new ways of collecting knowledge, which can be exemplified using approaches to support crisis management. Starbird et al. (2010) described how people affected by a flood used social media not only to communicate, but also to organize information flows, and suggested an approach to make use of these dynamics systematically (Vieweg et al. 2010). Vieweg (2012) then presented a framework to identify tweets that carry information about an ongoing crisis situation. Forming a volatile repository of very situated pieces of information, tweets may be analyzed to generate a rich picture of ongoing events while still providing a link back to the authors as original experts about local situations, strategies, and activities.

At the fringes of these developments in social media, other forms of involving human expertise in problem solving have emerged that make use of often partial expertise of often anonymous experts. Crowdsourcing (e.g., for sorting images, in von Ahn and Dabbish 2004) and microtask approaches (e.g., for language translation, in Ambati et al. 2012) offer new ways of generating knowledge that have proved helpful for a range of problem structures. Still, relying on knowledge aggregated from an anonymous crowd can be problematic (Kittur et al. 2008). There is ongoing work about supporting more complex task structures (e.g., Ambati et al. 2012; Dow et al. 2012; Kittur et al. 2012; Kulkarni et al. 2012), which may

alleviate the problem. (See Kittur et al. 2013 for a review of 'human computation' or 'social computation'), Other solutions are to not choose 'arbitrary' experts, but 'crowds' defined e.g., by location (Heimerl et al. 2012), or to combine expert and crowd results (Kandasamy et al. 2012).

CSCW research will need to investigate whether and how each field of practice we encounter is changed by these new opportunities. Reconsidering the medical domain for example, it might be interesting to look for changes in physician/nurse or clinician/patient cooperation through these emerging expertise resources.

5. Conclusion

CSCW has contributed many findings to a better understanding of knowledge and expertise sharing. In this paper, we separated the studies and their findings into two 'generations', noting that the distinction between the two is often only rhetorical, with recent studies (and many older ones) including elements of both. We used 'knowledge sharing' and 'expertise sharing' to separate out the two.

In what we termed the first generation, the knowledge sharing studies were artifact-centered, based in a repository model. These studies contributed many findings about the social nature of information production, storage, retrieval, and reuse that have been useful. CSCW researchers, and those in adjacent research fields, contributed to a general understanding that information use is heavily situated and socially contextualized. For practitioners, these studies served as critical reminders that one could not focus on the technical alone. Even technical studies were, necessarily, socio-technical.

In these studies, CSCW used and extended the concepts of boundary objects, common information spaces, and assemblies in order to understand and account for the element of social use. The concept of 'boundary objects' allows researchers to focus on the information objects that traverse organizational and social boundaries. Boundary objects were revealed to have a secret life where their ambiguity proved to be an asset, not an obstacle for knowledge sharing. CSCW researchers examined how they are used, and extensions have been offered so as to understand the social and communicative processes in which they are embedded. Other types of information objects have also come under investigation. 'Common information space' has been used to highlight the contingent understanding of information within social units, and it continues to enable researchers to define, analyze and design environments of knowledge exchange beyond organizational borders. Recently, concepts like 'assemblies' and 'ordering systems' have been used to investigate how information artifacts have been used in combination. It remains an important task for CSCW research to uncover commonalities and differences of ordering systems, especially with different stakeholder groups working with shared repositories.

In what we called a second generation of studies, the communication-centered expertise sharing studies focused attention on how knowledge is shared among knowledgeable actors. CSCW researchers, and those in adjacent areas, examined 'Communities of Practice' and 'social capital' in order to understand how and why knowledge is provided by others. The Communities of Practice concept established a perspective that embedded knowledge exchange not only into a practice of doing, but also in to a practice of learning (in and through practice). Social capital drew attention to relationship aspects in networks that guide the search for and exchange of knowledge. These concepts have been brought into expertise sharing systems, and researchers have produced a number of expertise-finding

and expertise-locating systems and studied them extensively. Expertise sharing, between coordinated and knowledgeable actors, has been found to be an often critical, if not required, part of knowledge sharing.

As time goes on, knowledge and expertise sharing are becoming standardized and routinized in many organizations and other collectivities. At the same time, new technologies, data collection mechanisms, analysis techniques, and fields of application have created many new research and design opportunities and challenges for the CSCW viewpoint to be extended anew.

The last 20 years of research into knowledge and expertise sharing has illuminated the enormous issues in dealing with information, knowledge, and expertise. It has challenged the normative, prescriptive understandings in other literatures. As a field, CSCW has created and examined new applications and systems over and over again, but more importantly, constantly attempted to view those applications and systems in terms of their social possibilities, issues, and challenges. In short, we have confronted the opportunities of 'knowledge management' socio-technically and in practice. As one might hope, the elaborated understanding has evolved into an accumulated literature, intellectual viewpoint, and epistemological stance. We fully expect the next twenty years, with its new technologies, interesting designs, and collaborative opportunities, to be as fruitful.

6. Addendum

Subsequent corrections, modifications, and additions to this document will be found at <http://www.socialworldsresearch.org/sharing-expertise-review>. The citation list, in bibtex format, will be available there as well.

7. Acknowledgements

To craft this paper, we acknowledge the support of the German Science Foundation's Graduate School 1769 'Locating Media'. The research on the Siegen side has been supported by various grants by the German Ministry of Education and Research, the German Ministry of Economics and Technology, and the State of North Rhine Westphalia over a period of more than 15 years.

This work was also supported, in part, by the U.S. National Science Foundation under grants IIS 0905460, IIS 0948639, and ACI 1035162.

We are indebted to many, many people who have helped us and engaged with us around these issues over many years. They are far too numerous to list here, but we are deeply appreciative. We do, however, want to acknowledge Kjeld Schmidt for his continued patience and interest, as well as the three anonymous reviewers who greatly strengthened this paper. We also thank all of current and past students and research staff for the enormous help they have provided.

8. Bibliography

- Ackerman, Mark S. (1993). Answer Garden: A Tool for Growing Organizational Memory of Work. Massachusetts Institute of Technology, Ph.D. Thesis.
- Ackerman, Mark S. (1994a). Augmenting the Organizational Memory: A Field Study of Answer Garden. *Proceedings of the Conference on Computer-Supported Cooperative Work (CSCW'94)*, pp. 243-252.

- Ackerman, Mark S. (1994b). Definitional and Contextual Issues in Organizational and Group Memories. *Proceedings of the Twenty-seventh IEEE Hawaii International Conference of System Sciences (HICSS 94)*, pp. 191-200.
- Ackerman, Mark S. (1998). Augmenting Organizational Memory: A Field Study of Answer Garden. *ACM Transactions on Information Systems*, vol. 16, no. 3, pp. 203-224.
- Ackerman, Mark S., James S. Boster, Wayne G. Lutters, and David W. McDonald (2002a). Who's There? The Knowledge Mapping Approximation Project. In Mark S. Ackerman, Volkmar Pipek and Volker Wulf (eds.), *Sharing Expertise: Beyond Knowledge Management*. Cambridge, MA: MIT Press, pp. 159-178.
- Ackerman, Mark S., and Christine Halverson (1998). Considering an Organization's Memory. *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW'98)*, pp. 39-48.
- Ackerman, Mark S., and Christine Halverson (2004a). Organizational Memory: Processes, Boundary Objects, and Trajectories. *Computer Supported Cooperative Work: The Journal of Collaborative Computing*, vol. 13, no. 2, pp. 155-189.
- Ackerman, Mark S., and Christine Halverson (2004b). Sharing Expertise: The Next Step for Knowledge Management. In Volker Wulf and Marylene Huysman (eds.), *Social Capital and Information Technology*. Cambridge, MA: MIT Press.
- Ackerman, Mark S., and Thomas W. Malone (1990). Answer Garden: A Tool for Growing Organizational Memory. *Proceedings of the ACM Conference on Office Information Systems*, pp. 31-39.
- Ackerman, Mark S., and Eric Mandel (1995). Memory in the Small: An Application to Provide Task-Based Organizational Memory for a Scientific Community. *Proceedings of the Hawaii International Conference of System Sciences (HICSS 95)*, pp. 323-332.
- Ackerman, Mark S., and Eric Mandel (1999). Memory in the Small: Combining Collective Memory and Task Support for a Scientific Community. *Journal of Organizational Computing and Electronic Commerce*, vol. 9, no. 2-3, pp. 105-127.
- Ackerman, Mark S., and David W. McDonald (1996). Answer Garden 2: Merging Organizational Memory with Collective Help. *Proceedings of the ACM Conference on Computer-Supported Cooperative Work (CSCW'96)*, pp. 97-105.
- Ackerman, Mark S., and Leysia Palen (1996). The Zephyr Help Instance: Promoting Ongoing Activity in a CSCW System. *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI'96)*, pp. 268-275.
- Ackerman, Mark S., Volkmar Pipek, and Volker Wulf (eds.) (2002b). *Sharing Expertise: Beyond Knowledge Management*. Cambridge, MA: MIT Press.
- Adler, Paul S, and Seok-Woo Kwon (2002). Social capital: prospects for a new concept. *Academy of Management Review*, vol. 27, no. 1, pp. 17-40.
- Alvesson, M., and D. Kärreman (2001). Odd couple: making sense of the curious concept of knowledge management. *Journal of Management Studies*, vol. 38, no. 7, pp. 995-1018.
- Ambati, Vamshi, Stephan Vogel, and Jaime Carbonell (2012). Collaborative workflow for crowdsourcing translation. *Proceedings of the ACM CSCW'12 Conference on Computer-Supported Cooperative Work*, pp. 1191-1194.
- Anderson, Stuart, Gillian Hardstone, Rob Procter, and Robin Williams (2007). Down in the (Data)base(ment) : Supporting Configuration in Organisational Information Systems. In Mark S. Ackerman, Christine A. Halverson, Thomas Erickson and Wendy A. Kellogg (eds.), *Resources, Co-Evolution, and Artifacts: Theory in CSCW*. New York: Springer, pp. 221-254.
- Andriessen, J.H. Erik (2005). Archetypes of Knowledge Communities. *Proceedings of the International Conference on Communities and Technologies 2005*, pp. 191-213.

- Auramäki, Esa, Mike Robinson, Anne Aaltonen, Mikko Kovalainen, Arja Liinamaa, and Taina Tuuna-Väiskä (1996). Paperwork At 78 kph. *Proceedings of the ACM Conference on Computer-Supported Cooperative Work (CSCW'96)*, pp. 370-379.
- Bandini, Stefania, Ettore Colombo, Gianluca Colombo, Fabio Sartori, and Carla Simone (2003). The role of knowledge artifacts in innovation management: the case of a chemical compound designer CoP. *Proceedings of the International Conference on Communities and Technologies*, pp. 327-345.
- Bannon, Liam, and Susan Bødker (1997). Constructing Common Information Spaces. *Proceedings of the European Computer Supported Cooperative Work Conference (E-CSCW'97)*, pp. 81-96.
- Bannon, Liam, and Kari Kuutti (1996). Shifting Perspectives on Organizational Memory: From Storage to Active Remembering. *Proceedings of the Hawaii International Conference on System Sciences (HICSS-29)*, pp. 156-167.
- Berg, Marc, and Geoffrey C. Bowker (1997). The Multiple Bodies of the Medical Record: Towards a Sociology of an Artifact. *Sociological Quarterly*, vol. 38, no. 3, pp. 513-537.
- Bergman, Mark, Gloria Mark, and Kalle Lyytinen (2004). Redefining Boundary Objects: An Examination of System Design. *Proceedings of the 6th International Conference on Organizational Discourse*.
- Berlin, Lucy M., Robin Jeffries, Vicki L. O'Day, Andreas Paepke, and Cathleen Wharton (1993). Where Did You Put It? Issues in the Design and Use of a Group Memory. *Proceedings of the ACM InterCHI'93 Conference on Human Factors in Computing Systems*, pp. 23-30.
- Bietz, Matthew J., and Charlotte P. Lee (2009). Collaboration in Metagenomics: Sequence Databases and the Organization of Scientific Work. *Proceedings of the the 11th European Conference on Computer-Supported Cooperative Work (E-CSCW'09)*, pp. 243-262.
- Blumer, Herbert (1986). *Symbolic Interactionism: Perspective and Method*. Berkeley: University of California Press.
- Boland, Jr., Richard J., Anil K. Maheshwari, Dov Te'eni, David G. Schwartz, and Ramkrishnan V. Tenkasi (1992). Sharing Perspectives in Distributed Decision Making. *Proceedings of the ACM CSCW'92 Conference on Computer-Supported Cooperative Work*, pp. 306-313.
- Bossen, Claus (2002). The parameters of common information spaces: the heterogeneity of cooperative work at a hospital ward. *Proceedings of the ACM CSCW'02 Conference on Computer-Supported Cooperative Work*, pp. 176-185.
- Bossen, Claus, Lotte Groth Jensen, and Flemming Witt (2012). Medical secretaries' care of records: the cooperative work of a non-clinical group. *Proceedings of the ACM CSCW'12 Conference on Computer-Supported Cooperative Work*, pp. 921-930.
- Boujut, Jean-François, and Eric Blanco (2003). Intermediary Objects as a Means to Foster Co-operation in Engineering Design. *Computer Supported Cooperative Work (CSCW)*, vol. 12, no. 2, pp. 205-219.
- Bourdieu, Pierre (1977). *Outline of a Theory of Practice*. Cambridge: Cambridge University Press.
- Bowker, Geoffrey C., and Susan L. Star (1999). *Sorting Things Out: Classification and Its Consequences*. Cambridge, MA: MIT Press.
- Bradner, E., W. Kellogg, and T. Erickson (1999). The Adoption and Use of "Babble": A Field Study of Chat in the Workplace. *Proceedings of the 6th European Conference on Computer Supported Cooperative Work (ECSCW'99)*, pp. 139 - 158.

- Brown, John S. , and P. Duguid (1991). Organizational Learning and Communities of Practice: Toward a Unified View of Working, Learning and Innovation. *Organization Science*, vol. 2, pp. 40-57.
- Brown, John Seeley, and Paul Duguid (2000). *The Social Life of Information*. Cambridge: Harvard Business Review Press.
- Cabitza, Federico, Gianluca Colombo, and Carla Simone (2013). Leveraging underspecification in knowledge artifacts to foster collaborative activities in professional communities. *International Journal of Human-Computer Studies*, vol. 71, no. 1, pp. 24-45.
- Cabitza, Federico, and Carla Simone (2012). Affording Mechanisms: An Integrated View of Coordination and Knowledge Management. *Computer Supported Cooperative Work (CSCW)*, vol. 21, no. 2-3, pp. 227-260.
- Cabitza, Federico, Carla Simone, and Marcello Sarini (2008). Knowledge Artifacts as Bridges between Theory and Practice: The Clinical Pathway Case. In Mark Ackerman, Rose Dieng-Kuntz, Carla Simone and Volker Wulf (eds.), *Knowledge Management in Action*. Boston: International Federation for Information Processing, pp. 37-50.
- Cabitza, Federico, Carla Simone, and Marcello Sarini (2009). Leveraging Coordinative Conventions to Promote Collaboration Awareness: The case of Clinical Records. *Computer Supported Cooperative Work (CSCW) Journal*, vol. 18, no. 4, pp. 301-330.
- Cohen, Don, and Laurence Prusak (2001). *In Good Company: How Social Capital Makes Organizations Work*. Boston: Harvard Business School Press.
- Constant, David, Sara Kiesler, and Lee Sproull (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.
- Cranshaw, Justin, and Aniket Kittur (2011). The polymath project: lessons from a successful online collaboration in mathematics. *Proceedings of the ACM CHI 2011 Conference on Human Factors in Computing Systems*, pp. 1865-1874.
- Dabbish, Laura, Colleen Stuart, Jason Tsay, and Jim Herbsleb (2012). Social coding in GitHub: transparency and collaboration in an open software repository. *Proceedings of the ACM CSCW'12 Conference on Computer-Supported Cooperative Work*, pp. 1277-1286.
- Dearman, David, and Khai N. Truong (2010). Why users of yahoo!: answers do not answer questions. *Proceedings of the ACM CHI 2010 Conference on Human Factors in Computing Systems*, pp. 329-332.
- Dörner, Christian, Volkmar Pipek, and Markus Won (2007). Supporting expertise awareness: finding out what others know. *Proceedings of the Symposium on Computer Human Interaction for the Management of Information Technology*.
- Dourish, Paul (2003). The Appropriation of Interactive Technologies: Some Lessons from Placeless Documents. *Computer Supported Cooperative Work (CSCW) Journal*, vol. 12, no. 4, pp. 465-490.
- Dourish, Paul, Victoria Bellotti, Wendy Mackay, and Chao-Ying Ma (1993). Information and context: lessons from the study of two shared information systems. *Proceedings of the ACM Conference on Organizational Computing*, pp. 42-51.
- Dourish, Paul, W. Keith Edwards, Anthony LaMarca, John Lamping, Karin Petersen, Michael Salisbury, Douglas B. Terry, and James Thornton (2000). Extending document management systems with user-specific active properties. *ACM Transactions on Information Systems*, vol. 18, no. 2, pp. 140-170.
- Dourish, Paul, W. Keith Edwards, Anthony LaMarca, and Michael Salisbury (1999). Presto: an experimental architecture for fluid interactive document spaces. *ACM Transactions on Computer-Human Interaction*, vol. 6, no. 2, pp. 133-161.

- Dow, Steven, Anand Kulkarni, Scott Klemmer, Bj Hartmann, ouml, and rn (2012). Shepherding the crowd yields better work. *Proceedings of the ACM CSCW'12 Conference on Computer-Supported Cooperative Work*, pp. 1013-1022.
- Ehrlich, Kate (2003). Locating Expertise: Design Issues for an Expertise Locator System. In Mark S. Ackerman, Volkmar Pipek and Volker Wulf (eds.), *Beyond Knowledge Management. Sharing Expertise*. Cambridge, USA: MIT Press, pp. 137-158.
- Ehrlich, Kate, C-Y. Lin, and V. Griffiths-Fisher (2007). Searching for Experts in the Enterprise: Combining Text and Social Network Analysis. *Proceedings of the ACM GROUP 2007 Conference on Supporting Group Work*, pp. 117-126.
- Faridani, Siamak, Ephrat Bitton, Kimiko Ryokai, and Ken Goldberg (2010). Opinion space: a scalable tool for browsing online comments. *Proceedings of the ACM CHI 2010 Conference on Human Factors in Computing Systems*, pp. 1175-1184.
- Fields, Bob, Paola Amaldi, and Antonello Tassi (2005). Representing collaborative work: the airport as common information space. *Cognition, Technology & Work*, vol. 7, no. 2, pp. 119-133.
- Finholt, Tom, and Lee S. Sproull (1990). Electronic Groups at Work. *Organization Science*, vol. 1, no. 1, pp. 41-64.
- Fiore, Andrew T., Scott Lee Tiernan, and Marc A. Smith (2002). Observed behavior and perceived value of authors in usenet newsgroups: bridging the gap. *Proceedings of the ACM CHI 2002 Conference on Human Factors in Computing Systems*, pp. 323-330.
- Fischer, Gerhard (2001). Communities of Interest: Learning through the Interaction of Multiple Knowledge Systems. *Proceedings of the 24th Information Systems Research Seminar in Scandinavia*, pp. 1-14.
- Fischer, Gerhard (2007). Designing socio-technical environments in support of meta-design and social creativity. *Proceedings of the 8th International Conference on Computer Supported Collaborative Learning (CSCL'07)*, pp. 2-11.
- Fisher, D., M. Smith, and H. T. Welsler (2006). You Are Who You Talk To: Detecting Roles in Usenet Newsgroups. *Proceedings of the 39th Annual Hawaii International Conference on System Sciences, (HICSS '06)*.
- Fitzpatrick, Geraldine (2003). Emergent Expertise Sharing in a New Community. In Mark S. Ackerman, Volkmar Pipek and Volker Wulf (eds.), *Sharing Expertise: Beyond Knowledge Management*. Cambridge, MA: MIT Press, pp. 81-110.
- Forte, Andrea, Niki Kittur, Vanessa Larco, Haiyi Zhu, Amy Bruckman, and Robert E. Kraut (2012). Coordination and beyond: social functions of groups in open content production. *Proceedings of the ACM CSCW'12 Conference on Computer-Supported Cooperative Work*, pp. 417-426.
- Gazan, Rich (2011). Social Q&A. *Journal of the American Society for Information Science and Technology*, vol. 62, no. 12, pp. 2301-2312.
- Gilbert, Eric, and Karrie Karahalios (2010). Understanding deja reviewers. *Proceedings of the ACM CSCW'10 Conference on Computer-Supported Cooperative Work*, pp. 225-228.
- Goggins, Sean , Isa Jahnke, and Volker Wulf (eds.) (2013). *Computer-Supported Collaborative Learning at the Workplace*. London: Springer.
- Goldberg, Yaron, Marilyn Safran, and Ehud Shapiro (1992). Active Mail -- A Framework for Implementing Groupware. *Proceedings of the ACM CSCW'92 Conference on Computer-Supported Cooperative Work*, pp. 75-83.
- Groth, K., and J. Bowers (2001). On finding things out: Situating organisational knowledge in CSCW. *Proceedings of the Seventh European Conference on Computer-Supported Cooperative Work*, pp. 279-298.

- Guy, I, M. Jacovi, N. Mushulam, I. Ronen, and E. Shahr (2008a). Public vs. private: comparing public social network information with email. *Proceedings of the ACM CHI 2008 Conference on Human Factors in Computing Systems*, pp. 393-402.
- Guy, I, S. Ur, and I. Ronen (2012). Best Faces Forward: A Large-scale Study of People Search in the Enterprise. *Proceedings of the ACM CHI 2012 Conference on Human Factors in Computing Systems*, pp. 1775-1784.
- Guy, I., M. Jacovi, E. Shahr, N. Meshulam, V. Soroka, and S. Farrell (2008b). Harvesting with SONAR: the value of aggregating social network information. *Proceedings of the ACM CHI 2008 Conference on Human Factors in Computing Systems*, pp. 1017-1026.
- Halverson, Christine A., and Mark S. Ackerman (2003). Yeah, the Rush ain't here yet -- Take a break: Creation and use of an artifact as organizational memory. *Proceedings of the IEEE Hawaii International Conference of System Sciences (HICSS'03)*.
- Halverson, Christine A., and Mark S. Ackerman (2007). The Birth of an Organizational Resource: The Surprising Life of a Cheat Sheet. In Mark S. Ackerman, Christine A. Halverson, Thomas Erickson and Wendy A. Kellogg (eds.), *Resources, Co-Evolution, and Artifacts: Theory in CSCW*. London: Springer, pp. 9-36.
- Halverson, Christine A., Thomas Erickson, and Mark S. Ackerman (2004). Organizational issues: Behind the help desk: evolution of a knowledge management system in a large organization. *Proceedings of the ACM Conference on Computer Supported Cooperative Work*, pp. 304-313.
- Hansen, Derek L., Mark S. Ackerman, Paul J. Resnick, and Sean Munson (2007). Virtual Community Maintenance with a Collaborative Repository. *Proceedings of the Annual Meeting of the American Society for Information Science and Technology (ASIS&T)*, pp. 1-20.
- Hardstone, Gillian, Mark Hartswood, Rob Procter, Roger Slack, Alex Voss, and Gwyneth Rees (2004). Supporting informality: team working and integrated care records. *Proceedings of the ACM CSCW'04 Conference on Computer-Supported Cooperative Work*, pp. 142-151.
- Harper, Richard (1998). *Inside the IMF*. New York: Academic Press.
- Harper, Richard, Rob Procter, Dave Randall, and Mark Rouncefield (2001). 'Safety in numbers': calculation and document re-use in knowledge work. *Proceedings of the International Conference on Supporting Group Work (GROUP'01)*, pp. 242-251.
- Harper, Richard, and Abigail Sellen (1995). Collaborative Tools and the Practicalities of Professional Work at the International Monetary Fund. *Proceedings of the ACM CHI'95 Conference on Human Factors in Computing Systems*, pp. 122-129.
- Heath, Christian, and Paul Luff (1996). Documents and Professional Practice: "Bad" Organizational Reasons for "Good" Clinical Records. *Proceedings of the ACM Conference on Computer-Supported Cooperative Work (CSCW'96)*, pp. 354-363.
- Heimerl, Kurtis, Brian Gawalt, Kuang Chen, Tapan Parikh, and Bjorn Hartmann (2012). CommunitySourcing: engaging local crowds to perform expert work via physical kiosks. *Proceedings of the ACM CHI 2012 Conference on Human Factors in Computing Systems*, pp. 1539-1548.
- Hiltz, Starr Roxanne, and Murray Turoff (1993). *The network nation: human communication via computer*. Cambridge, MA: MIT Press.
- Hinds, Pam, and Jeffrey Pfeffer (2003). Why organizations don't 'know what they know': Cognitive and motivational factors affecting the transfer of expertise. In Mark Ackerman, Volkmar Pipek and Volker Wulf (eds.), *Sharing Expertise. Beyond Knowledge Management*. Cambridge, MA: MIT Press, pp. 3-26.

- Hinrichs, Joachim, Volkmar Pipek, and Volker Wulf (2005). Context grabbing: Assigning metadata in large document collections. *Proceedings of the Ninth European Conference on Computer-Supported Cooperative Work*, pp. 367-386.
- Horowitz, Damon, and Sepandar D. Kamvar (2012). Searching the village: models and methods for social search. *Communications of the ACM*, vol. 55, no. 4, pp. 111-118.
- Hsieh, Gary, Scott E. Hudson, and Robert E. Kraut (2011). Donate for credibility: how contribution incentives can improve credibility. *Proceedings of the ACM CHI 2011 Conference on Human Factors in Computing Systems*, pp. 3435-3438.
- Hsieh, Gary, Robert E. Kraut, and Scott E. Hudson (2010). Why pay?: exploring how financial incentives are used for question & answer. *Proceedings of the ACM CHI 2010 Conference on Human Factors in Computing Systems*, pp. 305-314.
- Huysman, M.H., and Volker Wulf (2006). IT to support knowledge sharing in communities, towards a social capital analysis. *Journal of Information Technology*, vol. 21, pp. 40-51.
- Huysman, Marleen H., and Volker Wulf (2005). The role of information technology in building and sustaining the relational base of communities. *The Information Society*, vol. 21, no. 2, pp. 81-89.
- Huysman, Marleen, and Volker Wulf (eds.) (2004). *Social Capital & Information Technology*. Cambridge, MA: MIT Press.
- Jeon, Grace YoungJoo, Yong-Mi Kim, and Yan Chen (2010). Re-examining price as a predictor of answer quality in an online q&a site. *Proceedings of the ACM CHI 2010 Conference on Human Factors in Computing Systems*, pp. 325-328.
- Kamiya, Kenichi, Martin Roscheisen, and Terry Winograd (1996). Grassroots: Providing a Uniform Framework for Communicating, Sharing Information, and Organizing People. *Proceedings of the ACM CHI 96 Conference on Human Factors in Computing Systems*, pp. 239-240.
- Kandasamy, Durga M., Kristal Curtis, Armando Fox, and David Patterson (2012). Diversity within the crowd. *Proceedings of the ACM CSCW'12 Conference on Computer-Supported Cooperative Work, Companion* pp. 115-118.
- Katzenberg, Barbara, Fred Pickard, and John McDermott (1996). Computer Support for Clinical Practice: Embedding and Evolving Protocols of Care. *Proceedings of the ACM Conference on Computer-Supported Cooperative Work (CSCW'96)*, pp. 364-369.
- Kittur, Aniket, Ed H. Chi, and Bongwon Suh (2008). Crowdsourcing user studies with Mechanical Turk. *Proceedings of the ACM CHI'08 Conference on Human Factors in Computing Systems*, pp. 453-456.
- Kittur, Aniket, Susheel Khamkar, Paul Andreacuta, and Robert Kraut (2012). CrowdWeaver: visually managing complex crowd work. *Proceedings of the ACM CSCW'12 Conference on Computer-Supported Cooperative Work*, pp. 1033-1036.
- Kittur, Aniket, and Robert E. Kraut (2010). Beyond Wikipedia: coordination and conflict in online production groups. *Proceedings of the ACM CSCW'10 Conference on Computer-Supported Cooperative Work*, pp. 215-224.
- Kittur, Aniket, Jeffrey V. Nickerson, Michael Bernstein, Elizabeth Gerber, Aaron Shaw, John Zimmerman, Matt Lease, and John Horton (2013). The future of crowd work. *Proceedings of the ACM CSCW 2013 Conference on Computer Supported Cooperative Work*, pp. 1301-1318.
- Kling, Rob (1991). Value Conflicts in the Design and Organization of EFT Systems. In Charles Dunlop and Rob Kling (eds.), *Computerization and Controversy*. first ed. Boston: Academic Press, pp. 421-435.

- Ko, Andrew J., and Parmit K. Chilana (2010). How power users help and hinder open bug reporting. *Proceedings of the ACM CHI 2010 Conference on Human Factors in Computing Systems*, pp. 1665-1674.
- Kovalainen, Mikko, Mike Robinson, and Esa Auramaki (1998). Diaries at Work. *Proceedings of the ACM CSCW'98 Conference on Computer-Supported Cooperative Work*, pp. 49-58.
- Kulkarni, Anand, Matthew Can, and Bjorn Hartmann (2012). Collaboratively crowdsourcing workflows with Turkomatic. *Proceedings of the ACM CSCW'12 Conference on Computer-Supported Cooperative Work*, pp. 1003-1012.
- Lampe, Cliff, Nicole Ellison, and Charles Steinfield (2006). A face(book) in the crowd: social Searching vs. social browsing. *Proceedings of the ACM CSCW'06 Conference on Computer-Supported Cooperative Work*, pp. 167-170.
- Laqua, Sven, M. Angela Sasse, Steven Greenspan, and Carrie Gates (2011). Do you know dis?: a user study of a knowledge discovery tool for organizations. *Proceedings of the ACM CHI 2011 Conference on Human Factors in Computing Systems*, pp. 2887-2896.
- Lave, Jean, and Etienne Wenger (1991). *Situated Learning: Legitimate Peripheral Participation*. Cambridge, USA: Cambridge University Press.
- Lee, Charlotte P. (2007). Boundary Negotiating Artifacts: Unbinding the Routine of Boundary Objects and Embracing Chaos in Collaborative Work. *Computer-Supported Cooperative Work*, vol. 16, no. 3, pp. 307-339.
- Lesser, Eric L., Michael A. Fontaine, and Jason A. Slusher (eds.) (2000). *Knowledge and Communities*. Boston: Butterworth-Heinemann.
- Lin, C-Y., V. Griffiths-Fisher, Kate Ehrlich, and C. Desforges (2008). SmallBlue: People Mining for Expertise Search and Social Network Analysis. *IEEE Multimedia Magazine*, vol. 15, pp. 78-84.
- Lindstaedt, Stefanie N. (1996). Towards organizational learning: growing group memories in the workplace. *Proceedings of the Conference Companion for ACM Conference on Human Factors in Computing Systems (CHI'96)*, pp. 53-54.
- Lutters, Wayne G., and Mark S. Ackerman (2002). Achieving Safety: A Field Study of Boundary Objects in Aircraft Technical Support. *Proceedings of the ACM Conference on Computer-Supported Cooperative Work (CSCW'02)*, pp. 266-275.
- Lutters, Wayne G., and Mark S. Ackerman (2007). Beyond Boundary Objects: Collaborative Reuse in Aircraft Technical Support. *Computer-Supported Cooperative Work*, vol. 16, no. 3, pp. 341-372.
- Malone, Thomas W., Ken R. Grant, Franklyn A. Turbak, Stephen A. Brobst, and Michael D. Cohen (1987). Intelligent Information Sharing Systems. *Communications of the ACM*, vol. 30, no. 5, pp. 390-402.
- Malone, Thomas W., Kum-Yew Lai, and Christopher Fry (1995). Experiments with Oval: a radically tailorable tool for cooperative work. *ACM Transactions on Information Systems*, vol. 13, no. 2, pp. 177-205.
- Mambrey, P. , and M. Robinson (1997). Understanding the Role of Documents in a Hierarchical Flow of Work. *Proceedings of the ACM Group Conference (GROUP'97)*, pp. 119-127.
- Markus, M. Lynne (1983). Power, Politics, and MIS Implementation. *Communications of the ACM*, vol. 26, no. 6, pp. 430-444.
- McDonald, David W. (2000). *Supporting Nuance in Groupware Design: Moving from Naturalistic Expertise Location to Expertise Recommendation*.
- McDonald, David W. (2001). Evaluating Expertise Recommendations. *Proceedings of the ACM Conference on Supporting Group Work (Group'01)*, pp. 214-223.

- McDonald, David W. (2003). Recommending collaboration with social networks: a comparative evaluation. *Proceedings of the ACM CHI 2003 Conference on Human Factors in Computing Systems*, pp. 593-600.
- McDonald, David W., and Mark S. Ackerman (1998). Just Talk to Me: A Field Study of Expertise Location. *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW '98)*, pp. 315-324.
- McDonald, David W., and Mark S. Ackerman (2000). Expertise Recommender: A Flexible Recommendation System Architecture. *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW'2000)*, pp. 231-240.
- Mockus, Audris, and James D. Herbsleb (2002). Expertise Browser: A Quantitative Approach to Identifying Expertise. *Proceedings of the International Conference on Software Engineering (ICSE'02)*, pp. 503-512.
- Morris, Meredith Ringel, Jaime Teevan, and Katrina Panovich (2010a). A Comparison of Information Seeking Using Search Engines and Social Networks. *Proceedings of the Fourth International AAI Conference on Weblogs and Social Media (ICWSM'10)*, pp. 291-294.
- Morris, Meredith Ringel, Jaime Teevan, and Katrina Panovich (2010b). What do people ask their social networks, and why?: a survey study of status message & behavior. *Proceedings of the ACM CHI 2010 Conference on Human Factors in Computing Systems*, pp. 1739-1748.
- Morrison, J (1993). Team Memory: Information Management for Business Teams. *Proceedings of the 26th Hawaii International Conference on System Sciences*, pp. 122-131.
- Munkvold, Glenn, and Gunnar Ellingsen (2007). Common Information Spaces along the illness trajectories of chronic patients. *Proceedings of the Tenth European Conference on Computer-Supported Cooperative Work*, pp. 291-310.
- Nagar, Yiftach (2012). What do you think?: the structuring of an online community as a collective-sensemaking process. *Proceedings of the ACM CSCW'12 Conference on Computer-Supported Cooperative Work*, pp. 393-402.
- Nahapiet, J., and S. Ghoshal (1998). Social Capital, Intellectual Capital and the Organizational Advantage. *Academy of Management Review*, vol. 23, no. 2, pp. 242-266.
- Nam, Kevin K., and Mark S. Ackerman (2007). Arkose: Reusing Informal Information in Online Communities. *Proceedings of the ACM Conference on Supporting Group Work (Group'07)*, pp. 137-146.
- Nichols, Jeffrey, and Jeon-Hyung Kang (2012). Asking questions of targeted strangers on social networks. *Proceedings of the ACM CSCW'12 Conference on Computer-Supported Cooperative Work*, pp. 999-1002.
- Nonaka, Ikujiro, and Hirotaka Takeuchi (1995). *The knowledge-creating company: How Japanese companies create the dynamics of innovation*. New York: Oxford University Press.
- Normark, Maria, and Dave Randall (2005). Local expertise at an emergency call centre. *Proceedings of the Ninth European Conference on Computer-Supported Cooperative Work*, pp. 347-366.
- O'Day, V. , A. Adler, A. Kuchinsky, and A. Bouch (2001). When worlds collide: Molecular biology as interdisciplinary collaboration. *Proceedings of the the Seventh European Conference on Computer-Supported Cooperative Work*, pp. 399-418.
- Orlikowski, Wanda J. (1992). Learning from Notes: Organizational Issues in Groupware Implementation. *Proceedings of the Computer Supported Cooperative Work (CSCW'92)*, pp. 362-369.

- Orr, Julian (1986). Narratives at Work: Storytelling as Cooperative Diagnostic Activity. *Proceedings of the Computer Supported Cooperative Work (CSCW 86)*, pp. 62-72.
- Orr, Julian (1996). *Talking About Machines: an Ethnography of a Modern Job*. Ithaca: ILR/Cornell University Press.
- Pal, Aditya, James Margatan, and Joseph Konstan (2012). Question temporality: identification and uses. *Proceedings of the ACM CSCW'12 Conference on Computer-Supported Cooperative Work*, pp. 257-260.
- Panovich, Katrina, Rob Miller, and David Karger (2012). Tie strength in question & answer on social network sites. *Proceedings of the ACM CSCW'12 Conference on Computer-Supported Cooperative Work*, pp. 1057-1066.
- Pentland, Brian T (1992). Organizing moves in software support hot lines. *Administrative Science Quarterly*, vol. 37, no. 4, pp. 527-548.
- Pipek, Volkmar, and Volker Wulf (2003). Pruning the Answer Garden: Knowledge sharing in maintenance engineering. *Proceedings of the Eighth European Conference on Computer-Supported Cooperative Work*, pp. 1-20.
- Pipek, Volkmar, Volker Wulf, and A. Johri (2012). Bridging Artifacts and Actors: Expertise Sharing in Organizational Ecosystems. *Computer Supported Cooperative Work (CSCW)*, vol. 21, no. 2-3, pp. 261-282.
- Pipek, Volkmar, Joachim Hinrichs, and Volker Wulf (2003). Sharing expertise: challenges for technical support. In Mark Ackerman, V. Pipek and Volker Wulf (eds.), *Sharing Expertise-Beyond Knowledge Management*, MIT Press, Cambridge. Cambridge, MA: MIT Press, pp. 111-136.
- Polanyi, Michael (1967). *The Tacit Dimension*. London: Routledge and Kegan Paul.
- Preece, Jenny (2000). *Online Communities*. New York: Wiley.
- Randall, Dave, Richard Harper, and Mark Rouncefield (2007). *Fieldwork for Design: Theory and Practice*. London: Springer.
- Reddy, Madhu, and Paul Dourish (2002). A Finger on the Pulse: Temporal Rhythms and Information Seeking in Medical Work. *Proceedings of the ACM CSCW'02 Conference on Computer Supported Cooperative Work*, pp. 344-353.
- Reddy, Madhu, Paul Dourish, and Wanda Pratt (2001). Coordinating Heterogeneous Work: Information and Representation in Medical Care. *Proceedings of the European Conference on Computer Supported Cooperative Work (ECSCW'01)*, pp. 239-258.
- Reichling, Tim, Kai Schubert, and Volker Wulf (2005). Matching Human Actors based on their Texts: Design and Evaluation of an Instance of the ExpertFinding Framework. *Proceedings of the International Conference on Supporting Group Work (Group'05)*, pp. 61-70.
- Reichling, Tim, and Michael Veith (2005). Expertise sharing in a heterogeneous organizational environment. *Proceedings of the European Conference on Computer-Supported Cooperative Work (E-CSCW'05)*, pp. 325-345.
- Reichling, Tim, Michael Veith, and Volker Wulf (2007). Expert Recommender: Designing for a Network Organization. *Computer Supported Cooperative Work*, vol. 16, pp. 431-465.
- Reichling, Tim, and Volker Wulf (2009). Expert Recommender System in Practice: Evaluating Semiautomatic Profile Generation. *Proceedings of the ACM CHI 2009 Conference on Human Factors in Computing Systems*, pp. 59-68.
- Ribak, Amnon, Michal Jacovi, and Vladimir Soroka (2002). "Ask before you search": peer support and community building with ReachOut. *Proceedings of the ACM CSCW'02 Conference on Computer-Supported Cooperative Work*, pp. 126-135.

- Rodden, Tom (1996). Populating the Application: A Model of Awareness for Cooperative Applications. *Proceedings of the ACM CSCW'96 Conference on Computer-Supported Cooperative Work*, pp. 87-96.
- Rolland, Knut H., Vidar Hepsø, and Eric Monteiro (2006). Conceptualizing common information spaces across heterogeneous contexts: mutable mobiles and side-effects of integration. *Proceedings of the ACM CSCW'06 Conference on Computer-Supported Cooperative Work*, pp. 493-500.
- Saeed, Saqib, Volkmar Pipek, Markus Rohde, and Volker Wulf (2010). Managing nomadic knowledge: a case study of the European social forum. *Proceedings of the ACM CHI 2010 Conference on Human Factors in Computing Systems*, pp. 537-546.
- Schein, Edgar H. (1978). *Career Dynamics*. Reading, MA: Addison-Wesley.
- Schmidt, Kjeld (1999). Of maps and scripts: The status of formal constructs in cooperative work. *Information and Software Technology*, vol. 41, pp. 319-329.
- Schmidt, Kjeld (2002). The problem with 'awareness'. *Computer Supported Cooperative Work*, vol. 11, pp. 285-298.
- Schmidt, Kjeld (2012). The trouble with tacit knowledge. *Computer Supported Cooperative Work*, vol. 21, no. 2-3, pp. 163-225.
- Schmidt, Kjeld, and Liam Bannon (1992). Taking CSCW Seriously: Supporting Articulation Work. *Computer Supported Cooperative Work*, vol. 1, no. 1-2, pp. 7-40.
- Schmidt, Kjeld, and Carla Simone (1996). Coordination Mechanisms: Towards a Conceptual Foundation of CSCW Systems Design. *Computer Supported Cooperative Work*, vol. 5, no. 2/3, pp. 155-200.
- Schmidt, Kjeld, and Ina Wagner (2003). Ordering systems: coordinative practices in architectural design and planning. *Proceedings of the International Conference on Supporting Group Work (GROUP'03)*, pp. 274-283.
- Schmidt, Kjeld, and Ina Wagner (2004). Ordering Systems: Coordinative Practices and Artifacts in Architectural Design and Planning. *Computer-Supported Cooperative Work*, vol. 13, no. 5/6, pp. 349-408.
- Schmidt, Kjeld, Ina Wagner, and Marianne Tolar (2007). Permutations of cooperative work practices: a study of two oncology clinics. *Proceedings of the International Conference on Supporting Group Work (GROUP'07)*, pp. 1-10.
- Shami, N. Sadat, Kate Ehrlich, and D. R. Millen (2008). Pick me! Link selection in expertise search. *Proceedings of the ACM CHI 2008 Conference on Human Factors in Computing Systems*, pp. 1089-1092.
- Shapiro, Dan (1994). The Limits of Ethnography: Combining Social Sciences for CSCW. *Proceedings of the ACM CSCW'94 Conference on Computer-Supported Cooperative Work*, pp. 417-428.
- Shipman, Frank, III, and Catherine C. Marshall (1999). Formality Considered Harmful: Experiences, Emerging Themes, and Directions on the Use of Formal Representations in Interactive Systems. *Computer Supported Cooperative Work*, vol. 8, no. 4, pp. 333-352.
- Shipman, Frank, III, Robert Airhart, Haowei Hsieh, Preetam Maloor, J. Michael Moore, and Divya Shah (2001). Visual and spatial communication and task organization using the visual knowledge builder. *Proceedings of the International Conference on Supporting Group Work (GROUP'01)*, pp. 260-269.
- Shipman, Frank, III, and Ray McCall (1999). Supporting Incremental Formalization with the Hyper-Object Substrate. *ACM Transactions on Information Systems*, vol. 17, no. 2, pp. 199-227.

- Shipman, Frank, III, and Raymond McCall (1994). Supporting Knowledge-Base Evolution with Incremental Formalization. *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI'94)*, pp. 285-291.
- Singh, Vandana, and Michael B. Twidale (2008). The confusion of crowds: non-dyadic help interactions. *Proceedings of the ACM CSCW'08 Conference on Computer-Supported Cooperative Work*, pp. 699-702.
- Star, Susan Leigh, and James R. Griesemer (1989). Institutional Ecology, "Translations" and Boundary Objects: Amateurs and Professionals in Berkeley's Museum of Vertebrate Zoology, 1907-39. *Social Studies of Science*, vol. 19, pp. 387-420.
- Starbird, Kate, Leysia Palen, Amanda L. Hughes, and Sarah Vieweg (2010). Chatter on the red: what hazards threat reveals about the social life of microblogged information. *Proceedings of the ACM CSCW'10 Conference on Computer-Supported Cooperative Work*, pp. 241-250.
- Stevens, Gunnar, Michael Veith, and Volker Wulf (2005). Bridging among Ethnic Communities by Cross-cultural Communities of Practice. *Proceedings of the Second International Conference on Communities and Technologies (C&T 2005)*, pp. 377-396.
- Stevens, Gunnar, and Volker Wulf (2002). A new dimension in access control: studying maintenance engineering across organizational boundaries. *Proceedings of the ACM Conference on Computer-Supported Cooperative Work (CSCW'02)*, pp. 196-205.
- Stewart, Thomas A. (1991). BrainPower. *Fortune*, June 3, 1991, 44-60.
- Strauss, Anselm L. (1993). *Continual Permutations of Action*. New York: Aldine de Gruyter.
- Streeter, Lynn A., and Karen E. Lochbaum (1988). Who Knows: A System Based on Automatic Representation of Semantic Structure. *Proceedings of the RIAO 88 Conference*, pp. 380-388.
- Su, Norman Makoto, and H.N. Wilensky (2011). Doing Business with Theory: Communities of Practice in Knowledge Management. *Computer Supported Cooperative Work*, vol. 21, pp. 1-50.
- Suchman, Lucy A. (1983). Office Procedure as Practical Action: Models of Work and System Design. *ACM Transactions on Office Information Systems*, vol. 1, no. 4, pp. 320-328.
- Suchman, Lucy A. (1987). *Plans and Situated Actions: The Problem of Human-Computer Communication*. New York: Cambridge University Press.
- Suchman, Lucy A. (1994). Do Categories Have Politics? *Computer Supported Cooperative Work*, vol. 2, pp. 177-190.
- Suchman, Lucy A., and Eleanor Wynn (1984). Procedures and Problems in the Office. *Office: Technology and People*, vol. 2, pp. 133-154.
- Tausczik, Yla R., and James W. Pennebaker (2012). Participation in an online mathematics community: differentiating motivations to add. *Proceedings of the ACM CSCW'12 Conference on Computer-Supported Cooperative Work*, pp. 207-216.
- Terveen, Loren, and David W. McDonald (2005). Social matching: A framework and research agenda. *ACM Transactions on Computer-Human Interaction*, vol. 12, no. 3, pp. 401-434.
- Terveen, Loren, Peter Selfridge, and M. David Long (1995). Living Design Memory: Framework, Implementation, Lessons Learned. *Human-Computer Interaction*, vol. 10, no. 1, pp. 1-38.
- Torrey, Cristen, David W. McDonald, Bill N. Schilit, and Sara Bly (2007). How-To Pages: Informal Systems of Expertise Sharing. *Proceedings of the Tenth European Conference on Computer-Supported Cooperative Work*, pp. 391-410.
- Turner, Tammara Combs, Marc A. Smith, Danyel Fisher, and Howard T. Welser (2005). Picturing Usenet: Mapping Computer-Mediated Collective Action. *Journal of Computer-Mediated Communication*, 10.

- Van House, Nancy A. , Mark H. Butler, and Lisa R. Schiff (1998). Cooperative Knowledge Work and Practices of Trust: Sharing Environmental Planning Data Sets. *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW'98)*, pp. 335-343.
- Vieweg, Sarah (2012). Twitter communications in mass emergency: contributions to situational awareness. *Proceedings of the ACM CSCW'12 Conference on Computer-Supported Cooperative Work Companion*, pp. 227-230.
- Vieweg, Sarah, Amanda L. Hughes, Kate Starbird, and Leysia Palen (2010). Microblogging during two natural hazards events: what twitter may contribute to situational awareness. *Proceedings of the ACM CHI 2010 Conference on Human Factors in Computing Systems*, pp. 1079-1088.
- von Ahn, Luis, and Laura Dabbish (2004). Labeling images with a computer game. *Proceedings of the ACM CHI 2004 Conference on Human Factors in Computing Systems*, pp. 319-326.
- Walsh, James P., and Gerardo Rivera Ungson (1991). Organizational Memory. *The Academy of Management Review*, vol. 16, no. 1, pp. 57-91.
- Wellman, Barry (1999). *Networks In The Global Village: Life In Contemporary Communities*. Boulder: Westview Press.
- Wenger, E., Richard McDermott, and William Synder (2002). *Cultivating Communities of Practice*. Boston: Harvard Business School Press.
- Wenger, Etienne (1998). *Communities of practice : learning, meaning, and identity*. New York: Cambridge University Press.
- White, Ryen W., Matthew Richardson, and Yandong Liu (2011). Effects of community size and contact rate in synchronous social q & a. *Proceedings of the ACM CHI 2011 Conference on Human Factors in Computing Systems*, pp. 2837-2846.
- Won, Markus, and Volkmar Pipek (2003). Sharing knowledge on knowledge - the exact peripheral expertise awareness system. *International Journal on Universal Computer Science*, vol. 9, pp. 1388-1397.
- Wulf, Volker (1997). Storing and Retrieving Documents in a Shared Workspace: Experiences from the Political Administration. *Proceedings of the Conference on Human Computer Interaction: INTERACT 97*, pp. 469-476.
- Wulf, Volker, and Markus Rohde (1995). Towards an integrated organization and technology development. *Proceedings of the Conference on Designing Interactive Systems (DIS 1995)*, pp. 55-64.
- Wulf, Volker, Markus Rohde, Volkmar Pipek, and Gunnar Stevens (2011). Engaging with practices: design case studies as a research framework in CSCW. *Proceedings of the ACM CSCW'11 Conference on Computer-Supported Cooperative Work*, pp. 505-512.
- Yamauchi, Yutaka, Makoto Yokozawa, Takeshi Shinohara, and Toru Ishida (2000). Collaboration with Lean Media: How Open-Source Software Succeeds. *Proceedings of the CSCW'00 Conference on Computer-Supported Cooperative Work*, pp. 329-338.
- Yang, J., X. Wei, M. Ackerman, and L. Adamic (2010). Activity lifespan: An analysis of user survival patterns in online knowledge sharing communities. *Proceedings of the International Conference on Weblogs and Social Media (ICWSM'10)*.
- Yarosh, S., T. Matthews, and M. Zhou (2012). Asking the right person: supporting expertise selection in the enterprise. *Proceedings of the ACM CHI 2012 Conference on Human Factors in Computing Systems*, pp. 2247-2256.
- Zhang, Jun, Mark S. Ackerman, and Lada Adamic (2007). QuME: a mechanism to support expertise finding in online help-seeking communities. *Proceedings of the ACM Symposium on User Interface Software and Technology (UIST 2007)*, pp. 111-114.

Zhang, Jun, Yan Qu, Jane Cody, and Yulingling Wu (2010). A case study of micro-blogging in the enterprise: use, value, and related issues. *Proceedings of the ACM CHI 2010 Conference on Human Factors in Computing Systems*, pp. 123-132.

Zhou, Xiaomu, Mark Ackerman, and Kai Zheng (2011). CPOE workarounds, boundary objects, and assemblages. *Proceedings of the ACM CHI 2011 Conference on Human Factors in Computing Systems*, pp. 3353-3362.

1 In this paper, we will not differentiate between knowledge and information: We could easily spend several lifetimes teasing the two apart, and colloquial uses are sufficient (as argued in Normark and Randall 2005).

2 These generations are not completely distinct - some studies may have foreshadowed a generation, and later studies sometimes combined them. It is difficult to cover so much intellectual ground without being reduced to making seemingly arbitrary categorizations, and we beg the reader's indulgence on this point.