Collaborative Tabletops for Blind People: The Effect of Auditory Design on Workspace Awareness

DANIEL MENDES, INESC-ID Lisboa, Instituto Superior Técnico, Universidade de Lisboa / INESC TEC, Faculdade de Engenharia, Universidade do Porto

SOFIA REIS, Instituto Superior Técnico, Universidade de Lisboa

JOÃO GUERREIRO, INESC-ID Lisboa, Instituto Superior Técnico, Universidade de Lisboa / LASIGE, Faculdade de Ciências, Universidade de Lisboa

HUGO NICOLAU, INESC-ID Lisboa, Instituto Superior Técnico, Universidade de Lisboa

Fig. 1. Blind users collaborating on an interactive tabletop. Left: after receiving spatial audio feedback about her partners' findings, a user drags her finger to that location. Middle: a user found a relevant object for her partner, proceeding to grab and guide his hand to the object's position. Right: without knowing where their partners are exploring, the users experience unintentional physical contact.

Interactive tabletops offer unique collaborative features, particularly their size, geometry, orientation and, more importantly, the ability to support multi-user interaction. Although previous efforts were made to make interactive tabletops accessible to blind people, the potential to use them in collaborative activities remains unexplored. In this paper, we present the design and implementation of a multi-user auditory display for interactive tabletops, supporting three feedback modes that vary on how much information about the partners' actions is conveyed. We conducted a user study with ten blind people to assess the effect of feedback modes on workspace awareness and task performance. Furthermore, we analyze the type of awareness information exchanged and the emergent collaboration strategies. Finally, we provide implications for the design of future tabletop collaborative tools for blind users.

CCS Concepts: • **Human-centered computing** → **Accessibility**; **Empirical studies in accessibility**; **User studies**.

Authors' addresses: Daniel Mendes, danielmendes@fe.up.pt, INESC-ID Lisboa, Instituto Superior Técnico, Universidade de Lisboa / INESC TEC, Faculdade de Engenharia, Universidade do Porto; Sofia Reis, sofia.reis@tecnico.ulisboa.pt, Instituto Superior Técnico, Universidade de Lisboa; João Guerreiro, jpguerreiro@ciencias.ulisboa.pt, INESC-ID Lisboa, Instituto Superior Técnico, Universidade de Lisboa / LASIGE, Faculdade de Ciências, Universidade de Lisboa; Hugo Nicolau, hugo.nicolau@tecnico.ulisboa.pt, INESC-ID Lisboa, Instituto Superior Técnico, Universidade de Lisboa.

1 INTRODUCTION

In the past decade, interactive tabletops have been successfully used in multiple domains, including healthcare, education, entertainment, and cultural exhibitions [2, 18, 34, 36]. These devices show significant advantages in supporting multi-user interaction, namely due to their large shared surface. Also, tabletop geometry has been shown to encourage equitable participation among users as well as simultaneous individual interactions with the digital content [32, 33]. Previous research highlights several benefits of using interactive tabletops: improving collaborative learning [37, 55], supporting reflection-type conversations [32], enhancing social interaction [18], fostering creativity as well as engagement [10, 17]. However, people with visual impairments can struggle to engage with such large surfaces [23, 39]. We argue that not being able to use interactive tabletops, for example, due to visual impairments, and consequently participate in these group activities, can be a vehicle of social exclusion.

Despite being an inherently visual technology, touchscreen devices can be used by blind users. Screen readers, such as Apple's VoiceOver [3] or Android's Talkback [20], allow users to explore and control the device by providing audio feedback for touch actions. However, when considering the spatial awareness of interface elements, these accessibility services are mostly designed for smaller form-factors, rather than large collaborative surfaces where the ability to locate items and establish relationships between them is more challenging. Examples of applications for such devices include exploring maps [39] and anatomic models in educational settings, or mind maps in brainstorming sessions [53]. In these, the ability to locate artifacts without losing spatial awareness and to relate them is relevant. While there is a large body of work on touchscreen accessibility, research has been restricted to single-user interaction [15, 21, 23, 39].

This raises the question of how accessibility services can support nonvisual collaborative activities on large touchscreen devices. Particularly, how can we inform users about the actions of others? This notion of monitoring the activity of others, which provides context for one's activities is defined as *workspace awareness* [14, 28]. Workspace awareness is a crucial aspect of collaboration [28] and often relies on visual display techniques [30]. However, little is known about using audio for tabletop multi-user interaction, particularly on how to provide workspace awareness and support blind people in joint activities.

We present a study that contributes to the understanding of auditory design in enabling nonvisual collaboration in large touchscreen surfaces. We developed a nonvisual tabletop prototype, shown in action in Figure 1, with three feedback modes. We examined workspace awareness in co-located collaborative activities to answer: is auditory feedback effective in supporting nonvisual collaboration in large touchscreens? How does the amount of auditory information delivered influence patterns of awareness information exchange? How do blind users leverage audio feedback to engage with one another in co-located collaborative tasks? How is task performance in large touchscreens affected by the amount of auditory information displayed?

The main contributions of the paper are: first, the design of multi-user auditory feedback for nonvisual tabletop collaboration; second, analysis of the effects of auditory feedback on blind users' workspace awareness and task performance when using interactive tabletops; third, we describe emergent nonverbal collaboration behaviors when blind users engage in joint activities; and finally, we propose a set of implications for the design of multi-user auditory displays.

2 RELATED WORK

The related work reviewed in this section is two-fold: first, we discuss research on touchscreen accessibility for visually impaired users, including mobile devices and large surfaces; second, we examine previous attempts to create collaborative systems that leverage audio output as a feedback modality.

2.1 Touchscreen Accessibility

Mobile touchscreen devices have built-in screen readers that allow blind people to use their devices. They enable users to explore the screen by dragging their fingers and having the interface elements (e.g., buttons, labels, images) read aloud as they touch them (i.e., Explore by Touch). Alternatively, users can perform directional swipes to cycle through interface elements sequentially. A double-tap or split-tap selects the last heard element [38]. Over the last decade, numerous projects aimed at improving touchscreen interfaces, including understanding the performance of gestures [41], proposing new text input techniques [6, 24, 51, 52], and leveraging Braille knowledge [4, 58]. Other solutions took advantage of concurrent speech feedback [23], augmented touchscreens with physical interfaces [40, 60, 62] or haptic feedback [19, 46]. Touchscreen devices have also been combined with computer vision-based solutions to create novel assistive technologies [26, 27]. For more examples, Grussenmeyer and Folmer [21] provide an extensive survey of accessible touchscreen technologies.

Regarding large touch surfaces, Guerreiro et al. [25] investigated the exploration strategies employed by blind users when two-hand interaction is supported in a tabletop. Bardot et al. [5] also studied two-hand exploration with bi-lateral audio and vibrotactile feedback using a smartwatch on each wrist. Kane et al. [39] proposed three input techniques that enhanced performance with these devices by removing the need to perform absolute spatial explorations. These techniques were evaluated in an interactive map application. Indeed, this is the most frequent application associated with large touch surfaces for visually impaired people; several projects relied on large tactile maps augmented with auditory feedback and projections [1]. A detailed review was conducted by Brock et al. [7]. More recently, Ducasse et al. proposed the use of tangibles over interactive tabletops to ease blind users' spatial explorations [15]. The authors proposed a new type of tangible that uses a retractable reel to render digital links into physical. Another existing alternative for making large interactive surfaces accessible to the blind are large tactile displays [59], which can help preserve users' spatial memory. However, these do not scale properly for multiple users.

Overall, although efforts have been made to make large touch surfaces accessible to blind users, the potential to use these devices in collaborative activities with blind people has remained mostly unexplored.

2.2 Collaboration and Audio

According to Gutwin and Greenberg [28], it is important to consider workspace awareness when designing systems to support collaborative work over a shared workspace. This includes providing knowledge about where the other users are and what actions they are performing. Naturally, this involves monitoring both the shared space and others, as well as displaying one's activities to the group [14]. Audio has been shown to support awareness in groupware systems when combined with graphics to overcome limitations associated with visual displays [30, 31]. More recently, it has been shown that it can be used as a sole means for supporting nonvisual collaboration between sighted users in accessing and editing shared hierarchical menu structures [48].

One question that arises in auditory collaborative systems is the effect of the different means of delivering audio feedback on collaborative work. Hancock et al. [31] assigned a different timbre to

each user, showing increased group awareness but at the cost of decreased individual performances. Others [54] investigated the use of spatialized speech feedback, showing that users leverage the spatial position of their peers' speech to maintain awareness of their interactions with the shared space. Morris et al. [50] compared the use of headphones and shared speakers to deliver auditory feedback to a group of tabletop sighted users. They found that collaboration strategies changed when users wore headphones and that using headphones does not impede group communication. Although numerous projects have studied the effect of auditory designs on collaborative work, these studies focus on combining it with graphical interfaces.

In a seminal work, Winberg and Bowers [61] showed that visually impaired users could inter-weave their active manipulations of collaborative interfaces with talk; however, missing content changes from peers could result in disengagement from the task. McGookin and Brewster [47] investigated users collaborating to create graphs with an audio-haptic tool. The use of shared audio enabled awareness; however, the amount and type of audio needed seemed dependent on the strategy adopted (divide and conquer vs. turn-taking). Buzzi et al. [9] examined the accessibility challenges of an online text editing tool, highlighting the need for "accessible awareness" where the status and actions of users should be visible. On a similar topic, Das et al. [13] conducted interviews with 20 visually impaired professionals that use collaborative writing tools. They highlight the need for better auditory designs and the potential of using spatial audio, concurrent feedback, and multiple voices to enhance the visibility of collaborative actions. Kunz et al. [42] proposed to make pointing gestures above tabletop devices accessible to blind users as nonvisual cues are crucial in collaborative work. More recently, Pölzer et al. [53] presented a system to support blind people to participate in co-located brainstorm sessions with mix-ability teams by detecting pointing gestures and transcoding visual information to audio feedback. More broadly, Luque et al. [45] proposed a framework to analyse collaborative settings involving people with disabilities in order to identify potential challenges and coping mechanism. All these studies emphasize the importance of designing for workspace awareness in collaborative systems and the challenges of doing so through auditory feedback.

3 USER STUDY

Research on accessible large touchscreen surfaces is largely limited to single-user interaction. To the best of our knowledge, we are the first to investigate the effect of auditory design on tabletop awareness and how such information is then used by blind users in co-located collaborative activities.

Our main goal in this user study was to understand how the amount of auditory information affects both workspace awareness information exchange and task performance. Furthermore, we investigated how the existence of onscreen distractors – as a variable of task difficulty – affects each feedback mode. To this end, we designed and implemented a nonvisual tabletop prototype that can be customized to provide different auditory feedback modes.

This study aims to answer four main research questions:

(1) Is auditory feedback effective in supporting nonvisual collaboration in large touchscreens?
(2) How is task performance in large touchscreens affected by the amount of auditory information displayed?
(3) How does the amount of auditory information delivered influence patterns of awareness information exchange?
(4) How do blind users leverage audio feedback to engage with one another in co-located collaborative tasks?

3.1 Collaborative Task

Tabletops excel in conveying 2D spatial relations, due to their geometry and dimensions. We chose the fundamental tasks of exploring the screen and selecting targets. It is essential to highlight that given the size of a tabletop device, such a task can be challenging to complete using a screen reader [23, 39]. Furthermore, note that although sequential navigation – using directional swipes – can be used independently of the screen size, it does not convey information about the location of items and, therefore, is not appropriate for applications such as maps, diagrams, or 2D models, where the spatial relationship between interface elements is relevant.

We had three types of targets: squares, circles, and triangles. We asked pairs of participants to select all squares and circles. The task consisted of selecting five squares and five circles. Triangles were considered distractors and simulate irrelevant interface elements. Participants were informed that they could only select one type of target: either squares or circles; however, the overall goal was to finish selecting both types of targets as quickly as possible.

We developed a stimuli application responsible for generating target positions within a 9 x 5 grid (45 positions), as exemplified in Figure 2. All artifacts on the screen had an associated label that described its shape; i.e., "square", "circle", or "triangle". We used a sparse target layout to give the illusion of a freeform layout and avoid placing all targets on a limited screen area. For that purpose, we used three rules that needed to be met for a partially-randomized generation of artifacts' location: (1) each artifact needs to be at least one cell apart from the others (both horizontally and vertically); (2) there should be precisely one artifact type per grid line so that they are equidistant from both participants; (3) there should be at least two artifacts of each type on each side of the tabletop, to avoid clustering all targets on the same side. Participants were not informed that targets were distributed in a grid layout. All targets were 10cm x 10cm. The stimuli application recorded all screen layouts, including the type of targets, positions, and all users' touch interactions. For this evaluation, the visual display was removed from the tabletop.

3.2 Auditory Feedback Modes

Our prototype is based on the traditional explore-by-touch approach used in mainstream screen readers where users can drag a finger to explore onscreen content while interface elements are read aloud. Double-tapping selects the target.


Fig. 2. Targets distributed in accordance with our position generation rules.

Our multi-user nonvisual interactive tabletop provides users with a unique text-to-speech voice to support collaboration. For two people, each user is assigned a different voice – either male or female – as audio frequency is a good parameter to maximize speech discrimination [22]. For more than two users, voice variations can be used. Such approach has been previously used to support two-handed interaction [25].

Furthermore, we use spatialized speech to display target positions on the tabletop. Spatialization is relative to the user's location, which means that users can freely move around the tabletop, and auditory feedback dynamically adapts to their physical location relative to the target position. In terms of collaborative work, the spatialized feedback implicitly illustrates the screen areas being explored by partners, thus increasing the potential to increase workspace awareness.

We investigated three auditory feedback modes, which employ different trade-offs between privacy and visibility [16] of the users' information:

(1) *Private* - behaves similarly to a traditional screen reader as speech feedback is restricted to users' own actions;
(2) *Public* - users receive audio feedback about their own actions as well as all the actions of partners;
(3) *Task-Dependent* - in addition to their own actions, users receive feedback about their partners' actions that are relevant for the collaborative task. What constitutes a relevant action needs to be defined beforehand.

For the *task-dependent* feedback mode, we considered the type of target as the relevant attribute to filter information about the partner's actions. For instance, if users could select squares, only when their partner touched a square would they receive the respective spatialized speech feedback.

We emphasize that the partners' actions are read with a different voice (the one assigned to the partner) to enable distinguishing each other's actions.

3.3 Prototype

To test our multi-user nonvisual interactive tabletop and assess the effect of different feedback modes, we developed a prototype using a custom setup and the Unity engine.

3.3.1 Hardware setup. The hardware setup used for our prototype is shown in Figure 3. It consists of an interactive tabletop, two Bluetooth stereo headphones, and a pair of depth cameras, connected to two laptop computers. The tabletop has an interactive surface with 55" diagonal (1.21m x 0.69m). An infra-red based multi-touch frame placed on top of the surface is used to detect users' touches. The Bluetooth headphones allow us to give audio feedback to each user individually. The chosen headphones were intentionally over-ear, in opposing to in-ear alternatives. We chose models that did not isolate external audio so that verbal communication between users was not affected.

For tracking people around the tabletop, we use a non-invasive solution with two Microsoft Kinect v2 depth cameras, connected to our prototype through the Creepy Tracker toolkit [57]. Each camera is pointed at one of the participants' side of the tabletop to optimize tracking. The toolkit provides information about the participants' skeletons and allows the configuration of the position and orientation of the interactive surface. Therefore, both touch and user information can be matched to a single coordinate system.

The main computer runs our prototype and the tracker hub while connected to the multi-touch frame, a depth camera, and a headphone. The second computer connects to the second depth camera and the remaining headphone.

To map a touchscreen event to a single user, we relate the spatial location of both the touch event and the users' hands. For simplification of audio feedback, we only allow one touch per user,

Fig. 3. Prototype setup: an interactive tabletop, two Bluetooth stereo headphones and two Microsoft Kinect v2 depth cameras, connected to two laptop computers.

being considered the one that started earlier. When a new touch is detected, we associated it with the closest hand (i.e., hand tip joint from the Kinect skeletons).

3.3.2 Auditory feedback. After determining users' touch positions, our gesture recognizer detects both users' actions to give appropriate feedback. Users can either drag a finger on the surface, triggering feedback about the objects they intersect, or double-tap to select the last intersected object. If users attempt to select an object that cannot be selected, the prototype plays an error message; otherwise, it plays a successful selection.

We used two voices (one for each user), which consisted of two recordings of the same clip by a male and a female interlocutor. As for sound spatialization, we used the spatial audio capabilities of Unity: users' positions are used for audio listeners, and touch locations are used for audio sources. We do not consider head orientation as we intend the audio feedback to be provided considering the user's position in relation to the tabletop; i.e. if users hear the audio corresponding to an artifact coming from the left, it is because the artifact is, in fact, to their left. Note that the volume of the left and right channels is not binary. For instance, if an artifact is in front of the users, but more to the left, users will hear sound from both channels, but the volume of the left channel will be higher.

Both audio sources' and audio listeners' positions are projected to the same horizontal plane to amplify the spatialization effect. This makes the maximum volume to be played for an artifact close to the user, whereas if the original 3D positions were used, the height difference between the users and the tabletop would have a significant effect attenuating the sound, thus rendering the distance between artifacts less perceptible. To attenuate the feedback volume according to the distance between the artifacts and the users, we use a custom quadratic roll-off function where a zero volume corresponds to the furthest distance a target can be from users (i.e., the diagonal of the tabletop).

To deliver sound to both users, we use two satellite applications running on different computers that communicate with the main application, each connected to a Bluetooth headphone. Whenever a user touches an artifact or performs a selection, a message is broadcasted to these satellite applications, containing information about the action performed, the user who performed it, the position of the sound relative to both users, and the active feedback mode.

3.4 Participants

We recruited 10 participants (7 males), five pairs. Participants were recruited from a local social institution. Their ages ranged from 20 to 59 (M=46, SD=11.86) years old. All participants were legally blind: three of the participants were congenitally blind and the remaining onset blindness ranged from 13 to 56 years (M=27.3, SD=13.8). Eight participants were fully blind and 2 had residual vision. None reported having motor or hearing impairments. All participants had prior social relationships as they attended the same institution. All participants were right-handed and had daily experience with touchscreen mobile phones (performing tasks such as calls, text messages, radio apps, read email, access weather information, and social media) and consequently had experience with VoiceOver or Talkback for at least 2 months (M=5.22, SD=3.38 in years). None of the participants had previous experience with tabletop displays.

3.5 Procedure

Before each evaluation session, we calibrated the headphones' volume for each participant, until they confirmed that both audio feedback and external audio sources were heard. The two screen-reader voices (male and female) were randomly assigned to participants. We then instructed participants on how to use the interactive tabletop and invited them to explore our prototype until they were familiar with the three feedback modes. They were presented with tasks, similar to those used during the test, for each feedback mode. This training stage lasted for 10 minutes as participants were given time to interact with the screen while dissipating all doubts until they felt comfortable with both the feedback mode and the task. Participants stood facing each other, and the tabletop, wearing headphones.

We invited participants to perform the experimental tasks and instructed them to select all targets (i.e., circles and squares) as quickly as possible. They were reminded of the type of target they were able to select and to which voices both users' actions corresponded. Participants selected targets by performing a double-tap. The swipe to navigate gesture, common in mainstream screen readers, was not available to encourage users to explore the screen. The order of feedback conditions (*private, task-dependent, public*) was randomized. For each feedback condition, participants performed two tasks: one with five additional distractors on the screen (triangles) and another without distractors. Participants were given a maximum of 7 and 5 minutes to complete the tasks with distractors and without distractors, respectively.

After performing both tasks for each feedback mode, the pairs were asked to complete an individual questionnaire. The questionnaires consisted of six Likert items (1 - strongly disagree to 5 - strongly agree) with the following statements: "Tasks were easy to complete", "I helped my partner", "My partner helped me", "Too much information was being read aloud", "Tasks were frustrating", "Tasks were finished quickly". After finishing all experimental conditions, we interviewed participants to gather general opinions about the auditory feedback, the task, and the collaborative experience. The total time for the session was between 50 and 90 minutes. Each participant was compensated for their time with a €15 gift card.

3.6 Study Design and Data Analysis

We used a within-subjects design where each participant tested all feedback conditions: 3 feedback modes x 2 distractor conditions. All evaluation sessions were video recorded, and all interactions with the tabletop were logged.

Task performance was measured by task success rate, task completion time, selection attempts, and error rate (ratio between failed selections and total attempts). Moreover, we investigate touch exploration behaviors by measuring touch time, grid coverage, and distance covered.

Type	Coding	Description
Supply	What I did	Participant informs what action s/he just performed, without a previous request from the partner.
	What I am doing	Participant informs what action s/he is performing, without a previous request from the partner.
	What I will do	Participant informs what action s/he is about to perform, without a previous request from the partner.
	Where I am	Participant informs her/his location, without a previous request from the partner. Can be verbal or physical.
	Completion Status	Participant informs if s/he completed a task or an action, or the situation towards completion, without a previous request from the partner.
Request	What did you do?	Participant asks what action her/his partner just performed, without a previous supply from the partner.
	What are you doing?	Participant asks what action her/his partner is performing, without a previous supply from the partner.
	What will you do?	Participant asks what action her/his partner is about to perform, without a previous supply from the partner.
	Where are you?	Participant asks what her/his partner location is, without a previous supply from the partner.
	Completion Status	Participant asks if her/his partner completed a task or an action, or about the situation towards completion, without a previous supply from the partner.

Table 1. Coding scheme used to perform the video analysis.

We ran a two-way repeated-measures ANOVA with feedback mode (3 levels) and distractor condition (2 levels) as factors. Quantitative values were tested for normal distributions (Shapiro-Wilk's test). In case data failed the assumption, we applied a data transformation (e.g., log, square root) to guarantee that the data fit the repeated measures ANOVA assumptions. Finally, we tested for sphericity (Mauchly's test) and used the Greenhouse-Geisser correction when the assumption was not met. When statistically significant interactions were found, we used a one-way repeated measures ANOVA to assess differences between feedback modes, and a paired t-test to identify differences between distractor condition. Bonferroni corrections were used for post-hoc tests. For the questionnaire data, we ran a Friedman test.

Regarding the video analysis, we coded all footage to capture the content and patterns of workspace awareness information exchange during the collaborative activities. We used a coding scheme based on Gutwin and Greenberg's [28] framework for workspace awareness (WA) and later refined by Metatla et al. [48], as illustrated in Table 1. We divided the coding of WA information exchange into *Requested* and *Supplied* types. The requested type refers to occasions where an exchange is triggered by participants explicitly asking their partners for information, whereas the supplied type refers to information provided by participants to their partners without the latter having asked for it. We focused the coding scheme on *Intentions*, *Actions* or *Location*. Intentions and actions relate to explicit references to past, current, and future activities as well as indications of completion statuses. Concerning location, the coding scheme captured verbal and physical information exchanges about the position of targets on the touchscreen.

4 EFFECT OF FEEDBACK MODES

In this section, we present results related to the effect of feedback modes on task performance and WA information exchanges.

4.1 Task Performance Results

All pairs successfully completed all tasks within the time limit, resulting in 100% success rate. The statistical analysis of task completion time revealed no significant differences between conditions

(private M=88.6s, SD=21.3s; task-dependent M=102.9s, SD=32.5s; public M=99.4s, SD=32.2s), suggesting that although public mode delivers irrelevant auditory information, users are able to filter it out and maintain task efficiency.

When analyzing touch time, a statistically significant interaction was found for the time when the two participants were touching the tabletop at the same time ($F(2,8)=7.830$, $p<.05$), i.e. working simultaneously. Post-hoc tests revealed that in the *private* feedback mode, both participants explored the tabletop simultaneously for a significantly longer period of time in the distractors condition ($t(4)=-2.814$, $p<.05$, M=65.8%, SD=10.5%) than when they were not present (M=55.6%, SD=16.0%). This result is likely related to the increased number of targets (distractors) and users spending more time touching the screen to hear the private feedback.

The presence of distractors also had a significant effect on the total number of selections ($F(1,9)=9.223$, $p<.05$) and error rate ($F(1,9)=5.974$, $p<.05$). The distractor's condition had a higher selection count (M=8.5, SD=3.8) than the no distractors condition (M=6.7, SD=2.1), which led to a higher rate of wrong selections (distractors M=16.5%, SD=16.4% and no-distractors M=6.6%, SD=12.4%). These wrong selections happened because participants tried to select distractors, which is confirmed by the fact that there was no significant difference in tentative selections of other participants' targets ($F(1,9)=0.052$, $p=0.825$).

We also found a main effect of distractors on grid coverage ($F(1,9)=5.128$, $p<0.05$). In the distractors condition (M=84.2%, SD=10.9%) participants explored more of the grid than in the no-distractors condition (M=77.9%, SD=13.5%). For the total distance covered by each participant, we found an interaction between factors ($F(2,18)=5.292$, $p<.05$). The *public* feedback mode was significantly affected by the distractors ($t(9)=-3.829$, $p<.01$). The distance covered was shorter in the no-distractors condition (M=4.3m, SD=1.8) than in the other (M=6.62m, SD=2.2).

4.2 Workspace Awareness Information Exchanges

In this section, we analyze the patterns of WA information exchange between pairs and how feedback modes and distractors influenced them.

On average, we found 4.8 (SD=2.5), 5.5 (SD=3.5), and 6.6 (SD=2.8) WA information exchanges in the private, task-dependent, and public modes, respectively. We did not find a significant main effect of either feedback mode ($F(2,18)=2.357$, $p=.123$) or distractors ($F(1,9)=.95$, $p=.355$) on the number
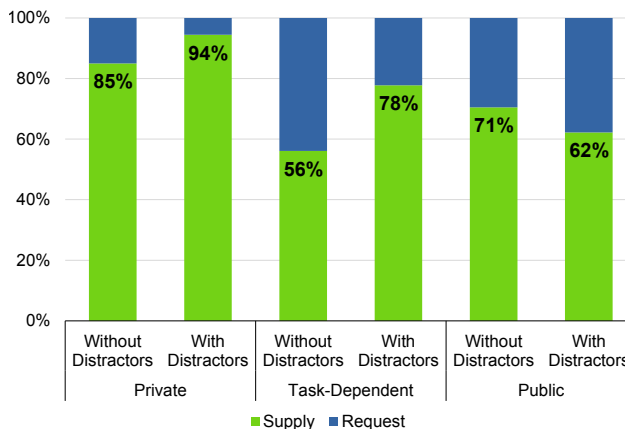

Fig. 4. Rate of supplies and request for each condition.

of WA information exchanges. However, we did find a significant interaction between factors (F(2,18)=4.489, p<.05), suggesting that the public feedback mode is more sensitive to distractors than private mode (F(2,18)=7.573, p<.05), increasing total information exchanges.

As shown in Figure 4, participants supplied proportionally significantly more WA information to each other than requested from one another in private (90% supplied vs. 10% requested, Z=-2.807, p<.01), and public conditions (66% supplied vs. 34% requested, t(9)=2.386, p<.05). However, we did not find significant differences in the task-dependent feedback mode (65% supplied vs. 35% requested, t(9)=1.477, p=.174). A separate comparison of the supplied and requested types of WA information across conditions revealed that the supplied type of exchanges was significantly higher when using the private mode (M=94%, SD=4.7%) in comparison to the public (M=71.5%, SD=5.1%, F(2,14)=6.020, p<.05).

One particular type of information was exchanged at significantly higher rates ($\chi^2(2)$=7.00, p<.05): "Where Are You?" (Figure 5). Significant differences occurred between private and task-dependent modes (p<.05), and at marginal significance between private and public feedback modes (p=.06). These results suggest that location requests were triggered when partners explored relevant targets. This could only happen in task-dependent and public conditions. Moreover, the result shows that participants were aware of relevant information that originated from partners' actions. This is a manifestation of an ability to move between focused individual work to collaborative work.

Results also show a significantly higher rate of "Where I am" supplies (F(9,81)=16.487, p<.001), which could suggest that the spatialized speech feedback was not effective. However, a detailed analysis of these supplies shows that some participants are highly proactive in providing target location information to their partners as soon as they find it (without an explicit request).
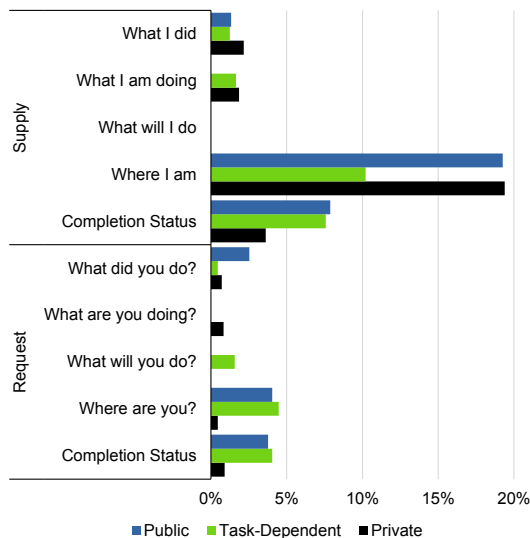

Fig. 5. Proportion of each type of workspace awareness information exchange regarding explicit references to past, current and future actions, location and completion statuses.

4.3 Subjective Ratings

Overall, all feedback modes were well perceived and rated similarly for all six Likert items. This holds true even for the amount of information made available simultaneously ("Too much information was being read aloud"), which we expected for the *public* feedback mode to be rated worse than the others. However, the median (inter-quartile rate) attributed by participants was similar: 1 (2), 1 (4), and 1.5 (3) for private, task-dependent, and public, respectively ($\chi^2(2)=2.33$, p=.311).

5 NONVISUAL COLLABORATION BEHAVIORS

In this section, we highlight examples of co-located nonvisual collaboration behaviors, namely how WA information was both supplied and extracted, how participants coordinated strategies, and discuss the relevance of physical exchanges.

5.1 Supplying Information about Location

Participants often supplied relevant target location information to their partners as soon as they found it and used various nonverbal behaviors that supplemented verbal instructions.

Deictic references. Blind participants often used referential communication to indicate location. Deictic references such as pointing or gesturing are widely used by sighted users in shared workspaces [28, 35]. We observed that it is also common in nonvisual collaboration, as all participants pointed at targets while making comments such as "there is a square here". Interestingly, deictic references themselves are inaccessible to blind participants due to their visual nature (e.g., pointing at a target). Participants thus used additional nonverbal/physical mechanisms to inform partners of target locations.

Audibly tapping on the tabletop. Rather than simply pointing at a target, 40% of the participants repeatedly tapped on the screen to make its position audible to their partners.

Leaving the finger as reference point. Another strategy to complement deictic references was leaving their finger on the target (67.1% of all location supplies) while their partners searched for it. In these cases, 90% of participants always dragged their finger on the tabletop until reaching them, while partners searched over the tabletop for their hand.

Retrieving partners' hand. Participants also used more intrusive approaches that consisted of performing wide sweep movements over the tabletop to eventually find their partners' hands, physically grabbing them, and guiding their movement to the target (Figure 1 middle). We observed this behavior in all physical location supplies. In this case, the helpers used the second hand as a reference point to facilitate re-locating the target.

Verbal instructions. One could expect that detailed verbal instructions about target locations would be common in nonvisual collaboration. However, it is worth highlighting that none of the verbal WA information exchanges contained absolute references to target locations, such as "bottom left corner" or "center of the tabletop". Instead, verbal instructions were often vague, such as "there are circles on this side" or used the partner or the tabletop as references: "there is one square here, to your right side, my left side" and "there is a square here, over this corner".

5.2 Extracting Information about Location

Participants used two main strategies to retrieve location information.

Reacting to spatialized feedback. The multi-user nonvisual interactive tabletop provided spatialized speech feedback on a given target, which could potentially help users infer the location of relevant targets being explored by their partners. In both task-dependent and public feedback modes, we observed 5 and 8 instances, respectively, where participants immediately reacted to speech feedback by moving their hand in the direction of the intended target without further instructions (Figure 1

left). Although such behavior suggests that spatialized feedback is effective, further research is needed to isolate this effect from, for example, the audible physical interactions of partners or knowledge of partners' position.

Offering hand. In 49% of all location requests, participants combined verbal requests with physical behaviors that included making their hands available to their partners, thus allowing them to be guided to the target. These occurred after hearing a relevant target triggered by the partners' actions.

5.3 Coordinating Strategies

Only one pair prearranged a strategy to find all targets efficiently. The strategy was based on divide-and-conquer, where each participant was responsible for exploring their half of the tabletop. The remaining strategies emerged from individual behaviors.

Individualist. Three participants from three different pairs adopted a strategy of prioritizing selecting their own targets over sharing relevant information with their partners. Only when all their targets were selected did they start supplying information about target locations. On the other hand, their partners assumed the role of supplying most of the information, accounting for 72% of exchanges.

Over-helper. One of the participants opted for sharing all information without much care for its relevance, which sometimes seemed to have a negative effect on his partner. For example, while both participants were exploring, the over-helper decided to supply information about the targets he was exploring, immediately grabbing his partner's hand without consent, and thus disrupting the partner's exploration.

5.4 Physical Information Exchanges

Supplying and requesting location information. Physical contact was widely advantageous to exchange information about target location and complement deictic references. Participants used a myriad of physical behaviors as described in the previous section. However, physical exchanges also came with some issues.

Unintentional contact. Overall, we observed several cases where participants' screen explorations resulted in unintentional physical contact (51% of all instances of physical contact). Although one could argue that such occurrences can benefit WA as participants are made aware of each other's hand positions while swiftly negotiating further exploration areas, it can also encompass negative consequences. First, it can interrupt individual screen exploration strategies [23], affecting users' task performance and experience. Second, depending on the relationship between participants (or the lack of it), touching each other's hands can be a socially uncomfortable experience [49]. We notice such effect in Pair 5, where the intentional physical exchanges were lower than the average (31% vs. 49%) and participants were visibly conscious of invading each other's personal space.

6 DISCUSSION

In this section, we answer our four research questions, summarizing the major results attained, derive design implications, and report limitations of this work.

6.1 Answering the Research Questions

1. Is auditory feedback effective in supporting nonvisual collaboration in large touchscreens? Results showed that all participants completed all tasks within the time limit. Participants were able to engage in parallel collaborative activities, independently of feedback mode. Furthermore, the multi-user nonvisual interactive tabletop fostered collaborative behaviors, particularly in the task-dependent and public feedback modes, shown by the increase of information requests. This means

that these feedback modes effectively promoted workspace awareness and enabled participants to move between individual and joint work.

2. How is task performance in large touchscreens affected by the amount of auditory information displayed? Results show that feedback mode did not have a significant influence on task completion time. This suggests that users were able to tease out irrelevant information in the *public* feedback mode without a decrease in task performance. Such effect is widely known as the cocktail party effect [11]. However, the existence of distractors significantly increased the number of errors and grid coverage, suggesting an impact of task difficulty on user performance.

3. How does the amount of auditory information delivered influence patterns of awareness information exchange? Neither feedback modes nor distractors revealed to have a significant effect on the total number of information exchanges. Nevertheless, feedback modes that provide higher amounts of information (*public* and *task-dependent*) revealed to have a better balance between the number of supplies and requests. In contrast, the *private* mode had a significantly higher rate of supplies that result from the lesser amount of information received about the partner's actions, which reduces awareness and, consequently, increases the need to share more information explicitly.

4. How do blind users leverage audio feedback to engage with one another in co-located collaborative tasks? Common behaviors included both verbal communication and physical interactions. Supplying location information, often entailed a verbal announcement complemented with a nonvisual deictic gesture. To make such nonverbal gestures accessible, blind users adopted multiple strategies, including audibly tapping on the screen, waiting for the partner to find their hand, or actively guiding their partners' hands to target. On the other hand, they spontaneously obtained location information either through the spatialized feedback or by making their hands available to their partners to be guided.

We used a framework for workspace awareness [29, 48] to analyze the data as it clearly relates to the context of our work. Still, links can be made to frameworks that approach coordination or collaboration more broadly or in different scenarios. To cite one example, Butchibabu et al. [8] define implicit coordination – providing anticipatory information or status updates – similarly to what we consider a supply information exchange. On the other hand, explicit coordination can be included in what we consider a request – prompting or requesting information from others.

6.2 Design Implications

Based on our results and observations, we were able to devise a set of implications that can be useful for interaction designers targeting multi-user screen readers, particularly for tabletop interaction.

Focus on information sensitivity. Since task performance is not significantly affected by the amount of information conveyed, interaction designers can focus instead on information sensitivity. For instance, if no private information is displayed, and designers want to promote workspace awareness, then the *public* feedback mode can be used without overloading users. On the other hand, if sensitive information exists, a *private* feedback mode would be a better option, letting users expressly exchange the information they intend, knowing that it would not significantly hinder task performance. Still, the collaborative system should be designed to allow users to request "Where are you" type of WA information.

Reduce on-screen content. The existence of unhelpful content in the scene can have a severe impact on task execution. Whenever possible, artifacts that are not useful for the task at hand should be omitted. This can help users focus on more relevant content, reduce cognitive load, and prevent errors.

Consider interaction on and above the tabletop. Information exchange between users happens not only while touching the interactive surface, but also in the space above it. Besides verbal exchanges, users often engage in physical communication, mainly to aid in location sharing. As

such, interaction design should target not only the interaction of the users with the tabletop but also the interaction between them. Furthermore, this physical communication means that solutions to support co-located and remote collaboration between blind tabletop users should not be addressed similarly.

Provide awareness about the other user. An approach based solely on a screen reader provides feedback when a user touches an artifact, remaining silent the rest of the time. This makes it more challenging to be aware of the partner's location, which can lead to undesired hand collisions. Future research in multi-user nonvisual tabletops should be sensitive to such social aspects and strive to minimize interaction that can be socially awkward. For example, users could be informed when they are getting closer to their partner's hand. Also, as a side effect, by helping users know where their partners are located, WA information exchanges can be made more accessible.

Make deictic expressions meaningful. Blind people also point. However, if their partner is also blind, the location they are pointing at will not be perceived without additional information exchanges. Strategies to make these deictic expressions more meaningful should be investigated to help guide users to the intended target without requiring their partners to hold the position and/or physically navigate them. For instance, such features could be triggered by an interface that recognizes multiple taps on the same location, holding a touch for a predetermined time period or verbal commands.

Use multiple taps wisely. Since users tend to perform multiple taps to convey location to their partners, designers should take extra care to prevent them from inadvertently invoking other functions. As an alternative, such functions could be activated by using touches with different signatures [43].

6.3 Limitations

Our study included 10 blind participants, which resulted in five pairs. Although this is a small number of pairs, we applied the statistical tests to assess significant variance between participants whenever possible. A higher number of pairs could have led to results with larger effect sizes. Also, it is known that prior social relationships can have an impact on collaboration [12, 44]. However, since all participants knew their partners somehow as they all attended the same institution, we could not assess this. Still, the observed WA information exchanges and emergent collaboration behaviors can be applied to other nonvisual collaborative use cases.

We only allowed one-handed interaction, which can be somewhat restrictive, considering the large touchscreen area available. However, to support two-hand interaction, we would have to provide more feedback channels so that it would be possible to distinguish content intersected by the left and right hands, which could make this first attempt at a multi-user nonvisual prototype overly complex. Further research should leverage this work and enable two-hand interaction.

The Task-Dependent mode requires defining in advance what constitutes a relevant action, which can be difficult (or sometimes impossible) in more complex collaborative tasks. In such cases, users may be required to explicitly define what is relevant to them. Such definition can be made, for instance, in terms of content type (as in our task), location of events on the tabletop, or (a)synchrony with the user's actions.

Finally, the Bluetooth headphones introduced a small latency, which could be perceived when participants dragged the finger quickly over an object; the headphones would play the corresponding sound slightly off location. We noticed this effect in some instances where participants were guiding partners' hands to a recently discovered target.

7 CONCLUSIONS

We presented the design and implementation of a multi-user auditory display for interactive tabletops. We then reported on a user study that examined its effectiveness on promoting nonvisual collaborative work as well as the effect of different feedback modes on task performance and workspace awareness information exchanges. Results show that nonvisual interactive tabletops are effective in allowing multiple users to simultaneously interact with content. Moreover, the amount of auditory provided to users did not have significant effects on task performance. On the other hand, task complexity made tasks slower and more erroneous. We further identified which type of workspace awareness information was exchanged under different feedback conditions, showing what information is relevant and should be conveyed by future screen reader design that aim to support multiple users. We also analyzed emergent collaboration behaviors, namely concerning the supply and extraction of location information, coordinating strategies, and physical interactions. Our findings provide empirical evidence that deictic gestures play a major role in nonvisual co-located collaboration, along with several physical coping strategies.

Our work is the first step towards inclusive tabletop collaboration and the presented results can help researchers and developers creating future co-located collaborative tools for blind people. Considering, for example, an application for map exploration, our guidelines can contribute to the development of interfaces that can successfully aid users in finding and discussing over points-of-interest without losing spatial awareness and being able to relate such points. Possible avenues for future research include the exploration of more complex collaboration mechanics, such as the support for shared resources, and the development of approaches towards inclusive tabletops that support users with mixed-visual abilities. Such developments would present an opportunity to contextualize natural interaction practices and emergent territories [56] within inclusive tabletop interfaces.

ACKNOWLEDGMENTS

REFERENCES

[1] Jérémy Albouys-Perrois, Jérémy Laviole, Carine Briant, and Anke M Brock. 2018. Towards a Multisensory Augmented Reality Map for Blind and Low Vision People: A Participatory Design Approach. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, USA, 629:1—-629:14. https://doi.org/10.1145/3173574.3174203

[2] Edward Anstead, Abigail Durrant, Steve Benford, and David Kirk. 2012. Tabletop Games for Photo Consumption at Theme Parks. In *Proceedings of the 2012 ACM International Conference on Interactive Tabletops and Surfaces (ITS '12)*. ACM, New York, NY, USA, 61–70. https://doi.org/10.1145/2396636.2396646

[3] Apple. 2019. VoiceOver. https://www.apple.com/lae/accessibility/iphone/vision/. [Online; accessed 10-September-2019].

[4] Shiri Azenkot, JAcob O. Wobbrock, Sanjana Prasain, and Richard E. Ladner. 2012. Input finger detection for nonvisual touch screen text entry in Perkinput. In *Proceedings of Graphics Interface*. Canadian Information Processing Society, Toronto, Ont., Canada, Canada, 121–129. https://doi.org/2305276.2305297

[5] Sandra Bardot, Marcos Serrano, Simon Perrault, Shengdong Zhao, and Christophe Jouffrais. 2019. Investigating Feedback for Two-Handed Exploration of Digital Maps Without Vision. In *Human-Computer Interaction – INTERACT 2019*, David Lamas, Fernando Loizides, Lennart Nacke, Helen Petrie, Marco Winckler, and Panayiotis Zaphiris (Eds.). Springer International Publishing, Cham, 305–324.

[6] Syed Masum Billah, Yu-Jung Ko, Vikas Ashok, Xiaojun Bi, and I V Ramakrishnan. 2019. Accessible Gesture Typing for Non-Visual Text Entry on Smartphones. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. ACM, New York, NY, USA, 376:1—-376:12. https://doi.org/10.1145/3290605.3300606

[7] Anke Brock, Bernard Oriola, Philippe Truillet, Christophe Jouffrais, and Delphine Picard. 2013. Map design for visually impaired people: past, present, and future research. *Médiation et Information* 36 (dec 2013), 117–129. http://oatao.univ-

toulouse.fr/12339/

[8] Abhizna Butchibabu, Christopher Sparano-Huiban, Liz Sonenberg, and Julie Shah. 2016. Implicit coordination strategies for effective team communication. *Human factors* 58, 4 (2016), 595–610.

[9] Maria Claudia Buzzi, Marina Buzzi, Barbara Leporini, Giulio Mori, and Victor M R Penichet. 2014. Collaborative Editing: Collaboration, Awareness and Accessibility Issues for the Blind. In *On the Move to Meaningful Internet Systems: OTM 2014 Workshops*, Robert Meersman, Hervé Panetto, Alok Mishra, Rafael Valencia-García, António Lucas Soares, Ioana Ciuciu, Fernando Ferri, Georg Weichhart, Thomas Moser, Michele Bezzi, and Henry Chan (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 567–573.

[10] Xiang Cao, Siân E Lindley, John Helmes, and Abigail Sellen. 2010. Telling the Whole Story: Anticipation, Inspiration and Reputation in a Field Deployment of TellTable. In *Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work (CSCW '10)*. ACM, New York, NY, USA, 251–260. https://doi.org/10.1145/1718918.1718967

[11] E Colin Cherry. 1953. Some Experiments on the Recognition of Speech, with One and with Two Ears. *The Journal of the Acoustical Society of America* 25, 5 (1953), 975–979. https://doi.org/10.1121/1.1907229

[12] Jonathon N Cummings and Sara Kiesler. 2008. Who collaborates successfully? Prior experience reduces collaboration barriers in distributed interdisciplinary research. In *Proceedings of the 2008 ACM conference on Computer supported cooperative work*. 437–446.

[13] Maitraye Das, Darren Gergle, and Anne Marie Piper. 2019. "It Doesn't Win You Friends": Understanding Accessibility in Collaborative Writing for People with Vision Impairments. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW (nov 2019). https://doi.org/10.1145/3359293

[14] Paul Dourish and Victoria Bellotti. 1992. Awareness and coordination in shared workspaces.. In *CSCW*, Vol. 92. 107–114.

[15] Julie Ducasse, Marc Macé, Marcos Serrano, and Christophe Jouffrais. 2016. Tangible Reels : Construction and Exploration of Tangible Maps by Visually Impaired Users. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2016). https://doi.org/10.1016/j.brainresbull.2008.09.003

[16] Thomas Erickson and Wendy A. Kellogg. 2000. Social Translucence: An Approach to Designing Systems That Support Social Processes. *ACM Trans. Comput.-Hum. Interact.* 7, 1 (March 2000), 59–83. https://doi.org/10.1145/344949.345004

[17] Taciana Pontual Falcão and Sara Price. 2009. What Have You Done! The Role of 'Interference' in Tangible Environments for Supporting Collaborative Learning. In *Proceedings of the 9th International Conference on Computer Supported Collaborative Learning - Volume 1 (CSCL'09)*. International Society of the Learning Sciences, 325–334. http://dl.acm.org/citation.cfm?id=1600053.1600103

[18] Eynat Gal, Liron Lamash, Nirit Bauminger-Zviely, Massimo Zancanaro, and Patrice L (Tamar) Weiss. 2016. Using Multitouch Collaboration Technology to Enhance Social Interaction of Children with High-Functioning Autism. *Physical & Occupational Therapy In Pediatrics* 36, 1 (2016), 46–58. https://doi.org/10.3109/01942638.2015.1040572

[19] Cagatay Goncu and Kim Marriott. 2011. GraVVITAS: generic multi-touch presentation of accessible graphics. *Human-computer interaction–INTERACT 2011* (2011), 30–48. http://link.springer.com/chapter/10.1007/978-3-642-23774-4{_}5

[20] Google. 2019. Talkback. https://play.google.com/store/apps/details?id=com.google.android.marvin.talkback. [Online; accessed 11-September-2019].

[21] William Grussenmeyer and Eelke Folmer. 2017. Accessible Touchscreen Technology for People with Visual Impairments: A Survey. *ACM Transactions on Accessible Computing (TACCESS)* 9, 2 (jan 2017), 6:1–-6:31. https://doi.org/10.1145/3022701

[22] João Guerreiro and Daniel Gonçalves. 2016. Scanning for digital content: How blind and sighted people perceive concurrent speech. *ACM Transactions on Accessible Computing (TACCESS)* 8, 1 (2016), 1–28.

[23] João Guerreiro, André Rodrigues, Kyle Montague, Tiago Guerreiro, Hugo Nicolau, and Daniel Gonçalves. 2015. TabLETS Get Physical: Non-Visual Text Entry on Tablet Devices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*.

[24] T Guerreiro, P Lagoá, H Nicolau, D Gonçalves, and J A Jorge. 2008. From tapping to touching: Making touch screens accessible to blind users. *IEEE MultiMedia* (2008), 48–50.

[25] Tiago Guerreiro, Kyle Montague, João Guerreiro, Rafael Nunes, Hugo Nicolau, and Daniel J V Gonçalves. 2015. Blind People Interacting with Large Touch Surfaces: Strategies for One-handed and Two-handed Exploration. In *Proceedings of the 2015 International Conference on Interactive Tabletops & Surfaces*. ACM, 25–34.

[26] Anhong Guo, Xiang'Anthony' Chen, Haoran Qi, Samuel White, Suman Ghosh, Chieko Asakawa, and Jeffrey P Bigham. 2016. Vizlens: A robust and interactive screen reader for interfaces in the real world. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. 651–664.

[27] Anhong Guo, Junhan Kong, Michael Rivera, Frank F Xu, and Jeffrey P Bigham. 2019. StateLens: A Reverse Engineering Solution for Making Existing Dynamic Touchscreens Accessible. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*. 371–385.

[28] Carl Gutwin and Saul Greenberg. 2002. A descriptive framework of workspace awareness for real-time groupware. *Computer Supported Cooperative Work* 11, 3-4 (2002), 411–446. https://doi.org/10.1023/A:1021271517844

[29] Carl Gutwin and Saul Greenberg. 2002. A descriptive framework of workspace awareness for real-time groupware. *Computer Supported Cooperative Work (CSCW)* 11, 3-4 (2002), 411–446.

[30] Carl Gutwin, Oliver Schneider, Robert Xiao, and Stephen Brewster. 2011. Chalk Sounds: The Effects of Dynamic Synthesized Audio on Workspace Awareness in Distributed Groupware. In *Proceedings of the ACM 2011 Conference on Computer Supported Cooperative Work (CSCW '11)*. ACM, New York, NY, USA, 85–94. https://doi.org/10.1145/1958824.1958838

[31] Mark S Hancock, Chia Shen, Clifton Forlines, and Kathy Ryall. 2005. Exploring Non-speech Auditory Feedback at an Interactive Multi-user Tabletop. In *Proceedings of Graphics Interface 2005 (GI '05)*. Canadian Human-Computer Communications Society, School of Computer Science, University of Waterloo, Waterloo, Ontario, Canada, 41–50. http://dl.acm.org/citation.cfm?id=1089508.1089516

[32] Amanda Harris, Jochen Rick, Victoria Bonnett, Nicola Yuill, Rowanne Fleck, Paul Marshall, and Yvonne Rogers. 2009. Around the Table: Are Multiple-touch Surfaces Better Than Single-touch for Children's Collaborative Interactions?. In *Proceedings of the 9th International Conference on Computer Supported Collaborative Learning - Volume 1 (CSCL'09)*. International Society of the Learning Sciences, 335–344. http://dl.acm.org/citation.cfm?id=1600053.1600104

[33] Steven E Higgins, Emma Mercier, Elizabeth Burd, and Andrew Hatch. 2011. Multi-touch tables and the relationship with collaborative classroom pedagogies: A synthetic review. *International Journal of Computer-Supported Collaborative Learning* 6, 4 (dec 2011), 515–538. https://doi.org/10.1007/s11412-011-9131-y

[34] Uta Hinrichs and Sheelagh Carpendale. 2011. Gestures in the Wild: Studying Multi-touch Gesture Sequences on Interactive Tabletop Exhibits. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. ACM, New York, NY, USA, 3023–3032. https://doi.org/10.1145/1978942.1979391

[35] Hiroshi Ishii and Minoru Kobayashi. 1992. ClearBoard: a seamless medium for shared drawing and conversation with eye contact. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 525–532.

[36] I. Jamil, C.S. Montero, M. Perry, K. O'Hara, A. Karnik, K. Pihlainen, M.T. Marshall, S. Jha, S. Gupta, and S. Subramanian. 2017. Collaborating around digital tabletops: Children's physical strategies from India, the UK and Finland. *ACM Transactions on Computer-Human Interaction* 24, 3 (2017), 1–30. https://doi.org/10.1145/3058551

[37] Izdihar Jamil, Kenton O'Hara, Mark Perry, Abhijit Karnik, and Sriram Subramanian. 2011. The Effects of Interaction Techniques on Talk Patterns in Collaborative Peer Learning Around Interactive Tables. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. ACM, New York, NY, USA, 3043–3052. https://doi.org/10.1145/1978942.1979393

[38] Shaun K. Kane, Jeffrey P. Bigham, and Jacob O. Wobbrock. 2008. Slide Rule: Making Mobile Touch Screens Accessible to Blind People Using Multi-Touch Interaction Techniques. In *Proceedings of the 10th International ACM SIGACCESS Conference on Computers and Accessibility* (Halifax, Nova Scotia, Canada) (*Assets '08*). Association for Computing Machinery, New York, NY, USA, 73–80. https://doi.org/10.1145/1414471.1414487

[39] Shaun K Kane, Meredith Ringel Morris, Annuska Z Perkins, Daniel Wigdor, Richard E Ladner, and Jacob O Wobbrock. 2011. Access Overlays: Improving Non-visual Access to Large Touch Screens for Blind Users. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology (UIST '11)*. ACM, New York, NY, USA, 273–282. https://doi.org/10.1145/2047196.2047232

[40] Shaun K Kane, Meredith Ringel Morris, and Jacob O Wobbrock. 2013. Touchplates: Low-cost Tactile Overlays for Visually Impaired Touch Screen Users. In *Proceedings of the 15th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '13)*. ACM, New York, NY, USA, 22:1–-22:8. https://doi.org/10.1145/2513383.2513442

[41] Shaun K Kane, Jacob O Wobbrock, and Richard E Ladner. 2011. Usable Gestures for Blind People: Understanding Preference and Performance. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. ACM, New York, NY, USA, 413–422. https://doi.org/10.1145/1978942.1979001

[42] Andreas Kunz, Dirk Schnelle-Walka, Ali Alavi, Stephan Pölzer, Max Mühlhäuser, and Klaus Miesenberger. 2014. Making Tabletop Interaction Accessible for Blind Users. *Proceedings of the Ninth ACM International Conference on Interactive Tabletops and Surfaces - ITS '14* (2014), 327–332. https://doi.org/10.1145/2669485.2669541

[43] Pedro Lopes, Ricardo Jota, and Joaquim A Jorge. 2011. Augmenting touch interaction through acoustic sensing. In *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces*. 53–56.

[44] Alina Lungeanu, Yun Huang, and Noshir S Contractor. 2014. Understanding the assembly of interdisciplinary teams and its impact on performance. *Journal of informetrics* 8, 1 (2014), 59–70.

[45] L Luque, L d. O. Brandão, and A A F Brandão. 2018. A Framework to Foster Diversity in Collaborative Activities. In *2018 IEEE Frontiers in Education Conference (FIE)*. 1–9. https://doi.org/10.1109/FIE.2018.8658438

[46] Muhanad S Manshad and Ahmad S Manshad. 2008. Multimodal Vision Glove for Touchscreens. In *Proceedings of the 10th International ACM SIGACCESS Conference on Computers and Accessibility (Assets '08)*. ACM, New York, NY, USA, 251–252. https://doi.org/10.1145/1414471.1414523

[47] David McGookin and Stephen Brewster. 2007. An Initial Investigation into Non-visual Computer Supported Collaboration. In *CHI '07 Extended Abstracts on Human Factors in Computing Systems (CHI EA '07)*. ACM, New York, NY, USA, 2573–2578. https://doi.org/10.1145/1240866.1241043

[48] Oussama Metatla, Nick Bryan-Kinns, and Tony Stockman. 2018. "I Hear You": Understanding Awareness Information Exchange in an Audio-only Workspace. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2018), 1–13. https://doi.org/10.1145/1235 arXiv:arXiv:1603.07016v1

[49] Meredith Ringel Morris, Anqi Huang, Andreas Paepcke, and Terry Winograd. 2006. Cooperative gestures: multi-user gestural interactions for co-located groupware. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*. 1201–1210.

[50] Meredith Ringel Morris, Dan Morris, and Terry Winograd. 2004. Individual Audio Channels with Single Display Groupware: Effects on Communication and Task Strategy. In *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work (CSCW '04)*. ACM, New York, NY, USA, 242–251. https://doi.org/10.1145/1031607.1031646

[51] Hugo Nicolau, Kyle Montague, Tiago Guerreiro, João Guerreiro, and Vicki L Hanson. 2014. B#: Chord-based Correction for Multitouch Braille Input. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'14)*. ACM, New York, NY, USA, 1705–1708. https://doi.org/10.1145/2556288.2557269

[52] João Oliveira, Tiago Guerreiro, Hugo Nicolau, Joaquim Jorge, and Daniel Gonçalves. 2011. Blind people and mobile touch-based text-entry: Acknowledging the Need for Different Flavors. In *The proceedings of the 13th international ACM SIGACCESS conference on Computers and accessibility - ASSETS '11*. ACM, 179. https://doi.org/10.1145/2049536.2049569

[53] Stephan Pölzer, Andreas Kunz, Ali Alavi, and Klaus Miesenberger. 2016. An Accessible Environment to Integrate Blind Participants into Brainstorming Sessions. In *Computers Helping People with Special Needs*, Klaus Miesenberger, Christian Bühler, and Petr Penaz (Eds.). Springer International Publishing, Cham, 587–593.

[54] R Ramloll and S Brewster. 2002. An environment for studying the impact of spatialising sonified graphs on data comprehension. In *Proceedings Sixth International Conference on Information Visualisation*. 167–174. https://doi.org/10.1109/IV.2002.1028773

[55] Jochen Rick, Paul Marshall, and Nicola Yuill. 2011. Beyond One-size-fits-all: How Interactive Tabletops Support Collaborative Learning. In *Proceedings of the 10th International Conference on Interaction Design and Children (IDC '11)*. ACM, New York, NY, USA, 109–117. https://doi.org/10.1145/1999030.1999043

[56] Stacey D. Scott, M. Sheelagh T. Carpendale, and Kori M. Inkpen. 2004. Territoriality in Collaborative Tabletop Workspaces. In *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work* (Chicago, Illinois, USA) *(CSCW '04)*. ACM, New York, NY, USA, 294–303. https://doi.org/10.1145/1031607.1031655

[57] Maurício Sousa, Daniel Mendes, Rafael Kuffner Dos Anjos, Daniel Medeiros, Alfredo Ferreira, Alberto Raposo, João Madeiras Pereira, and Joaquim Jorge. 2017. Creepy Tracker Toolkit for Context-aware Interfaces. In *Proceedings of the 2017 ACM International Conference on Interactive Surfaces and Spaces* (Brighton, United Kingdom) *(ISS '17)*. ACM, New York, NY, USA, 191–200. https://doi.org/10.1145/3132272.3134113

[58] Caleb Southern, James Clawson, Brian Frey, Gregory Abowd, and Mario Romero. 2012. An Evaluation of BrailleTouch: Mobile Touchscreen Text Entry for the Visually Impaired. In *Proceedings of the 14th International Conference on Human-computer Interaction with Mobile Devices and Services (MobileHCI '12)*. ACM, New York, NY, USA, 317–326. https://doi.org/10.1145/2371574.2371623

[59] Saiganesh Swaminathan, Thijs Roumen, Robert Kovacs, David Stangl, Stefanie Mueller, and Patrick Baudisch. 2016. Linespace: A sensemaking platform for the blind. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 2175–2185.

[60] Daniel Trindade, André Rodrigues, Tiago Guerreiro, and Hugo Nicolau. 2018. Hybrid-Brailler: Combining Physical and Gestural Interaction for Mobile Braille Input and Editing. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, USA, 27:1—-27:12. https://doi.org/10.1145/3173574.3173601

[61] Fredrik Winberg and John Bowers. 2004. Assembling the Senses: Towards the Design of Cooperative Interfaces for Visually Impaired Users. In *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work (CSCW '04)*. ACM, New York, NY, USA, 332–341. https://doi.org/10.1145/1031607.1031662

[62] Xiaoyi Zhang, Tracy Tran, Yuqian Sun, Ian Culhane, Shobhit Jain, James Fogarty, and Jennifer Mankoff. 2018. Interactiles: 3D Printed Tactile Interfaces to Enhance Mobile Touchscreen Accessibility. In *Proceedings of the 20th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '18)*. ACM, New York, NY, USA, 131–142. https://doi.org/10.1145/3234695.3236349