Handheld-Mediated Communication to Support the Effective Sharing of Meaning in Joint Activity

Kibum Kim, Deborah Tatar and Steve Harrison
Center for Human-Computer Interaction
Virginia Tech, Blacksburg, VA 24060, USA
{kikim, tatar, sharrison}@cs.vt.edu

Abstract

This paper reports the results of a laboratory study of a proposed solution to the problem of coordinating entry into collaborative activity while protecting the privacy and control of individuals over small-scale devices. The particular set of design requirements we work with are abstracted away from prior work on the constraints and affordances of a middle school teaching/learning setting. We implement on handheld devices a proposed functionality, Look*, intended to enable a side assistant to see what is happening in an activity without interrupting that activity. The question is whether Look is sufficient to allow effective comprehension of the situation and effective action based on that comprehension. To evaluate the success of this functionality, we draw on preexisting psycholinguistic/CSCL (computer supported collaborative learning) theory on (1) shared visual spaces and (2) overhearer comprehension. The associated experimental paradigm increases the occurrence of the interpersonal challenges we wish to study over their spontaneous rate in the classroom. Our variant of the paradigm provides preliminary evidence that our design is viable. It also extends the paradigm by casting light on which elements of the situation were responsible for results in prior experiments, and which may be further amenable to technological manipulation.*

1. Introduction

The problem of "side-assistant sufficiency" (SAS) occurs when two people are engaged in a primary task involving face-to-face focused interaction, using handhelds as well as speech, and a third person, the side assistant (SA) needs to monitor or join the interaction. In the classroom, this problem has two forms: "formative assessment", in which a teacher attempts to determine whether the activity is happening in an appropriate and sufficient manner and to intervene if necessary, and "peer-monitoring", in which a peer attempts to learn what is happening; however, the problem is not confined to education or to beginning with two-person activities. At its most general, the problem can arise whenever one or more people are focused on information that cannot be seen by a SA, making it difficult for the SA to gauge interruption and raising the effort involved in attaining sufficient common ground for informed participation.

We propose and evaluate a particular design solution, "Look." Within the constraints of the task and the media, Look makes physical co-presence of workspaces available to side assistants. Our research is motivated by situations that arose in real classroom use of the handhelds, that is, occasions on which a student had reason to want to know what other students were doing or when a teacher was trying to assess the utility of an on-going activity. Prior research on referential communication has often utilized experimental situations that create communication challenges for participants in a more condensed way than they typically occur spontaneously [4,7,8,11,12]. Therefore, we abstracted away from the original situation to create a more focused experimental investigation.

2. Shared visual workspaces

As Clark and Brennan long ago noted [6, 2, 1], different communication media put different constraints on the grounding of information. Clark and Brennan [6], Kraut, Fussell, Brennan and Siegel [11], Clark and Krych [3], and Kraut, Gergle and Fussell [12] provide both theory and experimental evidence that a shared visual space helps efficient communication and grounding of information. The underlying theory describes the need that participants have to create and maintain joint focus on the mental

Fourth IEEE International Workshop on
Wireless, Mobile and Ubiquitous Technology in Education (ICHIT'06)
0-7695-2723-X/06 $20.00 © 2006 IEEE

and physical objects of collaboration and the processes that they go through to maintain this state [5]. Clark and Krych [3] show that monitoring an addressees' workspace during a task involving referential communication is associated with an eighth the errors, and half the time needed for the work as not monitoring the workspace. In another task involving shared reference, Kraut, Gergle and Fussell conclude that:

> Delaying the visual update in the space reduces benefits and degrades performance. The shared visual space is more useful when tasks are visually complex or when actors have no simple vocabulary for describing their world. We find evidence for the ways in which participants adapt their discourse processes to their level of shared visual information. ([12], p. 31)

They conclude this from an experiment in which a Helper and a Worker communicate using a high quality audio link and varying degrees of shared visual space implemented on a desktop machine. A three second delay in the update of the visual space between helper and worker causes significant degradation in task completion time compared to immediate update. We note that, in a "foible-to-feature transform," the learning activity illustrated in previous projects [9,13, 14, 17, 19] utilizes differences between the views of two students to create a learning opportunity. The mental and verbal work that the students do to resolve reference to the functions on their screens and to make it clear to one another what they are talking about causes them to deepen their understanding of math representations.

However, the challenge for a side assistant (SA) is different. The SA needs to understand what participants are doing with as little impediment as possible. Furthermore, related research, motivated by the same theoretical basis in joint action, shows that there is also a cost to being an outsider or overhearer to an ongoing conversation. When two people are engaged in a referential communication tasks, a third connected by video is at a substantial disadvantage across repeated trials in decoding what the central participants are talking about [7].

Two gaps in theory follow from this prior work. The first has to do with the relationship between the cost of not sharing a visual space in focused interaction, and cost of being an overhearer as compared to a central participant. The research shows that both present problems, but leaves open how the problems related: are they (for example) cumulative or does one dwarf the other? The second question is the relationship between the findings about overhearers, who are completely isolated from participation, and SAs, who may participate but not interrupt or distract from the central learning goals. Together, the lack of theory in these areas leaves open design questions.

Our current work attempts to cast light on these questions by investigating a particular design solution. The proposed handheld network service, "Look," enables a coherent activity in which participants attempt to understand what others are talking about—in other words, what they are learning.

Figure 1. Pop-up windows notifying the user that IR communication has been initiated.

3. Handheld implementations

Some special conditions obtain as a result of using handhelds rather than other media for communication. Small screen size and glare mean that, unlike with desktop computers, looking at someone else's screen is difficult unless they interrupt what they are doing to hand it to you, or unless you are extremely close to them, closer than most people prefer to stand.

Another obvious way to find out what others are looking at is to gain a copy of it. This could be done unobtrusively, so that the SA would absorb the full cost of attempting to enter the conversation. However, manufacturers see handhelds as private machines and three major mechanisms for wireless connectivity on handhelds (i.e., infrared (IR) beaming, 802.11 wireless radio frequency communication (Wi-Fi), and Bluetooth communication) have low-level built-in "notification" conventions that prevent totally covert monitoring. As shown in Figure 1, notification windows appear on the screen informing the user of the operation. This means that any usual monitoring of what is happening on the handhelds is overt. By making confirmation windows automatically disappear, we developed techniques with covertly sharing screens to some degree.

An additional factor is an asymmetry between "pushing" and "pulling" information wirelessly. In IR communication particularly, only "pushing" information out from the PDA is supported directly, as when I beam my business card to you and you confirm by clicking on a pop-up window that you accept it. To

Fourth IEEE International Workshop on
Wireless, Mobile and Ubiquitous Technology in Education (ICHIT'06)
0-7695-2723-X/06 $20.00 © 2006 IEEE

create the illusion of a user's requesting and then getting information means, at a technical level, pushing the request to the machine with the information which must then pull the response back. Our "Look" network function supports a two-way roundtrip signal, a characteristic that provides for pulling information from remote companion handhelds to beaming handhelds.

4. Test system

To test whether we could come up with a design solution to the problem of side-assistant sufficiency within the conventions of handheld use, two conceptually separable systems were necessary, a primary activity system (the thing to be looked at) and the Look feature itself. Since the two systems are in fact integrated, we describe the general interaction factors first, then the primary activity, and finally the Look system. After this technical discussion, we present the experiment evaluating the utility of the Look feature.

4.1. System design considerations

Our system was implemented under the Palm OS 5.2 operating system to run on a Palm Tungsten E. Two general design considerations influenced us: input devices and limited screen real estate.

4.1.1. Input device. In our system, users use the stylus to select soft buttons on the screen and to rearrange screen items. Button selection is implemented through tapping with the stylus on the button. Stylus-based rearrangement is implemented with a drag-and-drop paradigm (see Figure 2).

Figure 2. A user drag-and-drops characters.

4.1.2. Limited screen real estate. Limited screen space also conditions our approach to design. Not only are overlapping windows impractical on the small-screen display of handhelds, but so are the tools for manipulating alternative displays (for an extensive discussion of the design tradeoffs involved, see [18]). Management of moving, resizing, scrolling, and even mode changing occupies space often needed for the main activities. Although the Look facility is a general one that does not need to be displayed on the same screen as the on-going task, we implemented it in an integrated way because we were imagining "Looking" as a first step towards entering into the activity. To make our tool flexible in this way, we tiled our display, with Look at the top of the screen and the primary activity at the bottom.

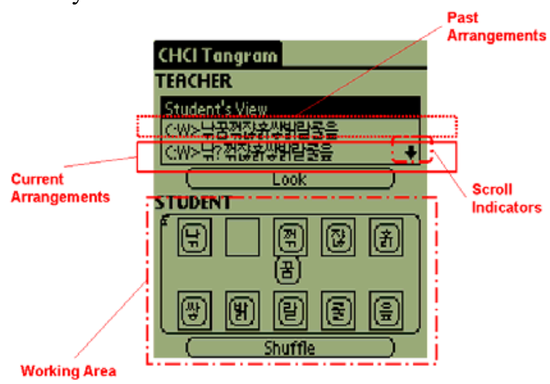

Figure 3. User interfaces of KCT game.

4.2. The primary activity: Korean character tangram

To evaluate the success of a SA in a complex situation involving shared reference, we created an electronic variant of the Tangram game which is called Korean Character Tangram (KCT) as the primary activity. Tangram is widely used to explore the creation and maintenance of common ground in computer supported cooperative work (CSCW) and psycholinguistic literature [5, 4]. In the Tangram game, there is a Matcher and a Director. Each has an identical set of cards with pictures made from Tangrams. The two participants are seated in such a way that they cannot see each other's cards, often by virtue of a physical barrier between them, but they can hear each other. Each set of cards is shuffled and laid out in a grid pattern in front of its owner. It does not matter exactly how many cards (usually between 12 and 24) or what the grid pattern is, as long as these are the same for the two participants. The two participants work together to put the Matcher's cards in the same order as the Director's. Over repeated trials, they become far more efficient [8]. This activity is simpler than the math or visual referential communication tasks discussed earlier in that all communication

Fourth IEEE International Workshop on
Wireless, Mobile and Ubiquitous Technology in Education (ICHIT'06)
0-7695-2723-X/06 $20.00 © 2006 IEEE

between Matcher and Director is accomplished verbally.

The images we used were not Tangrams, but rather Korean characters (see Figure 3). Like Tangram figures, they are unfamiliar and hard to name. However, the characters have the interesting educational property that "knowing" or being able to name them canonically has value. Cards were arranged in a 5 X 2 grid, and could be rearranged either with the "shuffle" button or by dragging and dropping.

4.3. The Look feature

"Looking" is thus the secondary task supported by the interface. When the SA taps the "Look" button, the most recent view of the student's characters appears in the "Student's View" window (Figure 3). After some informal experimentation, we decided to supplement the current view with a history of the rearrangement of characters as an ordered list. Each episode is marked with a prompt. Earlier episodes are accessible by scrolling. Notice that, in Figure 3, "?" mark shows up as the second element of the list because no matching Korean character is found inside the top second box in working area. This is because the Matcher was in the course of dragging that character elsewhere when the SA pressed Look.

Compared to the experience of an overhearer in experiments using the non-computational Tangram game, the SA here has on her screen more static information about the changes that the Matcher goes through from the time of one viewing to the next. On the other hand, the SA's view is punctuated rather than continuous.

In this first version of Look, we employ connectivity based on IR beaming. At the technical level, Look was implemented using *Exchange Manager* in the Palm OS API, which provides a high-level interface using the *exchange* socket structure. At a user level, this means that the SA specifies the Matcher by physically aligning her PDA with the Matcher's and tapping the Look button. The SA is relieved from the need to know or specify who the Matcher is in any more abstract terms such as by machine name, as she might need to with an RF (radio frequency) form of connection. The Matcher is notified by four messages in quick succession: "*Receiving: Chat*", "*Searching…*", "*Transferring: Tangram…*", and "*Disconnecting*". The SA also receives four messages: "*Searching…*", "*Sending: Chat*", "*Disconnecting*" and "*Receiving: Tangram…*" Both must maintain the alignment of machines until the transfer is complete, around 3 seconds. Because of the relationship between pushing and pulling, the SA sees a "*Disconnecting*" message before the transfer is actually complete, and must learn that this will be followed by the "*Receiving: Tangram…*" message.

Thus, compared to the non-mechanized version of the Tangram game, this mechanized version involves a little longer delay in utilizing shared visual spaces. On the other hand, it also involves a mutuality of the interaction between SA and central participants, which might cause subtle alignment in their behavior, especially compared to the complete disassociation between overhearer and matcher seen in the overhearer experiments [7].

5. Experiment

A preliminary experiment was conducted to evaluate whether Look is sufficient to allow effective comprehension of the situation and effective action based on that comprehension. It was a direct, between-subjects contrast of two conditions; that in which the SA had the Look facility and that in which the SA did not have Look.

The SA was conceptualized as a kind of facilitator. The prediction was that SA whose handheld was equipped with Look functionality should better understand the conversation and thus should be better able to provide improved facilitation compared to whom did not have Look. Outcomes centered on SA performance, and consisted of the SA ability to participate effectively and two quantitative learning measures: recognition of the characters and ability to name the characters after the episode.

5.1. Participant

Via announcement on mailing lists, sixty-three students and human-computer interaction (HCI) researchers were recruited to participate in the experiment yielding 21 groups of three. Most participants were volunteers, while some received *Professionalism in Computing* course credit. Reflecting the diverse ethnic constitution of the setting at Virginia Tech, participants' ethnic backgrounds were multi-cultural, included participants from fifteen countries around the world. Three people turned out to have prior exposure to Korean characters, due to having been stationed in Korea in the military. Their groups were excluded from analysis, yielding a net of 20 groups. Participant ages ranged from 20 to 56, with a mean age of 28 (SD: 7.19). 24% of the groups were female and 76% male. Most of the participants were graduate students studying computer science but other majors also took part, including students of industrial system engineering, educational technology, sociology, finance, and agriculture. All of the participants used computers in their everyday lives but just a few of them reported prior experience using handheld computers.

Figure 4. A side assistant was using Look to capture matcher's screen.

5.2. Procedure

After responding to our emailed requests, participants were scheduled in groups of three at a mutually agreeable time. After coming into the lab and discussing/signing the informed consent agreement, students were randomly assigned to one of the three conditions: Matcher, Director, or Side Assistant.

All players were introduced to the functionality of the handheld system. The roles of Matcher and Director were described as follows: "The director will hit the shuffle button to put his/her characters in a specific order. The game is for the matcher and director to put the matcher's characters into the same order as the director's without one person ever looking at the other's screen. The matcher may use the stylus to drag and drop characters on her screen to rearrange them. The director can use words to describe the characters, but he also has their names and we would like you to use the name of each character." They were also told "The third person will be the teacher. The teacher is like the judge or coach."

The group engaged in a round of ordering the cards until the Matcher and Director were sure they had the cards in the same order as one another. Half of the SAs had a handheld machine with Look on it and a picture of the Korean characters with their names taped to the folder containing the handhelds. The other half had only the picture of the Korean characters with their names. The entire process was video and audio-taped.

After the experiments, a questionnaire was administered that gathered two kinds of quantitative learning measurements: from a list of the twenty Korean characters, participants were asked to pick out ten characters with which they had worked and to match these characters with their English (transliterated) names. Finally, to seek as much detailed information as possible, a number of open-ended positive/negative comments were requested.

6. Results

6.1. Quantitative results

Two quantitative outcome measures were used to assess the benefits of having as compared to not having Look functionality; errors in recognizing and in naming the Korean characters. There was no significant difference in the number of recognition errors for SAs in the Look and No Look conditions, $F(1, 18) = 0.12$; p = n.s. However, the number of errors recorded by SAs in the naming of Korean characters was significantly smaller for the groups who had access to Look (mean = 3.0) when compared to the results of those who did not (mean = 6.3), ($F(1, 18) = 7.65, p < .015$).

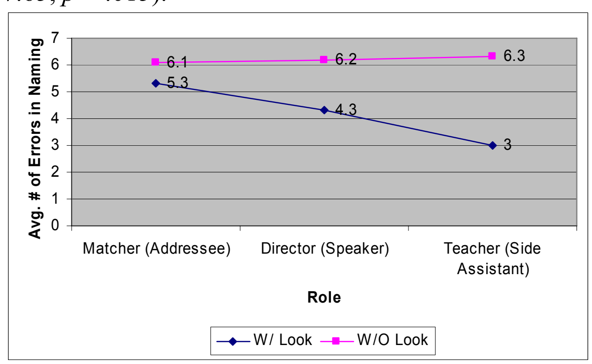

Figure 5. Average number of errors in naming the Korean characters. The reduction in errors for the SA with Look as compared to without Look is significant ($F(1, 18) = 7.65, p < .015$).

6.2. Questionnaire results

6.2.1. Advantages for SAs having "Look". Most SAs whose handhelds were equipped with Look realized that they were at an advantage due to the fact that they could easily share visual indicatory context: "Good thing about the teacher's action is that the teacher can help student to identify the mistakes without sitting beside.", "I like that the teacher can easily communicate with both the director and the matcher.", "It (Look) helps to see both people's work.", "Teacher let us know when we were doing something right. Teacher let us know when we're done or successful. Teacher did not interrupt our interactions."

6.2.2. Disadvantages for SAs not having "Look". On the other hand, SAs working under the "no-Look" condition pointed out the limitation of their interactions with students: "Students did not give enough feedback.", "Not sure if the student got the correct placement.", "Teacher might take more functionalities to help communication.", "I used my own memory and walked back and forth." Another argued that "It was difficult to go back and forth and try to look over their shoulders and to remember the shapes." Another participant claimed that "It can be difficult to express visual data in a form comprehensible to someone who is not also seeing the (same) visual data.", "Associating pictographs with names is not as easy as it looks, especially when pictographs have not seen before." Another echoed a comment often heard from students in project-based classes: "Teacher did not have to help in solving the puzzle. We just did it ourselves."

To determine whether the nature of participants' experiences—whether positive or negative—was related to the availability of Look, a coder blind to our hypotheses categorized each comment as positive, negative or neutral. It shows that significantly more positive comments were made in the Look condition and more negative comments in the No-Look condition ($\chi^2(2) = 6.568$, $\alpha = .05$, $p < .05$).

6.3. Video data: *Are you talking about the one with like…?*

Videotapes recorded during the experiment provide insight into the processes group engaged in. The following transcript (Table 1) from the experiment reveals such an occurrence of seamless entry into a conversation using shared visual clues. In this episode, when the director and the matcher had a hard time in matching their characters, the SA was able to comprehend the situation quickly and well enough to provide useful advice.

#	Role	Transcript
1	M	Oh, uh huh, where should I put it?
2	D	Umm. That goes in the bottom left corner.
3	S	(Look by beaming to Matcher)
4	M	(Hesitantly) I am not so sure about this one (being perplexed).
5	D	(being puzzled)
6	S	(with looking down her screen and raising her head to the director) Are you talking about the one with like a circle on the bottom and then two lines between the top part and the bottom part? There are two lines between the shoulder line and the bottom line (with hand gesture). Is the character you are talking about? *the character should be...*
7	D	*The one I am talking about* doesn't have a line separating the top and bottom half of the character. Umm…
8	S	You said there is a circle and then what's on top? ...
9	D	The circle is kind of wearing… The circle has connected to it a line and a little tick mark in the center… It kind of looks like the circle is wearing a hat…
10	S	Umm. I didn't see that character (with leaning to her own palm). Can you describe that for me, again? (silently)
11	D	Umm. Well, the top half of the character looks like… it is the Greek character μ. A μ, with little tails on the side.
12	S	Oh, okay, so, it kind of looks like a letter T? (with watching director)
13	D	Umm…, there is a sideways T! Yeah!
14	S	There is a sideways T. Um, like tail. Well, yeah, it is like a T and a U, share *a line*
15	M	*They share a line.*
16	S	So that's the character you're talking about?
17	D	T… and U... share a line? Oh! Are both on the top half of the same line of the character?
18	S	*A-ha.*
19	M	*Yes.*
20	D	Okay, yes, yes! (nods)
21	S	And then, the bottom there is actually a square from the formal writing, It is…*yeah, it is a square or, Umm, a circle, so yeah?*
22	M	*Yes, it is rectangular.* (with watching director)
23	D	I think that could be a square.
24	S	Yes, I think that's why she (Matcher) and I have misunderstood. Yeah, it looks like a circle, but actually, from the formal writing here I got it looks like a square
25	D	Okay!

Table 1. Transcript of an episode wherein Look provided assistance to the SA.
(Overlapping speech is marked with pairs of asterisks. Gestures are enclosed in parenthesis. D=director, M=matcher, S=side-assistant)

This interaction shows how the SA helped the director and matcher work through their confusion, which originated from the fact that they had slightly different views of the same object. As indicated in turn 21 and 24, the director had his own slightly different interpretation than either the matcher or the SA for the shape of the character 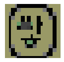 . This confusion began in the conversation that occurred between turns 1 and 5. In turn 6, the SA entered the conversation easily and appropriately. She used the evidence of shared visual clues to propose assistance (# 3). Turns between 7 and 14 show the interactive cycles of conversational turn-taking between the SA and the director. They acted jointly toward a convergence of meaning [15]. When the director indicated "the Greek character μ" (# 11), the SA added new symbolic explanation, "like a letter T?" (# 12), as a scaffold for director's understanding. The director's expression in turn 13 demonstrated that a conceptual change had been initiated.

Between turns 14 and 22, we find another example of how visual co-presence of objects can easily create common ground among participants. Turns 15, 19 and 22 show that the matcher was following correctly the discussion between the SA and the director. She could do so because she was watching the same object, and could even validate SA's description by displaying her understanding: "They share a line."(# 15) and "Yes, it is rectangular" (# 22). Finally, in turns between 23 and 25, the director was successful in creating convergent concepts and was able to complete the task.

7. Discussion

The experiment reported here investigated the adequacy of a particular implementation of "Look" to help with the problem of *side assistant sufficiency* (SAS): the problem of allowing a SA to gain adequate understanding of a shared visual workspace to enable meaningful monitoring of on-going conversation about that workspace or even entry into the conversation. Our experiment highlights a situation involving difficult shared reference. While our experiment, possibly due to lack of power, did not show that SAs with Look had a superior recognition of the referred to objects than those without Look in the post-test, we did provide preliminary evidence that they associated better the names of those objects with the objects. Moreover, we see evidence that side assistants with Look could enter into the conversation smoothly, while those without Look experience more difficulty. We also see that participants in groups with Look report a more positive experience about the experience compared to those without.

The experiment here is abstracted away from a situation found in the classroom use of handhelds to aid learning. However, it has implications beyond the use of handhelds in the classroom. In particular, it highlights a gap in predictive theory about the nature of joint action. Current theory does not accommodate the range of situations that we find in ubiquitous and pervasive computing. Although Kraut, Gergle and Fussell [12] have examined the cost of delay in sharing visual workspaces, they have not examined punctuated sharing such as obtained here. While other prior work has examined the information available to the overhearer, it has not examined either side assistance or the movement from side to central participant.

Two kinds of future work follow from this. Following from our interest in enabling classroom interactions, we plan to create experimental situations that more closely resemble the "formative assessment" and "peer sufficiency" problems we started with. In particular, the teacher or peer typically is also a latecomer to the on-going interaction, and does not witness, unlike our SAs did, the creation of common ground between Matcher and Director.

Additionally, the current work is focused on a particular implementation with special, even idiosyncratic, features. Three factors require further exploration. First, our infrared-based system costs the SA only public and emergent actions in defining the target of a Look action. That is, no private or previous work is required to name or locate a machine or a user. In a radio-frequency implementation, the SA would have to use her screen to somehow specify the IP address of the machine she wanted to look at, or to have done so previously. Second, our implementation supports only overt monitoring by the SA, enforced by the need to align the machines physically and the repeated notification messages. This might seem like a simple design flaw. However, while the act of looking covertly would interrupt less, it also might make entry into the conversation more difficult for the SA than it is currently. In the current version, the SA takes physical action similar to adjusting body position as one enters a hallway conversation. As the SA's access to information becomes more covert, the Matcher and Director gain less information about the SA's range of intentions. Third, our system provides only punctuated rather than continuous "looking" incidents. One way of thinking is that the SA has less information about what is happening than if the SA had continuous monitoring. However, another point of view is that the SA has an amount and kind of information that is controlled by herself. SA's control may arguably prove more important to understanding than continuous visual information.

Fourth IEEE International Workshop on
Wireless, Mobile and Ubiquitous Technology in Education (ICHIT'06)
0-7695-2723-X/06 $20.00 © 2006 IEEE

8. Conclusion

Previous research suggests that the sharing of computer screens provides certain advantages for collaborative learning endeavors [10]. Sharing a single display provides a shared artifact, which between collaborators—here central participants and side assistant—initiates discussions and enhances attention. However, handhelds do not naturally support the sharing of a single display among collaborators. As a result, collaborative activities that rely upon handhelds face greater challenges with regard to the maintenance of shared attention [16]. The Look feature makes it possible to capture the referential context during handheld mediated collaboration enabling participants to focus on the topic rather than the technology.

9. Acknowledgements

This work was supported by NSF grant REC2343442342.

10. References

[1] S. E. Brennan, *The grounding problem in conversation with and through computers*, in S. R. Fussell and R. J. Kreuz, eds., *Social and cognitive psychological approaches to interpersonal communication*, Lawrence Erlbaum, Hillsdale, NJ, 1998, pp. 201-225.

[2] S. E. Brennan, *Seeking and providing evidence for mutual understanding*, Stanford University, Stanford, CA, 1990.

[3] H. Clark and M. Krych, *Speaking while monitoring addressees for understanding*, Journal of Memory and Language, 50 (2004), pp. 62-81.

[4] H. H. Clark, *Plenary Session: Working Together at a Distance*, *CSCW*, ACM Press, Cambridge, MA, 1996.

[5] H. H. Clark, *Using Language*, Cambridge University Press, Cambridge, 1996.

[6] H. H. Clark and S. E. Brennan, *Grounding in Communication*, in L. B. Resnick, J. Levine and S. D. Behrens, eds., *Perspectives on Socially Shared Cognition*, 1991, pp. 127-149.

[7] H. H. Clark and M. F. Schober, *Understanding by addressees and overhearers*, in H. H. Clark, ed., *Arenas of Language Use*, University of Chicago Press, Chicago IL, 1992, pp. 176-197.

[8] H. H. Clark and D. Wilkes-Gibbs, *Referring as a collaborative process*, Cognition, 22 (1986), pp. 1-39.

[9] P. Vahey, D. Tatar and J. Roschelle, *Leveraging Handhelds to Increase Student Learning: Engaging Middle School Students with the Mathematics of Change*, Proceedings of ICLS 2004, Santa Monica, CA, 2004.

[10] K. Inkpen, K. Booth, M. Klawe and R. Upitis, *Playing Together Beats Playing Apart, Especially for Girls*, Proceedings of CSCL '95, 1995, pp. 177-181.

[11] R. E. Kraut, S. R. Fussell, S. E. Brennan and J. Siegel, *Understanding effects of proximity on collaboration: Implications for technologies to support remote collaborative work*, in P. Hinds and S. Kiesler, eds., *Distributed work* MIT Press, Cambridge, MA, 2002, pp. 137-162.

[12] R. E. Kraut, D. Gergle and S. R. Fussell, *The use of visual information in shared visual spaces: informing the development of virtual co-presence*, *CSCW 2002*, ACM Press, 2002, pp. 31-40.

[13] J. Roschelle, D. Tatar and J. Kaput, *Getting to scale with innovations that deeply restructure how students come to know mathematics*, in R. Lesh and A. E. Kelly, eds., *Design Experiments in Education*, 2005.

[14] D. Tatar, J. Roschelle, N. Shechtman, J. Kaput and D. Hopkins, *Scaling Up Innovative Technology-Based Math with a Wide Variety of 7th Grade Teachers*, Journal of the Learning Sciences, In press (2005).

[15] J. Roschelle, *Learning by Collaborating: Convergent Conceptual Change*, Journal of the Learning Sciences, 2 (1992), pp. 235-276.

[16] J. Roschelle and R. Pea, *A Walk on the WILD Side: How Wireless Handhelds May Change Computer-Supported Collaborative Learning*, International Journal of Cognition and Technology, 1 (2002), pp. 145-168.

[17] D. Tatar, J. Roschelle, P. Vahey and W. R. Penuel, *Handhelds Go to School*, IEEE Computer, 36 (2003), pp. 30-37.

[18] D. Tatar, *Design Tensions in Creating Wirelessly Networked Handheld Tools for Teaching/Learning Math* Submitted, International Journal of Human-Computer Interaction (IJHCI) (2006).

[19] C. DiGiano, L. Yarnall, C. Patton, J. Roschelle, D. Tatar, and M. Manley, *Collaboration design patterns: Conceptual tools for planning for the wireless classroom*, In Proceedings of WMTE 2002 (2002) pp. 39-47.