The Effects of Co-Present Embodiments on Awareness and Collaboration in Tabletop Groupware

Author

Organization or School

ABSTRACT

Most current tabletop groupware systems use direct touch, where people manipulate objects by touching them with a pen or a fingertip. The use of people's real arms and hands provides obvious awareness information, but workspace access is limited by the user's reach. Relative input techniques, where users manipulate a cursor rather than touching objects directly, allow users to reach all areas of the table. However, the only available awareness information comes from the virtual embodiment of the user (e.g., their cursor). This presents designers with a tradeoff: direct-touch techniques have advantages for group awareness; relative input techniques offer additional power but less awareness information. In this paper, we explore this tradeoff, and we explore the design space of virtual embodiments to determine whether factors such as size, realism, and visibility can improve awareness and coordination. We conducted a study in which seven groups carried out a picture-categorizing task using seven techniques: direct touch and relative input with six different virtual embodiments. Our results provide both valuable information to designers of tabletop groupware, and a number of new directions for future research.

KEYWORDS: Tabletop groupware, embodiments, interaction techniques.

INDEX TERMS: H.5.3 [Information Interfaces]: CSCW

1 INTRODUCTION

Tabletop groupware systems allow co-present collaborators to work together over a shared horizontal display. Tables are a natural site for group work, both because of their ubiquity in the real world, and because their physical characteristics support coordination and communication, such as the face-to-face orientation of people around the table, the central location of work artifacts, and the use of direct touch to manipulate objects on the work surface.

Direct touch – where people manipulate objects by touching them with a pen or a fingertip – provides a number of benefits for collaboration. In particular, the use of people's real arms and hands provides awareness information about 'who is working where' on the table, and makes it easy to watch other people work [16]. However, direct touch also has disadvantages: it can be difficult for people to reach objects that are far away; arms and bodies can get in the way of each other, preventing people from working in the same space at the same time; and it can be awkward or uncomfortable to work close to another person.

One way to deal with these problems is to use relative rather than direct input techniques – that is, techniques where each person manipulates a cursor rather than touching objects directly. Relative input techniques allow reaching to any part of the table and allow people to work in the same place, but since they do not use physical arms, this source of awareness information is lost.

The only awareness information produced by a relative input technique comes from the virtual embodiment of the user (e.g., their cursor) on the table. This visible representation provides each user with feedback about their own actions, but as a side effect, also provides awareness to other members of the group. Although this is the same mechanism by which real arms convey awareness, virtual embodiments are much less obvious than physical arms.

Tabletop designers are thus presented with a tradeoff: direct-touch techniques have advantages for group awareness but limitations for individual work; relative input techniques offer additional power but provide less awareness information. Most existing table systems have chosen direct touch input [1,4,14], but in fact little is known about how different techniques affect groupware usability.

In this paper, we explore the tradeoff between individual power and group awareness, and investigate the design of virtual embodiments for tabletop groupware. We built several different tabletop embodiments, and studied their use in a realistic collaborative task. We were interested in three main questions:

- What are the differences between real and virtual embodiments in terms of collaboration factors such as awareness, distraction, territoriality, and movement patterns?
- Does the appearance of a virtual embodiment (e.g., its size and shape) affect these factors, and can a virtual embodiment approach the awareness of a physical arm?
- Overall, do users prefer direct touch and physical embodiment, or relative input and virtual embodiment?

To investigate these questions, we carried out a study in which seven groups of three people carried out an open-ended picture-categorizing task on a digital table. We gathered movement data, tested people during the task on their awareness of other's locations, and asked several survey questions about perceived awareness, distraction, social distance, and overall preference. The study provided some results that confirmed our expectations (e.g., people were significantly more accurate in guessing others' locations with direct touch), but also many that were surprising. For example:

- type of virtual embodiment had no effect on awareness accuracy – even though some of the embodiments were much larger and more obvious than others;
- virtual embodiments did not appear to carry the same social constraints as did direct touch: people were more likely to cross embodiments, and they were much more likely to move into another user's personal territory;
- although there were few differences between the virtual embodiments in objective measures, there were strong subjective differences; in general, people preferred larger and more realistic embodiments, and felt that these provided much better awareness;

Street Address and Electronic Mail Address

- most surprising of all, a majority preferred the virtual embodiments over direct touch with physical arms, which goes against current practice in tabletop design.

Our study is the first exploration of the awareness differences between direct and relative interaction techniques. It shows that visual representation can have strong effects on subjective measures (even though there are few objective differences), and that relative input and virtual embodiments can actually be preferable to direct touch. Although exploratory, our results provide both valuable information to designers of tabletop groupware, and a number of new directions for future research.

2 GROUPWARE EMBODIMENTS

Embodiments are visual representations of users in groupware applications [2,26] that help people stay aware of others' presence, location, and movement [10].

2.1 Real Embodiments

Real embodiments make use of the actual physical body of the user [11], and they are used in co-present situations. Real embodiment techniques use direct touch (using hands, pens, or tangible blocks). For example, techniques such as pick-and-drop [24], drag-and-drop [3], and media blocks [30] do not require virtual embodiments because the user's body and tools allow others to track their actions.

Real embodiments have been used extensively in tabletop groupware because many current systems use multi-touch displays (e.g., [1,4,14). One of the main advantages of using touch input is that people can easily stay aware of others' interactions with the table since they can see them physically reaching into the workspace [11].

2.2 Virtual Embodiments

Virtual embodiments are digital representations of users. They are common in distributed groupware, where people must rely on computer mediated information to track others' actions. Telepointers, the simplest form of virtual embodiment, are widely used in shared-workspace groupware to represent each person's mouse cursor [9,26].

Virtual embodiments, however, are not frequently used in table systems. They are usually used only when relative input techniques are necessary, in order to show each person's cursor in the workspace. There are several tabletop techniques that support relative input, including standard mouse-based drag-and-drop [3], hyperdrag [23], cursor-extension techniques [8,12,17], laser pointing [19,15] and portal approaches such as radar views [27,17] and cutouts [21]. Many relative input techniques use simple pointers as embodiments and provide little information about who is controlling each cursor [3,19,16,21].

Recent comparative studies show that relative input techniques have several advantages over direct touch, such as improved access, less fatigue, and fewer occlusion problems [11,7]. However, studies have also shown that people have difficulties with staying aware of others since they cannot observe their actions directly [16,11].

Virtual embodiments have the potential to help people track others' actions more effectively when relative techniques are used; however, the design space for tabletop embodiments has not been explored in detail. In the next section, we propose a set of design dimensions for virtual embodiments in tabletop groupware.

3 VIRTUAL EMBODIMENT DIMENSIONS

The physical presence of people around the table makes it possible to develop new mappings between cursor and user that are not possible in distributed groupware. Lines can be drawn from the user to their cursor, or cursor arrows can point directly at the user that controls them. In this section, we propose visual design dimensions that can be used in virtual embodiments for tables. Little work has been done in this area previously, so these dimensions are not necessarily complete, and they are partially based on our experience with working with tabletop embodiments. We discuss five main dimensions in this section: size, realism, line, identity indicator, and profile.

Size. Telepointers used in distributed groupware are small so that they do not occlude content. However, when horizontal displays are used, it can be difficult to track small cursors since they are far away from people in some areas of the table [16].

Realism. Virtual embodiments can be abstract (telepointers or avatars), or they can be more realistic and can resemble people's bodies or arms. Video embodiments, such as those that use video of people's arms (e.g. [28,29]), have been used on distributed tables (e.g., [20]). However, video approaches are not useful with relative techniques because people's physical movements give limited information about where their cursor is located.

Line. User identity can be made more explicit in tabletop embodiments by drawing a line from the user's position to their cursor. This approach has been used in a version of the pantograph technique (described by Hascoët [12]) that was applied to tabletop systems by Nacenta et al. [16,17].

Identity indicator. Several visual variables, such as color, texture, orientation, and shape can be used to help identify embodiments [26]. Many of these variables have been previously used in distributed groupware systems (e.g., [9]). However, they can be combined in new ways on tables. For example, the orientation of an embodiment can dynamically change so that it always points to the user that controls it.

Profile. The visibility of virtual embodiments can change based on differences in several visual variables. Embodiments can have different levels of transparency and opacity, or they can be painted on top of or behind digital artifacts that are shown on the table. Reduced visibility can help to minimize occlusion and distraction caused by embodiments, but it also has the potential to make it more difficult for people to stay aware of others' actions.

4 EXPLORATORY STUDY

We conducted an exploratory study that compared seven embodiment techniques—one was a direct touch technique and six were relative input techniques that used virtual embodiments. Our three general research questions were to determine what the differences are between physical and virtual embodiments, to determine whether changing the visual representation of an embodiment makes a difference to collaboration, and to find out what people prefer.

The six virtual embodiment techniques used in the study cover the embodiment dimensions discussed in the previous section. They were implemented as part of a relative cursor extension technique, where users can access the table by using a stylus in a local input region. The virtual embodiments included: an arrow, a circle, a narrow pantograph, a wide pantograph, a wide pantograph drawn behind objects on the table, and an arm.

We tested seven groups of three participants while they carried out a photo grouping task; the task was chosen since grouping is a part of many collaborative activities [20,16], allowing us to assess embodiments during a realistic task. Grouping also contains elements that are common in most shared tabletop activities, including moving, passing, reorienting, and rearranging items; communicating to make decisions about the activity; and staying aware of others' actions to allow coordinated action.

4.1 Participants, Apparatus, and Procedure

21 paid participants, 8 female and 13 male, were recruited from a local university and organized into groups of three. Participants

ranged in age from 18 to 36 years (mean of 25). Groups performed the task with all seven techniques.

The experiment was conducted on a top-projected tabletop system with a display size of 1024x1536 pixels. The table surface measured 125x185cm with a display area of 118x178cm. Custom software was written in Java, and ran on a P4 Windows PC and a Polhemus Liberty magnetic tracker with three pens (see Figure 1).

Groups began the study by completing a training session, where they used each technique for two minutes. They then completed a 6 minute trial using each technique. Technique order was distributed among groups using a Latin square. During the trials, one participant sat at the head of the table and the others sat at the adjacent sides (Figure 1); after each trial they switched places. After the session, participants completed a preference questionnaire. Additional data were collected using system logs and observations.

Figure 1. Tabletop setting for the experiment. The wide pantograph in back embodiment is shown.

4.2 Embodiment Techniques

We implemented six virtual embodiment techniques and direct touch (Figure 2). In each virtual embodiment, the user moves the pen in their local input space (shown as a white rectangle). The user's cursor position on the table is proportional to the position of the pen (5X) so that the whole table can be reached with little movement. All virtual embodiments are drawn using a color that corresponds to the user's location at the table (red, blue, or green). There were six types of virtual embodiment:

- *Circle*. Cursors are represented using opaque circles.
- *Arrow*. Cursors are represented using an opaque arrow that always points to the user that controls it. The arrow has the same area as the circle.
- *Narrow pantograph*. A translucent line is drawn from the table edge where the user is sitting to the cursor position.
- *Wide pantograph in front*. This embodiment is the same as the narrow pantograph, except ten times wider.
- *Wide pantograph in back*. This embodiment is the same as wide pantograph in front, except the line is painted behind the objects on the table. The tip of the line, which is shown using a small triangle, is painted on top of the objects so that users can always track the cursor position.
- *Arm*. The arm embodiment is similar to the pantograph techniques, except an arm is used to represent the user's cursor. The arm width is the same as that of the wide pantograph techniques, and it is always painted in front of objects. The arm length changes as a user reaches across the table (the width is static), and it uses the same transparency level used in the pantograph techniques.
- *Direct touch*. Users select an object by touching it with the pen and clicking their button, and they can move objects by dragging the tip of their pen across the table. In this technique, the user's body acts as the embodiment.

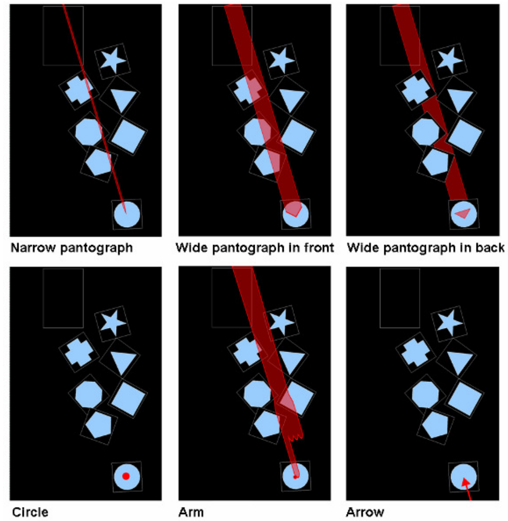

Figure 2. Virtual embodiments. User controlling embodiment sits at the top, where the input rectangle is shown.

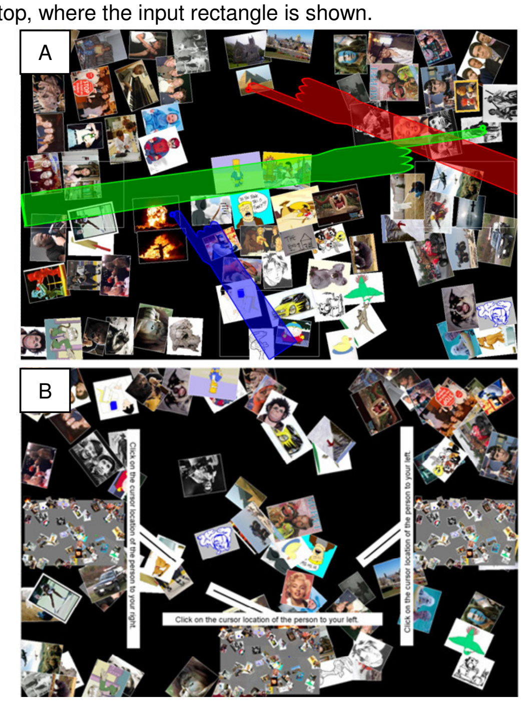

Figure 3. Task overview. (A) completed task with sorted piles of images, arm shown. (B) freeze mode where each user clicks on the location of the cursor for the specified user.

4.3 Task

Participants were asked to carry out a grouping task, where they sorted 80 pictures into three different categories. At the beginning of the task, the pictures were placed on the table in random locations and with random orientations. They were asked to work together, positioning the pictures so that they did not overlap, and so that they were grouped with other pictures in the category (see Figure 3A). They carried out a 6 minute trial with each technique. Participants were given a different set of categories and a different set of images for each trial, but they were not given any guidance on the strategies that they should use when carrying out the task.

The categories were intentionally ambiguous, so that pictures would not clearly fit into any of the categories. Examples include: "animal, vegetable, mineral", "liquid, solid, gas", and "happy, sad, angry." Our goal was to facilitate communication within the group about how pictures should be categorized so that the task would not turn into a strict divide and conquer activity.

We assessed participants' awareness of others' cursor position by freezing the task at random intervals and asking them to specify the table position where each user was working. Our approach was based on the situation awareness global assessment technique (SAGAT) that was described and validated by Endsley [5,6]. While participants were engaged in the activity, the screen cleared and the embodiments were removed from the tabletop. A miniature of the workspace was displayed in front of each person, and users were asked to click on the view to specify the positions of the other people's cursors. After each freeze, the task resumed normally The freeze technique was used 8 times for each trial. Freezes did not occur in the first 30 seconds of the task, and there was at least 20 seconds between each freeze. The time spent in freeze mode did not count toward overall task time.

5 RESULTS

Our three research questions were to determine the main differences between real and virtual embodiment types, to determine whether changing the visual representation of an embodiment makes a difference to collaboration, and to find out what people prefer. We organize the results of our investigations by our main collaboration issues: awareness and distraction, movement patterns and territoriality on the table, and overall subjective preferences.

5.1 Awareness, Distraction, and Interference

We gathered two kinds of data about people's ability to maintain awareness: the 'freeze test' data, in which people were periodically asked to mark others' locations on a map of the workspace; and the subjective questionnaire, which asked people to rank the techniques according to how well the embodiment helped them stay aware of others' actions.

Our hypothesis was that people would be better able to stay aware of others' locations with direct touch, followed by the large virtual embodiments, then the small virtual representations; and that people's subjective assessments would line up with the objective results.

5.1.1 Freeze Test: Awareness of Location

The freeze test measured error in terms of pixel distance from people's true locations (see Figure 4). A factorial ANOVA with embodiment and table location as fixed factors and group as random factor showed a significant main effect of embodiment type on error amount ($F_{6,34}=6.46$, $p<0.001$). Follow up Tukey tests show that direct touch was significantly different from all virtual techniques ($p<0.001$ in all cases), but that there were no differences between the virtual embodiments.

5.1.2 Subjective Awareness Ranking

We asked people to rank the techniques according to how well the embodiment helped them stay aware of others' actions (see Figure 5). A Friedman test on the rankings showed significant differences between techniques ($\chi2(6)=49.69$, $p<0.001$). Post-hoc Wilcoxon Signed Rank tests with a Bonferroni correction ($\alpha=.0024$) identified several pairwise differences: Circle was significantly different from all other techniques; Arrow was different from Arm and Wide Front; and Narrow Pantograph was different from Wide Front (Arm and Narrow Pantograph were not significantly different, $p=.0037$).

The three large virtual embodiments had the highest rankings, and were all ranked higher than the physical arms of direct touch (Figure 5). The small embodiments had the lowest rankings, with Circle ranked last by 13 of 14 participants. The low average ranking of direct touch was unexpected. Although six people ranked direct touch highest, eight others ranked Arm as the best technique.

Participant comments provide some insight into these rankings. The size of the large embodiments, and their connection to their 'owner' appeared to be a factor in their high rankings: for example, one participant stated "the size of the arm pointer made it easier to follow…also, the arm extended directly from the person so was easier to track."

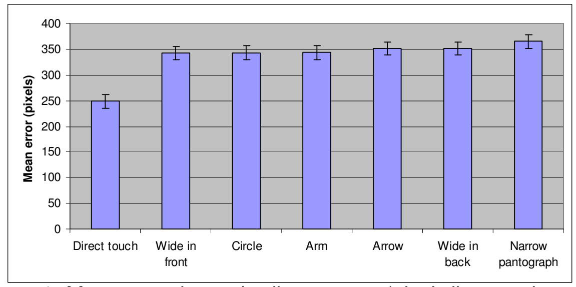

Figure 4. Mean error by embodiment type (pixel distance between guess and true location)

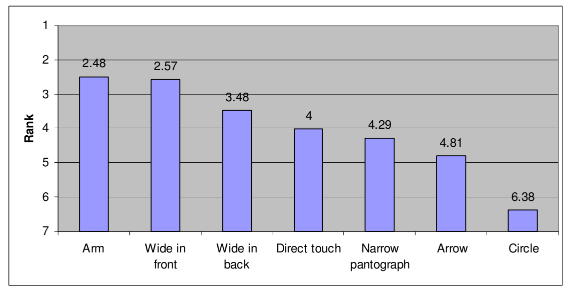

Figure 5. Subjective awareness (taller bars mean better rank)

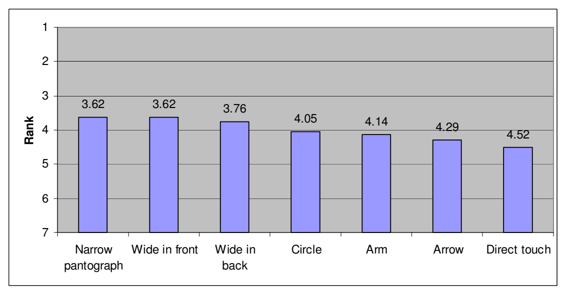

Figure 6. Subjective distraction ranking (taller bars indicate less distraction).

Comments also suggest that the arrow and the circle were difficult to track due to their small size. However, even though the area of both embodiments was the same, people's comments show that they found the arrow slightly easier to track than the circle: one participant said "the floating arrow was larger lengthwise so it was easier to track others' movements."

Participants also stated that identity – determining to whom the cursor belonged – was a problem for the circle embodiment. One person stated "the circle's size and lack of connection…made it the least easy to track;" another said "it was not connected and did not give any hint on where it was coming from."

5.1.3 Distraction and Interference

A questionnaire recorded subjective distraction rankings (Figure 6). We expected roughly the inverse of the subjective awareness rankings (Figure 5): that is, that people would feel the large virtual embodiments were the most distracting and interfered most with their work. However, there was no clear trend in the rankings, and a Friedman test showed no effect of embodiment type ($\chi2(6)=3.27$, $p=.775$). As can be seen from Figure 6, participants ranked the different techniques in several different ways, leading to a narrow range between the means (3.62 to 4.52).

Comments from the open-ended questions showed that most people did not find that large embodiments caused more distraction and occlusion problems. Five participants stated that they did not feel that there was any difference at all between the different techniques.

5.2 Movement Patterns

We gathered three measures to examine the effect embodiment and interaction technique have on the way that people move their cursors or hands around the table. We looked first at whether the techniques altered the territories where people did their work, and second, at issues of social distance (how often did people move into another person's territory, and how often did people cross arms or cursors).

5.2.1 Territoriality

We recorded all cursor movements, and used this data to calculate the areas of the table where users were active. The movement maps in Figure 7 show examples of this data. As a measure of the spread of user activity across the table, we performed a test based on O'Brien's variance comparison procedure [18,22]. The test measures whether the variance (spread) of two sets of data points are significantly different from one another.

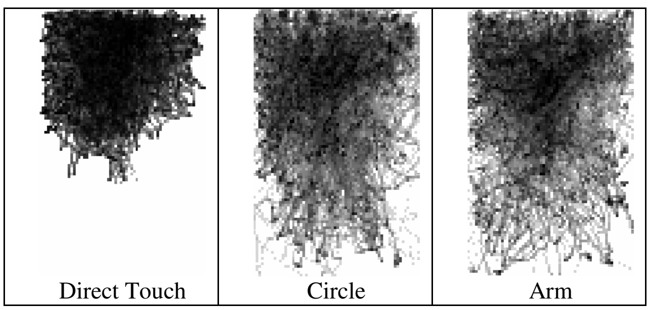

Figure 7. Example movement maps (includes all participants sitting at top position; darker areas mean more movement).

We expected to see clear differences between the direct input technique and the relative techniques because of the obvious reach limitations of physical arms. In addition, we were interested in whether any of the virtual embodiments would result in different movement maps.

We found a significant effect of embodiment type on the spread of activity ($F_{6,120}$ = 10.749, p<0.001), and as expected, pair-wise follow-up tests showed that activity with the direct touch technique was much less spread than with all of the relative techniques (all t<-4.5, p<0.001). However, there were no differences found between any of the relative techniques (all |t|<1.96, p>0.068), and no differences between the large and small embodiment types either. As suggested by Figure 7, there are no obvious differences between the Circle and Arm movement maps.

There were two characteristics of the virtual embodiments' maps. First, even though they could reach anywhere on the table, people still worked primarily in the area within arms' reach. Second, when people did move outside their home area, they worked over most of the table, even areas at the far edge (which was 175 cm away from participants seated at the ends of the table)

5.2.2 Cursor Crossings

To explore differences in the way that people reached across each others' arms and embodiments, we calculated the number of times that people's cursors (or pens) crossed (that is, where crossings occurred in lines drawn from the cursor or pen back to the person's seat location).

Crossings are one indication of whether people think of the virtual embodiments in the same way as real arms. Assuming that people do not wish to reach across another person's real arm, we hypothesized that people would also be less likely to cross more realistic virtual embodiments (i.e., Arm) than the disconnected embodiments.

Figure 8 shows the average number of crossings for each embodiment. A repeated-measures ANOVA showed a significant effect of embodiment type on number of crossings ($F_{2.99,17.98}$=11.749, p<0.001); the Greenhouse-Geisser correction was applied because of the lack of sphericity. Follow-up comparisons showed that Direct Touch was different from all virtual embodiments (all |t|>5.3, p<0.002). As shown in Figure 8, there were fewer crossings with direct touch (mean of 11) than with any of the other techniques (mean 53 to 79). There were no differences found between any of the virtual embodiments (all |t|<2.99, p>0.02, no Type I error correction).

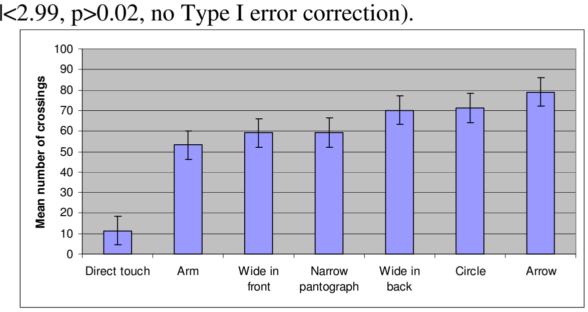

Figure 8. Mean number of cursor crossings by technique (error bars show standard error)

5.2.3 Reaching into Others' Personal Territories

Our second measure of social distance effects was the percentage of time that people spent reaching into another person's 'personal territory.' Scott and colleagues [25] suggest that the area within a person's immediate reach is recognized as belonging to that person, which should lead to fewer intrusions by other people.

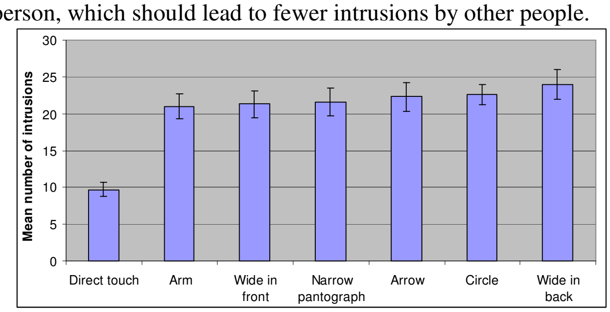

Figure 9. Percentage of time cursor was in another user's personal territory (error bars show standard error)

We defined personal territory for the analysis as the area within a radius of 65cm of the seating point of the user. As with crossings, a repeated-measures ANOVA showed a main effect of embodiment type ($F_{6,120}$=11.22, p<0.001, no correction); follow-up paired-samples t-tests showed differences between direct touch to be significantly different from all other techniques (all t<-6.2, p<0.001). Participants had their cursors in somebody else's territory an average 9.6% of the time; with the rest of techniques the percentages ranged from 21% with Narrow Pantograph to 24% with Wide Pantograph Back (see Figure 9). As with the crossings, differences between any of the virtual embodiments were non-significant (all |t|<1.7, p>0.10).

5.2.4 Comments about Social Distance

We asked participants whether any of the input techniques or embodiment types affected their comfort levels in working close to another person. There were a variety of responses, indicating that people have widely different views on social distance issues.

Eight of the 21 participants stated that the embodiment type made no difference to their comfort level, and that they were willing to reach anywhere on the table that they needed to. In contrast, six other participants stated that they were uncomfortable working close to other people when using direct touch. Two of these indicated that physical crowding between group members was the main problem, and four others stated that they felt uneasy working too close to others. For example, one participant stated that "direct touch made me uncomfortable with going into someone else's space."

Six people stated that they preferred using smaller embodiments when working close to another person. One participant stated that "a somewhat smaller cursor was more comfortable to use when accessing objects near other users." A main reason for this attitude was people's concern about blocking others' views with their embodiment. For example, participants stated "[embodiments] that didn't obscure the users' view worked best", and "I preferred the [Circle] because it didn't feel like I was getting a huge line in front of the person's work area."

5.3 Overall Preferences

After the session, we asked participants to rank the techniques according to overall preference. A Friedman test found a significant effect of embodiment type ($\chi2(6)=28.57$, $p<0.001$); as can be seen in Figure 10, there are substantial differences in the mean ranks given to the different techniques, with Arm ranked highest (mean rank 2.7 of 7) and direct touch ranked lowest (mean rank 5.6 of 7).

We performed post-hoc Wilcoxon Signed Rank tests to identify pairwise differences between techniques. Wide Front and Circle had similar means and standard deviations with Wide Back and Arrow, so we removed them from posthoc testing to reduce the number of comparisons. Direct touch was significantly different from all other techniques ($p<0.005$) except Arrow; and Arrow was significantly different from Arm ($p=.001$).

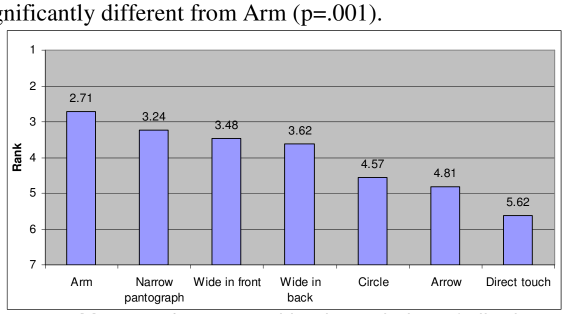

Figure 10. Mean preference ranking by technique (taller bars mean better ranking).

Arm had the highest average ranking, and most (14/21) people preferred it over the wide pantograph techniques, which had the same width and transparency levels. People's comments showed that they liked the Arm embodiment because it was more natural, and felt more like an extension of their real arm: for example, one participant stated that the Arm was "more personalized…I liked the arm because it was more fun and natural;" another said "I felt I was directly accessing the object. It was 'my hand' not just the cursor;" and a third stated "I liked the arm better because it seemed like a more natural extension of a person's reach."

The narrow pantograph also had a high ranking; several participants stated that they liked the balance between identity awareness (provided by the line back to the person's location) and low distraction compared to the wider embodiments. However, three people felt that the narrow line was difficult to track and did not promote as good a sense of awareness as the other techniques.

The disconnected techniques (Circle and Arrow) both had low average rankings. Several participants stated that they found it difficult to keep track of others with these cursors because there was no visible connection to the person, and that the color mappings were difficult to remember. In addition, four people found the variable orientation used in the Arrow confusing; for example, one participant stated that "I found the arrow difficult to use because it was pointed at me…it was actually more confusing with the arrow, which pointed towards me and, therefore, did not indicate the direction of the object."

Direct touch was ranked the lowest by 14 of 21 participants. Comments suggested that people disliked the technique for two main reasons: the limitations on reaching distance, and the physical effort that was required to access objects.

6 DISCUSSION

Our exploratory study raised several issues – some of which we expected, and others that were unexpected.

- The study found substantial differences in all objective measures (awareness freeze test, crossings, intrusions) between direct touch with physical arms, and relative input with virtual embodiments.
- We found no differences between virtual embodiments on any of the objective measures, even though some embodiments are much more visually obvious.
- The study did find differences between virtual embodiments in subjective perception of awareness, but no difference in perception of distraction.
- We found that most participants preferred the relative/virtual techniques over direct touch.

In the remainder of this section we consider explanations for our results – in particular, why there were subjective awareness differences but not objective differences, why different virtual embodiments did not cause different levels of subjective distraction or concern about social distance, and why people preferred the virtual embodiments. Last, we consider how our results can be generalized, and what issues must be considered further in future work.

6.1 Explanation of Results

6.1.1 Objective differences in location awareness

People were significantly more accurate in marking others' locations with direct touch than any virtual embodiment. The accuracy difference between direct touch and the virtual embodiments was approximately 100 pixels, which corresponds to approximately 12 cm on the table.

One possible reason for this difference is that physical arms do, as expected, provide better awareness than virtual embodiments. This is a reasonable explanation, both since people are experienced at interpreting the motions of others' arms, and because arms are obvious physical features that exist in the 3D space above the table (rather than 2D images on the table surface).

However, another explanation suggests that the increased accuracy could be due simply to people's understanding of how far people can reach with direct touch. That is, it was possible for people to constrain their guesses with direct touch to the area within arms' reach of a person, whereas this strategy was not possible with the virtual embodiments, since people could easily be anywhere on the table.

Although this may seem to be just an artifact of the freeze test method, it highlights an important difference between direct and virtual techniques: that the natural restrictions of a technique can be used as implicit aids to location awareness. Assumptions (if correct) can be as useful as new information, and when people are restricted to their home territory, it becomes easier to assume that they are close to home. In addition, moving outside the home territory requires a large and obvious action with direct touch, but looks similar to local actions when using relative input.

6.1.2 Subjective awareness differences

The difference just described is also important when considering people's subjective awareness ratings of the different techniques. In contrast to the freeze data, people rated all three of the large virtual embodiments as easier to keep track of than the physical arms of direct touch.

Why did people rate these techniques highly, when the freeze-test data shows a different story? One explanation of this seeming contradiction relates to the relative visibility of the different embodiments during the task. Since the relative input techniques could move anywhere on the table, the virtual embodiments were much more likely to be in view of others; in addition, since people could reach for things anywhere on the table, they were also much more likely to be looking around the entire table, and therefore more likely to see others' embodiments. This increased visibility may have given participants the sense that they had up-to-date information about others' locations. In contrast, direct touch forced people to work closer to their home positions, which meant that others' actions were not as often in their field of view. This relative lack of visibility may have convinced participants that they did not know where others were working, even though they could make assumptions based on reach limits (described above).

The visibility of the different virtual embodiments could also explain the discrepancy between the low subjective rankings for the small embodiments (line, arrow, and circle) and their equally-accurate freeze test results.

6.1.3 Subjective distraction

Visibility, however, does not explain why there was so little difference in people's distraction rankings. We were surprised that people did not find the larger embodiments more distracting, given that they showed dramatically more visual information on the table. For example, when reaching across the table, the Arm had an average visible area of 1500cm2, whereas the Circle had a visible area of only 9cm2; despite this 150-times enlargement, the two embodiments were ranked almost equally (Figure 6).

We suggest three explanations for these results. First, it is possible that distraction has more to do with interruption in terms of the task, and less to do with simple visual information. That is, even though some of the embodiments were large, they did not break people's focus on their individual work (the large representations were either translucent, or were shown behind the objects). Similarly, it may be that large objects are not necessarily distracting if they behave in predictable ways, allowing people to accommodate to, and therefore ignore, the incoming information. This is certainly true with real arms; they are large, but were seen in our study as least distracting. Third, our findings suggest that the identity problems of the disconnected embodiments caused distraction – people had to think about who belonged to the Circle and Arrow – that did not occur with the connected embodiments.

6.1.4 Differences in social distance measures

The study showed large differences between real and virtual embodiments in the number of crossings and the amount of time spent in others' personal territories. As with the objective awareness measure, the simple physical limits on reach are a main reason for this difference – for example, people cannot be in parts of the table that they cannot reach. The simple fact of reach could be the reason why people spent more time in others' territories with virtual embodiments, and it is clear that when people have a virtual embodiment, they do spend more time intruding.

However, there were also indications from participants that reach was not everything, and that social distance was at least a minor factor in these results. For example, several participants commented that they felt uncomfortable being too close to another person when using direct touch. Further study is needed to determine whether different embodiments lead to different feelings of social distance. Based on an informal post-hoc analysis of the movement maps (e.g., Figure 7), it does not appear that the participants seated at the ends of the table differed greatly in the amount of time they spent in left or right side of the table. Since only one side was occupied by another person, this suggests that the virtual embodiments may reduce people's discomfort in working close to another person. This issue should be looked at more carefully in future studies.

6.1.5 Differences in preference

The low ranking for the direct touch technique is an interesting result primarily because current design practice for digital tabletops has almost always used direct touch for input. From comments and observations, it seems reasonably clear why relative techniques were preferred; most participants felt that the reach limitations of direct touch were too much of a problem in this task, despite the technique's naturalness and simplicity. We observed numerous situations where an object was just out of reach, and the participant would have to stretch or half rise from their chair in order to obtain it. In contrast, the virtual techniques allowed easy reaching to any part of the table.

It is interesting to note that our task did not explicitly require any participant to reach across the table – they could have completed a great deal of the task working only within their own territories. However, the open-ended nature of the task meant that people could use any objects from anywhere on the table – and it was clear that when they were able to reach, they chose to do so, and to use the entire table as their work surface.

6.2 Generalizing the Results for Designers

Even though this was an exploratory study, we believe that several of our findings are robust enough to generalize to other tables and other tasks. First, generalizing our results to other tables focuses on the issues of table size and participants' reach. On small surfaces where everyone can reach all parts of the table, it is clear that relative input methods are not required. The larger that tables grow, however, the more important relative input becomes, and the more relevant our results become as well.

Second, it is clear that some of our results are strongly related to characteristics of the task. In particular, we believe that in tasks that have stronger requirements for awareness, we might see objective differences between the different virtual embodiments. For example, in a task where people had to hand off objects to another person's cursor, it is possible that the larger embodiments would lead to better performance (this will be tested in a future study). However, we chose our open-ended task because it was more like the casual and unrestricted tasks that we see on real-world tables. In these tasks, there are no strict requirements on maintaining awareness, and therefore it is reasonable to expect our results to transfer to other situations as well.

We do not have enough evidence yet to state our findings as design principles. Nevertheless, there are strong indications in our work that can already be valuable to designers. Therefore, we present a list of issues that designers of tabletop groupware should consider when building interaction support into their applications (at least when working with open-ended tasks and small groups).

- Direct touch may not be the best choice. If the table is large enough to make reaching difficult, then the limits of direct touch could outweigh its advantages.
- Relative pointing can be successful on tables; our participants did not have any problems with accuracy or errors, and the techniques were easy to learn.
- The type of virtual embodiment used with a relative input technique does affect perception of awareness and overall

preference, even though there do not seem to be any differences in objective measures.

- The best virtual embodiment that we tested was the Arm: it was strongly preferred, and several participants liked the way that the natural representation felt like an extension of their own arm. In addition, even though it was large, participants did not find it more distracting than smaller embodiments.

6.3 Future Work

The results raise several questions that warrant further investigation. Since our study was exploratory in nature, we took objective measures that cover several areas of tabletop work. Our task was open-ended and not designed to emphasize any specific dimension of group activity. We plan to carry out additional studies, where we use tasks with stricter requirements for awareness and distraction, to look for performance differences in virtual embodiments.

We believe the limitations that are seen in both direct touch and relative input can be improved by developing new hybrid interactions. New techniques can use direct touch when people work in personal space, but can also allow users to switch to relative input when they need to access distant parts of the table. We are currently working on the design of such techniques.

People ranked the Arm technique highest in terms of preference and awareness. This seems to be partially due to the size and naturalism of the Arm, leading to perceptions that it felt like an extension of users' real arms. The issue of naturalism needs to be investigated in further detail, and future work will compare computer-modeled arms and more realistic arms that are based on real photos.

The issue of social distance with virtual embodiments also deserves further consideration. Future research directions include investigating variable-embodiment techniques that change their size or transparency as they moves closer to another person's work area. We also plan to study proxemics and territoriality issues in other tasks where people are forced to manage shared access to resources.

7 CONCLUSION

In this paper, we explore the tradeoff between real and virtual embodiments, and we explore the design space of virtual embodiments to determine whether factors such as size, realism, and visibility can be used to improve group awareness and coordination. We conducted a study in which seven groups of three people carried out a picture-categorizing task using seven techniques: direct touch and relative input with six different virtual embodiments. The study found substantial differences in all objective measures (awareness freeze test, crossings, and intrusions) between direct touch with physical arms, and relative input with virtual embodiments—direct touch showed more localized movements and higher levels of group awareness than the virtual embodiment techniques. We found no differences between virtual embodiments on any of the objective measures, even though some embodiments are much more visually obvious. The study did find differences between virtual embodiments in subjective perception of awareness, but no difference in perception of distraction. We also found that most participants preferred the relative/virtual techniques over direct touch, and showed a higher preference for virtual embodiments that connect the cursor to the user. Our results provide valuable information to designers of tabletop groupware, and suggest several new directions for future research.

REFERENCES

[1] Bakker, S., Vorstenbosch, D., van den Hoven, E., Hollemans, G., Bergman, T. Weathergods: tangible interaction in a digital tabletop game. *Proc. Tangible & Embedded Interaction 2007*, pp. 151-152.

[2] Benford, S., Bowers, J., Fahlen, L., Greenhalgh, C., and Snowdon, C., User Embodiment in Collaborative Virtual Environments. *CHI 1995*, 242-249.

[3] Brignull, H., Izadi, S., Fitzpatrick, G., Rogers, Y., Rodden, T., The Introduction of a Shared Interactive Surface into a Communal Space. *Proc. CSCW'04*, 49-58.

[4] Dietz, P.H., Leigh, D.L. DiamondTouch: A Multi-User Touch Technology, *Proc. UIST 2001*, pp. 219-226.

[5] Endsley, M.R. Measurement of situation awareness in dynamic systems. *Human Factors* 37, 1 (1995), 65-84.

[6] Endsley, M.R., Kiris, E.O. *Situation awareness global assessment technique (SAGAT) TRACON air traffic control version user guide*. Tech. report TTU-IE-95-02, Industrial Eng., Texas Tech, 1995.

[7] Forlines, C., Wigdor, D., Shen, C., Balakrishnan, R. Direct-touch vs. mouse input for tabletop displays. *Proc. CHI 2007*, 647-656.

[8] Geiβler, J. Shuffle, throw or take it! Working Efficiently with an Interactive Wall, *Proc. CHI 1998*, 265-266.

[9] Greenberg, S., Gutwin, C., and Roseman, M., Semantic Telepointers for Groupware. *Proc. OzCHI 1996*, 54-61.

[10] Gutwin, C., and Greenberg, S., A Descriptive Framework of Workspace Awareness for Real-Time Groupware. *JCSCW*, 11(3), 2002, 411-446.

[11] Ha, V., Inkpen, K., Mandryk, R., Whalen, T., Direct Intentions: The Effects of Input Devices on Collaboration around a Tabletop Display. *Proc. IEEE TableTop 2006*, 177-184.

[12] Hascoët, M., Throwing Models for Large Displays. *Proc. HCI 2003*, British HCI Group, 73-77.

[13] Ishii, H., Kobayashi, M., Grudin, J. Integration of inter-personal space and shared workspace: ClearBoard design and experiments. *Proc CSCW '92*.

[14] Microsoft Surface Computing. http://www.surface.com.

[15] Myers, B., Bhatnagar, R., Nichols, J., Peck, C., Kong, D., Miller, R., and Long, A., Interacting At a Distance: Measuring the Performance of Laser Pointers and Other Devices, *Proc. CHI 2002*, 33-40.

[16] Nacenta, M.A., Pinelle, D., Stuckel, D., Gutwin, C. The Effect of Interaction Technique on Coordination in Tabletop Groupware, *Proc. GI 2007*, 191-198.

[17] Nacenta, M., Aliakseyeu, D., Subramanian, S., and Gutwin, C. A Comparison of Techniques for Multi-Display Reaching, *Proc. of CHI 2005*, 371-380.

[18] 18. O'Brien, R.G. A General ANOVA Method for Robust Tests of Additive Models for Variances, *Journal of the American Statistical Assoc*., 74(368), 1979, 877-880.

[19] Parker, J., Mandryke, R., and Inkpen, K., TractorBeam: seamless integration of local and remote pointing for tabletop displays, *Proc GI 2005*, 133-140.

[20] Pauchet, A., Coldefy, F., Lefebvre, L., Louis dit Picard, S., Perron, L., Guérin, J. TableTops: worthwhile experience of collocated and remote collaboration. *Proceedings of Tabletop 2007*.

[21] Pinelle, D., Dyck, J., Gutwin, C., Stach, T. (2006) Cutouts: Multiple Views for Tabletop Groupware, Tech Report HCI-TR-06-04.

[22] Ramsey, P.H. Testing Variances in Psychological and Educational Research, *Journal of Educational Statistics* 19(1), 1994, 23-42.

[23] Rekimoto, J. and Saitoh, M. Augmented Surfaces: A Spatially Continuous Work Space for Hybrid Computing Environments, *Proc. CHI 1999*, 378-385.

[24] Rekimoto, J. Pick-and-Drop A Direct Manipulation Technique for Multiple Computer Environments, *Proc. UIST 1997*, 31-39.

[25] Scott, S., Carpendale, S., and Inkpen, K., Territoriality in collaborative tabletop workspaces. *Proc. CSCW 2004*, 294-303.

[26] Stach, T., Gutwin, C., Pinelle, D., Irani, P., Improving Recognition and Characterization in Groupware with Rich Embodiments. *CHI 2007*, 11-20.

[27] Swaminathan, K. and Sato, S. Interaction design for large displays. *Interactions*, 4, 1, January 1997, 15-24.

[28] Tang, J. C., and Minneman, S., VideoWhiteboard: Video Shadows to Support Remote Collaboration. *Proc. CHI 1991*, 315-322.

[29] Tang, A., Neustaedter, C., Greenberg, S., VideoArms: Embodiments for Mixed Presence Groupware. *Proc HCI 2006*, 85-102.

[30] Ullmer, B., Ishii, H., Glas, D., mediaBlocks: Physical Containers, Transport, and Controls for Online Media. *SIGGraph '98*, 379-386.