

Mediating Group Dynamics through Tabletop Interface Design

Meredith Ringel Morris, Anthony Cassanego,
Andreas Paepcke, and Terry Winograd
Stanford University

Anne Marie Piper
Microsoft

Anqi Huang
Harvard University

We encounter tables in a variety of situations in our everyday lives—at work, school, and home as well as in restaurants, libraries, and other public venues. The ubiquity of tables results from the utility of their affordances. Tables' horizontal surfaces permit the placement of objects, and their large surface area affords the spreading, piling, and organization of these items. Chairs afford sitting and relaxing, making work around tables leisurely and comfortable. Perhaps most importantly, tables allow face-to-face collaboration among a small group of colocated individuals. One of the primary reasons people perform tasks at tables is because of the social affordances they provide. Consequently, when designing next-generation interactive table technology, this technology's impact on group dynamics is a key issue. The effect of group dynamics on the use of the technology likewise has important bearing on interactive table design.

Our tabletop research efforts at Stanford University have focused on how tabletop user interfaces (UIs) might respond to and even influence a user group's social dynamics. In this article, we provide an overview of four projects:

- Multi-User Coordination Policies (a joint research effort with Mitsubishi Electric Research Laboratories),
- cooperative gesturing,
- educational tabletop interfaces, and
- Shared Interfaces for Developing Effective Social Skills (Sides).

Multi-User Coordination Policies explore how tabletop interfaces can react to breakdowns in group social dynamics to avoid resulting workflow interruptions.

Cooperative gesturing is an interaction technique that can mitigate negative social dynamics (for example, by drawing attention to actions that potentially disrupt other group members' activities) and encourage positive group dynamics, such as participation and socialization.

Our work on educational tabletop interfaces focuses on how subtle variations in tabletop UIs can achieve specific pedagogical goals, such as minimizing the free-rider problem during small-group activities and facilitating collaborative problem solving. We have evaluated several design variations related to these educational concerns.

Additionally, we discuss our work on Sides, which explores how the design of an interactive tabletop computer game can promote desired group interactions among adolescents diagnosed with Asperger's syndrome.


Through all these projects, we learned that many aspects of group dynamics, including conflict, awareness, participation, and communication, can influence and be influenced by interactions with a shared tabletop display.

All of the projects described use Mitsubishi's DiamondTouch,¹ an 85.6-cm by 64.2-cm touch-sensitive surface. Up to four users sit on conductive seat pads. The device identifies which user touched at which point based on capacitive coupling. All four users may interact simultaneously if they wish. DiamondTouch's user-identification feature is key to our interface prototypes' functionality. However, even without a DiamondTouch surface, user identification can occur using a variety of current and emerging technologies, including computer vision and face recognition, biometrics, or fingertip proxies (such as styli or mice) with user-specific signals or tags. All of the projects we describe also use Mitsubishi's DiamondSpin tabletop interface toolkit.² DiamondSpin is a Java toolkit that simplifies developing tabletop UIs by providing a polar-coordinate-based programming model, which

A Stanford University research group explores how design alternatives for tabletop interfaces can impact group dynamics to promote effective teamwork. They built and evaluated a series of novel prototypes that explore Multi-User Coordination Policies and cooperative gesturing, encouraging equitable participation in educational tasks and supporting social skills development for special-needs populations.



1 The duplicate coordination policy reacts to a whole-element conflict by replicating the contested item.



2 The tear coordination policy reactions to a whole-element conflict by mimicking paper-based interactions to draw attention to social breakdowns.

allows flexible specification of objects' orientations on the horizontal display.

Coordination policies

Our exploration of the interplay of tabletop UI design and group dynamics began with our observations of undesirable group dynamics during tabletop use. The status quo for colocated groupware is to assume that social protocols (standards of polite behavior) are sufficient to coordinate a group of users' actions. However, prior studies of groupware use have found that social protocols do not always provide sufficient mediation. For instance, Greenberg and Marwood³ observed that social protocols cannot prevent many classes of conflicts, including those caused by accident or confusion, unanticipated side effects of a user's action, and interruptions or power struggles. Examples of such breakdowns were reported in Izadi et al.'s evaluation of Dynamo,⁴ a large, shared wall display. Studies of Dynamo revealed a problem with *overlaps*—situations where one user's interactions interfered with another's, such as when a user closed a document that belonged to someone else without asking first, or when users were concerned that group members might steal copies of their work without permission. Our own observations of a tabletop document-sharing and annotation application, Mitsubishi's Table-for-N,² yielded additional observations of conflict behaviors, such as a user clearing the screen or quitting the application while other group members were working, and users grabbing digital documents away from others who were using them.


These observations prompted us to conduct two surveys to better understand users' expectations regarding the outcomes of conflicts when using a shared tabletop display.⁵ These surveys presented users with a conflict scenario: Users A and B are sitting across from each

other at a table, and B has a document. User A then reaches across and grabs B's document. In one survey, the table was a traditional piece of furniture, and the document a piece of paper. In the other, the table was a DiamondTouch, and the document was virtual. In either case, respondents were asked to list all the outcomes they could imagine for the conflict described. The survey responses (20 people completed the survey for the paper scenario, and 27 for the digital scenario) revealed several popular envisioned outcomes to the conflicts. We embodied these popular suggestions in our Multi-User Coordination Policies.

Multi-User Coordination Policies⁶ are software mechanisms for avoiding and/or mitigating the effects of conflicts that result from breakdowns of social protocols. Having tabletop interfaces respond to social breakdowns in reasonable, deterministic ways can reduce workflow interruptions and allow users to focus on the task at hand. We developed two classes of coordination policies: policies for global conflicts and for whole-element conflicts. Global conflicts involve changes that impact the application as a whole (such as when one user clears the screen while others are still working). Whole-element conflicts involve disputed access to a single object. The aforementioned survey results influenced this latter class of coordination policies. We created test implementations of each of our proposed policies as an extension of the DiamondSpin toolkit.

Example global coordination policies include these:

- *No holding documents.* Changes that would impact global state can only succeed if no other group members are currently touching digital documents on the tabletop.
- *Voting.* Voting widgets appear in response to proposed global state changes, and group members can vote for or against the change.



3 The relocate interaction technique is one means of implementing a sharing coordination policy. Users can transition a document from a privately accessible state to public availability by moving it from their personal region into the table’s shared, central region.

■ *Privileged objects.* Global changes can only be made from a special menu. The use of a special interface mechanism for these changes is intended to draw attention to the potential for these reserved actions to impact other group members.

Example whole-element coordination policies include the following:

- *Duplicate.* When two users both try to grab the same virtual document, it automatically duplicates itself (see Figure 1).
- *Tear.* A contested document breaks into two pieces, drawing attention to a social breakdown and encouraging negotiation for reassembly (see Figure 2).
- *Rank.* A higher ranking user can always take documents from lower ranking users (for example, a teacher with students).

Duplicate, tear, and rank are examples of reactive policies, which respond to a breakdown with a deterministic behavior. Our policies also include proactive options, which let users keep their documents in a private state to prevent conflicts (such as document grabbing) from occurring in the first place. One example of a proactive whole-element policy is sharing, which lets users dynamically transition a document between a public state (where there are no limits on who can access the item) and a private state (where tabletop identification features enforce access only by the owning user).

We have implemented and evaluated four interaction techniques to allow fluid transitions between public and private document access: release, relocate, reorient, and resize.⁷ The release technique uses a hand-off gesture to transfer document access permissions, while the relocate technique ties access to tabletop location: documents located in the center are publicly accessible, while documents located in a user’s personal space are private (see Figure 3). Reorient uses a document’s orientation to determine access rights (for example, documents facing their owner are private, but turning them to face others makes them accessible). With the resize interaction, when a user enlarges documents, those documents then become public.

Coordination mechanisms beyond social protocols can play a valuable role in tabletop groupware. In addition to preventing conflicts that might arise due to confusion or discord, such policies also help ensure that

groupware has deterministic, predictable responses to multiuser interactions.

Cooperative gestures

Our work on Multi-User Coordination Policies led us to consider other interaction techniques that could potentially mitigate undesired aspects of group behavior by increasing group awareness of important interactions and encouraging a sense of involvement and togetherness. This line of inquiry led us to design cooperative gestures: interactions where a tabletop system interprets the gestures of more than one user as contributing to a single, combined command.⁸ This interaction technique trades off some performance efficiency for the benefits of enhanced collaboration, communication, awareness, and/or fun. These benefits might indirectly improve efficiency by reducing errors or miscommunications, although we have not yet formally explored this possibility.

Several motivating scenarios exist for cooperative gesturing techniques. Interactions that require explicit coordination among multiple users can increase participation or collaboration among group members—this might be particularly relevant for educational activities for children or for special-needs groups (see the “Educational interfaces” and “Interfaces for special-needs populations” sections for more detailed discussions of these issues).


Another application of invoking a group gestural command is for potentially disruptive system commands. In this scenario, the cooperative gesture serves to increase group awareness about irreversible program actions.

Cooperative gestures can also facilitate reach on large surfaces, provide a means for implicit access control of virtual tabletop items, or provide a fun and sociable feel to entertainment-oriented applications. Some of these applications of cooperative gesturing relate to our previous work on Multi-User Coordination Policies—using such a gesture to increase awareness of important system events could be one means of implementing a global coordination policy, while cooperative gestures for access control are a possible mechanism for a whole-element policy.

To explore the properties of cooperative gesture interaction techniques, we developed CollabDraw. CollabDraw is a tabletop art application that lets groups of two to four users create drawings and photo collages. A set of 16 gestures controls CollabDraw: 5 single-user gestures and 11 cooperative gestures. CollabDraw recog-



4 When all four group members simultaneously perform the “neaten” gesture, the result is the cooperative gesture for “organize table.” All photos on the table are swept into a single, central pile. Because the same action has valid single-user and multiuser interpretations in CollabDraw, it exemplifies an additive-effect gesture.



5 CollabDraw adds a creative aspect to drawing by making it a cooperative action, as shown in this modify ink gesture. The left-side user’s finger specifies the (x, y) coordinates where ink will appear, and the right-side user’s fingers specify the resulting stroke’s width.

nizes gestures through a combination of machine-learning techniques and heuristic rules.

One gesture that CollabDraw recognizes is a piling motion made with two hands (see Figure 4), which collects any photos between those hands into a single pile. However, the effect is amplified when all four members of a group simultaneously perform this gesture—the system combines all photos on the table, regardless of whether they are between anyone’s hands, into a single pile in the table’s center. This illustrates one class of cooperative gestures, additive gestures, which have both single-user and cooperative interpretations.

Another additive gesture in CollabDraw’s repertoire is the action for erasing digital ink, which involves rubbing the open palm of the hand back and forth over the marks to be erased. When all four group members simul-

taneously perform the erase motion, the result is to clear all ink on the entire tabletop, rather than only the ink directly underneath the users’ hands. In both these cases, the additive effect (organizing or clearing the table’s contents) could potentially disrupt users’ work activities if invoked unexpectedly by a teammate. Thus, we require the group’s input to collectively trigger those events.

The modify ink gesture interprets the actions of more than one user to determine the appearance of the digital ink that can be drawn on the tabletop. For example, one user’s finger determines where the ink is drawn, while another user can control the thickness of the drawn line by placing two fingers on the table’s surface and varying the distance between them (see Figure 5). This gesture’s design is intended to increase users’ sense of team involvement during the activity by coupling users’ input in a creative and entertaining manner.

Exiting CollabDraw also requires a cooperative action, to prevent accidental invocation of this command while other group members are still drawing. The exit gesture requires all group members to hold hands, and the user at the end of this chain touches the table’s surface (see Figure 6). CollabDraw recognizes hand-holding by exploiting DiamondTouch’s special properties: When any one member of a hand-chain touches the table’s surface, the electrical coupling between the users causes the device to sense that all four group members have simultaneously pressed at a single point.

Our evaluation of cooperative gesturing involved seven pairs of subjects, who were trained on each of the gestures. Subjects then had to recreate a target drawing using the gestures they had learned and provided questionnaire feedback about their experience. This evaluation highlighted several lessons relevant for designing future iterations of cooperative gesture interfaces:

- the need to develop techniques to avoid accidental invocation of cooperative actions by coincidental, simultaneous performance of single-user gestures with additive meanings;
- the impact of proxemics (intimacy levels) on gesture design; and
- the potential tedium of using cooperative interactions for frequently invoked system commands.

A more complete description of the other gestures that CollabDraw recognizes as well as details on our evaluation procedure and results are available elsewhere.⁸

Based on our experiences designing, implementing, and evaluating an initial set of cooperative gestures, we articulated some important axes of the design space for these interactions. The seven design axes we identified are symmetry, parallelism, proxemic distance, additivity, identity awareness, number of users, and number of devices. An interesting avenue for future work is to analyze the impact of these design axes on cooperative gestures’ learnability, memorability, usability, and naturalness.

Our initial experiences with cooperative gesturing suggest that integrating this interaction technique into tabletop systems can emphasize particularly important application events while promoting a highly social interactive experience.

Educational interfaces

Coordination policies and cooperative gestures aim to mitigate the effects of poor team coordination. Our next project tackled the issue of group dynamics and tabletop displays with a more ambitious goal: to actively encourage desired group work styles through UI design. The importance of collaboration in small-group work and methods for facilitating effective group work, specifically through group problem-solving tasks, is a prominent research topic in the education field. Small group work is particularly valuable in the foreign-language-learning domain, where peer-to-peer interaction provides important opportunities to practice conversational skills. Interactive table technology is an exciting new platform for educational activities because it combines face-to-face small group work with the benefits of digital media for providing students with immediate feedback about their progress. Such technology also allows flexible adaptation of content for different topics and skill levels.

As part of our efforts to understand how we can design interactive table applications to support educational group work, we conducted interviews with foreign language teachers at our institution. These teachers indicated that the free-rider problem—someone not participating as much as other group members—is a key concern when allowing students to work in groups. Thus, we chose to explore how we could design tabletop UIs to reduce free riders by increasing participation equity among group members. In the previous section, we discussed cooperative gesturing, which can lend a social and participatory feel to an application by requiring the active involvement of all group members to execute certain actions. Adding cooperative gestures to an educational activity could be one means of ensuring participation by all students in a group. However, our evaluation of the CollabDraw system found that excessive use of cooperative gesturing could make an activity tedious. Consequently, we decided to investigate additional, more lightweight UI designs that might also encourage more participation in tabletop activities for educational tasks.


We created three tabletop language-learning applications with flexible, adaptable structures to fit varying content:

- MatchingTable, which lets students match words and phrases with images (see Figure 7);
- PoetryTable, which lets students create free-form sentences from word tiles; and
- ClassificationTable, which lets students sort vocabulary words or sentences into one of four corners of the table based on various properties (see Figure 8 on the next page).

We explored four design variations of these applications, to assess the impact on participation equity: feedback modality, feedback privacy, interaction visualizations, and spatial configuration.

Feedback modality

To explore the impact of feedback modality on group participation during an educational tabletop activity,



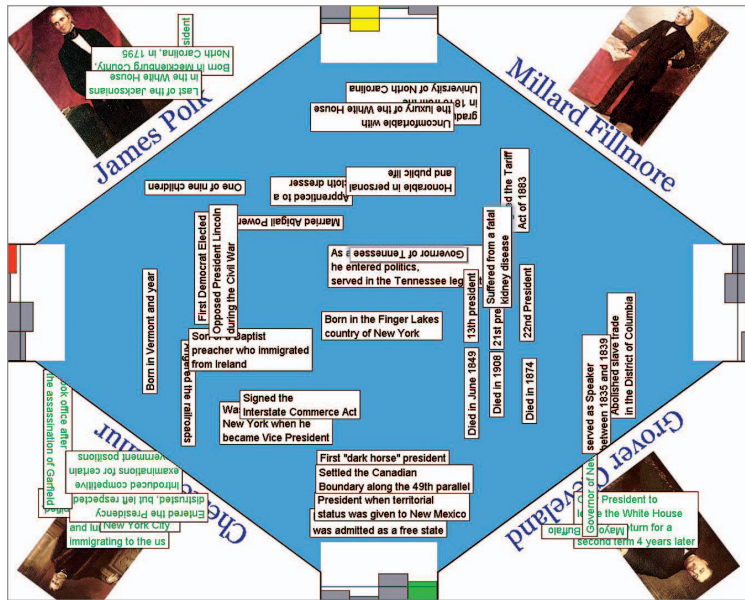
6 Exiting CollabDraw requires a cooperative gesture, to prevent accidental invocation of this potentially disruptive command. To perform the exit gesture, all users hold hands in a chain and the last user touches the tabletop.



7 Students in one of Stanford University's English as a Second Language courses interact with the MatchingTable, combining English descriptions with photos of their classmates. The one-eared headsets they wear serve several purposes: during some activities, touching words on the tabletop causes them to be pronounced through the earpiece. The earpiece can also provide individually targeted audio feedback regarding the correctness of an action. The microphones on each headset gather speaker-volume data that can optionally be displayed via interaction visualizations.

we had 32 paid subjects (divided into eight groups of four each) who had taken Spanish-language courses complete a series of ClassificationTable activities. These activities used Spanish-language clues describing sets of Latin American countries, to simulate a foreign-language-learning activity. Subjects were asked to speak in Spanish during the activity.

One variation we explored was the modality through which we provided feedback regarding the correctness of clue classifications in the Spanish ClassificationTable application. We explored two alternatives: visual feedback (the clues' text turned either green or red to indicate correct or incorrect placement) and audio feedback (either an upbeat or discordant tone was played to indicate correct or incorrect placement). We hypothesized that audio feedback would increase the amount of con-



8 This screen shot shows a ClassificationTable activity. Groups of students work together to assign each fact to the US president that it describes. Interaction visualizations in front of each user’s seat reflect how many clues that student has classified relative to other group members, with the goal of encouraging more equitable distributions of work among group members. In this example, the yellow user (top) has been most active in classifying clues, while the blue user (right) has interacted the least.

versation among group members (important for rehearsing a foreign language) by breaking the ice and inserting noise into the environment, making it less awkward for students to generate their own noise by talking. Audio feedback should also increase students’ awareness of their participation levels because it pushes information to them.

By comparing subjects’ self-assessments of their contribution to the activity (in terms of number of clues classified) with logs of actual tabletop interaction, we confirmed that audio feedback increased subjects’ accuracy for these estimates as compared to visual feedback. Subjects also conversed more in the presence of audio, rather than visual, feedback, speaking in Spanish during 74.3 percent of the session (rather than only 60.8 percent with visual feedback).

Feedback privacy

Another interface variation we explored was whether feedback regarding the correctness of clue placement in the ClassificationTable activity was conveyed publicly (to the entire group) or privately (only to the group member who moved a particular clue). We explored these conditions as part of the same experimental sessions as the feedback modality variants, with the Spanish-language-learning content. To explore private feedback in the shared environment context, we used individually targeted audio feedback via one-eared headsets. We hypothesized that private feedback would increase participation equity by reducing the potential for embarrassment over incorrect answers, and thereby encouraging shy and underperforming students to contribute more to the activity. We also hypothesized

that private feedback would increase the accuracy of students’ self-assessments of performance by drawing more attention to their individual contributions.

We found that when subjects received private feedback, they tended to clearly emphasize positive results to the group, via thumbs-up gestures or exclamations of “Sí,” “Bueno,” and “Yeah!”. Negative feedback was typically acknowledged more subtly, by a slight head shake or the user moving the incorrect clue back into the central region without commenting. This lack of drawing attention to the negative feedback supports our design motivation for providing such feedback privately—to reduce potential embarrassment over incorrect answers by not pointing them out to the entire group. Likert-scale questionnaire responses confirmed that users felt less self-conscious about their performance when using the private audio feedback, and their comments, such as “I prefer private audio, so I can know what I got correct or incorrect with a sense of confidentiality,” echoed this trend. However, subjects’ sentiments were not strongly reflected in their actions in terms of clear changes in patterns of participation equity. We suspect that the artificial experimental situation might have impacted this. That is, subjects’ behavior was not strongly influenced by feelings of embarrassment since they were not in a real classroom situation, and were not likely to encounter their coparticipants in the future.

Interaction visualizations

DiMicco et al. explored the impact of a real-time visualization on group participation in a planning task (where participation referred to the amount each person contributed to a conversation).⁹ For the planning task, DiMicco et al. found that overcontributors spoke less in the presence of the visualization, but that undercontributors did not increase their participation because they did not believe the display was accurate.

Inspired by this study, we integrated real-time histograms into some versions of our educational tabletop activities (see Figure 8). The histograms appear on the table in the region directly in front of each user and can reflect either the number of answers contributed by each group member based on tracking touch interactions with the table or the amount of time each user has spent speaking relative to other group members.

We hypothesized that the histograms in our application would increase participation equity and would have greater impact than in the setting studied by DiMicco et al. because the histograms should have more credibility to users in the context of a computer-mediated activity. We also altered the design of the visualizations that DiMicco et al. described to make them more appropriate for a collaborative educational activity by having a customized visualization for each group member (instead of a single visualization for viewing by the entire group). We highlighted the current user’s bar in color and grayed out the bars representing the other three users, so that students would feel they were comparing themselves more to a group average than competing directly against individual group members.

We evaluated the impact of interaction visualizations in a separate study, completed by 40 subjects divided into

three- and four-person groups. Subjects used three variants of the ClassificationTable. Before the experiment, they read articles about past US presidents, first ladies, and Supreme Court justices. Facts from these articles were then presented as part of the ClassificationTable exercises, to simulate a classroom experience. In one variant, interaction visualizations showed each user how many clues they had classified relative to other group members. In another, visualizations showed the amount each user spoke during the activity relative to everyone else in the group, and another variant had no visualizations at all. We hypothesized that the presence of visualizations showing speaking contributions would result in more participation equity in terms of amount of conversation contributed by each user, while the presence of visualizations showing the amount of clues classified would increase participation equity in that domain.

We found that the visualizations reflecting speech contributions increased participation equity among group members in terms of amount of conversation contribution by each person ($F(1, 2) = 3.93, p < 0.05$). The classification visualizations, on the other hand, did not produce a change in classification participation equity relative to the other two conditions. Our observations during the sessions offer a possible explanation for this difference in impact. During the sessions with speech visualizations, we noticed that many users would engage in filler speech—speech that was related to the activity (such as reading the clues aloud or asking other group members for their opinions), but which did not offer a concrete contribution to the activity by suggesting an answer or offering information. However, these types of filler actions were not possible for the clue classification, where any action would involve a concrete attempt to make a statement about the answers. We need to do further work in this area to determine whether awareness visualizations can impact contribution in an educationally beneficial way, rather than only motivating the production of filler behaviors.

Spatial configuration

We altered the initial configuration of clues and photos in our three tabletop activities to explore the impact of initial layout on participation. In the “four piles” design, we initially placed all objects (clues, word tiles, photos, and so on) into four random, equally sized virtual piles near the four users’ seats around the table’s borders. In the “central pile” design, we initially placed all objects into a single virtual pile in the center of the table.

We informally explored these spatial configuration variations as part of the use of our tabletop applications in an English as a Second Language course at Stanford. About 20 students from this class used the table during two sessions over the course of two weeks, completing activities that tied in to the current lessons in their course.

We hypothesized that the four piles design would increase participation equity as compared to the cen-

tralized design by making undercontributors feel more responsibility for the items that started out nearest them and by making overparticipants hesitant to reach out and take responsibility for objects that originated near others. This hypothesis is in accordance with Scott et al.’s studies of tabletop group work that have found that a table’s central area is considered a group-owned, public space, while the areas directly in front of each user are considered personal or private zones.¹⁰

Trends in the data collected during the ESL class’s use of the ClassificationTable, MatchingTable, and Poetry-Table applications support our hypothesis that laying out information in four piles—one near each group member, rather than in the center of the table—encouraged more equitable participation. For each student group, we calculated the standard deviation of the percent of

touch events contributed by each group member, and of the percent of talking time contributed by each group member. Lower standard deviations reflect more equitable contributions among group members. For each of the three tabletop activities completed, groups had lower standard deviations for both the percent of touch interactions contributed and the percent of conversation contributed under the four piles condition. However, students’ comments on questionnaires distributed after the activities reflected a

potential drawback of the four piles layout, suggesting that it detracted from the activity’s collaborative feel.

More details about our educational software, testing scenarios, and statistical analyses are available elsewhere.¹¹ Overall, our experiences in designing and evaluating educational tabletop groupware indicate that relatively small UI variations, such as the modality or privacy of feedback, the layout of objects on the shared surface, or the presence of explicit awareness information, can produce observable impacts on students’ team work styles. The ability of UI design to facilitate specific pedagogical goals presents an exciting, and potentially high-impact, area for further research.

Interfaces for special-needs populations

Continuing our theme of exploring the potential for tabletop interface design to promote positive group dynamics, we created a tabletop interface for a user population for whom social interactions are a particularly relevant topic: youths diagnosed with Asperger’s syndrome (AS). AS is a pervasive developmental disorder and is considered an autism spectrum disorder. Individuals with AS have difficulty understanding accepted social conventions, reading facial expressions, interpreting body language, and understanding social protocols. These social deficits can lead to challenges in learning effective group work skills, including negotiation, perspective taking, and active listening.

Most computer programs for social skills development are designed for one user working directly with the application, and lack the face-to-face interaction found in

**The ability of UI design
to facilitate specific
pedagogical goals presents
an exciting, and potentially
high-impact, area for
further research.**



9 The Sides tabletop computer game requires a group of children to form a path enabling a frog to travel from the start lily pad to the end lily pad. Along the way, if this path intersects with insects, the group scores bonus points. The distribution of arrow pieces forces all four users to cooperate in order to form a valid path. Buttons along each side of the table enable users to vote on key game interactions (for example, testing the path or quitting).

authentic social situations. Our contextual inquiry at a local middle school special education classroom revealed discontent with current classroom board games and group activities for social skills therapy. Tabletop technologies are more compelling for this user population than traditional board games for several reasons:

- many individuals with AS describe an affinity for technology,
- computer technology has the ability to enforce basic game rules (freeing up therapists’ time to deal with higher level group issues), and
- computer games have the flexibility to adapt content and difficulty level.

Tabletop technology is a unique platform for multiplayer gaming, combining the benefits of computer games with the affordances of face-to-face interaction. We were particularly interested in exploring how tabletop interface design could impact group dynamics for this unique user population, who stand to benefit from such systems.

Based on the results of our contextual inquiry process, which involved observing and interviewing students aged 11 to 14 with AS in a middle school social skills therapy class and their therapist, we developed the Sides game.¹² Sides is played by four students with AS who are supervised by their social skills therapist (so that he or she can help the students discuss and understand issues that arise during gameplay). The game aims to create a path enabling a frog to travel between two specific lily pads (see Figure 9). If this path intersects a variety of insects located on the table, players earn extra

points. The path is constructed from directional arrow pieces; the distribution of different types of arrow pieces among the four players ensures that they cooperate and negotiate with each other to construct a valid path.

Sides incorporates several features designed with the particular social needs of AS students in mind. The social coordination challenges characteristic of AS heighten the need of this user group for some of the control mechanisms, such as coordination policies and cooperative gestures, discussed earlier. Thus, we designed Sides to include system-enforced ownership of game pieces, an example of the “private” whole-element coordination policy. The DiamondTouch table’s user-identification features enforce turn-taking behaviors and prevent users from interacting with arrow pieces that belong to other group members. These restrictions necessitate reliance on social interactions (such as offering suggestions or asking questions) to complete the game, rather than allowing a student to bypass conversation with group members by completing the entire puzzle independently. Also, Sides includes buttons that must be pressed by all four players to achieve certain key game events, an example of a cooperative gesture for achieving global coordination. These special voting buttons are located in front of each user’s seat and must be activated by all four group members to test the current path for correctness or to quit the game. This feature encourages negotiation and coordination among the students, requiring them to strategize and make requests of their teammates.

Three groups of students from a local special-needs school played Sides under the guidance of their social skills therapist. Initial reactions from students, the therapist, and the students’ parents have been encouraging. Our work on Sides provides a starting point for thinking more broadly about user populations and computing scenarios that can benefit from the social computing experience that tabletop technology provides.

Discussion and conclusion

The projects described in this article represent an initial step toward considering social dynamics when designing software for multiuser shared-surface settings. We found that providing users with the ability to dynamically change document accessibility on a tabletop can facilitate shared document inspection and alteration, while also letting users guard their data against intentional or inadvertent tendencies of users to interfere with other group members’ content on a shared tabletop. Providing this type of built-in structure for an activity (such as computer-enforced access control, turn-taking, and so on) can also make tabletop software more accessible to special-needs groups, such as adolescents with AS. For all user groups, however, making items’ current accessibility status visible is important for preventing confusion regarding accessibility.

In addition, cooperative gesturing is one means of facilitating active engagement of all group members with a tabletop application. Explicit awareness information that a tabletop system provides is another means to help users regulate their participation levels. This information can be subtle (for example, the use of audio feedback attached to particular interactions can keep users

aware of group members' activity levels) or more explicit (for example, individually targeted interaction visualizations can convince users to change their levels of activity). We also found that the location of key interaction items on a tabletop display impacts participation patterns. Next-generation tabletop systems could monitor patterns of interaction and proactively move data objects near particular users' seats to encourage them to become more involved in an activity.

By considering tabletop design holistically, including both the human-computer and human-human interactions that take place during tabletop activities, we hope to ultimately create more usable and useful tabletop groupware. ■

Acknowledgments

We thank Mitsubishi Electric Research Labs for donating a DiamondTouch table and for their contributions to the coordination policies work. We also acknowledge Eileen O'Brien for her contributions to Sides.

References

1. P. Dietz and D. Leigh, "DiamondTouch: A Multi-User Touch Technology," *Proc. User Interface Software and Technology (UIST)*, ACM Press, 2001, pp. 219-226.
2. C. Shen et al., "DiamondSpin: An Extensible Toolkit for Around-the-Table Interaction," *Proc. Conf. Human Factors in Computing Systems (CHI)*, ACM Press, 2004, pp. 167-174.
3. S. Greenberg and D. Marwood, "Real Time Groupware as a Distributed System: Concurrency Control and its Effect on the Interface," *Proc. Conf. Computer-Supported Cooperative Work (CSCW)*, ACM Press, 1994, pp. 207-217.
4. S. Izadi et al., "A Public Interactive Surface Supporting the Cooperative Sharing and Exchange of Media," *Proc. UIST*, ACM Press, 2003, pp. 159-168.
5. M.R. Morris et al., "Conflict Resolution in Paper and Digital Worlds: Two Surveys of User Expectations," *CSCW Conf. Supplement*, ACM Press, 2004.
6. M.R. Morris et al., "Beyond Social Protocols: Multi-User Coordination Policies for Co-located Groupware," *Proc. CSCW*, ACM Press, 2004, pp. 262-265.
7. M. Ringel et al., "Release, Relocate, Reorient, Resize: Fluid Interaction Techniques for Document Sharing on Multi-User Interactive Tables," *Proc. CHI Extended Abstracts*, ACM Press, 2004, pp. 1441-1444.
8. M.R. Morris et al., "Cooperative Gestures: Multi-User Gestural Interactions for Co-located Groupware," *Proc. CHI*, ACM Press, 2006, pp. 1201-1210.
9. J. DiMicco, A. Pandolfo, and W. Bender, "Influencing Group Participation with a Shared Display," *Proc. CSCW*, ACM Press, 2004, pp. 614-623.
10. S.D. Scott, M.S.T. Carpendale, and K. Inkpen, "Territoriality in Collaborative Tabletop Workspaces," *Proc. CSCW*, ACM Press, 2004, pp. 294-303.
11. M.R. Morris et al., *Supporting Cooperative Language Learning: Interface Design for an Interactive Table*, tech. report, Computer Science Dept., Stanford Univ., 2005.
12. A.M. Piper et al., "SIDES: A Cooperative Tabletop Computer Game for Social Skills Development," *Proc. CSCW*, ACM Press, 2006, in press.



Meredith Ringel Morris is a researcher in Microsoft Research's Adaptive Systems and Interaction group. Her research interests include human-computer interaction, computer-supported cooperative work, and educational technologies. Morris has a PhD in computer science from Stanford University. Contact her at merrie@stanfordalumni.org.




Anne Marie Piper is a user researcher in Microsoft's Entertainment and Devices Division. Her research interests include the design and evaluation of techniques to support colocated cooperative technology use. Piper has an MA in learning, design, and technology from Stanford University. Contact her at ampiper@microsoft.com.



Anthony Cassanego is an undergraduate student in computer science at Stanford University. Cassanego's research interests include computer-mediated interaction and interaction visualizations. Contact him at tonyc@cs.stanford.edu.



Andreas Paepcke is a senior research scientist and director of the Digital Library and BioACT Projects at Stanford University. His research interests include interfaces for small devices, Web search, and browsing facilities for difficult-to-index artifacts. Paepcke has a BS and an MS in applied mathematics from Harvard University and a PhD in computer science from the University of Karlsruhe, Germany. Contact him at paepcke@cs.stanford.edu.



Terry Winograd is a professor of computer science at Stanford University, where he codirects the Human-Computer Interaction Group and the teaching and research programs in HCI design. He is also a founding faculty member of the Stanford Institute of Design. Contact him at winograd@cs.stanford.edu.



Anqi Huang is an undergraduate in computer science at Harvard University. His research interests include cooperative gestures and ways to design graphical interfaces to deter phishing. He is a recipient of the Arthur D. Little Innovations Fellowship. Contact him at huang3@fas.harvard.edu.