Computer Supported Cooperative Work **11:** 411–446, 2002.

A Descriptive Framework of Workspace Awareness for Real-Time Groupware

CARL GUTWIN[1] & SAUL GREENBERG[2]
[1]*Department of Computer Science, University of Saskatchewan, 57 Campus Drive, Saskatoon, Canada, S7N 5A9 (E-mail: gutwin@cs.usask.ca);* [2]*Department of Computer Science, University of Calgary, 2500 University Drive NW, Calgary, Canada, T2N 1N4 (E-mail: saul@cpsc.ucalgary.ca)*

Abstract. Supporting awareness of others is an idea that holds promise for improving the usability of real-time distributed groupware. However, there is little principled information available about awareness that can be used by groupware designers. In this article, we develop a descriptive theory of awareness for the purpose of aiding groupware design, focusing on one kind of group awareness called *workspace awareness*. We focus on how small groups perform generation and execution tasks in medium-sized shared workspaces – tasks where group members frequently shift between individual and shared activities during the work session. We have built a three-part framework that examines the concept of workspace awareness and that helps designers understand the concept for purposes of designing awareness support in groupware. The framework sets out elements of knowledge that make up workspace awareness, perceptual mechanisms used to maintain awareness, and the ways that people use workspace awareness in collaboration. The framework also organizes previous research on awareness and extends it to provide designers with a vocabulary and a set of ground rules for analysing work situations, for comparing awareness devices, and for explaining evaluation results. The basic structure of the theory can be used to describe other kinds of awareness that are important to the usability of groupware.

Key words: awareness, groupware design, groupware usability, real-time distributed groupware, situation awareness, shared workspaces, workspace awareness

1. Introduction

Awareness has recently begun to receive considerable attention in CSCW and groupware research (e.g. Dourish and Bellotti, 1992; McDaniel and Brinck, 1997; Rodden, 1996; Gutwin and Greenberg, 1998a). While staying aware of others is something that we take for granted in the everyday world, maintaining this awareness has proven to be difficult in real-time distributed systems where information resources are poor and interaction mechanisms are foreign. As a result, working together through a groupware system often seems inefficient and clumsy compared with face-to-face work. It is becoming more and more apparent that being able to stay aware of others plays an important role in the fluidity and naturalness of collaboration, and supporting awareness of others is looked on as one way of reducing the characteristic awkwardness of remote collaboration. Awareness is a design concept that holds promise for significantly improving the usability of real-time distributed groupware.

Despite this attention, no clear overall picture of awareness has yet emerged from the CSCW community. With a few exceptions, awareness support presented to date involves localized solutions to specific domain problems, and isolated approaches and principles that are difficult to generalize to other situations. Most importantly, this void means that groupware designers have little principled information available to them about how to support awareness in other domains and new systems. Faced with a blank slate for each new application, designers must reinvent awareness from their own experience of what it is, how it works, and how it is used in the task at hand.

Our goal in this article is to develop a descriptive theory of awareness for the purpose of aiding groupware design. We synthesize and organize existing research on awareness, and extend this work through a conceptual framework. Our motivation is the observation that current groupware systems are not particularly usable – and here we are more concerned with how well a system supports activities of collaboration like communication, coordination, and assistance, than we are with how well the system supports the domain task (Salas, 1995). Our overall research hypothesis is that helping people to stay aware in groupware workspaces will improve a groupware system's usability.

Our conceptual framework differs from previous work on groupware awareness in three ways:

- it integrates and expands upon a variety of observations and previous theories of awareness;
- it addresses a particular type of situation – small groups working over medium sized shared workspaces; and
- it is intended to assist the iterative design of real-time distributed groupware.

We examine one kind of awareness in collaboration – called *workspace awareness* because of its intimate relationship with shared workspaces – and construct a framework that describes the concept for use in groupware design. Workspace awareness is the up-to-the-moment understanding of another person's interaction with a shared workspace (Gutwin and Greenberg, 1996). Workspace awareness (WA) involves knowledge about where others are working, what they are doing, and what they are going to do next. This information is useful for many of the activities of collaboration – for coordinating action, managing coupling, talking about the task, anticipating others' actions, and finding opportunities to assist one another.

Starting from recent human factors research on awareness and from Neisser's (1976) cognitive model of how awareness is maintained, our WA framework is organized around three issues:

- what kinds of information people keep track of in shared workspaces;
- how people gather workspace awareness information; and
- how people use workspace awareness information in collaboration.

These three areas inform three problems faced by groupware designers setting out to support awareness: what information to gather and distribute, how to present the information to the group, and when the information will be most useful. The framework provides designers with a structure to organize thinking about awareness support, a vocabulary for analysing collaborative activity and for comparing solutions, and a set of starting points for more specific design work. We do not give prescriptive rules and guidelines, however, since each groupware application will have to operate within particular awareness requirements dictated by the task and the group situation.

The framework was developed iteratively over several years (e.g. see Gutwin and Greenberg, 1996; Gutwin, Greenberg, and Roseman, 1996; Gutwin and Greenberg, 1998a) and is derived from a variety of sources:

- observations and insights of other groupware developers on issues concerning awareness (e.g. Stefik et al., 1987a; Tang, 1991; Beaudouin-Lafon and Karsenty, 1992; Dourish and Bellotti, 1992; Dix et al., 1993);
- theories developed by psychologists, linguists, ethnographers and human factors researchers on awareness (e.g., Clark, 1996; Brennan, 1990; Heath and Luff, 1995; Endsley, 1995);
- our own observational studies of face to face groups performing tasks over shared work surfaces (see Appendix A);
- our own iterative development and testing of many awareness widgets and displays, where we analyzed reasons for success and failure (e.g. Gutwin and Greenberg, 1996b, 1998a).

In this article, we explore workspace awareness and detail the three parts of the conceptual framework. To begin, we outline the concepts that underlie and bound the research, such as real-time distributed groupware, shared workspaces, and workspace awareness. Next, we give more detail on why awareness is a problem in groupware, and on the difficulty of supporting workspace awareness in a distributed computational setting. Third, we discuss human factors research into what awareness is and how it works, research that underlies the conceptual framework. We then introduce the three-part framework itself.

2. Setting the scene

There are bounds on the collaborative situations that we consider in this research. Our boundaries involve the kinds of groups we are trying to support, the workspace environment where collaboration takes place, the kinds of tasks that groups will undertake, and the kinds of groupware that will be used.

Systems: Real-time distributed groupware. Real-time distributed groupware systems allow people to work or play together at the same time, but from different places (e.g. Ellis et al., 1991). Although many kinds of group activity can be supported with real-time distributed groupware, we are particularly interested in applications that provide a shared workspace.

Environment: Shared workspaces. Many real-time groupware systems provide a bounded space where people can see and manipulate artifacts related to their activities. We concentrate on flat, medium-sized surfaces upon which objects can be placed and manipulated, and around which a small group of people can collaborate. In these spaces, the focus of the activity is on the visible and manipulable objects through which the task is carried out. The combination of physical space and artifacts makes a shared workspace an external representation of the joint activity (Clark, 1996; Norman, 1993; Hutchins, 1990).

Tasks: Generation and execution. Primary task types in shared workspaces are *generation* and *execution* activities (McGrath, 1984). In particular, these tasks tend to involve creation of new artifacts, navigation through a space of objects, or performance of physical manipulation on existing artifacts. Examples include activities such as construction (page layout, diagram assembly), organization (arranging, ordering, or sorting artifacts), design (drawing, generating an outline), or exploration (finding certain types of artifacts in the space). Other types of tasks (e.g. decision-making) also involve workspace awareness, but as these types involve less interaction with the artifacts, we do not consider them as primary for the framework.

Groups: Small groups and mixed-focus collaboration. Small groups of between two and five people primarily carry out tasks in these medium-sized workspaces. These groups often engage in mixed-focus collaboration, where people shift frequently between individual and shared activities during a work session (e.g. Dourish and Bellotti, 1992; Salvador et al., 1995). Although larger groups may also engage in tasks that require workspace awareness, it is less common for large groups to work synchronously over a shared workspace (because of space limitations), and we take small group activity as our primary focus.

Within these boundaries, a rich variety of small-group collaboration is possible. For example, typical examples might include two people organizing slides on a light table, a research group generating ideas on a whiteboard, or the managers of a project planning a task timeline. These and all the other group activities within our boundaries share a common problem when they take place in a groupware setting: it is difficult to maintain awareness of others in the workspace.

3. The awareness problem in groupware workspaces

In a face-to-face workspace, awareness of one another is relatively easy to maintain, and the mechanics of collaboration are natural, spontaneous, and unforced. Unfortunately, workspace awareness is much harder to maintain in groupware workspaces than in face-to-face environments, and it is often difficult or impossible to determine who else is in the workspace, where they are working, and what they are doing.

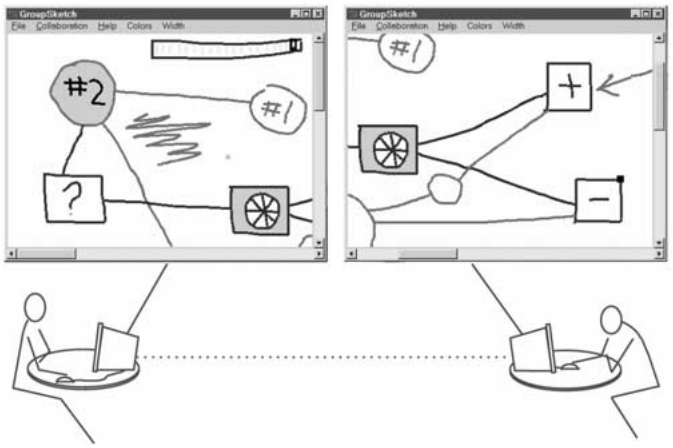

Figure 1. GroupSketchpad, a relaxed-WYSIWIS shared whiteboard.

There are three main reasons why this is so. First, the input and output devices used in groupware systems generate only a fraction of the perceptual information that is available in a face-to-face workspace. Second, a user's interaction with a computational workspace generates much less information than actions in a physical workspace. Third, groupware systems often do not present even the limited awareness information that is available to the system.

As an example, consider a basic shared whiteboard such as the GroupSketchpad system from the GroupKit toolkit (Roseman and Greenberg, 1995), seen in Figure 1. As each person draws, their actions are communicated to the other machine, so both participants' workspaces contain the same objects. At this moment in their task, the participants have scrolled their viewports to different parts of the workspace, and only a portion of their views overlap.

Systems like this one show almost none of the awareness information that would be available to a group working with a physical whiteboard. People's hands and bodies are reduced to simple telepointers, there is no sound, and only a small piece of the entire drawing can be seen at one time. In this situation, it will be difficult or impossible for the two participants to discuss particular objects, provide timely assistance, monitor the other person's activities, or anticipate their actions. In short, lack of information about others means that many of the little things that contribute to smooth and natural collaboration will be missing from the interaction.

In relaxed-WYSIWIS systems like this one, the awareness problem is particularly severe. When different people can scroll to different parts of the workspace (as in mixed-focus collaboration), they still need to maintain awareness of others; however, any information about where the other person is working or what they are doing can only be gathered through verbal communication. Once a person loses track of their partner, collaborating with them in real time becomes much more difficult.

How can groupware designers address the awareness problem? Part of the solution is to provide people with more information about their collaborators. As it is infeasible to replicate the detail and size of real-world workspaces, however, designers must carefully determine what information is most important, and how it can be put to best advantage in the system. The framework of workspace awareness is intended to provide designers with assistance in making these decisions. The first step involves setting out more precisely what workspace awareness is, and the process by which people manage to maintain it.

4. Awareness

In this section, we outline characteristics of awareness that are relevant to group work, describe prior research in awareness, describe the concept of workspace awareness in more detail, and set out a model of how awareness is maintained.

4.1. CHARACTERISTICS OF AWARENESS

Previous researchers have defined awareness as knowledge created through interaction between an agent and its environment – in simple terms, "knowing what is going on" (Endsley, 1995, p. 36). This conception of awareness involves states of knowledge as well as dynamic processes of perception and action. Four basic characteristics run through prior work on awareness (e.g. Adams et al., 1995; Norman, 1993; Endsley, 1995).

1. Awareness is knowledge about the state of an environment bounded in time and space.
2. Environments change over time, so awareness is knowledge that must be maintained and kept up to date.
3. People interact with and explore the environment, and the maintenance of awareness is accomplished through this interaction.
4. Awareness is a secondary goal in the task – that is, the overall goal is not simply to maintain awareness but to complete some task in the environment.

Everyone has experienced this kind of awareness; at its most basic, it is what allows us to walk around without bumping into things. As situations and environments become more complex, however, information demands sometimes outstrip our ability to attend, and awareness becomes more noticeable. In these contexts, previous researchers have explored what they call *situation awareness*, a concept that underlies the idea of workspace awareness in groupware.

4.2. SITUATION AWARENESS (SA)

Research into awareness as we describe it above originated in the study of military aviation, where pilots interact with highly dynamic, information-rich environments. In recent years, researchers have expanded their focus to other environments where situation awareness plays a major role, such as commercial aviation (Sarter and Woods, 1995), air traffic control (Smith and Hancock, 1995), and anesthesiology (Gaba and Howard, 1995). These environments all share the characteristics of "dynamism, complexity, high information load, variable workload, and risk" (Gaba and Howard, 1995).

The human factors community has not settled on a single definition of situation awareness, but most researchers include aspects of product (i.e. knowledge that an actor can make use of), and process (i.e. how that knowledge is created through interaction with the environment). A good general definition of SA is as "the up-to-the minute cognizance required to operate or maintain a system" (Adams et al., 1995, p. 85). Endsley (1995) focuses more on the process, proposing a three stage definition:

> Level 1: *perception of relevant elements of the environment*. An actor must first be able to gather perceptual information from the environment, and be able to selectively attend to those elements that are most relevant for the task at hand.
>
> Level 2: *comprehension of those elements*. An actor must be able to integrate the incoming perceptual information with existing knowledge, and make sense of the information in light of the current situation.
>
> Level 3: *prediction of the states of those elements in the near future*. To perform well in a situation, an actor must also be able to anticipate changes to the environment and be able to predict how incoming information will change.

The characteristics of awareness as introduced above also apply to workspace awareness: it is knowledge of a dynamic environment, it is maintained through perceptual information gathered from the environment, and it is peripheral to the primary group activity. We view workspace awareness as a specialization of situation awareness, one that is tied to the specific setting of the shared workspace.

4.3. WORKSPACE AWARENESS

We define workspace awareness as the up-to-the-moment understanding of another person's interaction with the shared workspace. This definition bounds the concept in two ways. First, workspace awareness is awareness of people and how they interact with the workspace, rather than just awareness of the workspace itself. Second, workspace awareness is limited to events happening in the workspace – inside the temporal and physical bounds of the task that the group is carrying out. This means that workspace awareness differs from informal awareness of who is around and available for collaboration, and from awareness of cues and turns in verbal conversation, both of which have been studied previously in CSCW (e.g. Borning and Travers, 1991; Dourish and Bly, 1992; Greenberg, 1996) and in linguistics (e.g. Clark, 1996; Goodwin, 1981).

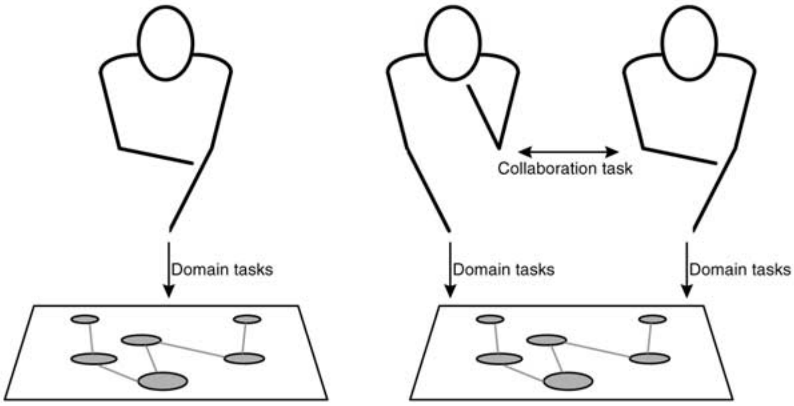

Figure 2. Domain and collaboration tasks.

The shared workspace setting makes workspace awareness a specialized kind of situation awareness. When someone works alone in a workspace, their activities and their SA involve only the workspace and the domain task (see Figure 2). In a collaborative situation, however, people must undertake another task, that of collaboration, and therefore their situation awareness must involve both the domain and the collaboration.

A second apparent difference between workspace awareness and situation awareness is that collaborating in most shared workspaces often does not involve high information load or extreme dynamism.[1] That is, it is not generally difficult to maintain workspace awareness in the real world: sorting slides on a table does not seem very similar to air combat in a jet fighter. However, the two types of situations do share an important characteristic, : that people are unable to gather the information that they need from the environment. In the jet aircraft, the information load exceeds the pilot's ability to take it all in. In the slide-sorting task, although the participants' perception would normally be perfectly adequate, a groupware system has artificially reduced their abilities to gather awareness information.

This means that the initial problems of maintaining WA in groupware revolve around obtaining useful information, rather than around what people make of the information. In the situations that SA research currently studies, problems can occur at any of Endsley's three levels: people can fail to gather important information from the environment, they may fail to understand what that gathered information means to the activity, or they may fail to predict what that information means for future events. In workspace situations, all of these can also occur, but we must focus first on the lack of information at the first and second levels. People's perception is artificially hampered by the technological constraints of a groupware system: information may be unavailable, or it may be presented in a form that makes the information unusable for maintaining up-to-the-moment awareness. The designer's task and our conceptual framework concentrate on these two levels: on

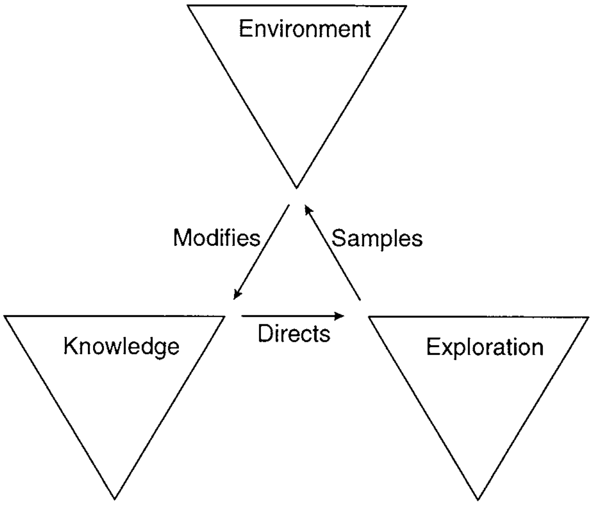

Figure 3. The perception-action cycle (Neisser, 1976).

determining what information to present, and on presenting that information so that people can maintain awareness easily and naturally.

4.4. MAINTAINING AWARENESS

Understanding how people maintain awareness is crucial if we are to design systems that support workspace awareness. Adams et al. (1995) suggest a cognitive model that shows how awareness is maintained in dynamic environments, a model that also draws together the process and product aspects of different definitions of SA. The model is Neisser's (1976) perception-action cycle, a "cognitive framework for the interdependence of memory, perception, and action" (Adams et al., 1995, p. 88). Neisser's model, shown in Figure 3, captures the interaction between the agent and the environment, and incorporates relationships between a person's knowledge and their information-gathering activity. It differs from linear models of information processing by recognizing that perception is influenced and directed by existing knowledge.

Awareness of an environment is created and sustained through the perception-action cycle. When a person enters an environment to do a particular task, they bring with them a general understanding of the situation and a basic idea of what to look for. The information that they then pick up from the environment can be interpreted in light of existing knowledge to help the person determine the current state of the environment – that is, what is happening – and also help them to predict what will happen next. These expectations lead to a further refinement in perceptual sensitivity, as when the expectation of seeing another aircraft sensitizes a pilot to subtle variations in the visual field (Adams et al., 1995, p. 89). The perception-action cycle combines both product and process aspects of awareness. Product is captured by the active knowledge created by previous cycles, and process is captured by the movement around the cycle.

To summarize thus far, Neisser's cycle and the research into situation awareness provide us with a foundation for a conceptual framework of workspace awareness. We have established that workspace awareness is a specialization of SA, where the 'situation' is well-defined – others' interactions with a shared workspace. Workspace awareness is maintained through a perception-action cycle, in which awareness knowledge both directs, and is updated by, perceptual exploration of the workspace environment. Finally, the initial problem in maintaining workspace awareness in distributed settings is that groupware technology limits what people can perceive of others in the workspace, hindering their ability to gather WA information from the environment.

We now turn to the conceptual framework itself. Part one involves the types of information that make up WA, Part two involves the mechanisms people use to gather WA information, and Part three involves the ways that people use WA information in collaboration. The contents of the framework come from existing research in CSCW, HCI, and human factors, and from our own observations of simple tabletop tasks and of real world group work in offices and control rooms.[2] After the three sections dealing with the framework, we discuss ways in which the knowledge of the framework can be used in the design of interface widgets and interaction techniques.

5. Framework part one: What information makes up workspace awareness?

Workspace awareness is made up of many kinds of knowledge, and the first part of the framework divides the concept into components. This part of the framework gives designers a basic idea of what information to capture and distribute in a groupware system. Even though a person can keep track of many things in a shared workspace, elements from a basic set make repeated appearances in research literature (e.g. Dourish and Bellotti, 1992; Sohlenkamp and Chwelos, 1994; McDaniel and Brinck, 1997). The basic set is the elements that answer "who, what, where, when, and how" questions. That is, when we work with others in a physical shared space, we know who we are working with, what they are doing, where they are working, when various events happen, and how those events occur. People keep track of these things in all kinds of collaborative work, and these are the kinds of information that should be considered first by designers.

Within these basic categories, we have identified specific elements of knowledge that make up the core of workspace awareness. Tables I and II show these elements and list the questions that each element can answer. Table I contains those elements that relate to the present, and Table II contains those that relate to the past. The elements are all commonsense things that deal with interactions between a person and the environment. Awareness of presence and identity is simply the

Table I. Elements of workspace awareness relating to the present

Category	Element	Specific questions
Who	Presence	Is anyone in the workspace?
	Identity	Who is participating? Who is that?
	Authorship	Who is doing that?
What	Action	What are they doing?
	Intention	What goal is that action part of?
	Artifact	What object are they working on?
Where	Location	Where are they working?
	Gaze	Where are they looking?
	View	Where can they see?
	Reach	Where can they reach?

knowledge that there are others in the workspace and who they are, and authorship involves the mapping between an action and the person carrying it out. Awareness of actions and intentions is the understanding of what another person is doing, either in detail or at a general level. Awareness of artifact means knowledge about what object a person is working on. Location, gaze, and view relate to where the person is working, where they are looking, and what they can see. Awareness of reach involves understanding the area of the workspace where a person can change things, since sometimes a person's reach can exceed their view.

Awareness of the past involves several additional elements. Action and artifact history concern the details of events that have already occurred, and event history concerns the timing of when things happened. The remaining three elements deal with the historical side of presence, location, and action. We do not include elements relating to the future in our framework, because designers are unlikely to be able to support maintenance of those elements. This is because past and present information can be determined from raw perceptual information, whereas belief about the future involves inference, extrapolation, and prediction.

Workspace awareness knowledge will be made up of these elements in some combination, and participants in a face-to-face group activity will generally know the basic elements (consciously or unconsciously). This does not mean, however, that the designer should support all elements equally in the interface. Two factors are critical in determining how the designer should treat each element. First, the degree of interaction between the participants in the activity indicates how specific or general the information in the interface should be. Second, the dynamism of the element – how often the information changes – indicates how often the interface will need to be updated. In some situations, certain elements never change, and so do not require explicit support in the interface. For example, if the participants in

Table II. Elements of workspace awareness relating to the past

Category	Element	Specific questions
How	Action history	How did that operation happen?
	Artifact history	How did this artifact come to be in this state?
When	Event history	When did that event happen?
Who (past)	Presence history	Who was here, and when?
Where (past)	Location history	Where has a person been?
What (past)	Action history	What has a person been doing?

an activity are always assigned to particular areas of the workspace, there is little need for the system to gather and distribute location information.

Although there will also be additional kinds of information specific to the task or the work setting, these basic elements provide a high-level organization of workspace awareness. The elements are a starting point for thinking about the awareness requirements of particular task situations, and provide a vocabulary for describing and comparing awareness support in groupware applications.

6. Framework part two: How is workspace awareness information gathered?

The groupware designer must attempt to present awareness information in ways that make the maintenance of workspace awareness simple and straightforward. We believe that this will be easier if people can gather information in familiar ways, even though the actual interface devices in a groupware system may not be familiar. This means understanding the mechanisms people use to gather workspace awareness information from the workspace environment – basically, how people find the answers to the who, what, where, when, and how questions listed in Tables I and II. In this section, we outline some of the ways that people find those answers.

Prior research suggests three main sources of workspace awareness information, and three corresponding mechanisms that people use to gather it (Segal, 1994; Norman, 1993; Dix et al., 1993; Hutchins, 1990). People obtain information that is produced by people's bodies in the workspace, from workspace artifacts, and from conversations and gestures. The mechanisms that they use to gather it are called consequential communication, feedthrough, and intentional communication.

6.1. BODIES AND CONSEQUENTIAL COMMUNICATION

The first information source is the other person's body in the workspace (e.g. Segal, 1994; Norman, 1993; Benford et al., 1995). Since most things that people do in a workspace are done through some bodily action, the position, posture, and movement of heads, arms, eyes, and hands provide a wealth of information about what's going on. Therefore, watching other people work is a primary mechanism for gathering awareness information: "whenever activity is visible, it becomes an essential part of the flow of information fundamental for creating and sustaining teamwork" (Segal, 1994, p. 24). Although people also contribute to the auditory environment, much of the perception of a body in a workspace is visual. In all of the tabletop tasks that we observed, for example, participants would regularly turn their heads to watch their partners work.

The mechanism of seeing and hearing other people active in the workspace is called *consequential communication*: information transfer that emerges as a consequence of a person's activity within an environment (Segal, 1994). This kind of bodily communication, however, is not intentional in the way that explicit gestures are: the producer of the information does not intentionally undertake actions to inform the other person, and the perceiver merely picks up what is available. Nevertheless, consequential communication provides a great deal of information. In a study of piloting teams, Segal reports that:

> [Pilots] spent most of their time – over 60% – looking across at their [partner's] display while it was being manipulated. This suggests that beyond the information provided by the display itself, these pilots were specifically looking for information provided by the dynamic interaction between their crewmembers and that display. (p. 24)

This study also suggests that movement is particularly important in consequential communication, since our attention is naturally drawn to motion. Norman (1993) gives an example, when he relates the value of "obvious actions" in aircraft cockpits:

> When the captain reaches across the cockpit over to the first officer's side and lowers the landing-gear lever, the motion is obvious: the first officer can see it even without paying conscious attention. The motion not only controls the landing gear, but just as important, it acts as a natural communication between the two pilots, letting both know the action has been done. (p. 142)

6.2. ARTIFACTS AND FEEDTHROUGH

The artifacts in the workspace are a second source of awareness information (e.g. Dix et al., 1993; Gaver, 1991). Artifacts provide several sorts of visual information: they are physical objects, they form spatial relationships to other objects, they contain visual symbols like words, pictures, and numbers, and their states are often shown in their physical representation. Artifacts also contribute to the acoustic environment, making characteristic sounds when they are created, destroyed, moved, stacked, divided, or manipulated in other ways (Gaver, 1991). Tools in particular have signature sounds, such as the snip of scissors or the scratch of a pencil. By seeing or hearing the ways that an artifact changes, it is often possible to determine what is being done to it.

This mechanism is *feedthrough* (Dix et al., 1993): when artifacts are manipulated, they give off information, and what would normally be feedback to the person performing the action can also inform others who are watching. When both the artifact and the actor can be seen, feedthrough is coupled with consequential communication; at other times, there may be a spatial or temporal separation between the artifact and the actor, leaving feedthrough as the only vehicle for information. For example, in our observations of the Calgary air traffic control centre (Appendix 1), the departures controller cannot monitor all of the arrival controller's actions, but can see the status of arriving aircraft on their display change from "approaching" to "landed." When they see this change in the artifact, they can also infer the activities of the arrivals controller.

6.3. CONVERSATION, GESTURE, AND INTENTIONAL COMMUNICATION

A third source of information that is ubiquitous in collaboration is conversation and gesture, and their mechanism is intentional communication (e.g. Clark, 1996; Heath and Luff, 1995; Birdwhistell, 1952). Verbal conversations are the prevalent form of communication in most groups, and there are three ways in which awareness information can be picked up from verbal exchanges. First, people may explicitly talk about awareness elements with their partners, and simply state where they are working and what they are doing. Our observations of shared-workspace tasks suggest that these direct discussions happen primarily when someone asks a specific question such as "what are you doing?" or when the group is planning or replanning the division of labour.

Second, people can gather awareness information by overhearing others' conversations. Although a conversation between two people may not explicitly include a third person, it is understood that the exchange is public information that others can pick up. For example, navigation teams on navy ships talk on an open circuit, which means that everyone can hear each others' conversations. Hutchins (1990) details how members of the team listen in on these conversations, either to monitor the actions of a junior member, or to learn from more experienced members. For this reason, voice loops – audio channels that allow directed and overheard communication among spatially separate sub-groups of people – have evolved as standard practice in mission control domains such as air traffic management, aircraft carrier operations, and space mission control (Watts et al., 1996).

Third, people can pick up others' verbal shadowing, the running commentary that people commonly produce alongside their actions, spoken to no one in particular. Heath and Luff (1995) have observed this behaviour, which they call "outlouds." They note that although these "outlouds . . . might be thought relatively incursive, potentially interrupting activities being undertaken by [others] in the room, [they are] perhaps less obtrusive than actually informing particular persons" (p. 157).

The style of verbal shadowing can be explicit or highly indirect. In the our observations of a newspaper-layout task (Appendix 1), participants regularly stated exactly what they were doing, saying things like "I'm going to cut this article," or "I'll move this over here." In other work situations like the London Underground (Heath and Luff, 1992), controllers talk more to themselves and use oblique references like curses or song phrases, but are nevertheless able to convey information to others in the control room.

Gestures and other visual actions can also be used to carry out intentional communication. These differ from consequential communication in that they are intended, and are often used alongside verbal productions. Short, Williams, and Christie (1976) note two forms of visual communication used to convey task information. First is illustration, where speech is illustrated, acted out, or emphasized. For example, people often illustrate distances by showing a gap between fingers or hands. The second form is the emblem, where words are replaced by actions: for example, a nod or shake of the head indicates 'yes' or 'no' (p. 45). These types of gestures have also been observed in CSCW studies (e.g. Ishii and Kobayashi, 1992; Tang, 1991).

7. Framework part three: How is workspace awareness used in collaboration?

A groupware designer needs to know the situations and activities where workspace awareness will be used, to better analyze collaborative tasks and to better determine when groupware support is called for. Workspace awareness is used for many things in collaboration. Awareness can reduce effort, increase efficiency, and reduce errors for the activities of collaboration. This section describes five types of activity, reported in the literature and as seen in our observational studies (Appendix 1), that are aided by workspace awareness (e.g. Tatar et al., 1991; Clark, 1996; Tang, 1991; Salvador et al., 1996). These provide a basic set of collaborative activities that designers can look for as they analyse work situations. The five activities are: management of coupling, simplification of verbal communication, coordination, anticipation, and assistance.

7.1. MANAGEMENT OF COUPLING

Several researchers have recognized that when people collaborate, they shift back and forth between individual and shared work, and that awareness of others is important for managing these transitions. For example, Dourish and Bellotti (1992) observed that people involved in a shared editing task "continually moved between concurrent, but more or less independent, work ... to very tightly focused group consideration of single items. These movements were opportunistic and unpredictable, relying on awareness of the state of the rest of the group" (p. 111). Gaver (1991) adds that "people shift from working alone to working together, even when joined on a shared task. Building systems that support these transitions is important, if difficult" (p. 295).

Salvador et al. (1996) call the degree to which people are working together *coupling*.[3] In general terms, coupling is the amount of work that one person can do before they require discussion, instruction, action, information, or consultation with another person. Some of the reasons that people may move from loose to tight coupling are that they see an opportunity to collaborate, that they need to come together to discuss or decide something, that they need to plan their next activity, or that they have reached a stage of their task that requires another person's involvement. A sense of awareness about what another person is doing makes each of these situations more feasible, by allowing people to recognize when tighter coupling could be appropriate. Heath and Luff (1995) give the example of a financial dealing office where dealers manage coupling by carefully monitoring their colleagues' activities:

> . . . though dealers may be engaged in an individual task, they remain sensitive to the conduct of colleagues and the possibility of collaboration . . . 'Peripheral' monitoring or participation is an essential feature of both individual and collaborative work within these environments. (p. 156)
>
> So, for example, it is not unusual in the dealing room for individuals to time, with precision, an utterance which engenders collaboration, so that it coincides with a colleague finishing writing out a ticket or swallowing a mouthful of lunch. By monitoring the course of action in this way and by prospectively identifying its upcoming boundaries, individuals can successfully initiate collaboration so that it does not interrupt an activity in which a colleague is engaged. (p. 152)

Although these examples deal with a wider environment than a flat shared workspace, the idea is the same – that people keep track of others' activities when they are working in a loosely coupled manner, for the express purpose of determining appropriate times to initiate closer coupling. Without workspace awareness information, people will miss opportunities to collaborate, and will often interrupt the other person inappropriately.

7.2. SIMPLIFICATION OF COMMUNICATION

Workspace awareness lets people use the workspace and the artifacts in it to simplify their verbal communication and make interaction more efficient. When discussion involves task artifacts, the workspace can be used as an external representation of the task that allows efficient nonverbal communication (Hutchins, 1990; Clark, 1996). That is, the artifacts act as conversational props (Brinck and Gomez, 1992) that let people mix verbal and visual communication.

Workspace awareness is important because interpreting the visual signals depends on knowledge of where in the workspace they occur, what objects they relate to, and what the sender is doing. The nonverbal actions simplify dialogue by reducing the length and complexity of utterances. Four kinds of these communicative actions have been previously observed in studies of face-to-face collaboration: deictic reference, demonstration, manifesting actions, and visual evidence.

Deictic references. Referential communication involves composing a message that will allow another person to choose a thing from a set of objects (Krauss and Fussell, 1990). When transcripts of a collaborative activity are reviewed, however, many of these messages are almost unintelligible without knowledge of what was going on in the workspace at the time. For example, consider a fragment from a transcript of a puzzle task (Appendix 1):

> A: How about this thing ... <points to diagram> ... the tail? The only thing that can be is ...
> B: <holds up a piece> No, not that.
> B: <holds up another piece> This thing? It could be that thing <points to diagram> ...
> A: Yeah, could be that thing ...
> A: <holds up another piece> Could be that thing ...

The verbal communication does not convey what people are pointing at or indicating when they say "this," "that," "here," or "there." The practice of pointing or gesturing to indicate a noun used in conversation is called deictic reference, and is ubiquitous in shared workspaces (e.g. Segal, 1995; Tatar et al., 1991; Tang, 1991). For example, in a flight simulation experiment with two pilots, Segal (1994) found that many of the transcribed utterances could not be interpreted without reference to a videotape of the cockpit displays. Deictic reference is a crucial part of the way we communicate in a shared space. As Seely Brown and colleagues (1989) state:

> Perhaps the best way to discover the importance and efficiency of indexical terms and their embedding context is to imagine discourse without them. Authors of a collaborative work will recognize the problem if they have ever discussed the paper over the phone. "What you say here" is not a very useful remark. Here in this setting needs an elaborate description (such as "page 3, second full paragraph, fifth sentence, beginning ...") and can often lead to conversations at cross purposes. The problem gets harder in conferences calls when *you* becomes as ambiguous as *here* ... The contents of a shared environment make a central contribution to conversation. (p. 36)

Demonstrations. In addition to gestures used to illustrate conversation (e.g. Clark, 1996), people use gestures in workspaces to demonstrate actions or the behaviour of artifacts. As Tang (1989) states, "ideas are often enacted gesturally in order to express them effectively to others, especially if they involve a dynamic sequence of actions" (p. 76). Common demonstrations include tracing a path in the workspace with a finger or illustrating how an artifact operates. For example, Tang (1989) observed a participant in a design session turning her hand over to demonstrate how a card would flip back and forth (p. 76).

Manifesting actions. Actions in the workspace can also replace verbal communication entirely. When people replace an explicit verbal utterance with an action in the shared workspace, they are performing a manifesting action (Clark, 1996). Placing my groceries on the counter tells the clerk "I wish to purchase these items" without me having to say so. However, manifesting actions must be carried out carefully to prevent them being mistaken as ordinary actions: the action must be stylized, exaggerated, or conspicuous enough that the "listener" will not mistake it (Clark, p. 169). Therefore, I must place my groceries on the counter in such a way that the clerk realizes I am making a purchase request and not just resting my arms.

Visual evidence. When people converse, they require evidence that their utterances have been understood. In verbal communication, a common form of this evidence is back-channel feedback. In shared workspaces, however, visual actions can also provide evidence of understanding or misunderstanding. Clark (1996) provides an example from an everyday setting, where Ben is getting Charlotte to center a candlestick in a display:

> Ben: Okay, now, push it farther – farther – a little more – right there. Good. (p. 326)

Charlotte moves the candlestick after each of Ben's utterances, providing visual evidence that she has understood his instructions and has carried them out to the best of her interpretation. This kind of evidence can be used whenever people carry out joint projects involving the artifacts in a shared workspace.

The success of these four kinds of nonverbal communication depends on two aspects of workspace awareness. First, and most obvious, the communicative action must be perceived before it can be understood; if the action is invisible, it is impossible to interpret. For example, if I cannot see that you are pointing, or what you are pointing at, I cannot ground your deictic reference. Second, the receiver needs to have an idea of the workspace context in which the visible actions occur, since the meaning of the action may be ambiguous without certain information. For example, if there are several green blocks in the workspace, seeing only that you are pointing to a green block may not be enough information to correctly ground the reference. Or, if you hand me an object in a way that appears to be a request, I may need knowledge of your current activities before I can determine your expectations.

The important thing here is that the sender has to understand what the receiver can see in order to construct useful non-verbal communications. This means that workspace awareness is part of conversational common ground in a shared workspace. Common ground is the mutual knowledge that people take advantage of to increase their communicative efficiency (Clark, 1996). The principle of least collaborative effort suggests that people expend only the minimum effort in composing an utterance that they believe is necessary for their message to get across to the hearer (Clark and Brennan, 1991). If they can exploit common ground, they can reduce the work that goes into communication. Without common ground, people must do more work to compose exact, complete, and literal utterances. Workspace awareness as common ground means that people can further simplify their communication even without visual productions. They do this by assuming that the other person's awareness will help them correctly interpret highly under-specified utterances. For example, if I believe that you know where I am and what I'm working on, I can say something like "do you think that it will fit?" instead of "do you think that the smaller of the two arches will fit at the top of the tower that's at the right side of the picture?," a much more complicated and exact utterance.

7.3. COORDINATION OF ACTIONS

Coordinating actions in a collaborative activity means making them happen in the right order, at the right time, and generally, making them meet the constraints of the task. Coordination is necessary at several levels of granularity, from small hand movements to large-scale divisions of labour. In addition, certain kinds of joint activities require the concerted action of two people.

Coordination can be accomplished in two ways in a shared workspace: "one is by explicit communication about how the work is to be performed ... another is less explicit, mediated by the shared material used in the work process" (Robinson, 1991, p. 42). This second, less explicit way uses workspace awareness. Awareness aids both fine and coarse-grained coordination, since it informs participants about the temporal and spatial boundaries of others' actions, and since it helps them fit the next action into the stream. Workspace awareness is particularly evident in continuous action where people are working with the same objects. For example, CSCW researchers have noted that concurrency locks are less important or even unnecessary when participants have adequate information about what objects others are currently using; when the awareness information is available, people can use social protocols to coordinate access to objects (Greenberg and Marwood, 1994). Another example is the way that people manage to avoid bumping into each others' hands in a confined space. Tang (1989) saw this kind of coordination in design activity:

> the physical closeness among the participants ... allows a peripheral awareness of the other participants and their actions, as evidenced in the many 'coordinated dances' observed among the hands of the collaborators in the workspace. There were many episodes of intricate coordinated hand motions, such as getting out of the way of an approaching hand or avoiding collisions with other hands. These coordinated actions indicate a keen peripheral awareness of the other participants ... (p. 95)

Workspace awareness is also useful in the coordination and division of labour and in the planning and replanning of the activity. As the task progresses, groups regularly reorganize what each person will do next. These decisions depend in part on elements of workspace awareness – what the other participants have done, what they are still going to do, and what is left to do in the task. Based on another person's activities, I may decide to begin a complementary task, to assist them with their job, or to move to a different area of the workspace to avoid a conflict. It may be more efficient to have the members of the group do work that is near in proximity or in nature to what they are currently doing or have done in the past. Knowing activities and locations, therefore, can help in determining who should do what task next. For example, in one of the puzzle tasks we observed, the structure was symmetric, and people would regularly choose to do the symmetrical complement to their partner's action immediately after the partner had completed it.

7.4. ANTICIPATION

Another common behaviour in collaboration is anticipation, where people take action based on their expectations or predictions of what others will do in the future (Tang, 1989; Hall, 1959). People anticipate others in several ways. They can prepare for their next action in a concerted activity, they can avoid conflicts, or they can provide materials, resources, or tools before they are needed.

Anticipation is based on prediction, and people can predict workspace actions at both small and large time scales. First, people can predict some types of events by extrapolating forward from the immediate past. For example, if I see someone reaching towards a pair of scissors, I might predict that they are going to grab them. This prediction allows me to anticipate the event: I might pick up the scissors and pass them to the reacher, I might replan my own movements to avoid a collision, or I might reach for them myself to grab them before the other person gets them. This kind of anticipation is integral to the fine-grained coordination discussed above. Although ordinary, anticipation is difficult without workspace awareness – in the scissors example, without up-to-the-moment knowledge of where the other person's hand is moving, and of their location in relation to the scissors. In addition to this information, my prediction could have also taken into account other workspace awareness knowledge, such as their current activities and whether they were doing something that required scissors.

When prediction happens at a larger time scale, people learn which elements of situations and tasks are repeated and invariant. People are experts at recognizing patterns in events, and quickly begin to predict what will come next in situations that they have been in before. Workspace awareness is again important, but this time provides people with the information they need to determine whether others' behaviour match the patterns that they have learned. For example, in air traffic control, regional controllers hand flights off to the Calgary controllers when they come within 35 miles of the city. The transfer is done entirely through the shared workspace. The regional controller tags the aircraft's icon, and the Calgary controller must acknowledge the handoff by pressing a command key while their trackball cursor is overtop the aircraft. This handoff procedure is done for each flight, so the controllers are extremely familiar with it. Accordingly, the Calgary controllers anticipate the handoff, based on the information available in the workspace and their experience of what the regional controllers do in this situation. When a Calgary controller sees an incoming aircraft appear on the edge of the radar screen, they will often move their cursor over the aircraft, waiting for the handoff indicator from the regional controller to appear.

7.5. ASSISTANCE

Assisting others with their local tasks is an integral part of collaboration, and one that also benefits from workspace awareness. Assistance was extremely common in the tasks we observed, but not usually explicit. Often, one participant would make some indirect statement indicating that they wanted assistance, and their partner would look over and leave their tasks for a few moments to help out, and then return to what they were doing. For example, one participant was unable to find a piece that she needed for the cathedral puzzle task (Appendix 1), and so indirectly asked her partner for assistance:

A: Do you have another one of these guys here? <holds up piece>
B: They're, uh, red?
A: Yeah.
B: Yep, there's one . . . <hands piece to A>

People were also able to provide assistance without a prior request. In the same task, one participant simply reached over and placed a piece for the other:

A: Oh, and I found another triangle thing for you . . . here. <places piece>

Awareness in these situations is useful because it helps people determine what assistance is required and what is appropriate. In order to assist someone with their tasks, you need to know what they are doing, what their goals are, what stage they are at in their tasks, and the state of their work area. In the second example above, the helper knew what their partner had already completed; in particular, that she had not yet found all of the needed "triangle things," and that adding one to the cathedral would be beneficial.

This section has outlined five kinds of collaborative activity that are aided by greater workspace awareness; these are summarized in Table III. Groupware designers can use this part of the framework in two ways: first, as an analysis tool to help them determine the degree of awareness support that is needed for a particular work situation (since different collaborative situations involve these activities in different amounts); and second, as a guide to determining where in the interface that awareness support should be provided. In the next section, we discuss some of the ways in which awareness support can be provided in the interface.

Table III. Summary of the activities in which workspace awareness is used

Activity	Benefit of workspace awareness
Management of coupling	Assists people in noticing and managing transitions between individual and shared work.
Simplification of communication	Allows people to the use of the workspace and artifacts as conversational props, including mechanisms of deixis, demonstrations, and visual evidence.
Coordination of action	Assists people in planning and executing low-level workspace actions to mesh seamlessly with others.
Anticipation	Allows people to predict others' actions and activity at several time scales.
Assistance	Assists people in understanding the context where help is to be provided.

8. Supporting examples: Applying the workspace awareness framework to interface design

The framework describes what the elements of workspace awareness are, what mechanisms are used to maintain it, and when it is useful in collaborative work situations. In this section, we look at ways designers can apply the knowledge of the framework to the design of groupware interfaces, and review a set of techniques that can be used to provide different elements of workspace awareness information. We give examples of how a designer can use the framework: to think about the representation and placement of awareness information within the interface; to analyze and categorize existing awareness techniques, displays and widgets; and to inform the design evolution of a particular awareness widget.

We caution that these are representative and illustrative examples, rather than as an exhaustive list of previous work. We do not attempt to explain the details of approaches or interface widgets. Also, the awareness displays and widgets used in these examples are oriented towards only one part of the process of maintaining awareness – making information available – so designers must also consider whether people interpret the information correctly, and whether their resulting actions are appropriate. Our examples are also biased towards our own experiences: many of the techniques presented arise from our work with the GroupKit groupware toolkit (Roseman and Greenberg, 1996).

8.1. ORGANIZING DISPLAY SPACE

Our first example suggests how a designer can think about the general representation and placement of how awareness information is presented within the interface. A designer faces basic questions of where and how to display workspace awareness

Table IV. Presentation and placement of awareness display techniques

		Placement	
		Situated	Separate
Presentation	Literal		
	Symbolic		

information in a groupware interface. We have determined two basic dimensions that provide boundaries for some of these questions. First, when considering where information will be displayed, the dimension of *placement* draws a basic distinction between information that is *situated* within the workspace and information that is presented *separate* from it. Situated placement implies that the information is displayed at the workspace location where it originated, and separate placement means displaying the information outside the workspace in a separate part of the interface. Second, the issue of how information will be displayed suggests the dimension of *presentation*: a display can be either be *literal* or *symbolic.* Literal presentation implies that the information is shown in the same form that it is gathered, and includes low-level movement and feedback. Symbolic presentations extract particular information from the original data stream and display it explicitly. These two dimensions combine to form the matrix shown in Table IV.

Of these divisions, the approach that holds perhaps the most promise for natural and effective awareness support is the situated-literal approach. Here, awareness information is integrated into the workspace's existing representation, and is shown in the same form that it was produced by another person. This approach is the closest approximation of how awareness information appears in the real world, and it is the only one that allows people to use their existing skills with the mechanisms of feedthrough, consequential communication, and gestural communication. In addition, situated and literal information best supports the three activities in the perception-action cycle that people use to maintain awareness: it is available in the environment but need not be attended to all the time; it provides low-level information that can be interpreted in light of other existing knowledge; and it allows further exploration or action to be taken in the same context in which the information was gathered. Situating awareness information, however, raises the possibility that people may not notice important events; furthermore, there is no guarantee with any awareness technique that people are going to interpret the information correctly or use it effectively.

Two critical design elements of the situated-literal approach are *embodiment* and *expressive artifacts*. Embodiments are visible representation of each person's body in the workspace – representations that have been used include telepointers (Hayne, Pendergast and Greenberg, 1993), view rectangles (Beaudoin-Lafon and Karsenty, 1992), avatars (Benford et al., 1995), and video images (Tang and Minneman, 1991; Ishii and Kobayashi, 1992). Depending upon its expressiveness, a workspace embodiment can provide information about who is in the workspace, where they are, and what they are doing, and can afford both consequential and gestural communication. The third mechanism, feedthrough, is provided by expressive artifacts – artifacts that maximize the amount of usable awareness information produced for the group. Although the design of specific artifacts cannot be predetermined, there are general strategies for designing and displaying common types of manipulations that increase expressiveness, such as action indicators or action animations (Gutwin and Greenberg, 1998).

One particular drawback to the situated-literal approach is that of visibility – when information is situated in the workspace, others have to be looking at the appropriate part of the workspace in order to see the information. This can be a severe problem in relaxed-WYSIWIS systems that allow people to scroll to entirely different parts of the workspace. A solution to this problem is to provide multiple views that offer visibility to awareness information in unseen parts of the workspace; for example, radar views (Smith et al., 1998) show information in the entire workspace. In the next section, we map techniques from both the situated-literal approach and other parts of the design space to the elements of workspace awareness.

8.2. TECHNIQUES, DISPLAYS, AND WIDGETS

Our second example shows how a designer can use the workspace awareness framework to analyze and categorize existing techniques, displays and widgets. Tables V, VI, and VII below use the elements of workspace awareness to organize a variety of awareness displays and techniques that have appeared in previous literature. We concentrate here on real-time aspects of workspace awareness – elements that answer the *who*, *what*, and *where* questions. The techniques are grouped according to what workspace awareness elements are supported; some displays appear several times since they support more than one element of awareness. This review is intended as an illustrative and representative list rather than an exhaustive one, and due to our familiarity with GroupKit, is slightly skewed towards solutions that have been built with that toolkit.

8.3. CASE STUDY – EVOLUTION OF A RADAR VIEW

As an example of how the knowledge in the framework can be used, we review the design evolution of a radar view built for the GroupKit toolkit. Radar views are secondary windows used with a detailed view of the shared workspace; they show miniatures of the artifacts in a shared workspace, and can also be used to show awareness information about the participants in the session. Our original radar view showed only the movement of workspace objects (Figure 4a). As we worked with the display in a newspaper-layout domain, it became apparent that several aspects of awareness were not well supported.

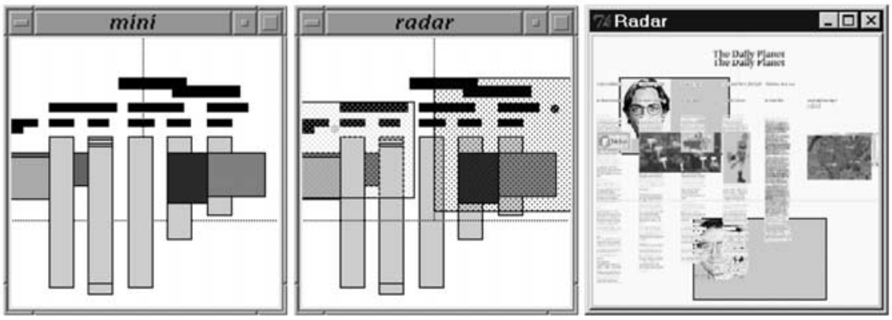

Figure 4. Three versions of the GroupKit radar view. Version 4a shows object movement only; 4b adds location information by showing each person's main view as a shaded rectangle; 4c adds photographs for participant identification.

Table V. Workspace awareness techniques for "who" questions

WA element (Who)	Example interface techniques
Presence (*Is anyone there?*)	• Participant list (e.g. Sohlenkamp and Chwelos, 1994). The most basic awareness display, the participant list shows who is currently logged in to the system (although several other types of awareness information can be added to this basic idea). Presence is shown by presence in the list. • Embodiment solutions (telepointers, view rectangles, avatars, video images). Since an embodiment is a representation of an actual person, presence is shown by the existence of the embodiment. In some cases, presence can also be heard if embodiments emit sound as they interact with the workspace (Gaver, 1991).
Identity (*Who is that?*)	• Participant list identifies participants with a name or picture. • Embodiments show identity through visual characteristics of the representation, such as colour (telepointers or view rectangles), shape and appearance (avatars), or actual images (video techniques).
Authorship (*Who is doing that?*)	• Creation colouring (e.g. Mitchell, 1996). When activities involve the creation of new artifacts, the objects (such as characters in a text window) can be coloured to indicate authorship. • Embodiment proximity. The proximity of a person's representation to an action is a strong authorship clue in direct-manipulation environments. • Authorship lines (e.g. Sohlenkamp and Chwelos, 1994). Lines drawn from actions or artifacts to a participant list to indicate authorship.

Table VI. Awareness techniques for “what” questions

WA element (What)	Example interface techniques
Action (*Is anything happening? What is she doing?*)	• Activity and change indicators (e.g. Ackerman and Starr, 1995). “Change meters” placed in the interface to indicate the occurrence or rate of activity or edits in the workspace. • Consequential communication through embodiment. People’s workspace representations convey both that actions are happening, and also what actions are occurring through characteristic motions. • Mode indicators. Representations of the mode in which each person is working. Modes can be shown separately (in a participant list) or can be situated. For example, telepointers can show each person’s mode in a drawing program (e.g. Greenberg and Bohnet, 1991). • Action indicators and animations. Actions that are hard to see can be made artificially more perceptible with visible indicators; actions that are instantaneous can be lengthened with animations (e.g. Gutwin and Greenberg, 1998). • Visibility of actions (Smith et al., 1998; Gutwin, Greenberg and Roseman, 1995). Separate views of the workspace provide visibility to actions that are in other parts of the workspace. Radar views show the entire workspace. Over-the-shoulder views show a miniature version of another person’s main view. Cursor’s-eye views show the area immediately around another person’s cursor in full detail. • Audible actions. Others’ actions can be represented with sound to show both existence and type of activity (e.g. Gaver, 1991).
Intention (*What is she going to do?*)	• Embodiment frame rate. Showing embodiments at a real-time frame rate allows observers to accurately predict movements and anticipate actions (e.g. Gutwin, 2000). • Marking artifacts. Explicit notification of future intentions by visibly marking workspace artifacts (e.g. Gutwin, Roseman and Greenberg, 1996).
Artifact (*What object is she using?*)	• Embodiment proximity. The proximity of embodiment to an artifact is a strong clue in direct-manipulation environments. • Artifact indicators. Artifacts that are currently being edited can be represented on a separate display such as a participant list. • Characteristic sounds. Different objects can produce different types of sounds, giving some indication of which artifact is in use (e.g. Gaver, 1991).

Table VII. Workspace awareness techniques for "where" questions

WA element (Where)	Example interface techniques
Location (*Where is she working?*)	• Embodiment techniques show location by the position of the person's representation. Outside the main workspace view, visibility techniques such as radar views are required. • Radar or gestalt views (Smith et al., 1998; Baecker et al., 1993). These show location using view rectangles and telepointers on a miniature of a two-dimensional workspace. • Multi-user scrollbars (Baecker et al., 1993). These show location using view bars in one-dimensional workspaces. • Distortion-oriented workspace representations. The visibility problem can also be addressed by always showing the entire workspace in the main view, and then using magnification techniques to show detail (e.g. Greenberg, Gutwin and Cockburn, 1995). • Sound distance. Activity sounds can indicate distance and location of activity by changes in volume and direction (e.g. Smith, 1999). • Location indicators. In structured environments (such as rooms-based systems), indications of location can be placed on a separate display such as a participant list (e.g. Roseman and Greenberg, 1996).
Gaze (*Where is she looking?*)	• Eye-contact video. Certain types of video embodiments show gaze direction accurately (Ishii and Kobayashi, 1992). • Embodiment position. The position of the control part of an embodiment (e.g. the telepointer or the hand of an avatar) is often a reasonable clue as to a person's gaze direction.
View (*What can she see?*)	• View rectangles. Explicit representations of another person's view show what they can see in detail (e.g. Beaudoin-Lafon and Karsenty, 1995). • Duplicate views. The over-the-shoulder view provides a miniature of another person's detail view (Gutwin, Greenberg and Roseman, 1995). • View slaving. Being able to temporarily switch to another person's view shows what they can see in full detail (Gutwin, Roseman and Greenberg, 1996).
Reach (*What can she manipulate*)	• View rectangles. Representations of a person's detail view indicate what a person can reach for detailed work; overviews show what can be reached for large-scale manipulation (often the entire workspace)

Table VIII. Awareness elements supported by features of the radar view

Feature added	Awareness elements	Mechanism
Object movement	Action, artifact	Feedthrough
View rectangles	Identity, location, view, reach	Consequential communication
Radar telepointers	Identity, location, action, intention	Consequential communication
Participant photos	Identity	Embodiment

In analysing the drawbacks of the device for the tasks being carried out, the WA framework was used as an analysis tool to help identify which elements of awareness should be better supported. We determined that more information about location, activity, and identity was required for some tasks. This led to two redesigns. To the version in Figure 4b, we added location information with shaded viewport rectangles and miniature telepointers, and to the version in Figure 4c, we added portraits for participant identification.

We also used the idea perception-action cycle to change the way that the radar works. The first two versions of the radar are display only, and we found that people were having difficulty acting on information that they gathered from the window (Gutwin, Roseman and Greenberg, 1996). Therefore, the third version of the radar was made into a fully interactive secondary workspace rather than a view-only display: people can interact with its objects, and moving the telepointer over it lets people gesture anywhere within it. Our evaluations confirm that users do find the later devices more useful for some kinds of collaborative tasks (e.g. Gutwin, Roseman, and Greenberg, 1996; Gutwin and Greenberg, 1998).

The features added to the radar view, and the elements of workspace awareness that they support, are shown in Table VIII.

9. Summary of the workspace awareness framework

Workspace awareness is the up-to-the-moment understanding of another person's interaction with the shared workspace. The conceptual framework sets out basic issues that designers need to consider when building workspace awareness support into groupware systems. The framework describes three aspects of workspace awareness: its component elements, the mechanisms used to maintain it, and its uses in collaboration. These parts correspond to three tasks that the groupware designer must undertake in supporting workspace awareness: understand what information to provide, determine how the knowledge will be gathered, and determine when and where the knowledge will be used. The framework is illustrated in Figure 5, overlaid on Neisser's original perception-action cycle. In addition, we add a new link to the cycle (action) to indicate that people take action based on their knowledge as well as exploring the environment.

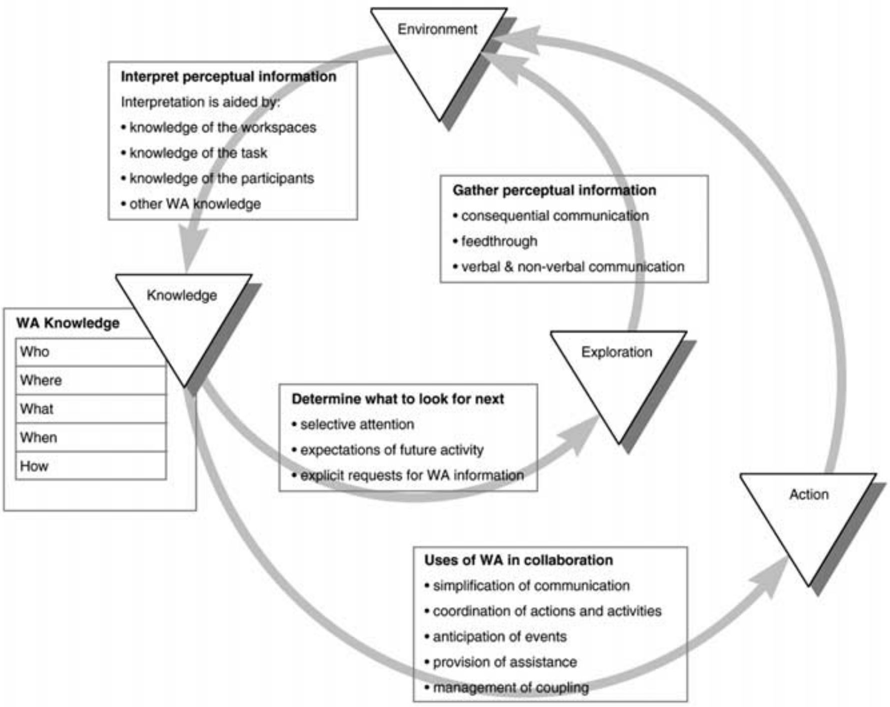

Figure 5. The workspace awareness framework.

The elements of workspace awareness answer who, where, when, how, and what questions. They deal with issues like who is present and who is responsible for actions, where people are working and where they can see, and what actions they are performing and what their intentions are. Other elements of workspace awareness consider awareness of history and past events. The elements are a starting point for thinking about the awareness requirements of particular task situations, and provide a vocabulary for describing and comparing awareness support in groupware applications.

Workspace awareness is maintained through a perception-action cycle in which people gather perceptual information from the environment, integrate it with what they already know, and use it to look for more information in the workspace. Information is gathered primarily through three mechanisms. First, the presence and movement of hands and bodies in the workspace provide consequential communication. Second, movement and changes to artifacts in the workspace provides feedthrough information. Third, information is gathered through intentional communication, which can be either verbal or gestural. People are already familiar with these three ways of gathering workspace awareness information, from their experiences in face-to-face workspaces. In groupware, designers can simplify information-gathering by using these mechanisms in their awareness displays, even though the displays themselves will likely bear little resemblance to face-to-face environments.

Workspace awareness is useful for making collaborative interaction more efficient, less effortful, and less error-prone. There are several activities of collaboration where the benefits of workspace awareness are evident: in helping people to recognize opportunities for closer coupling, in reducing the effort needed for verbal communication, in simplifying coordination, in allowing people to act in anticipation of others, and in providing context for appropriate help and assistance. Designers can use this part of the framework as an analysis tool to help them determine the awareness support that is needed for a particular work situation, and as a guide to determining where in the interface that awareness support should be provided.

The role of the framework in the groupware design process is not as a prescriptive design guide, but rather as a structured collection of knowledge that can assist the iterative development of awareness support. The framework identifies three steps that designers should undertake – think about what information to provide, what perceptual mechanisms to use to convey the information, and when and where in the interface to provide the information – and provides a set of alternatives and possibilities for each step.

The knowledge in the conceptual framework will allow designers to build more usable groupware, and this knowledge has not previously been available to groupware designers in one place. However, workspace awareness is only one type of group awareness, and the knowledge in our framework must be used along with other tools. For example, another model of awareness in collaborative virtual environments is the focus/nimbus model (e.g. Benford et al., 1995; Rodden, 1996). The model offers a way to determine what the level of awareness *should* be for two actors in a shared space. The actors' physical location and the distance between them are two important factors in the model, and states an inverse relationship between distance and awareness – the farther you are from someone, the less aware you should be of them. In addition, the model incorporates the possibility that actors can affect their own degree of awareness: these capabilities are represented in the concepts of *focus* and *nimbus*. The focus/nimbus model is concerned with large spaces that can contain many people, and hence the focus on determining how much awareness information should be provided. Our framework, in contrast, is oriented towards small groups in medium-sized workspaces where it is more likely that participants are always interested in maintaining awareness of all the members of the group. Therefore, we see the focus/nimbus model as a higher-level complement to our framework. The two models can work together in environments where people can work together at both a large and a small scale – the focus/nimbus model would operate in the large, and the workspace awareness model in the small.

10. Conclusion

In this paper we have presented a descriptive theory of awareness for small groups in shared-workspace groupware. Our motivation for the research is that although the idea of group awareness shows great promise for improving groupware usability, groupware designers do not have access to principled information about how to support it in their interfaces. Our goal, therefore, was to provide developers with useful knowledge about how to design for awareness in multi-user systems, and in particular, how to design for one kind of awareness called workspace awareness. The main structure of the descriptive theory is a framework of workspace awareness that organizes the concept and that informs designers as they analyse work situations and consider the design of awareness support. The framework is based on sound psychological principles of what awareness is and how people maintain it in dynamic environments. The framework can both educate designers about the importance of awareness in groupware and help to improve the quality of the systems that are built.

We believe that the foundations and basic structure of the framework can be used to characterize and describe other types of awareness that affect distributed group work. First, the perception-action cycle is a general model that can be used to explain how people keep track of a wide variety of information in a collaborative situation. Second, the three design issues of what information to present, how to present it, and where and when to present it apply equally well to supporting (for example) informal awareness and conversational awareness in groupware. Since workspace awareness is not independent of these other types, a more comprehensive theory that integrates several different aspects of group awareness is needed. Extending the framework is one of our current ongoing projects. Other current work includes assessing the effects of awareness support on groupware usability (Gutwin and Greenberg, 1998a) and developing new awareness displays and devices (Gutwin and Greenberg, 1998b).

Acknowledgements

This research was supported in part by the Natural Sciences and Engineering Research Council of Canada, by Intel Corporation, and by Microsoft Research. Our thanks to the anonymous referees for their valuable comments and recommendations.

Notes

1. However, these qualities could easily be part of collaborative work: for example, in a fast-paced multiplayer video game.
2. Appendix A briefly describes our own observational studies used in the conceptual framework.
3. This is different from Dewan and Choudhary's (1991) earlier notion of coupling, which involves coupling between interface elements in shared interfaces.

Appendix A: Observational studies used in the framework

We observed several groups performing simple tasks in physical shared workspaces, in order to gather basic information about the uses and mechanisms of workspace awareness, and to gain first-hand experience with phenomena described in research literature. Findings from these studies contribute to the structure and content of the conceptual framework. The studies were informal and varied widely in task, group structure, setting, and realism; in some cases, we even participated as part of the group. We did not consistently employ one particular methodology, but in all cases we observed the collaboration and recorded our observations. In some sessions, the collaboration was videotaped for later review.

Below, we introduce each session to give an idea of the settings and the tasks that were observed. The first five tasks were completed in a laboratory setting, and the final two were visits to real work environments. In the laboratory tasks, people were allowed to organize their collaboration however they saw fit. All of the laboratory tasks were made-up activities, while the two real work visits involved people's normal work activities.

Blocks and puzzles. We began our observations by asking people to complete simple tabletop tasks with one of us as a partner. Three people each completed three different tasks. The first task was a jigsaw puzzle, the second was a puzzle with pentominoes pieces, and in the third, we built a house out of toy blocks. All three tasks were carried out at an ordinary table. These tasks took approximately 10 minutes each to complete.

String. Three dyads were asked to measure the distance between several pairs of points on a whiteboard, using a long piece of string as a measuring tool. The points were far enough apart that each person had to hold one end of the string. The participants did the task in two settings: first, in front of a normal whiteboard, and second, with a divider that prevented them from seeing one another's work areas. The tasks took about 20 minutes in total.

Cathedral. Two pairs completed a more complicated construction task, that of building a two-dimensional plan of a cathedral using a variety of cardboard pieces. The task included constraints (such as keeping the colours symmetrical) to encourage more interaction between the two participants. The task took place on a large table, and participants were allowed to move where they wished around the workspace. The cathedral task took about 40 minutes to complete.

Concept map. Three pairs were asked to complete a half-finished concept map using a written paragraph as their guide to the entities and relationships in the map. Again, the materials were paper and pencils, and the workspace was a large table. Pairs had to organize a set of existing objects and relations, and then add to the diagram until the paragraph was fully represented by the map. The concept map tasks took people about 50 minutes to finish.

Newspaper layout. Nine pairs completed a newspaper layout task. Groups were asked to put together a two-page spread of a fictional newspaper, using paper articles, pictures, and headlines supplied to them. Groups were allowed to lay out the pages as they wished, as long as the paper had a roughly consistent style. These tasks required about 40 minutes.

Results of this study were reported in (Gutwin, Roseman, and Greenberg 1996).

Newsroom. A visit to the student newspaper offices on production day was one of two observations of real work situations. We spent approximately six hours in the production room of the Gauntlet, the University of Calgary student newspaper, watching activities that ranged from story composition to page layout. In the part of the office we observed, five writers and two editors worked on the paper.

Air traffic control. The second real work situation that we visited was the air traffic control centre at the Calgary airport. We spent about four hours observing three collaborating controllers who supervise the airspace in a 35-mile radius around Calgary. A controller is in charge of one of three stations: commercial arrivals, commercial departures, or small private aircraft that operate under visual flight rules. Controllers sit in front of large radar screens that show all flight activity within an adjustable radius from the airport. Therefore, controllers see one another's aircraft on their screens. The controllers interact with each other, with the tower operators who supervise takeoffs and landings, and with regional controllers who supervise the airspace beyond the 35-mile radius. A typical high-level task for the arrivals controller would be to accept an aircraft from the regional controllers, guide it into its final approach, and hand it off to the tower controllers (*cf.* Heath and Luff 1992).

References

Ackerman, M. and B. Starr (1995): Social Activity Indicators: Interface Components for CSCW Systems. *Proceedings of the ACM Symposium on User Interface Software and Technology*, pp. 159–168.

Adams, M., Y. Tenney and R. Pew (1995): Situation Awareness and the Cognitive Management of Complex Systems. *Human Factors*, vol. 37, no. 1, pp. 85–104.

Baecker, R., D. Nastos, I. Posner and K. Mawby (1993): The User-Centred Iterative Design of Collaborative Writing Software. *Proceedings of the Conference on Human Factors in Computing Systems*, Amsterdam, pp. 399–405.

Beaudouin-Lafon, M. and A. Karsenty (1992): Transparency and Awareness in a Real-Time Groupware System. *Proceedings of the Conference on User Interface and Software Technology*, Monterey, CA, pp. 171–180.

Benford, S., J. Bowers, L. Fahlen, C. Greenhalgh and D. Snowdon (1995): User Embodiment in Collaborative Virtual Environments. *Proceedings of the Conference on Human Factors in Computing Systems (CHI'95)*, pp. 242–249.

Birdwhistell, Ray L. (1952): *Introduction to Kinesics: An Annotation System for Analysis of Body Motion and Gesture*. University of Kentucky Press.

Borning, A. and M. Travers (1991): Two Approaches to Casual Interaction over Computer and Video Networks. *Proceedings of the Conference on Human Factors in Computing Systems*, New Orleans, LA, pp. 13–19.

Brennan, S. (1990): *Seeking and Providing Evidence for Mutual Understanding*. Ph.D. thesis, Stanford University, Stanford, CA.

Brinck, T. and L.M. Gomez (1992): A Collaborative Medium for the Support of Conversational Props. *Proceedings of Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW'92)*, Toronto, Ontario, pp. 171–178.

Clark, H. (1996): *Using Language*. Cambridge: Cambridge University Press.

Clark, H.H. and S.E. Brennan (1991): Grounding in Communication. In R.M. Baecker (ed.): *Readings in Groupware and Computer Supported Cooperative Work: Assisting Human-Human Collaboration*. Mountain View, CA: Morgan-Kaufmann Publishers, pp. 222–233.

Dix, A., J. Finlay, G. Abowd and R. Beale (1993): *Human-Computer Interaction*. Prentice Hall.

Dewan, P. and R. Choudhary (1991): Flexible User Interface Coupling in a Collaborative System. *Proceedings of CHI'91*, pp. 41–48.

Dourish, P. and V. Bellotti (1992): Awareness and Coordination in Shared Workspaces. *Proceedings of the Conference on Computer-Supported Cooperative Work*, Toronto, pp. 107–114.

Dourish, P. and S. Bly (1992): Portholes: Supporting Awareness in a Distributed Work Group. *Proceedings of the Conference on Human Factors in Computing Systems*, Monterey, CA, pp. 541–547.

Ellis, C., S. Gibbs and G. Rein (1991): Groupware: Some Issues and Experiences. *Communications of the ACM*, vol. 34, no. 1, pp. 38–58.

Endsley, M. (1995): Toward a Theory of Situation Awareness in Dynamic Systems. *Human Factors*, vol. 37, no. 1, pp. 32–64.

Gaba, D., S. Howard and S. Small (1995): Situation Awareness in Anesthesiology. *Human Factors*, vol. 37, no. 1, pp. 20–31.

Gaver, W. (1991): Sound Support for Collaboration. *Proceedings of the Second European Conference on Computer Supported Cooperative Work*, pp. 293–308.

Gilson, R.D. (1995): Introduction to the Special Issue on Situation Awareness. *Human Factors*, vol. 37, no. 1, pp. 3–4.

Goodwin, C. (1981): *Conversational Organization: Interaction Between Speakers and Hearers*. New York: Academic Press.

Greenberg, S. and R. Bohnet (1991): GroupSketch: A Multi-User Sketchpad for Geographically-Distributed Small Groups. *Proceedings of Proceedings of Graphics Interface '91*. Calgary, Alberta.

Greenberg, S. (1996): Peepholes: Low Cost Awareness of One's Community. *Proceedings of the Conference on Human Factors in Computing Systems (Conference Companion)*. Vancouver, pp. 206–207.

Greenberg, S. and D. Marwood (1994): Real Time Groupware as a Distributed System: Concurrency Control and its Effect on the Interface. *Proceedings of the Conference on Computer-Supported Cooperative Work*. Chapel Hill NC, pp. 207–217.

Gutwin, C. (2000): *Slow and Sticky: Effects of Network Delay on Real-Time Groupware*. Technical Report 2000-02, Department of Computer Science, University of Saskatchewan.

Gutwin, C., S. Greenberg and A. Cockburn (1996): Using Distortion-Oriented Displays to Support Workspace Awareness. *Proceedings of People and Computers XI (BCSHCI'96)*. London, pp. 299–314.

Gutwin, C., S. Greenberg and M. Roseman (1996): Workspace Awareness in Real-Time Distributed Groupware: Framework, Widgets, and Evaluation. In R.J. Sasse, A. Cunningham and R. Winder (eds.): *People and Computers XI (Proceedings of the HCI'96)*. Springer-Verlag, pp. 281–298.

Gutwin, C. and S. Greenberg (1996): Workspace Awareness for Groupware. *Proceedings of the Conference on Human Factors in Computing Systems*. Vancouver, pp. 208–209.

Gutwin, C., S. Greenberg and M. Roseman (1996): Workspace Awareness in Real-Time Distributed Groupware: Framework, Widgets, and Evaluation. *People and Computers XI (Proceedings of HCI'96)*. Springer-Verlag, pp. 281–298.

Gutwin, C., M. Roseman and S. Greenberg (1996): A Usability Study of Awareness Widgets in a Shared Workspace Groupware System. *Proceedings of the Conference on Computer-Supported Cooperative Work*. Boston, pp. 258–267.

Gutwin, C. and S. Greenberg (1998a): Effects of Awareness Support on Groupware Usability. *Proceedings of ACM CHI'98*. ACM Pres: Los Angeles.

Gutwin, C. and S. Greenberg (1998b): Design for Individuals, Design for Groups: Tradeoffs between Power and Workspace Awareness. *Proceedings of ACM CSCW'98*. ACM Press: Seattle.
Hall, E. (1959): *The Silent Language*. Doubleday.
Hayne, S., M. Pendergast and S. Greenberg (1993): Gesturing Through Cursors: Implementing Multiple Pointers in Group Supports Systems. In *Proceedings of the HICSS Hawaii International Conference on System Sciences*. IEEE Press.
Heath, C., M. Jirotka, P. Luff and J. Hindmarsh (1995): Unpacking Collaboration: The Interactional Organisation of Trading in a City Dealing Room. *Computer Supported Cooperative Work*, vol. 3, no. 2, pp. 147–165.
Heath, C. and P. Luff (1992): Collaboration and Control: Crisis Management and Multimedia Technology in London Underground Line Control Rooms. *Computer-Supported Cooperative Work*, vol. 1, nos. 1–2, pp. 69–94.
Hutchins, E. (1990): The Technology of Team Navigation. In J. Galegher, R. Kraut and C. Egido (eds.): *Intellectual Teamwork: Social and Technological Foundations of Cooperative Work*. Hillsdale, NJ: Lawrence Erlbaum, pp. 191–220.
Ishii, H. and M. Kobayashi (1992): ClearBoard: A Seamless Medium for Shared Drawing and Conversation with Eye Contact. *Proceedings of the Conference on Human Factors in Computing Systems*. Monterey, CA, pp. 525–532.
James, W. (1981 [written 1890]): *The Principles of Psychology*. Cambridge, MA: Harvard University Press.
Krauss, R. and S. Fussell (1997): Mutual Knowledge and Communicative Effectiveness. In J. Galegher, R. Kraut and C. Egido (eds.): *Intellectual Teamwork: Social and Technological Foundations of Cooperative Work*. Hillsdale, NJ: Lawrence Erlbaum, pp. 111–145.
McDaniel, S.E. and T. Brinck (1997): Awareness in Collaborative Systems. Workshop Report. *SIGCHI Bulletin*, October.
McGrath, J., (1984): *Groups: Interaction and Performance*. Englewood Cliffs, NJ: Prentice-Hall.
Mitchell, A. (1996): *Communication and Shared Understanding in Collaborative Writing*. M.Sc. thesis, University of Toronto, Toronto.
Neisser, U. (1976): *Cognition and Reality*. San Fransisco: W.H. Freeman.
Norman, D. (1993): *Things That Make Us Smart*. Reading, MA: Addison-Wesley.
Robinson, M. (1991): Computer-Supported Cooperative Work: Cases and Concepts. *Proceedings of Groupware '91*, pp. 59–75.
Rodden, T. (1996): Populating the Application: A Model of Awareness for Cooperative Applications. *Proceedings of ACM CSCW'96 Conference on Computer-Supported Cooperative Work*, pp. 87–96.
Roseman, M. and S. Greenberg (1996): Building Real-Time Groupware with GroupKit, a Groupware Toolkit. *Transactions on Computer-Human Interaction*, vol. 3, no. 1, pp. 66–106.
Roseman, M. and S. Greenberg (1996): TeamRooms: Network Places for Collaboration. *Proceedings of the Conference on Computer-Supported Cooperative Work*. Boston.
Salas, E., C. Prince, D. Baker and L. Shrestha (1995): Situation Awareness in Team Performance: Implications for Measurement and Training. *Human Factors*, vol. 37, no. 1, pp. 123–136.
Salvador, T., J. Scholtz and J. Larson (1996): The Denver Model for Groupware Design. *SIGCHI Bulletin*, vol. 28, no. 1, pp. 52–58.
Sarter, N. and D. Woods (1995): How in the World Did We Ever Get into That Mode? Mode Error and Awareness in Supervisory Control. *Human Factors*, vol. 37, no. 1, pp. 5–19.
Seely Brown, J., A. Collins and P. Duguid (1989): Situated Cognition and the Culture of Learning. *Educational Researcher* (January–February), pp. 32–42.
Segal, L. (1994): Effects of Checklist Interface on Non-Verbal Crew Communications, NASA Ames Research Center, Contractor Report 177639.

Segal, L. (1995): Designing Team Workstations: The Choreography of Teamwork. In P. Hancock, J. Flach, J. Caird and K. Vicente (eds.): *Local Applications of the Ecological Approach to Human-Machine Systems*. Hillsdale, NJ: Lawrence Erlbaum, pp. 392–415.

Short, J., E. Williams and B. Christie (1976): Communication Modes and Task Performance. In R.M. Baecker (ed.): *Readings in Groupware and Computer Supported Cooperative Work: Assisting Human-Human Collaboration*. Mountain View, CA: Morgan-Kaufmann Publishers, pp. 169–176.

Smith, K. and P. Hancock (1995): Situation Awareness is Adaptive, Externally Directed Consciousness. *Human Factors*, vol. 37, no. 1, pp. 137–148.

Smith, R. (1999): *The Kansas Project*. http://www.sun.com/research/ics/kansas.html.

Smith, R., R. Hixon and B. Horan (1998): Supporting Flexible Roles in a Shared Space. *Proceedings of ACM CSCW'98 Conference on Computer-Supported Cooperative Work*, pp. 197–206.

Sohlenkamp, M. and G. Chwelos (1994): Integrating Communication, Cooperation and Awareness: The DIVA Virtual Office Environment. *Proceedings of ACM CSCW'94 Conference on Computer-Supported Cooperative Work*, pp. 331–343.

Stefik, M., G. Foster, D. Bobrow, K. Kahn, S. Lanning and L. Suchman (1987a): Beyond the Chalkboard: Computer Support for Collaboration and Problem Solving in Meetings. *Communications of the ACM*, vol. 30, no. 1, pp. 32–47.

Stefik, M., D. Bobrow, G. Foster, S. Lanning and D. Tatar (1987b): WYSIWIS Revised: Early Experiences with Multiuser Interfaces. *ACM Transactions on Office Information Systems*, vol. 5, no. 2, pp. 147–167.

Tang, J. (1989): *Listing, Drawing, and Gesturing in Design: A Study of the Use of Shared Workspaces by Design Teams*. Ph.D. thesis, Stanford University, Stanford, CA.

Tang, J. (1991): Findings from Observational Studies of Collaborative Work. *International Journal of Man-Machine Studies*, vol. 34, no. 2, pp. 143–160.

Tatar, D., G. Foster and D. Bobrow (1991): Design for Conversation: Lessons from Cognoter. *International Journal of Man-Machine Studies*, vol. 34, no. 2, pp. 185–210.

Tang, J.C. and S.L. Minneman (1991): VideoWhiteboard: Video shadows to support remote collaboration. *Proceedings of ACM SIGCHI Conference on Human Factors in Computing Systems, New Orleans*, pp. 315–322.

Watts, J., D. Woods, J. Corban, E. Patterson, R. Kerr and L. Hicks (1996): Voice Loops as Cooperative Aids in Space Shuttle Mission Control. *Proceedings of ACM CSCW'96 Conference on Computer-Supported Cooperative Work*, pp. 48–56.