

Consistent practices in artifact-mediated collaboration

Nathan Dwyer · Daniel D. Suthers

Received: 14 March 2006 / Revised: 18 October 2006 /

Accepted: 20 October 2006 / Published online: 30 November 2006

© International Society of the Learning Science, Inc.; Springer Science + Business Media, LLC 2006

Abstract The design of collaborative representations faces a challenge in integrating theoretical communication models with the context-sensitive and creative practices of human interaction. This paper presents results from a study that identified multiple, invariant communicative practices in how dyads appropriated flexible, paper-based media in discussions of wicked problems. These invariants, identified across media, participants and topics are a promising first step towards creating an abstract model for design that connects representational affordances and communicative functions. The authors identify areas where this model may challenge conventional design wisdom and discuss directions for further research.

Keywords Descriptive studies · Interactional practices · Representational affordances · Shared workspaces · Video analysis

Introduction

The mapping of communication models to user interfaces, to date, has not produced collaboration tools that approach the effectiveness of face-to-face (FTF) communication. In contrast to the concrete structures of theoretical communication models (Gerbner, 1956; Jakobson, 1960; Shannon, 1948), studies of spoken language have established that the meanings of utterances are contextual and negotiated only to the level of agreement needed to support action (Galantucci, 2005; Garfinkel & Sacks, 1970; Goodwin & Heritage, 1990). The work reported in this paper is based on the empirically grounded premise that the same is true of nonlinguistic representations. Meanings of representations are not fixed in advance, but change according to context. Efforts to provide users with a visual language

N. Dwyer (✉) · D. D. Suthers

Department of Information and Computer Sciences, University of Hawai'i at Manoa,
1680 East West Road, POST 317, Honolulu, HI 96822, USA
e-mail: ndwyer@hawaii.edu

D. D. Suthers

e-mail: suthers@hawaii.edu

for reasoning or argumentation have often encountered difficulty creating a functional notation (van Bruggen & Kirschner, 2003) and in getting people to adopt such a notation (Conklin, 2003; Selvin, 2003). Also, users resist the very idea of categorizing their thinking, although they do seek representational aids in organizing ideas (Shipman & McCall, 1994). People make flexible use of representations, and the affordances of representations are appropriated in sometimes-unexpected ways. For example, Dillenbourg & Traum (1999) had participants use synchronous chat and white-boards in a MOO environment while solving a murder mystery. The researchers expected that the white-board would be used for disambiguating spatial references through its two dimensionality and drawing affordances. Instead, the most important affordance of the white-board for participants turned out to be its persistence: information that had to be recorded permanently was written in the white-board.

Although there is substantial work on replicating the properties of FTF communication by using high bandwidth video and audio (Finn, Sellen, & Wilbur, 1997; G. M. Olson & J. S. Olson, 2002) and techniques such as clever placement of cameras and screens for accurate conveyance of gesture and gaze (e.g., Kato et al., 2001), others believe that collaborative technologies offer unique opportunities (Dillenbourg, 2005; Suthers, 2006) and should go “beyond being there” (Hollan & Stornetta, 1992) by exploiting the special properties of computational media in ways that make distance interaction more effective. This message resonates with educators’ recognition of the need to guide and scaffold learning. Whether the application is intended for learning or work, it is not enough to simply provide a channel of communication as rich as face-to-face communication. We truly realize the potential of these technologies only if we use them to guide and enable more effective learning and problem solving practices. This point applies equally well to face-to-face interaction. If the richness of FTF were sufficient to solve problems of learning and collaboration, we would not see interest in technologies that support these activities specifically in FTF contexts (Kaput & Hegedus, 2002; Lingnau, Hoppe, & Mannhaupt, 2003; Sugimoto, 2003).

This situation presents a dilemma for designers of computer-supported collaborative learning (CSCL) or collaborative work (CSCW) systems, or indeed computer-mediated communication (CMC) in general. (We will refer to these collectively as *collaborative technologies*.) Designers seek to build representational and interactional tools that guide and support cognitive and social activities while also allowing for flexible use of both linguistic and nonlinguistic representations, but user interface toolkits tend to define rigid mappings between graphical user interface (GUI) elements and functionality. User interfaces that are easy to build with typical GUI toolkits are not a good match to the flexible nature of human communication, nor do they adapt to the changing needs of the user. Furthermore, the needs of the user and the purpose of the interface change over time. Constrained representational tools may guide novices in their enculturation to a new field, but must be kept simple for this learning period and this simplicity may soon become too limiting. Conversely, a set of tools that is sufficiently complex for supporting experts in a field can be daunting, and dissuade novices.

The present study attempts to address this dilemma by identifying how people appropriate flexible representations to meet their needs—essentially, to identify what Garfinkel calls “member’s methods” (Garfinkel, 1967) for synchronous collaboration via written representations—so that we can build CMC tools with affordances (Gibson, 1977) that support the kinds of flexibilities observed. We want to answer such questions as: When people communicate via written means, what strategies do they use to manage the interaction? How do they appropriate the affordances of media to carry out these strategies?

How can we design user-interfaces for CMC tools that provide those affordances? This paper addresses the first two questions empirically and discusses implications for the third.

We cannot adequately answer such questions with studies of CMC, nor of FTF interaction alone. A third strategy is needed and taken by this study. We cannot effectively conduct the study with an existing CMC technology, since any CMC technology we chose would carry with it the very assumptions of collaborative technology design that we are questioning. Exploratory development of collaboration technology is an alternative, but comes with the high cost of implementing each iteration of the software. It would also be a mistake to conduct the study with unrestricted FTF interaction because there are too many differences between FTF and CMC to create a reasonable mapping from one to the other. FTF interaction includes many subtle cues that are difficult to replicate online (Clark & Brennan, 1991; Goodwin, 2000; G. M. Olson & J. S. Olson, 2000), and we want to find the special advantages of artifact-mediated communications that might not be evident in FTF interaction. Therefore, we take a middle road.

Our strategy is to start with FTF, but restrict or remove some of its features that are especially hard to replicate online. We require that people communicate with written representations, but to do so using very familiar and flexible tools—paper office supplies—so we can get a sense of which affordances of flexible representational tools participants take up, and for what communicative functions. Many attributes of the tools were varied so as to highlight invariance in communicative functions across different permutations. In this paper we report on the range of ways in which the tools were exploited, identify invariances in terms of what was being accomplished by these uses, and discuss ways in which these findings challenge conventional wisdom regarding CMC software design.

Materials and methods

Strategy

The purpose of this study was to discover strategies or methods that people use to collaborate through shared written representations. This understanding is sought *both independently of and in reference to the representations used*:

- We want to know what kinds of communicative or coordinative functions people consistently attempt to implement independent of the representation used.
- We want to know how people appropriate the affordances of specific representations for these purposes.

Our methodological strategy is to vary the representational tools provided and look for invariants across the different configurations. This differs from an experimental design, which attempts to control as many factors as possible and show there is a difference between experimental groups correlated with the one thing that varies between those groups. We are not making comparisons between experimental groups, and the variation in participants and materials between the sessions is desirable, as it strengthens claims of generality.

Like Sacks (1984a) we are “trying to find the machinery” of social interaction, but we are looking at a different level of the machinery than traditional ethnomethodological analysis.

We provide a constrained, and possibly odd-seeming, environment for people to interact in, and then examine how they make sense of each other. The participants are required to create shared methods by which they make their actions accountable to each other. This is the heart of Garfinkel's (1967) definition of ethnomethodology, in which he goes to some length to explain that "any occasion whatsoever" is an appropriate setting for analysis if the "communality of practical actions is a project of members' actions." We rely heavily on Sacks' (1984b) assertion of "order at all points," but we apply his logic at the level of media appropriation. Sacks observes that individuals experience only a small, random portion of their culture, yet these experiences consistently generalize to their culture as a whole. A parallel observation can be made that, confronted with new communication media, individuals are capable of rapidly generalizing methods of communication to make use of the media's particular capabilities. In our case, we are focusing on methods of successful collaboration via written media rather than the mechanisms of casual conversation. Face-to-face interactions use a broad array of semiotic resources (Goodwin, 2000), many of which are prohibitively difficult to re-create online. The environment designed for this study limits the available semiotic resources in ways that mimic online limitations. Our goal is to document the methods that unfold in this intentionally constrained environment in order to understand how we can re-create them in the more constrained online world.

The study design echoes the methods of Vygotsky and Garfinkel as well as some more recent research. Vygotsky's work started from the premise that higher forms of behavior should not be considered stable forms of interaction. "Any psychological process, whether the development of thought or voluntary behavior, is a process undergoing changes right before one's eyes" (Vygotsky, 1978, p. 61). The collaborative practices that support intersubjective meaning-making are, by nature, evolving interactive processes. Vygotsky makes a strong case for studying these changes as they unfold, and not the "fossilized" practices that are eventually produced by them. He goes on to argue that while an experimenter could wait for any process to be exhibited, he proposes a methodology that "artificially provokes or creates a process of psychological development" (Vygotsky, 1978, p. 61). This study echoes Vygotsky's methodology in that it is designed to provoke opportunistic and creative media appropriation so that that the process may be observed as it unfolds. Vygotsky calls his approach "the functional method of double stimulation" (Vygotsky, 1978, p. 74).

Garfinkel took a similar methodological stance in designing the "breaching" experiments he created early on in his construction of Ethnomethodology (Garfinkel, 1967). Garfinkel's primary interest focuses on "the familiar commonsense world of everyday life" (Garfinkel, 1967, p. 36). His premise is that it is a set of "expected, background features" of everyday activities that provides them with their commonsense nature, and his breaching experiments were attempts "to detect some expectancies that lend commonplace scenes their familiar, life-as-usual character, and to relate these to the stable social structures of everyday activities" (Garfinkel, 1967, p. 37). He goes on to summarize his methodology, "Procedurally it is my preference to start with familiar scenes and ask what can be done to make trouble" (Garfinkel, 1967, p. 36). Garfinkel directed his students to act in ways that challenged the background expectancies of those around them, provoking psychological and social processes of adaptation that revealed the "fossilized" (Vygotsky, 1978) background expectancies in which he was interested.

Recent research following in the same methodological vein examines the development of communicative practices when normal communication channels are limited or removed. Observations of deaf children (Goldin-Meadow & Feldman, 1977; Goldin-Meadow &

Mylander, 1998) documented the spontaneous development of communication systems in the absence of verbal interaction. Healey, Swoboda, Umata, and Katagiri (2002) describe a series of experiments in which participants were asked to communicate about pieces of music using only drawings but without using letters or numbers. Galantucci (2005) conducted a series of experiments in which pairs of participants were asked to play a collaborative logic game using only a severely restricted drawing tool that disallowed the use of recognizable symbol systems or complex, iconic representations. The methods utilized in these studies were effective at documenting the communication systems that developed and, in the latter two, the processes through which they developed. The present work takes a similar approach.

Participants

The study involved six pairs of friends recruited from a pool of community college and university students. Their ages ranged from 18 to 24, with an average age of 20.9. They had a diverse set of majors, including Digital Art, Nursing, Environmental Studies, and Computer Science. All pairs had consistent social contact and all but one had known each other for multiple years. An additional six pairs were recruited for pilot studies in order to test different discussion questions and configurations of the materials. Unless noted otherwise, the observations below do not include data from the pilot studies.

Environment

A large table was prepared with a screen suspended above it such that participants seated on opposite sides of the table could not see each others' faces (see Fig. 1). The seating was arranged such that participants could reach and therefore manipulate the entire workspace. This allowed them to see each others' arms and infer the positions of each others' torsos.

A video camera was placed above and to the side of the participants, positioned in the same plane as the screen to minimize the extent to which the screen blocked the camera's view of the workspace. The camera captured both participants and the entire workspace, except for a thin line blocked by the screen. Digital video output was streamed to a hard drive in real time.

Materials and methods

In all cases, the table was covered with a secured sheet of butcher paper. All pairs had access to tape, rulers, scissors, string, paperclips, and sticky labels, dots and stars. All pairs were provided with the same collection of various writing utensils, including red, green, blue, and black whiteboard markers and ballpoint pens, and blue and black permanent markers.

Different pairs were provided with different kinds of paper products (see Table 1). Three pairs of participants were given a variety of office supplies such as varying sizes and colors of paper, multiple sizes of index cards, and multiple sizes and colors of Post-It™ (sticky) notes. One pair was only given 3×5 in. (7.6×12.7 cm) index cards. Two pairs were given a single, large sheet of 2×3 foot (61×91 cm) unlined paper. The original intent was to have two pairs for each set of materials, but the first pair given unrestricted supplies used almost nothing but 3×5 index cards, so an additional session was done with the unrestricted set of supplies.

Fig. 1 Study environment and materials



Office supplies were chosen partially for their familiarity. It was assumed that participants would have used them at school or work and would not need time to learn their affordances. Also, office supplies have a demonstrated history of versatility and effectiveness. The three sets of materials were chosen for the specific affordances they provided (see Table 2). Individual pieces of paper and index cards can be moved around and repositioned in relation to each other, and Post-It™ notes can be attached and detached

Table 1 Session information

Session	Gender	Given	Used
S1	Males	3×5 index cards	Only 3×5 index cards
S2	Males	Unrestricted paper supplies	3×5 and 4×6 cards, two sheets of 8.5"×11" paper, and one post-it
S3	Males	Single large sheet of paper	Also wrote on table top
S4	Females	Single large sheet of paper	Also wrote on instruction sheet
S5	Females	Unrestricted paper supplies	All materials
S6	Males	Unrestricted paper supplies	All materials

Table 2 Selected affordances of writing surfaces

Paper medium	Reorientable	Repositionable	Space limited	Attachable
Butcher paper taped to table				
Large sheet of paper	Y			
Letter sized paper	Y	Y		
Index cards	Y	Y	Y	
Post-it notes	Y	Y	Y	Y

from other materials. The index card only condition removes the attachment affordance, and the large sheet of paper additionally disallows repositioning the participant's contributions.

Procedure

Pairs of participants were given a sequence of discussion topics and asked to brainstorm ideas, discuss them and come to some kind of final agreement. The pairs were given three topics chosen from a pool of five "wicked" (Rittel & Webber, 1973) problems, such as: How do we save Hawaii's environment? Space aliens are coming; how should we respond when they arrive? What is the appropriate relationship between science and religion? Wicked problems are typified by the lack of clear evaluation metrics for any answer as well as the lack of a well-specified process for approaching them. These problem features required the participants to collaboratively develop processes for collaboration and evaluation, and to negotiate when they had come to the end of the discussion. The order and selection of topics assigned to the pairs were permuted to minimize any confounding effects such as fatigue, learning, and familiarity.

Every effort was made to support the greatest amount of flexibility in the representational medium, but in order to approximate limitations of online communication several restrictions were placed on the participants. Since video-mediated communication can be problematic (Heath & Luff, 1991; Mark & Abrams, 2005) and since text is by far the predominant on-line communication medium, we chose to limit visual and verbal communication channels. Participants were required to communicate using only the pens and materials, and because of the screen, participants could not communicate using facial expressions. They were also asked to remain silent during the written discussion. The participants' hands and arms were visible to each other. Given that gesture is so fundamental to communication (Goldin-Meadow, 1999; Krauss, 1998), and since telepointers and avatars are viable CMC tools, this seemed to be a reasonable allowance.

Discussion problems were printed on sheets of paper and these 'problem sheets' were given to the pairs one at a time during the session. Pairs were allowed to ask questions about the problems before they began each discussion. They were told that they should discuss the problem in as much detail as possible, and that they needed to agree on a final conclusion. Pairs were given 30 min to discuss each problem. One of the authors recorded observations, notes, and comments during the sessions. After each problem, the pairs were interviewed on their conclusion, what they thought of the interaction, difficulties or issues with the procedure and any other reaction to the discussion. This time was also used to ask for clarification on any activities that had been observed during the discussion.

Analysis

Grounded Theory methods (Glaser, 1992; Glaser & Strauss, 1967) were used to analyze the data. In the first pass at the analysis, the goal was to cover the breadth of the data and identify its coarse structure. The video data was reviewed in multiple passes using a custom application that was developed to support variable speed review as well as annotation of the video. Initial review of the video was done at normal speed, to reinforce an awareness of the contents, and then at double or triple speed to get a sense of the larger-scale recurring patterns. Initial reviews of the video identified several of the most obvious interaction patterns—generally those that dealt with use of the space and discussion structure. Successive reviews looked specifically for these patterns and identified several more. The entire collection of video (approximately 8 h) has been reviewed multiple times, and several interesting segments have been reviewed at much slower playback speeds in order to study the fine details of the interaction. Consistent with grounded theory methods, the video was heavily annotated during the multiple analysis passes, and the constant comparative method was used to evaluate and refine our understanding of the contents. Extensive memos were kept and these notes along with the annotated video eventually became the basis of the following observations.

Observations

On the surface, each pair's interaction style was decidedly unique. Some created chaotic collages of text decorated with symbols and tied together with arrows, while others constrained themselves to linear contributions on individual index cards. Participants' interactions repeatedly demonstrated that the properties of FTF do not define the universe of all possible means of interaction. The study environments provided significantly different sets of representational affordances and constraints, yet the participants readily adapted to these environments. Participants interacted in ways that are not possible in spoken interaction, so their interactions cannot be considered solely in terms of adaptation of FTF practices.

Despite the multiple differences in materials, participants and topics and the different interaction styles these engendered, there was an underlying consistency to the structure of the interactions and the methods the pairs employed to define the content and the process of the collaboration. These consistent communicative functions are interesting in that they suggest the existence of foundational invariants of interaction that could form the basis for user-interface design principles. Some of these invariants suggest specific features that could benefit collaborative software, while others challenge common HCI design wisdom and point to a possible re-conceptualization of collaborative software design.

This section presents the consistencies that were observed and presents examples of the different practices by which the pairs enacted them. Examples are drawn from multiple interactions to show the diversity in the pairs' practices. One pair (S5¹) provides particularly striking examples, being the only pair to exhibit overt hostility during their interactions. Despite the antagonism of their interaction, the S5 pair's interactions provide evidence of the same communicative functions as the other pairs. Analysis of their conflict

¹ The sessions are designated S1–S5 and the participants are designated with L and R. For example, 'S5' refers to the fifth study session, and 'S5L' indicates the person on the left from the perspective of the observer.

provides useful, contrasting evidence for the existence of invariants. The S5 pair will be discussed separately, and unless otherwise noted, the generalizations about behavior made below do not include either the pilot study or the S5 pair. After briefly commenting on the content of participants' discussions, we examine interaction structure, use of materials, workspace organization, and interactionally negotiated conventions in turn, followed by a discussion of S5's conflict.

Discussion content

Neither the content of the discussions nor the conclusions reached in each problem were remarkable. The discussions all followed an informal style and displayed a topically episodic structure: topics and ideas introduced by each participant were explored only to the point where the pair agreed on either the topic's relevance or a general conclusion. Few topics were revisited, and usually only to verify the initial agreement before the information was incorporated into the final product. The topics of conversation were usually organized spatially, which is reflected in the resulting artifacts. On a single large sheet of paper, areas can be identified for each topic discussed. Pairs who used 8.5×11 pages generally confined each topic to a single page, and pairs who relied on smaller materials almost always introduced a new post-it or 3×5 card to start a new topic.

The physical nature of the materials was not exploited to investigate connections between ideas or to propose groupings or inclusions. In fact, there was no evidence of meta-level information management such as grouping ideas into categories or making explicit pro and con lists. The result was that despite the kinds of simultaneous contributions discussed below, the discussions were essentially linear. This suggests some important limitations of this study. The limited time frame, synchronous interaction, and the dependence on personal opinion and knowledge (rather than external data and/or formal evaluation criteria) allowed the participants to manage the salient elements of the conversation without writing many of them down. It is possible that a longer time frame or the requirement to use external information and formal criteria would have motivated participants to invent more organizational mechanisms. Also, writing all contributions by hand limited the amount of text that could reasonably be generated in 30 min. This seemed to dampen enthusiasm for extended debates or long explanations.

This study did not evaluate the participants' conclusions, and does not comment on what role the environment played in the quality of their discussions. The participants consistently arrived at conclusions to their discussions and reported being satisfied with their conclusions and their interactions. It is reasonable to conclude that the environment was neither unduly beneficial nor detrimental to the unfolding of the collaboration.

Interaction structure

In contrast to the content, the ways in which each pair interacted were remarkably complex, subtle, and multi-layered. Unlike the alternating exchange of contributions postulated by models of face-to-face conversation (Gerbner, 1956; Jakobson, 1960), partners frequently contributed to the workspace *simultaneously*. Their actions overlapped, and so they interrupted each other, and occasionally interrupted themselves, in order to draw *attention* to their work. Contributions to the workspace most often consisted of task-related information, but process negotiation and personal exchanges were also common. Overlaid on all the contributions was *affective information* that informed the progression of the

interaction and the appropriate roles of the participants. The following sections illustrate each of these phenomena with examples from the data.

Simultaneity

The nature of the material allowed participants to read and write independently of each other. Individual pieces of paper supported each participant having their own writing area, while the single sheet of paper was large enough for both participants to write concurrently. The ability to work simultaneously allowed the pairs to have multiple active discussion threads. Some pairs made considerable use of this ability. The participants in S1 used multiple 3×5 cards to represent conversational threads. The pair accomplished this by writing their contribution on the card and then passing it across the table to their partner. The S4 pair, and S3 to a lesser degree, used different areas of the single sheet of paper to contain different threads. The S4 pair often wrote simultaneously about different topics. As each finished their contribution, they would negotiate which thread to pursue. Often, while one partner was writing her reply, the other would start a new thread or contribute to one of the existing threads.

The S2 and S6 pairs had minimal simultaneous threads, but the members of these pairs often wrote simultaneously. In all the pairs, each participant could observe the production of their partner's contribution. Since reading the contribution took much less time than writing it, the S2 and S6 participants would frequently read their partner's partially completed contribution and then start writing a reply. The S2 pair was the most extreme in this regard. Even though they were using materials that could be repositioned on the tabletop, each maintained their own material to write on and would alternate between writing and then reading what their partner had written. They did not reorient their material so the other could read it—they consistently read their partner's contributions upside down. While this use of the materials limited their ability to carry on multiple threads of conversation, it gave them the most freedom to write whenever they wished.

In addition to writing, the participants spent time reading and re-reading text that had been produced during the session. Participants could read the contribution their partner was currently writing, old contributions from either participant or the problem sheet that had been handed out. Participants would occasionally interrupt their own writing to read what their partner was currently working on. When participants had come to the end of a contribution, it was not uncommon for participants to use the time while their partner was writing to review the problem statement.

Summary of simultaneity Unlike FTF, written media afford simultaneous production and comprehension. Participants readily appropriate these affordances to manage their awareness of multiple simultaneous discussion threads.

Attention

The multiple concurrent threads and complex written artifacts required that participants take a more active role in directing their partner's attention. Participants sometimes wanted their partner to take immediate notice of their contribution, and communicated the urgency of the desire through the way they requested the attention. At the lowest level of urgency, participants often indicated the desire for attention by tapping their partner on the arm or by audibly tapping their finger or pen several times on the table, but if they felt a greater need

for their partner's attention some employed more intrusive means such as repositioning their hands or materials to cover their partner's work. As the sessions progressed, some pairs developed well-defined protocols for interrupting each other and drawing attention to a specific piece of text.

The S1 pair based their attention management on exchanging 3×5 cards. At the beginning of problem one, each would write on a card, wait for the partner to finish writing, and then the two would exchange cards. As the session went on, this protocol changed. When a participant was done writing he would slide the card over to his partner's side of the table, usually positioning the card to the partner's left. When the partner finished writing he would slide his card across the table, and then move the card from his left to be directly in front of himself. As the pair became more comfortable interacting with the materials, this protocol became even less formal. Several times, one of the two would finish writing on a card and then pick the card up and casually drop it on top of the partner's hand, or slide it over the card on which the partner was writing. The body language used (e.g., after dropping a card on his partner's writing hand, S1L usually turned his attention to a new topic) indicated that this was meant good-naturedly, and didn't imply any special urgency or importance.

While the protocols for getting attention differed, all the pairs were able to manage pending interactions and to indicate the desire for attention without requiring an immediate response. The S4 pair had the most developed protocol for interruptions. These participants were often writing simultaneously in different places on a single large sheet of paper. When one wanted the other's attention, she would tap her partner on the arm, and then would often immediately go back to work. The partner might not respond to the interruption immediately, but would finish her current task (e.g., complete the sentence she was writing) before turning her attention to her partner's current work and interpreting it as the motivation for the interruption.

In addition, participants often interrupted themselves. In the S4 pair, the participant on the right would often start writing and then interrupt herself to get her partner's attention. Once she had had the other's attention she would continue writing. In the first problem, she interrupts herself as shown in Fig. 2.

S4R would interrupt herself in this way even if her contribution was not particularly urgent. It is possible that she wanted the social approval of having her partner's attention or that she was merely excited about her contribution.

Where S4R interrupted herself to get her partner's attention, members of the S3 pair would interrupt themselves to pay attention to what their partner was writing. For example, during the first problem, the S3 pair has the exchange shown in Fig. 3.

This pair's interactions became finer grained as the session progressed. In transcript 2, S3R briefly interrupts himself in the middle of a word to respond to S3L's jibe (see Fig. 4).

Summary of attention The persistent media enabled the participants to split their attention between their own and their partner's work as well as attract their partner's attention by indicating the desire for attention or using the media. The commonality amongst all these interactions is that the participants are able to change their attentional focus at a finer granularity than individual contributions, interrupting even the production of words. Existing turn-taking models of interaction do not adequately account for this data.

Affective information

Since the members of each pair already knew each other, they had already developed interaction patterns unique to their relationship. These patterns were expressed in the

S4L is writing <star> Water Usage/Electricity
 15:30.46 S4R: How do you want to
 15:38.05 S4L taps S4L on the arm and goes back to writing
 15:39.23 S4R: write this proposal
 15:45.28 S4R taps more insistently on S4L's arm until she responds
 15:47.49 S4R: ?
 15:48.56 S4R: outline form.
 15:55.13 S4L: Yes
 15:57.04 S4R: B/C it looks really messy but I don't know if
 that matters
 16:15.62 S4L: nah, doesn't matter!
 16:17.70 S4R: <laughs silently>
 S4L points with both hands at Water Usage, but
 S4R reaches across S4L's hands, blocking them,
 and writes:
 16:25.22 S4R: <smiley>
 The discussion then switches to S4L's
 contribution.



Almost four minutes later, S4L looks for this interaction and writes:
 21:00.41 S4L: Perhaps we should make a box in a corner for the
 proposal.

Fig. 2 S4R interrupts herself to get her partner's attention. S4L returns to this context 4 min later to append a suggestion

artifact. The S3 pair was moderately competitive in having the last word or finding errors in each other's contributions. They repeatedly wrote the word "owned" to indicate a victory of this type over their partner (see Fig. 5a). At one point in their interactions, participants in the S3 and S6 pairs, both male, drew outlines of a hand with the middle finger extended (see Fig. 5b). All of these interactions were done in the context of the pairs' friendships, and were not intended or taken as harsh or critical.

S3L has written the title, "What the gov't would do in case of an alien attack"

20:45.94 S3L: I. NO INFORMATION
 20:51.92 S3R: ALIEN ENCOUNTER
 L stops writing and reads R's contribution.
 21:22.42 S3R: WE DON'T KNOW I[F THE]
 21:25.51 S3L: [NO ATTACK]
 21:31.32 S3R: Y ARE A
 21:39.26 S3L: OK. I SEE THINK I

L crosses out the word "attack" in the title and writes "encounter" above it.

21:51.31 S3L: OK?
 21:55.85 S3R: OWNED
 21:55.91 S3L: <crosses out 'owned'>

L then completes his sentence from 20:45:94:
 22:11.31 S3L: WOULD BE RELEASED



Fig. 3 S3L interrupts himself to read what his partner is writing. Alignment of text indicates timing of simultaneous actions

8:38.65 S3R: NEED \$\$ TO PAY FOR [REN]
8:56.42 S3L:
8:59.55 S3R:

<crossing out> T




Fig. 4 The S3 pair’s interruption protocol becomes finer-grained

Some expressions of affect were explicit; the S6 pair drew stars and the S4 pair used thumbs up gestures to indicate assent, and the S4 and S5 pairs drew several happy faces at the point of an agreement. Interestingly, the only audible communication that the pairs seemed unable to constrain was laughter, and all of the pairs at some point exchanged some form of stifled mirth.

The nature of the environment also allowed the participants to indicate their affective state within the same actions they took to manipulate the physical environment, and this was the most frequent way in which the participants communicated their feelings. Most of the pairs altered the enthusiasm with which they would write or identify contributions, and changed the rate or intensity of how they tapped on the table to indicate levels of excitement or frustration. During problem two, S3L became more and more frustrated with S3R’s inability to generate new ideas, and repeatedly wrote “next topic” while tapping on the table with increasing intensity and volume. Exchanges such as this not only made the affect of each participant visible, but functioned to define the roles of each participant.

Summary of affect Expressions of affect were an integral part of individuals’ contributions that resulted in relational terms or (more rarely) symbolic representations of affect being incorporated directly into the task-related content. Participants also used dimensions of rate (e.g., underlining or tapping) or intensity (e.g., volume) afforded by the media to express affect.

Conclusions on interaction

The exchanges demonstrated repeatedly that task and affective information is communicated simultaneously and continuously. Each of the participants’ actions were extremely dense composites of multiple information types, and these actions overlapped and interacted with each other. The complexity of the interactions, and of the resulting artifact, required that the participants pay explicit attention to their partner’s contributions as well as that they explicitly direct their partner’s attention.




Fig. 5 S3 participants’ explicit expressions of affect

Use of materials

Interactions are expressed in and through the materials and physical space in which the activity takes place. The sets of materials were chosen for specific affordances they did and did not supply (see Table 1). The smaller paper materials could be moved around, turned to face any direction, and could be lined up next to each other or stacked. The large sheet of paper disallowed the ability to re-position or re-orient previous contributions, and the addition of post-its contributed the ability to attach one contribution to another. While the pairs' artifacts were obviously shaped by the materials being used, they rarely explored the unique capabilities of the media to any extent. Instead, the pairs used only a *limited set* of the materials that allowed them to manage their partner's *awareness* of their contributions, indicate *relationships* between individual contributions, create *groupings* of contributions and partition the workspace into *functional and conceptual spaces*. The following sections illustrate each of these phenomena with examples from the data.

Limited, polymorphic repertoire

Despite the availability of a wide variety of materials, every pair constrained themselves to a very *limited subset of the materials*. Even when the widest variety of materials was made available, all pairs tended to use only one or two material types almost exclusively. New types of materials were generally introduced to distinguish types of information. For example, in a discussion entirely on 3×5 cards a pink post-it was used to record conclusions.

In addition to limiting their choice of materials, all the participants made use of a remarkably *limited set of gestures and deixis*. Women were considerably more likely to use hand gestures than men. Still, these gestures were mostly limited to an approval gesture (e.g., thumbs up) and a questioning gesture (e.g., hands spread, palm up, like what usually accompanies a shrug). Actions that related to the artifact consisted of a variety of pointing gestures. Participants pointed at, tapped on, or ran a finger along (under) artifact elements. Gestures involving both hands were also used, e.g., indicating some written text with one hand, either pointing, tapping, or underlining, and then tapping the area being used to record conclusions with the other hand.

A similarly *limited number of symbols* were employed in the discussions. Symbols were almost never used independently, but almost always as an annotation or in reference to some other piece of text. The symbols used regularly were question marks (by far the most prevalent), arrows, smiles, stars, and check marks. Two different males each drew representations of a hand with the middle finger extended, but only once each. Despite a variety of labels and stickers, pairs only made use of stars and colored dots.

The limited repertoire required that actions, gestures, and symbols be *polymorphic*, meaning that individual actions or symbols carried multiple potential meanings and could only be understood in context. Pointing, for example, might be used to indicate suggested topics, related information, reminders, request for clarification, or illegible handwriting. Participants used underlining to indicate emphasis while they were writing, but also to indicate repetition (described below). The meanings of question marks and other symbols were similarly context-sensitive. A specific and recurring example was the use of tapping and underlining to draw the partner's attention to some text. When a participant wanted to start a new topic, the writer would produce the introductory sentence, get their partner's attention, and then tap or underline the new text. This was often followed by the writer running a finger under the text, indicating what the partner should read. When a participant

wanted to return to a previous topic, the order of actions was reversed. The participant would indicate previously written text by tapping on it or underlining it, and then write the related contribution. Remarkably, the participants seemed to have almost no problem correctly interpreting each others' actions, gestures, and symbols.

Summary of repertoire Participants consistently used a limited repertoire of simple, flexible tools and applied them in polymorphic ways to their collaborative activity.

Managing awareness

Collaboration requires managing mutual awareness of each other's contributions. Awareness was identified as a broader category than attention, since participants were observed to be peripherally aware of many actions and workspace elements that were outside their current focus of attention. For example, members of the S1 pair were aware of the index cards that had been placed near them while they were writing, and demonstrated this with the immediate transition to them once their writing was complete.

Awareness is a prerequisite for attention, and the participants' practices for increasing and decreasing awareness can be seen as intentionally altering the likelihood of something becoming the attentional focus. Participants took responsibility both for maintaining their awareness of their partner's contributions and making their partner aware of their own contributions. Three primary conventions were used for this purpose: (1) Using the movable materials, a participant would write a contribution and then move the material over to the partner's side of the table, near where the partner was attending. Participants maintained a peripheral awareness of these materials and consistently incorporated them into their activity. (2) Participants would make a contribution temporally proximal to their partner's previous contribution. When participants were both attending the same material the pair would alternate making contributions, reasonably certain that the partner was paying attention. Interestingly, these contributions did not have to be spatially proximal. Both pairs using the single sheet of paper would sometimes alternately write contributions on the far side of the paper from themselves (closest to their partner), seemingly to make it easy for their partner to read the contribution. (3) Participants would get their partner's attention before making a contribution. This was accomplished by tapping on the tabletop or the partner's arm and was indicated when the partner stopped writing and changed positions so their body language indicated they were paying attention.

Using the movable materials, position was also used to *reduce* awareness of texts. After a thread had come to conclusion, or the pair was otherwise done using a particular material, it would be moved off to the side where it was less likely to become the focus of attention. S2R kept the collection of 3×5 cards he'd written in a stack under the card he was currently writing on, making it difficult for his partner to refer to previous contributions.

Summary of awareness Participants used spatial and temporal affordances to make their partner aware of their contributions. If the media could not be used directly, or there was greater urgency, awareness was managed through attention-getting mechanisms.

Indicating relatedness

Making a partner aware of a contribution was always accompanied by indicating how the new contribution related to existing texts. By far the most common method for indicating

relatedness was to position the new text close to the text to which it related. Pairs S1 and S6 kept related contributions on a single piece of material and the S3, S4, and S5 pairs regularly wrote new contributions near related text, but with little regard for the resulting orientation of the text.

Since those in the S2 pair were each maintaining their own materials, they could not make use of spatial proximity, but instead relied on temporal sequence. In one case, S2R indicated agreement by repeatedly underlining the word “yes” (approximately every 2–3 s) until his partner noticed. In this case the creation of the underline response was insufficient on its own and required more synchronized timing of reading and writing between the two partners. The S3 pair also used temporal sequence to indicate the relationship between contributions made on opposite sides of the large sheet of paper.

When the simpler mechanisms of proximity or temporal sequence could not be used (e.g., if there was no space to add a proximal contribution), participants relied on symbolic indications. In problem one, when S1R ran out of room on a card he drew an arrow along the bottom of the card and continued his contribution on a second card. When he was finished writing he moved the two cards together across the table, maintaining the spatial relationship between the two. The S4 pair also occasionally used arrows and other types of connecting lines (see Fig. 6). Near the end of problem one, the S6 pair produced lists of issues and responses and made a point of lining the two lists up next to each other.




Fig. 6 S4 participants' use of space to group threads and their division of the worksheet into discussion area and conclusions

Summary of relatedness Indications of the existence of relationships were intrinsic to the act of contributing. Relatedness was almost always indicated by way of some kind of similarity, whether it was spatial or temporal proximity, or similar visual attribute (e.g., color, size, alignment), although symbolic representations such as arrows were sometimes used.

Grouping

In addition to indicating relationships between representational elements, the participants also created groups of representational elements. Groups differ from other relationships in that a group has an identity as a representational element that is in addition to the identities of its members. Some indications of groups were implicit in the use of materials. The S1 and S6 pairs, using moveable materials, used individual sheets of paper to represent groups of contributions. Other groupings were expressed more explicitly in how the participants' utilized the materials. On the single sheet of paper, the S4 pair left distinct spaces around discussion threads (see Fig. 6) and the S3 pair labeled their workspace and conclusions (see Fig. 7). The S3 pair created the most explicit expression of groups by paper-clipping together collections of index cards.

Closer analysis of the texts indicates that the use of grouping was pervasive at multiple levels of detail. Spatial proximity was used to group contributions into threads or to explicitly define them as a collection (e.g., conclusions). In addition to proximity, participants would indicate the association of a contribution to an existing text by mimicking visual attributes of the existing text. Lists and multi-line contributions usually maintained a consistent left margin. Conversely, the S1 pair, even though they were using small index cards, distinguished their contributions to a thread by maintaining different margins.

Grouping, and its counterpart individuation, were used pervasively and also at multiple levels. Pointing, tapping, and underlining were used equally often to indicate a symbol or




Fig. 7 S3 participants write over the top of existing text in the work area rather than re-designate part of the mostly blank conclusions area

word, a few words from a contribution, a whole contribution or thread, multiple threads, an individual piece of material including the whole large sheet of paper, the whole workspace or the tabletop. The participants exhibited the ability to identify elements from the workspace simultaneously as a whole and as a collection of other elements.

Summary of grouping Elements of the workspace were fluidly collected together and broken up into constituent elements. The workspace was not treated as a collection of atomic data items, but as a field of information that could be fluidly reconfigured using explicit borders, spatial proximity, and typographic similarity or continuity.

Functional/conceptual spaces

Some groupings emerged in every interaction of every pair to fulfill common functions. Every pair used some mechanism to separate the text of the discussion from the text of the conclusions (e.g., Fig. 6). Once these spaces were designated, participants resisted altering their purpose, and in several cases pairs continued to squeeze the discussion text into a constrained area long after this became problematic. The S3 pair, using a single large sheet of paper, split the paper into a discussion area (~30%) and a results (~70%) area. When the discussion area had been completely filled with text, the results area was still almost completely unused. Rather than re-designate space from the results area, the participants used larger markers to write over the top of the existing text in the discussion area (see Fig. 7). This behavior was not unique. Many pairs constrained their discussion to an impractical writing area while leaving large areas designated for results unused.

All the pairs with movable materials used the space directly between them as the active workspace and set old or discarded materials off to the side. This organization was echoed in how the space directly in front of each participant was managed. When participants slid a card over to their partner's side of the table it was often placed to the side of the partner's current work, which the partner correctly interpreted as "pending."

Summary of functional/conceptual spaces Many of the functional spaces emerged from the specific natures of the interactions and came and went as they were needed. On the other hand, a few recurring conceptual spaces (i.e., personal space, shared work area, conclusions, and discard) were present in every interaction of every pair. Different affordances were drawn on to create spaces—some were explicitly indicated with boundaries or labels but many were defined implicitly through the participants' use of the physical space.

Conclusions on material use

The participants used only a small number of materials, gestures, and symbols, but allowed the meaning of any tool or action be derived from the informational and social context of its use. Participants used media affordances to manage each other's awareness of contributions and the connections between them. All the sessions used materials to indicate when two contributions were related, to indicate groups of contributions, and to constitute functional and conceptual spaces in the work area. Very few of the conventions used were discussed in advance or explicated in the workspace. Nonetheless, common conceptual spaces emerged in all the sessions.

Workspace organization

Both the macro-structure of the discussions (i.e., a linear series of topics culminating in a conclusion) and the micro-structure of the interactions (i.e., the practical aspects of action and material use) demonstrated a great deal of regularity. On the other hand, the workspaces developed in decidedly different ways. In all cases, though, the workspaces became complex, messy collections of disconnected and overlapping contributions. These artifacts appear extremely *disorganized*, and contain a great deal of *incomplete and obsolete content*. Nonetheless, the participants reported no difficulty using them to collaborate, and reported that the workspace had adequately supported their exchanges. Each of these phenomena is illustrated below.

Apparent disorganization

No organizational scheme was suggested to the pairs, and while conclusions were explicitly organized, none of the pairs imposed much structure on the rest of the artifact beyond some simple grouping (discussed above as *Use of Materials*). The result was complex, messy workspaces that reflected the discussions' unstructured progression from topic to topic. The different materials led to different kinds of apparent disorganization. The S1 pair generally covered the table with multiple 3×5 cards; some scattered and some loosely grouped (see Fig. 8a). The S3 and S4 pairs produced sheets of paper with text running in all directions, and often text was written over the top of other text (e.g., Figs. 6 and 7). The S6 pair had the same disarray as S1, but since S6 used 8.5×11 " sheets of paper, there was much more overlapping and piling of materials. S6 also made contributions by drawing on the tabletop, so an additional layer of contributions covered the table (see Fig. 8b). The S2 pair made the most conservative use of the materials, with each maintaining separate materials for their own contributions. Even so, this practice produced materials on which the ordering of the exchange was extremely difficult to recognize.

Summary of apparent disorganization At the end of each session, the workspaces appeared complex and messy to an outside observer. The participants, however, did not demonstrate any difficulty navigating and making use of them. The participants chose to rely on their involvement in the interaction rather than using the affordances of the media to structure the workspace.



Fig. 8 a S1 participants' and b S6 participants' organization of workspaces

Incomplete and obsolete content

Two common practices increased the complexity of the task-related inscriptions: littering the text with out-of-date annotations and failing to record important information about the pair's interactions and agreements.

The great majority of the text recorded in the artifacts was task-related, consisting of contributions of information and responses to contributions. However, the workspace was made much more complex by the inclusion of text containing process-related discussion and social interactions. For example, when participants could not read their partner's handwriting, or did not understand their partner's contribution, they would often annotate their partner's contribution with a question mark or underline the text in question. Their partner would respond by retracing the writing or by adding clarifying text. In any case, the (now obsolete) question marks and underlines remained part of the artifact. In one exchange, S1R wrote on cards so that the text was oriented toward his partner (i.e., upside down), and S1L responded with the sardonic comment "Stop writing upside down, nobody thinks it's cool." Several of the questions required pairs to produce a list of responses. Both the S3 and S6 pairs were explicit about the change from one list item to the next, writing "cancel" or "next topic" to move the discussion along. In their first problem, the S4 pair had an interaction about how to record their conclusions (see Fig. 2) that concluded with S4R interrupting S4L's writing to draw a smiley face. Many of these annotations were eventually scribbled over or had lines drawn through them to indicate a kind of deletion. Nonetheless, all of these annotations and process negotiations became part of the workspace record, cluttering the workspace with symbols and text that no longer served an obvious purpose.

In addition to contributions that became out of date, participants used a variety of practices that left only a partial record, or even no record whatsoever. Many process and personal interactions took place in gestures and other attitudinal indicators. Thumbs-up gestures were used to indicate agreement. Other hand gestures indicated indecision, e.g., one hand turned palm upwards, or both palms upward with a shrugging motion), disagreement (e.g., palm forward 'stop' gesture, or waving a hand side to side over the text that was being disagreed with), or a restatement of the conclusion (e.g., in problem 3, after S4 began agreeing on an integration of science and religion, S4R repeatedly made a gesture where she held her hands in front of herself with the palms facing each other and moved her hands until they touched). To tie two ideas together, a participant might tap or point at one text and then the other, draw an arrow or line from one text to another, or add a new text in proximity to an existing text. Some relationship was indicated in all of these cases, but the nature of the relationship was very rarely specified or recorded. The contributor relied on the partner's ability to infer the relationship's intended meaning. This seemed consistently successful in that partners rarely asked for clarification and contributors rarely corrected their partner's interpretations. However, this practice left important information unrecorded in the workspace.

Summary of content management While some contributions were expressed in the persistent paper medium, and others were expressed in the ephemeral gestural medium, the choice between these media did not correlate with whether the contribution was potentially important to subsequent interaction or only transiently relevant. By the end of the session, the workspace included obsolete information and did not have a record of the gestural contributions. However, the pairs reported that the extra complexity of the persistent record did not impair their ability to communicate.

Multiple contexts

In the course of the sessions, the workspaces often represented multiple loci of interaction. The S1 pair put different threads on different 3×5 cards, so the physical switch from one card to another required the conceptual switch from one topic to another. The same was true for the S6 pair with 8.5×11" sheets of paper. The S4 pair was required to do similar contextual shifts as they moved their focus from one area of the sheet of paper to another. The pairs would occasionally nest new contexts in the current interaction. As problem one was wrapping up, S1R brought the current conversational thread to a halt and injected a side conversation, as shown in Fig. 9.

The S4 pair's negotiation of the conclusion space in Fig. 2 was similarly nested in S4L's introduction of the Water Usage/Electricity topic (at 15:35.62 and 16:27.13). Remarkably, this interaction created a significant enough context that almost 5 min later (21:00.41) S4L explicitly looked for this exchange and added the last comment in Fig. 2.

Summary of multiple contexts Participants fluidly shifted activity between physical locations or pieces of paper to manage multiple contexts, including nested and parallel conversations.

Conclusions on workspace organization

Despite the complexity and seeming disorganization of the workspace artifacts, the participants had no apparent difficulty navigating the information and reported that the workspaces had adequately supported their interactions. The different exchanges that took place during the interaction were often scattered across the workspace, but despite this, the participants smoothly transitioned from the context of one exchange to another.

Interactionally negotiated conventions

Negotiation was a foundational aspect of the interactions and, like affect and process information, the negotiation itself was intrinsically interwoven with the pursuit of the task's goals. In very few cases were practices explicitly proposed. Instead, the pairs' practices emerged from their negotiation of acceptable action. Almost always, a practice was introduced through one participant enacting it, and then affirmed through their partner's adoption of the practice.

The participants' interaction constituted contexts and objects in the workspace as well as acceptable practices and social roles. Almost universally, conventions were constructed within the first few minutes of a problem session and displayed remarkable stability over the course of the interaction. During the interactions, conventions did evolve somewhat. This seemed to be a reflection of the participants becoming more comfortable with the

Fig. 9 The S1 pair nests one interaction inside another

At 16:10.47, S1L has written a final proposal with a short questionnaire.
 16:41.20 S1R gets a blank card
 16:43.17 S1L points insistently at his questionnaire
 16:43.61 S1R makes a 'wait' gesture, and then slides S1L's card to the side
 16:47.10 S1R: Wait
 The pair negotiates the meaning of the word 'alien', after which
 18:35.88 S1R checks the 'yes' box on L's questionnaire

environment. Major changes to a convention, or the adoption of new conventions only happened at the beginning of a new problem session.

Access and ownership

Some of the most interesting emergent behaviors involved the negotiation of access to and ownership of both the information and the physical materials. Changes in position and orientation indicated changes in who was allowed to write on a particular material as well as the kinds of contributions that were permitted. The changes made by one participant were rarely disputed by the partner, and the specific ways in which access was managed/granted reflected the existing social relationship of the pair.

All of the pairs maintained some way of indicating who had made each contribution. The S2 pair was the most extreme in that each participant had separate materials for their contributions. In all the other pairs, the participants consistently, and intentionally, used different colored pens. Although this effort was made to distinguish each other's contributions, no observable use was made of this information: after the contributions had been made, there was no explicit reference to whose they were.

Pairs using movable materials positioned materials in the space between themselves so as to mediate access, and individual contributions were consistently made with the material close to the writer (e.g., both pairs in Fig. 8). A standard pattern was for writers to position a card nearby, write a contribution, and then move the card to their partner's side of the table. The position also regulated the contributions the partner was allowed to make. When the material was in the center of the table, both participants would regularly write words and complete sentences on it, but when the material was closer to one participant, the partner limited their contributions to pointing, tapping, or writing individual symbols—most often a question mark. Even though the S2 pair did not move materials across the table, they imposed these same limits on themselves.

More collaborative interactions took place nearer the center of the table (see Fig. 8). When the two participants were engaged in negotiating an agreement, the interaction often took place on a single piece of paper positioned in the center of the table. In Fig. 8a, S1L is proposing a final conclusion on a card positioned in the center of the table. The shared paper could be temporarily turned so a participant could write easily, but participants usually returned the paper to a sidewise orientation that allowed each participant the same ability to read and contribute, and rarely moved the paper closer to their own side of the table. Papers used to record conclusions or final results also often stayed near the center of the table and usually remained oriented sideways. The S6 pair displayed this behavior most dramatically, maintaining a single collaborative thread and keeping almost all of their materials oriented sideways in the center of the table (see Fig. 8b).

This sideways orientation was also used in one of the large sheet of paper conditions to record the shared conclusions. In every problem the S4 pair consistently oriented their conclusions sideways on the paper even though their other contributions were generally made with no apparent regard to orientation. Once S4 started writing conclusions, the artifact in Fig. 6 was repositioned to orient the conclusions area sideways, even though this caused the paper to hang over the sides of the table. The S3 pair was one of the most asymmetrical pairs in terms of social roles, and the conclusions were usually oriented toward S3L (see Fig. 7), who was the more dominant partner. In problem two, S3L demonstrated the use of orientation as a proxy for ownership by explicitly re-orienting the sheet of paper and indicating that S3R should contribute to the conclusions.

Conclusions on interactionally negotiated conventions

Contributions to a collaborative discourse are not just context-sensitive, but constitute an active negotiation of accepted processes and social structure. Process conventions emerge and stabilize quickly. While they do evolve over time, significant changes in practice seem to require disengagement from the task-focused interaction. Additionally, the more subtle conventions of access and ownership of information are negotiated within the same interaction. Emergent conventions are reflexive in that they are enacted in the interactions they help to structure and, being so, they resist being defined a priori or independently of the interaction.

Conflict

The sessions were remarkably consistent in their high degree of collaboration, success of arriving at shared conclusions and construction of working social structures. The S5 pair was a marked exception. They displayed a high level of hostility towards each other and they failed to arrive at a conclusion to any of their three topics—the only pair to fail to do so. The two women had very different demeanors; S5L was very aggressive while S5R tried to be more conciliatory. S5L tended to produce either long, uninterrupted blocks of text or short, negative contributions such as “no” or “who cares?” with a considerable amount of underlining. Nonetheless, it is striking how the structure of their interaction and the conventions they developed compare to those of the other pairs.

The S5 pair demonstrated a linear discussion structure similar to the other pairs, but the pair rarely came to a conclusion or agreement on any topic. Most topics lasted only a few exchanges before S5L forced a topic change or took over the exchange. The only major exception was the beginning of problem one when S5L spent 20 min writing an essay, after which the pair did very little work independently, hardly ever writing simultaneously. During S5L's essay writing, S5R tried to interrupt several times, once by scribbling on the top of the essay page and twice by writing short phrases (“too complex” and “for me”) on S5L's problem sheet. S5L did not respond to these requests for attention, but later on she would tap loudly on the table or interrupt her partner's work to make sure that her own requests for attention were responded to immediately.

The pair's conventions did not achieve any level of stability. Their contributions sprawled across multiple different kinds of paper, post-its, and the problem sheet, and the pair never settled on a specific place for conclusions. S5R tried to designate such a space in the second problem, but S5L did not adhere to the designation. S5R tried several times to change the confrontational nature of the discussion. At the beginning of the second problem she started writing her contributions on post-its, explaining later that she thought the smaller size would force her partner to make shorter responses. At one point in problem two she resorted to holding her non-writing hand over the material she was writing on to keep S5L from taking the paper or from writing on it. This was only temporarily successful.

The pair did make use of a limited and polymorphic repertoire like the other pairs, but they exhibited significant differences in how they employed the grouping of contributions and the management of access and ownership. A typical pattern of exchange was for S5R to position a piece of paper in the center of the table on which the pair would make several alternating contributions, at which point S5L would become agitated and act in one of two ways. One typical action was to turn the material over, or introduce new material and start a new thread. In addition to disallowing S5R access to the previous contributions, this

disallowed any subsequent contributions from being positioned in proximity to the previous ones, thereby ending the thread. S5L's other typical action was to pull the paper over to her side and write a long response that filled the remaining area on the paper. S5R was disallowed from contributing during this time. This had the same effect in that it made the thread inaccessible to S5R, and once the paper was filled with writing it disallowed future contributions from being positioned in proximity to the text on the paper. In the normal course of the interaction, S5L would often write overly large text with multiple underlines and circlings in what seemed like an attempt to fill as much space on the paper as possible.

While S5R followed the common convention of only making small contributions to material on her partner's side of the table, S5L would often reach across the table and write whole sentences. Several contentious exchanges in problem two take place entirely on S5R's side of the table with S5L stretching her torso across the table to write. S5L's violations of the spatial conventions observed in all other sessions were a direct expression of her aggressiveness.

In summary, some regularities and conventions observed in other sessions were seen in S5, and others were not. We can make sense of these observations by noting that those results consistent with other sessions (e.g., limited repertoire, polymorphism) pertain to basic mechanisms of communication in this medium, while those results that differ from other sessions (e.g., ignoring interruptions, violation of spatial conventions of ownership) are precisely how aggression and anger were expressed. The meanings of the conventions remained the same, but the S5 pair demonstrated their conflict by consistently acting in opposition to these conventions. In this sense, the exceptions of S5 "prove the rule."

Discussion

On the surface, the six pairs' interactions were very different. The interpersonal interactions varied from outright hostility to more subtle power dynamics to egalitarian collaboration. The artifacts that resulted were equally varied—a result of both the social dynamics and the materials provided. However, at a deeper level, the pair's interactions were remarkably similar. Rather than explore the unique capabilities of the workspace, the different pairs appropriated different social and physical affordances of their environment to enact functionally equivalent practices. The consistency in the pairs' practices indicates a common need for the functions these practices provided. Even in the conditions where the materials did not support the most effective of these practices, participants did their best to recreate the functions, accepting second-rate approximations (see Tables 3, 4, and 5). The left columns in Tables 3, 4, and 5 list the communicative functions that were common across all the variability introduced in the study. The right columns provide examples of the different practices by which each communicative function was enacted with the different sets of materials provided.

These observations might lead one to the conclusion that the specifics of the environment make little difference. It is possible that given any environment, people will create tools appropriate to the environment, adapting to the limitations of their tools and "making do" with whatever affordances are available in order to perform important communicative tasks. In this process, however, people will often settle on "good enough" tools that may in fact be suboptimal (Galantucci, 2005). Conversely, it might be argued that we could determine the "best" affordances, and simply ensure that any collaborative environment supported them. There are pitfalls to both of these approaches. The first may

Table 3 Consistent communicative functions (interaction)

Function	Materials	Enacted practices
Attention	Single large sheet of paper	Tap on their partner's arm Tap or point at text or materials Mime or draw an underline
	3×5 index cards and unrestricted materials	Move materials over the top of their partner's current work Hold materials up for their partner to read Make their partner aware of the material, but wait for the partner to shift attention
Awareness	Single large sheet of paper	Wait for their partner's attention before writing Write on the side of the sheet closest to their partner Tap or point at a contribution repeatedly
	3×5 index cards and unrestricted materials	Move materials closer to their partner Orient materials to face their partner Move materials off to the side Wad up materials and remove them from the workspace
Access and ownership	Single large sheet of paper	Orient the text of contributions (write upside down) Orient conclusions sideways Offer the pen for writing conclusions to their partners Reorient the sheet of paper
	3×5 index cards and unrestricted materials	Move materials closer to or farther from their partner Reorient materials to face their partner Maintain sideways orientation of the conclusions Square up a stack of cards.

Communicative functions were enacted with a variety of practices. The practices enacted with 3×5 index cards and unlimited materials were similar enough that they were consolidated

provide insufficient guidance or structure, while the second runs the risk of crippling the negotiation of practices and conventions. The two approaches must be balanced against each other.

Commonalities between the pairs point the way to invariant aspects of collaboration that might provide a simplified model of the complexities of interaction and so effectively inform software design. Too rapid a move to implications for design, however, dismisses much of the value of these observations as a lens through which to critique the current assumptions that underlie collaborative technology design (Dourish, 2006). Instead, the discussion below revisits topics from our observations and discusses how some of these commonalities provide new ways of defining design issues and possibly indicate new conceptualizations of the nature of collaborative technologies.

Content

Although our casual observation is that the content of the pairs' interactions and conclusions was unremarkable, the management of this content and the relationships between the content and the representation is representative of the remarkable practices of every collaborative interaction.

The term 'content' is a gloss for the complex collection of private knowledge, publicly stated information, external data, and the multiple interpretations that are applied to them. The short duration of this study allowed discussions to rely on the individual knowledge

Table 4 Consistent communicative functions continued (material use)

Function	Materials	Enacted practices
Relatedness	Single large sheet of paper	Add new contributions in proximity to previous contributions Draw arrows or lines between contributions
	3×5 index cards and unrestricted material	Put contributions on the same material Create numbered lists Use consistent left margins Align or reposition materials together
Grouping	Single large sheet of paper	Draw a boundary or circle contributions Add titles to areas Use the same writing tool for elements of the group
	3×5 index cards and unrestricted materials	Pile or stack materials Paperclip materials together Give members of the group similar annotations (stars, numbers) Use unique materials
Functional and conceptual spaces	Single large sheet of paper	Explicitly divide the workspace Use the problem sheet, workspace, and tabletop for different purposes
	3×5 index cards and unrestricted materials	Move materials to and from the center of the table Move materials off to the side
Multiple contexts	Single large sheet of paper	Group related contributions Use the problem sheet, workspace, and tabletop for different purposes
	3×5 index cards and unrestricted materials	Introduce new materials Add related contributions to previous materials

Table 5 Consistent communicative functions, continued (interactionally negotiated conventions)

Function	Materials	Enacted practices
Access and ownership	Single large sheet of paper	Orient the text of contributions (write upside down) Orient conclusions sideways Offer the pen for writing conclusions to their partners Reorient the sheet of paper
	3×5 index cards and unrestricted materials	Move materials closer to or farther from their partner Reorient materials to face their partner Maintain sideways orientation of the conclusions Square up a stack of cards.

that the participants brought with them. Longer sessions, or more complex requirements, may have required the explicit management of external representations of information.

How should software manage these different aspects of ‘content’? Most designs assume that the representation contains the entire content of the interaction. This study indicates that there is content that plays a part in interactions yet is not explicitly represented. Furthermore, new collaborative technologies might create new definitions of information that further extend our understanding of what constitutes the content of an interaction. This is one example of how collaboration technologies could go “beyond being there” (Hollan & Stornetta, 1992) and provide environments that exceed current interaction media.

Interaction

Like FTF conversation, artifact-mediated interaction is much finer grained than the individual contributions. The observations from this study challenge the atomicity of individual contributions. Interaction is less an exchange of contributions than an ongoing simultaneous production of them. Users need to be able to produce their own contributions simultaneously with their perception of others’ contributions. This work indicates that the production itself often carries important information about the meaning and purpose of the contribution.

The term *contribution* itself may imply an inappropriate degree of chunking. Study participants regularly interrupted themselves and each other to respond to something in the workspace. Users may be producing multiple contributions at any one time as well as responding to others’ partially complete contributions or providing an ongoing response to the production process (Sacks, Schegloff, & Jefferson, 1974). Participation in an interaction is a dense composite of task, affective, and social information (Bronckart, 1995; Whitworth, Gallupe, & McQueen, 2000). It may not be worthwhile, or even possible, to separate the channels of information since the meaning of the contribution is usually derived from the interpretation of all of the channels simultaneously.

The completion of a contribution is only a single point in the process of perceiving and responding, and yet many collaboration technologies rely on chunking contributions (e.g., e-mail messages, discussion postings, or instant messages). Is there a way to increase the granularity and density of users’ participatory actions? The complexity of the evolving interaction as well as its component contributions already requires more explicit management of awareness and attention, but people have demonstrated the ability to manage larger numbers of contexts in text chat or using e-mail (Herring, 1999; O’Neill & Martin, 2003). It may be that increasing the granularity of the interaction would situate participants more fully in the interaction and *increase* their ability to handle multiple contexts of interaction.

Representation

The design of the study was partially motivated by the pair of implicit assumptions that the participants would make full use of the wide range of tools provided to them and that their entire interaction would be recorded in the workspace. To the contrary, the participants consistently used only a few of the tools and the persistent artifacts are at best a partial record of complex and subtle interactions. The participants used the minimal record in the workspace to manage the much larger body of historical and current task, affect, process, and role information that informed their discussion.

Observations from this study challenge the need for explicit, concrete representational structures. Participants used only a small number of the materials provided, and even these in fairly limited ways. In general, people use representations to anchor their discourse, not mirror it, and representations may have multiple (even conflicting) meanings. People create their own rules for coherency, and they develop their own, context-sensitive meanings for the tools they use.

What is the correct level of specificity for software tools? These observations seem to imply that simple tools with multiple possible semantics might better serve users. Care has to be taken, however, that the user is not required to explicitly specify the semantics of the tools but is allowed to enact the semantics of the tool through its use. A tool's effectiveness in the interaction comes partially from its lack of specification and partially from its ability to carry meaning.

Emergent practices

Negotiation through interaction is a foundational component of collaborative discourse. Pairs interactionally negotiated conventions for using the materials, assigning and playing roles, and furthering the conversation. These interactions are difficult to predict or script—negotiation emerges through users' interactions and is implicitly proposed and taken up. What is the role of scaffolding or scripting in an environment that facilitates emergent practices? Too much structure can undermine the development of situated practices, but collaborators may adhere to "good enough" practices that can render their interaction less effective or more cumbersome.

Designing for semiotic resources

One approach to supporting open-ended, flexible negotiation practices would be to de-emphasize the computer's role as mediator and instead conceptualize the computer as a medium that facilitates the users' interactions with each other (Suthers, 2006). The users are fundamentally engaged in creating an intersubjective understanding with each other; a process that should be supported with adequate semiotic resources (Goodwin, 2000). Our analysis has identified several categories of semiotic resources, including:

Multiple loci for interaction in a persistent medium visible to all participants, enabling maintenance of multiple contexts, simultaneous production, and fine-grained interruptions and context switching.

Association of regions of the shared space with individuals, enabling management of awareness, attention, and ownership through orientation and placement with respect to "personal space."

The ability to vary rate and intensity in a manner integrated with production of contributions, enabling expression of affect, urgency, or other dimensions simultaneous with the literal content.

Variations in attributes of inscriptions (e.g., color, thickness, style, size, and typography), enabling expression of relatedness, grouping, ownership, and other distinctions through selection of similar and different attributes.

These resources provide an alternative factoring of the enacted practices shown in Tables 3, 4, and 5; one that groups practices by the semiotic value of affordances relied on rather than by communicative function. Our ongoing work is exploring the role of cate-

gories, such as the above, to design sets of affordances in new media that offer sufficient degrees of freedom for expression and negotiation of communicative conventions.

Conclusions

This study was conducted in an environment significantly different from either typical conversation or typical on-line communication. Face-to-face conversation is generally not persistent, and relies heavily on a wide range of non-verbal cues. Online environments provide a structured set of tools to facilitate communication, but are often over-constrained by the model of communications embedded in the tools. This study environment attempted to marry the limited communication channels of the online environment with the flexible representational abilities of pen and paper in order to uncover strategies participants used in appropriating affordances of the written media. Participants' level of engagement suggests that this marriage was successful.

Analyzing interaction from the bottom-up gives a sense of how people act "naturally" and provides a new lens through which to examine the assumptions at play in the design of collaborative systems. The data gathered from this study shows a great deal of consistency at the micro-level of artifact-mediated communication. As an alternative to creating more complex group cognition models, software design should instead focus on the invariant aspects of practices that emerge independent of the specifics of the interaction. This work demonstrates the generality and practical value of these invariants, and shows they provide an empirically grounded and tractable model of interaction. Furthermore, invariants can be discovered by studying the concrete and observable consistent practices of participants in a collaborative discourse. The question shifts from understanding why people did something to documenting how they appropriated the material environment to do it.

This study should not, however, be taken too literally as a design for an online environment. Some of the behaviors observed are deeply tied to the physicality of the situation, e.g., managing placement and orientation of materials relative to participants' locations in a shared space, using both hands for gesturing, or touching to get one's attention. A direct implementation of this environment would have difficulty reproducing this physicality and at the same time fail to take advantage of abilities afforded by the electronic medium (Dillenbourg & Traum, 1999; Hollan & Stornetta, 1992). Instead, we should recognize that collaborative practices produce powerful, context-specific mechanisms and we should create software environments that cultivate them. At a deeper level, this study points the way to studying interaction at the level of invariants of communicative practices, which is a promising re-conceptualization of how interaction is accomplished in terms of properties that span differing media, topics, and participants.

Acknowledgements This work was supported by the National Science Foundation under award #0093505.

References

- Bronckart, J. P. (1995). Theories of action, speech, natural language, and discourse. In J. V. Wertsch, P. D. Rio, & A. Alvarez (Eds.), *Sociocultural studies of mind* (pp. 75–91). New York: Cambridge University Press.

- Clark, H. H., & Brennan, S. E. (1991). Grounding in communication. In L. B. Resnick, J. M. Levine, & S. D. Teasley (Eds.), *Perspectives on socially shared cognition* (pp. 127–149). Washington, DC: American Psychological Association.
- Conklin, J. (2003). Dialog mapping: Reflections on an industrial strength case study. In P. A. Kirschner, S. J. B. Shum, & C. S. Carr (Eds.), *Visualizing argumentation: Software tools for collaborative and educational sense-making* (pp. 117–136). Berlin, Heidelberg, New York: Springer.
- Dillenbourg, P. (2005). Designing biases that augment socio-cognitive interactions. In R. Bromme, F. W. Hesse, & H. Spada (Eds.), *Barriers and biases in computer-mediated knowledge communication-and how they may be overcome* (pp. 243–264). Berlin, Heidelberg, New York: Springer.
- Dillenbourg, P., & Traum, D. (1999). The long road from a shared screen to a shared understanding. In C. Hoadly & J. Roschelle (Eds.), *Proceedings of the 3rd conference on computer-supported collaborative learning* (pp. 12–15). Stanford.
- Dourish, P. (2006). Implications for design. In *Proceedings of the SIGCHI conference on Human Factors in computing systems* (pp. 541–550). Montréal, Québec, Canada: ACM.
- Finn, K. E., Sellen, A. J., & Wilbur, S. B. (Eds.) (1997). Video-mediated communication. Mahwah, NJ: Erlbaum.
- Galantucci, B. (2005). An experimental study of the emergence of human communication systems. *Cognitive Science*, 29, 737–767.
- Garfinkel, H. (1967). *Studies in ethnomethodology*. Englewood Cliffs, New Jersey: Prentice-Hall.
- Garfinkel, H., & Sacks, H. (1970). On formal structures of practical actions. In J. C. McKinney & E. A. Tiruakian (Eds.), *Theoretical sociology: Perspectives and developments* (pp. 338–366). New York: Appleton, Educational Division.
- Gerbner, G. (1956). Toward a general model of communication. *Audio Visual Communication Review*, IV(3), 171–199.
- Gibson, J. J. (1977). The theory of affordances. In R. Shaw & J. Bransford (Eds.), *Perceiving, acting, and knowing*. Hillsdale, NJ: Erlbaum.
- Glaser, B. G. (1992). *Basics of grounded theory analysis: Emergence vs. forcing*. Mill Valley, CA: Sociology.
- Glaser, B. G., & Strauss, A. L. (1967). *The discovery of grounded theory: Strategies for qualitative research*. Chicago, IL: Aldine.
- Goldin-Meadow, S. (1999). The role of gesture in communication and thinking. *Trends in Cognitive Sciences*, 3(11), 419–429.
- Goldin-Meadow, S., & Feldman, H. (1977). The development of language-like communication without a language model. *Science*, 197(4301), 401–403.
- Goldin-Meadow, S., & Mylander, C. (1998). Spontaneous sign systems created by deaf children in two cultures. *Nature*, 391(15), 279–281.
- Goodwin, C. (2000). Action and embodiment within situated human interaction. *Journal of Pragmatics*, 32, 1489–1522.
- Goodwin, C., & Heritage, J. (1990). Conversation analysis. *Annual Review of Anthropology*, 19, 283–307.
- Healey, P. G. T., Swoboda, N., Umata, I., & Katagiri, Y. (2002). Graphic representation in graphical dialogue. *International Journal of Human-Computer Studies*, 57, 375–395.
- Heath, C., & Luff, P. (1991). Disembodied conduct: Communication through video in a multi-media office environment. In *Proceedings of the SIGCHI conference on human factors in computing systems: Reaching through technology* (pp. 99–103). New York: ACM.
- Herring, S. C. (1999). Interactional coherence in CMC. Paper presented at the International Conference on System Sciences, Maui, HI.
- Hollan, J., & Stornetta, S. (1992). Beyond being there. Paper presented at the Conference on Human Factors in Computing Systems, Monterey, CA.
- Jakobson, R. (1960). Closing statement: Linguistics and poetics. In T. Sebeok (Ed.), *Style and language*. Cambridge, MA: MIT.
- Kaput, J., & Hegedus, S. (2002). Exploring classroom connectivity by aggregating student constructions to create new learning opportunities. Paper presented at the 26th Annual Conference of the International Group for the Psychology of Mathematics Education, UK.
- Kato, H., Yamazaki, K., Suzuki, H., Kuzuoka, H., Miki, H., & Yamazaki, A. (2001). Designing a video-mediated collaboration system based on a body metaphor. In T. Koschmann, R. Hall, & N. Miyake (Eds.), *CSCL 2: Carrying forward the conversation* (pp. 409–423). Mahwah, NJ: Erlbaum.
- Krauss, R. M. (1998). Why do we gesture when we speak? *Current Directions in Psychological Science*, 7, 54–59.
- Lingnau, A., Hoppe, H. U., & Mannhaupt, G. (2003). Computer supported collaborative writing in an early learning classroom. *Journal of Computer Assisted Learning*, 19(2), 186–194.
- Mark, G., & Abrams, S. (2005). Differential interaction and attribution in collocated and distributed large-scale collaboration. In Proceedings of the 38th Hawai'i International Conference on System Sciences. Hawai'i: IEEE.

- Olson, G. M., & Olson, J. S. (2000). Distance matters. *Human-computer Interaction*, 15(2/3), 139–178.
- Olson, G. M., & Olson, J. S. (2002). Groupware and computer-supported cooperative work. In J. A. Jacko & A. Sears (Eds.), *The human-computer interaction handbook: Fundamentals, evolving technologies and emerging applications* (pp. 583–595). Mahwah, NJ: Erlbaum.
- O’Neill, J., & Martin, D. (2003). Text chat in action. In *Proceedings of the 2003 international ACM SIGGROUP conference on supporting group work* (pp. 40–49). Sanibel Island, FL, USA: ACM.
- Rittel, H., & Webber, M. (1973). Dilemmas in a general theory of planning. *Policy Sciences*, 4, 155–169.
- Sacks, H. (1984a). Notes on methodology. In J. M. Atkinson & J. Heritage (Eds.), *Structures of social action* (pp. 21–27). Cambridge: Cambridge University Press.
- Sacks, H. (1984b). On doing “being ordinary”. In J. M. Atkinson & J. Heritage (Eds.), *Structures of social action* (pp. 413–429). Cambridge: Cambridge University Press.
- Sacks, H., Schegloff, E. A., & Jefferson, G. (1974). A simplest systematics for the organization of turn-taking for conversation. *Language*, 50(4), 696–735.
- Selvin, A. M. (2003). Fostering collective intelligence: Helping groups use visualized argumentation. In P. A. Kirschner, S. J. B. Shum, & C. S. Carr (Eds.), *Visualizing argumentation: Software tools for collaborative and educational sense-making* (pp. 137–163). Berlin, Heidelberg, New York: Springer.
- Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27, 379–423.
- Shipman, F. M., III, & McCall, R. (1994). Supporting knowledge-base evolution with incremental formalization. In *Proceedings of the SIGCHI conference on Human factors in computing systems: celebrating interdependence* (pp. 285–291). Boston, MA: ACM.
- Sugimoto, M. (2003). How sensing and mobile technologies can enhance collaborative learning in classrooms and museums. Paper presented at CSCL 2003: Designing for Change in Networked Learning Environments, Bergen, Norway.
- Suthers, D. (2006). Technology affordances for intersubjective meaning-making: A research agenda for CSCL. *International Journal of Computers Supported Collaborative Learning*, 1(3), 315–337.
- van Bruggen, J. M., & Kirschner, P. A. (2003). Designing external representations to support solving wicked problems. In J. Andriessen, M. Baker, & D. Suthers (Eds.), *Arguing to learn: Confronting cognitions in computer-supported collaborative learning environments*. Dordrecht: Kluwer.
- Vygotsky, L. S. (1978). *Mind in society*. Cambridge, MA: Harvard University Press.
- Whitworth, B., Gallupe, B., & McQueen, R. (2000). A cognitive three-process model of computer-mediated group interaction. *Group Decision and Negotiation*, 9, 431–456.