

Visual Attention-Based Access: Granting Access Based on Users' Joint Attention on Shared Workspaces

BARIŞ SERİM, Helsinki Institute for Information Technology HIIT, Department of Computer Science, University of Helsinki, Finland and Department of Design, Aalto University, Finland

KEN PFEUFFER, Bundeswehr University Munich, Germany

HANS GELLERSEN, Lancaster University, United Kingdom

GIULIO JACUCCI, Helsinki Institute for Information Technology HIIT, Department of Computer Science, University of Helsinki, Finland

During collaboration, individual users' capacity to maintain awareness, avoid duplicate work and prevent conflicts depends on the extent to which they are able to monitor the workspace. Existing access control models disregard this contextual information by managing access strictly based on who performs the action. As an alternative approach, we propose managing access by taking the visual attention of collaborators into account. For example, actions that require consensus can be limited to collaborators' joint attention, editing another user's personal document can require her visual supervision and private information can become unavailable when another user is looking. We prototyped visual attention-based access for 3 collaboration scenarios on a large vertical display using head orientation input as a proxy for attention. The prototype was deployed for an exploratory user study, where participants in pairs were tasked to assign visual attention-based access to various actions. The results reveal distinct motivations for their use such as preventing accidents, maintaining individual control and facilitating group awareness. Visual attention-based access has been perceived as more convenient but also less certain when compared to traditional access control. We conclude that visual attention-based access can be a useful addition to groupware to flexibly facilitate awareness and prevent conflicts.

CCS Concepts: • **Human-centered computing** → **Collaborative interaction; Synchronous editors; Collaborative and social computing devices; Interactive whiteboards;**

Additional Key Words and Phrases: Visual attention, joint attention, access control, awareness, CSCW, single display groupware, design, qualitative study

ACM Reference Format:

Bariş Serim, Ken Pfeuffer, Hans Gellersen, and Giulio Jacucci. 2018. Visual Attention-Based Access: Granting Access Based on Users' Joint Attention on Shared Workspaces. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 2, 3, Article 133 (September 2018), 23 pages. <https://doi.org/10.1145/3264943>

Authors' addresses: Bariş Serim, Helsinki Institute for Information Technology HIIT, Department of Computer Science, University of Helsinki, Finland, Department of Design, Aalto University, Finland, baris.serim@helsinki.fi; Ken Pfeuffer, Bundeswehr University Munich, Germany, ken.pfeuffer@unibw.de; Hans Gellersen, Lancaster University, United Kingdom, h.gellersen@lancaster.ac.uk; Giulio Jacucci, Helsinki Institute for Information Technology HIIT, Department of Computer Science, University of Helsinki, Finland, giulio.jacucci@helsinki.fi.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2474-9567/2018/9-ART133 \$15.00

<https://doi.org/10.1145/3264943>

1 INTRODUCTION

Concurrent input of multiple users on shared workspaces comes with the benefit of task parallelization, but also introduces the challenge of maintaining coordination between multiple users, in particular, ensuring that the work done by individual users is relevant to the joint activity and that individual users' actions do not interrupt others. The need for coordination led to research on “awareness”, which emerged both as a conceptual tool to understand how multiple users coordinate their activities and as a design goal for groupware design [47]. A design implication of awareness is making the actions and the state of individual users publicly available [17, 23].

At the same time, *awareness depends on the active monitoring of actions by individual users* as much as it depends on their public availability [47]. Because human attention is limited, the public availability of information does not guarantee individual users' awareness of others' actions. The disparity between what is visible and what is monitored can be significant during episodes of attentional disconnect [21], or when the joint activity is conducted on larger workspaces, where users visually attend to different regions of the workspace at a given time [8, 32]. This has potential implications for design. Rather than assuming users' awareness of public information, the system can track users' locus of visual attention and adapt the interaction accordingly.

In this paper, we investigate adapting the access rights based on how multiple users visually attend to the interface and each other's actions on a shared display. During collaboration, users can switch between working on different tasks in parallel to working in tight coordination on the same screen region [29, 32, 53], leading to different visual attention configurations. In return, actions can require varying degrees of oversight or consensus based on their scope or reversibility. These demands have traditionally been satisfied by access control models that are based on *who performs the action* (such as when editing or viewing rights are restricted to particular users). In contrast, we investigate how traditional access control models can opportunistically be relaxed if users are visually attending to each other's actions. Consider a document that can only be edited by the document owner under traditional access control. Access rights for the document can be relaxed to enable editing by other users if the owner of the object is paying visual attention as visual monitoring increases owner's capacity to intervene and keep track of the changes. Similarly, certain commands that cause global changes in the workspace can be restricted to input with joint attention to prevent interruptions to other users.




Fig. 1. An example of visual attention-based access: Actions can be configured to require the joint attention of multiple users.

In the rest of the paper, we motivate visual attention-based access by reviewing previous work on shared workspaces and discuss how our work relates to the long-acknowledged ([17, 22]) trade-off between awareness and individual power in groupware design. Then, we present a framework for visual attention-based access and introduce 4 different access types based on their availability in solitary and joint attention situations. *Universal* actions are available under any attention situation, *consensual* actions require the joint attention of all the relevant users, *supervised* actions become available to other users during joint attention (under the supervision of a particular user such as owner/supervisor) and *private* actions become unavailable during joint attention. We conducted an exploratory study with 20 participants in pairs, where participants were instructed to assign these access types to various actions (such as editing, moving, deleting) in 3 different task scenarios on a large vertical

display that tracked their head orientation. The data from the study showed salient differences between user preferences based on the type and scope of an action, as well as distinct reasons for assigning the same access type. We draw on these findings to generate a list of design implications for future work. Overall, our contributions include the novel concept of visual attention-based access, an initial framework for visual attention-based access and empirical findings and design implications derived from the user study.

2 BACKGROUND

2.1 Visual Attention and Coordination on Shared Workspaces

To facilitate coordination, researchers often contrast two, sometimes competing approaches, namely allowing users to coordinate their actions through dynamic self-organization or structuring the interaction through access control, or “role-restriction”, mechanisms [17, 41]. Here, we outline these two approaches and discuss how they relate to visual attention.

As a design approach, self-organization aims to increase users' own ability to avoid conflicts and keep their actions relevant. Because users' capacity to self-organize depends on their knowledge about other users' actions and intentions, maintaining workspace awareness has been identified as an important goal [17, 22]. Most systems rely on “passive awareness” mechanisms [17] that automatically make actions available to other users. On the other hand, users' awareness of publicly available information depends on their ability to monitor the workspace, which is ultimately influenced by various design decisions. Here, designers of groupware systems are often confronted with a trade-off between individual power and group awareness: The more flexibility and power individual users have the less aware they are of each other's actions [17, 22]. For example, the ability of individual users to view different parts of the workspace at the same time (as in relaxed WYSIWIS [what you see is what I see] interfaces) provides flexibility, but potentially decreases their awareness of others' actions as they can visually attend to different parts of the workspace. This led to various design interventions to work around the trade-off and facilitate awareness while maintaining individual power [22]. Some of these interventions address the challenge of divergent visual attention by indicating the visual attention of other users. For example, the system can provide visual feedback about other users' field of view through view windows [7] or multi-user scroll bars [24]. These representations act as a rough estimate of other users' visual attention, but more recent work has employed eye tracking to make collaborators' gaze points available to each other for remote [11, 13, 14] and collocated (on a large shared display) collaboration [59]. Overall, these examples address the challenge of divergent visual attention by facilitating awareness about other users' attention, but do not alter the way user actions are handled by the system.

Another design approach for coordination has been access control. Access control for groupware determines the conditions under which computational resources become available. A well-known example of access control is edit/view permissions for specific users or roles in online collaboration tools (e.g., [1, 2]). Such “access-matrix” or “role-based” models, however, are not the only means for managing access (for a review, see [54]). For ubiquitous computing applications, role-based models have been extended to take various contextual information (such as the time of day or the location) into account, resulting in context-aware access models [12, 34]. Overall, an important motivation for utilizing access control has been addressing security concerns during collaboration [54], but as Dourish and Bellotti argue, access control also contributes to a heightened workspace awareness due to decreased uncertainty about other users' actions [17].

To summarize, while both design approaches foster awareness, they put different demands on users' visual attention. Relying on self organization, especially in the absence of established social protocols, introduces uncertainty and requires situational awareness, which for many interfaces relies on visual monitoring. In contrast,

awareness through access control can be achieved without users having to monitor each other's actions, since role-restriction rules out the possibility of another user performing unwanted actions or viewing private information. On the other hand, the certainty provided by access control comes at the expense of individual power.

With visual attention-based access we aim to work around the trade-off between awareness and power by relying on self-organization when users visually attend to an action and restricting access when they do not. In doing so, we address what we identify as a gap in research: While previous work acknowledged the role of access-control in fostering awareness [17], *how traditional access control models can opportunistically be relaxed if users are visually attending to each other's actions* has not been explored. Our approach is similar to context-aware access models [12], but instead of location, presence or time, which has been the focus of earlier work on ubiquitous environments, we focus on visual attention due to its significance for awareness.

2.2 Coordination on Single Displays

Single display groupware reflects the trade-off between awareness and power described above, but also comes with particular considerations for coordination. Single display applications are technically WYSIWIS, since the interface state is uniform across users. The uniformity of the interface and the visibility of off-screen actions provide a common reference for users [22, 27, 53]. Thus, awareness on single displays has been treated as less problematic when compared to remote collaboration on shared workspaces [22]. However, *what is actually seen* by different users at a given time can diverge significantly due to the larger display size and users working in a "loosely coupled" manner (i.e., independently on different interface regions without coordinating their actions [29, 32, 53]). Coupling has been observed through different metrics such as task focus [29], spatial arrangement around the display [53] and, most relevant for our study, visual attention area as measured by head orientation tracking [32]. A general observation from these studies is the ease with which users fluidly switch between loosely coupled individual work to tightly coupled joint work [29, 32, 53]. As with remote collaboration, different levels of coupling on single displays come with various trade-offs. While loose coupling can facilitate fast, concurrent work, it can also result in duplicate efforts due to lack of awareness [29] or decreased concern for building consensus [8]. In a competitive use scenario, Birnholz et al. observe that concurrent input by multiple users on a single display enabled fast task execution but also resulted in users acting more in their own interests (when compared to single input) [8]. They explain this through the decreased likelihood of users to scrutinize others' actions due to their preoccupation with their own input that has led to decreased awareness. Along the same lines, Mayer et al., observe that competitive tasks result in an increased effort by participants to monitor the whole display space [39]. The studies discussed so far focus on semi-public settings in a confined space within a small group, but similar considerations have informed the design of public displays in urban environments [19, 42, 44]. For example, Fischer and Hornecker regard public visibility of actions as one of the reasons that people refrain from posting inappropriate content when interacting with urban media facades, but otherwise do not report to what degree appropriate behavior relies on the presence, and thus the monitoring, of other people [19].

Lack of customization for individual users is an additional source of conflict for single display interactions. Previous work has documented instances when navigation and object manipulation actions performed by one user interrupted another or when users competed for interaction area on public or semi-public displays [30, 44]. Users can to a certain extent avoid conflicts by assigning territories, employing turn-taking procedures or other social protocols [19, 44, 48, 55]. Importantly, successful self-organization such as turn-taking assumes users' awareness of contextual factors (e.g., who else is interacting with display or whether someone else is waiting in the line [19, 44]). Even so, conflicts can still occur when users knowingly interrupt others [44].

Another line of research has taken a more structured approach to conflict avoidance on single displays. This led to work on interaction mechanisms for conflict resolution [41] and providing access control in a way that is similar

to access-matrix models in remote groupware but without having to rely on device-based login information. Thus, previous work utilized hand gestures [58], touch capacitance [15, 40], fingerprint patterns [26] and proximity to an interactive surface [4, 33, 56] to identify users and manage access. For example, manipulation of objects can be restricted to touch actions by authorized users [26] or global changes can be programmed to require the parallel input of all the users in the form of “cooperative gestures” [40]. Some previous work also utilize user input as a proxy for visual attention to control view access permissions. For example, personal information can be shown only when users gets close enough to the screen so that their body acts as a visual obstacle against information voyeurism [56]. Additionally, various view-dependent projection [38], shutter glass [3, 36, 51] and proximity aware [16] techniques have been developed to build multi-view tabletops that provide personalized views for individual users on the same display. However, while various access mechanisms have been devised to determine what items are editable or viewable by different users, managing access based on who *visually attends* to an action or an object at a given time has not been explored. The closest work is the cooperative game GazeArchers [45], which requires both users to look at moving targets when shooting. In this case, visual attention input is used to add additional challenge, rather than to address awareness and conflict issues.

3 VISUAL ATTENTION-BASED ACCESS

Visual attention-based access determines what actions are available based on who visually attends the action. In this section, we distinguish between solitary and joint attention situations during synchronous collaboration and describe 4 different access types that are generated from their combination.

3.1 Visual Attention Situations

In contrast to previous observational studies that categorized attention patterns at the workspace level (how different users visually attend to different areas and to each other) [29, 32, 53], our basic unit is a single visual target; we are interested in how a specific visual target or action is visually attended to at a given time. This leads to a basic distinction between solitary and joint attention situations (Figure 2).

Solitary attention refers to situations in which an item or action is visually attended to by a single user, while other users are away or attending to another area of the workspace. In some cases, the attending user can be the owner of the object, which requires special consideration for access management.

Joint attention refers to situations in which an item or an action is visually attended to by all the relevant users.




Fig. 2. Examples of solitary and joint attention situations on a visual target at a given time.

3.2 Visual Attention-based Access Types

Actions can be assigned 4 different access types based on their availability in solitary and joint attention situations:

Universal actions can be accomplished under both attention situations. This access type can also be defined as lack of any role restriction. Potential use cases are actions that are not very critical, that can be easily reversed without serious consequences or when users' self-organization alone is sufficient for coordination.

Consensual actions require joint attention to be accomplished. They expand on the concept of cooperative gestures ([9, 10, 37, 40]), which require coordinated input from multiple users for the realization of certain actions (such as actions that affect the whole workspace [40]). However, while cooperative gestures require synchronous

Table 1. Types of actions that are available (●), unavailable (-) or only available to a particular user (◐) under solitary and joint attention situations.

Action can be accomplished		Solitary Atten.	Joint Atten.
Universal	under any attention situation	●	●
Consensual	only under joint attention	-	●
Supervised	if object owner or supervisor is attending	◐	●
Private	only if the owner is attending and no one else	◐	-

manual input, consensual actions are based on visual attention. This is based on the assumption that monitoring alone increases the capacity of users to intervene or prevent conflicts. Consensual actions potentially require less effort when compared to manual consent mechanisms (e.g., cooperative gestures [40] or confirmation buttons), since users can already be visually attending to the action. However, they are more restrictive than universal access.

Supervised actions require the attention of a particular user (instead of all users as in consensual). They are enabled for joint attention situations and also for solitary attention situations as long as the attending user has special rights. Potential use cases are actions that can benefit from the awareness of a particular user such as the supervisor of a session or owner of a document. When compared to traditional access rights management, which strictly restricts actions to a specific user, supervised access relaxes access by enabling actions by others if the particular user is monitoring. Both consensual and supervised actions are positive access criteria, because they require the visual attention of certain user(s) for access.

Private actions can only be accomplished if a particular user (such as the owner) is attending and no one else. The private access category is the counterpart of consensual access; it is a negative access criterion as the action becomes unavailable during joint attention. While consensual and, to a certain extent, supervised actions enforce awareness, private actions enforce privacy. Private access aims to limit “information voyeurism” of private information or actions on public displays [52]. This makes it similar to other solutions that provide personalized views for individual users (e.g., [3, 36, 51]), but instead of dedicated hardware, such as shutter glasses, private access relies on visual attention information.

Table 1 summarizes the availability of each access type under different attention situations. Note that we so far defined solitary and joint attention situations, respectively, as involving single and all relevant users. The distinction is straightforward for two users, but it would be more appropriate to view solitary and joint attention as a continuum for larger groups. In intermediary situations, such as when a subset of users is attending to an action, the system can be either strict or flexible regarding how it grants access. For example, consensus can strictly be interpreted as requiring the attention of all relevant users or, flexibly, a subset of users. We discuss these different approaches later in the paper in terms of scalability.

3.3 Input Handling

When an initiated action is available for a given visual attention situation, the system grants immediate access by executing the action promptly. A mismatch between the visual attention situation and the access type, however, can be handled in different ways:

Restricting access entirely, without any further interaction, is the least complicated approach. It is also the handling method we employed in the user study due to its straightforwardness.

Deferring action execution until additional input is another method. For example, an action initiated by a solitary user can be later confirmed by another user. When compared to full access restriction, deferring allows for more individual power, but also comes at the expense of real-time awareness and associated risks such as duplicate or irrelevant contributions.

Notifying other users is a real-time alternative to deferring. If a consensual action is initiated under solitary attention, the system can evaluate whether other users are visually attending to the workspace. If so, the system can provide a visual notification of the action near the users' locus of visual attention to inform them about the ongoing action (similar to previous work on gaze-adaptive visual feedback [49, 50] for single user applications). Notifications work around the trade-off between power and awareness even further when compared to restricting the action. On the other hand, they can introduce another long-acknowledged trade-off in groupware design, namely the trade-off between awareness and disruption [28], especially when users would prefer to remain focused on their individual tasks.

4 STUDY: VISUAL ATTENTION-BASED ACCESS PREFERENCES

Having defined different visual attention-based access types, the question remains how the framework can be put into use in collaborative applications, specifically, which access type makes sense for a particular action (such as editing or deleting an item) and on what grounds. We conducted an exploratory study to answer this question and find out user preferences and motivations related to different access types. Participants (in pairs) were tasked to decide which actions should belong to universal, consensual, supervised and private access types as they completed three different use scenarios of project planning, brainstorming and document sharing. The scenarios and their applications have been prototyped for a large vertical interactive display that was situated in a meeting room.

While visual attention-based access can be prototyped using different hardware setups, we chose the interactive whiteboard setup due to a number of practical reasons. First, we expected the collocated, single-display setup to increase participants' awareness of whether their input is performed under joint or solitary attention. Secondly, the setup was chosen to facilitate verbal feedback and deictic references to elements on the workspace (e.g., pointing to different elements while explaining their preferences) for data gathering purposes. In addition to providing a clear solitary–joint attention distinction, the 2-user design ensured equal centrality for the participants, as participants could position themselves to the left and right of each other.

The assignment of access types had instant effect, allowing the participants to immediately observe and test the effects of their preference. Additionally, the joint nature of the assignment task enabled us to observe the agreement process between the participant pairs (i.e., the participants' externalized reasoning and discussions about why a certain action should be available or unavailable for different attention situations). Overall, the research questions that motivated the study were:

- (1) What are the access type preferences for different actions?
- (2) What are the motivations reported by participants as they assign access types to different actions?
- (3) How do the particular qualities of visual attention-based access manifest in participant preferences and interactions?

4.1 Apparatus

The study has been conducted on a $2,05 \times 1,20$ meter vertical interactive surface consisting of three adjacent displays, each with a resolution of 1080×1920 pixels. The displays were able to distinguish between touch and pen input using IR image recognition. In all of the applications, pen input was mapped to editing content while touch input was mapped to moving the elements on the screen. Both could be used for other actions that are accomplished by buttons. Head position and orientation of users were tracked (Figure 3), by an OpenCV

application that detects head-worn markers using a web camera (running at 640×480 pixel resolution) mounted at the ceiling. Both touch and head tracking data were transmitted to the web-based whiteboard applications through web sockets. We determined whether a visual area is attended by a participant by scoring the visual attention information using visual angle (θ) and distance (d) values between the head and the target on the screen (Figure 3). The participants' gaze points were made visible on the screen to indicate how the system sensed their visual attention.




Fig. 3. Detection of head position and orientation in relation to the vertical interactive screen. We determined whether a participant is attending to a target by scoring the visual angle (θ) and distance (d) values between the head and the target on the screen.

4.2 Scenarios

The multiple-application design of the user study was intended to observe the potential commonalities in user preferences across actions in 3 different scenarios (Table 2). The tasks have been created so that participants could accomplish them without any need for specific knowledge, by drawing on their daily experiences.

4.2.1 Project Planning (PP). The first scenario involved creating different “to-do” elements for a hypothetical project that aims to “decrease the energy consumption of households in Helsinki”. The project planning scenario and related actions was inspired by the use of whiteboards in agile and scrum software development methods, in which different tasks are represented as cards that can be assigned to different individuals [18, 20]. The application, in total, supported 11 actions to which participants were instructed to assign access types. 8 of these actions were element based including *adding*, *editing*, *viewing*, *moving*, *changing the owner of* and *deleting* to-do elements and *editing* and *viewing* a personal calendar that was assigned to each participant. The remaining 3 actions were global and included *aligning all to-do* elements (that vertically positioned them based on their owners), *deleting all to-do* elements and *exiting the session*.

4.2.2 Brainstorming (B). The second scenario required participants to brainstorm for content ideas for a website about “life in the city” by writing their ideas on post-its. In contrast to to-do elements in the project planning task, post-its did not have any owners. The application supported 7 actions. Five of these actions were element based including *adding*, *editing*, *viewing*, *moving* and *deleting* post-its. The remaining 2 actions were global and included *deleting all post-its* and *exiting the session*.

4.2.3 Document Sharing (DS). The last scenario involved sharing personal documents for a magazine layout project. This scenario builds on previous research that investigated access management for personal media on shared displays (e.g., [30, 40, 46]). The application allowed participants to place different personal elements from a menu including two article drafts (the participants were told that they had composed the articles), one personal

Table 2. Overview of use scenarios and the list of object-level and global actions for each application. In each of the use scenarios the participants were instructed to “Discuss and assign attention based access rights for different items and action”.

Task description	Action list
Project planning. Imagine that you are tasked to make a plan and different “to-do” items for a project that aims to decrease the energy consumption of households.	To-do (object): add, draw/edit, view, move, change owner, delete Personal calendar (object): edit, view Global: align all to-dos, delete all to-dos, exit session
Brainstorming. Imagine that you are in a brainstorming session about gathering new ideas to make a website about life in the city of X.	Post-it (object): add, draw/edit, view, move, delete Global: delete all post-its, exit session
Document sharing. Imagine that you are working on a layout project that require you to go over your own documents.	Document (object): annotate, view, move, remove Global: pile all media, remove all media, exit session

bookmark element and one personal note element. While the other two applications required participants to create content from scratch, all of the elements in the document sharing scenario had an owner and were pre-configured but could be annotated on. Overall the application supported 7 actions. Four of these actions were element-level including *annotating*, *viewing*, *moving* and *removing* documents. The remaining 3 actions were global and included *piling all media* (moving all the documents to the left-hand side of the screen), *removing all media* and *exiting the session*.

Note that, the three applications featured elements with different levels of ownership. To ensure the relevance of supervised and private access types for global actions and elements with no owners, the participants could assign a general “session master” who could act as a substitute for the owner. Overall, the actions across different applications can be analyzed through these two dimensions:

Element-level/Global: Previous work distinguished between element-level (e.g., editing a single element) and global actions (e.g., piling up all the elements on the workspace) on shared workspaces [40]. We were interested in observing whether participants would select higher attention demand access types (consensual and supervised) for actions with global scope, in other words, whether previous insights from cooperative gestures ([40]) would extend to visual attention-based access.

Action typology: While scenarios involved different elements, the actions they support can be grouped under different typologies such as editing, viewing, moving or deleting. Crucially, different actions types had different levels of reversibility; moving or creating new elements could easily be reversed, but deletion was irreversible. We expected irreversible actions such as deletion to be assigned more restrictive access types (instead of universal). Additionally, some actions such as moving or drawing result in gradual changes that enable the other participant to intervene, while others such as deletion result in discrete and sudden changes. To make these two types of actions comparable, discrete actions required a continuous press on the button for 1 seconds, accompanied by a horizontal progress bar for visual feedback.

4.3 Access Type Control Panel

Participants could select an access type for an action anytime during the session through a contextual control panel (Figure 4) by pointing at one of the 4 access types. We had initially implemented a central menu to control the access types of all the instances of the same action, but during pilot studies this proved to be too abstract and

not flexible enough to match the evolving and element-specific participant preferences. Thus, we opted for an object-based control interface that contextually appeared under the related interface element.

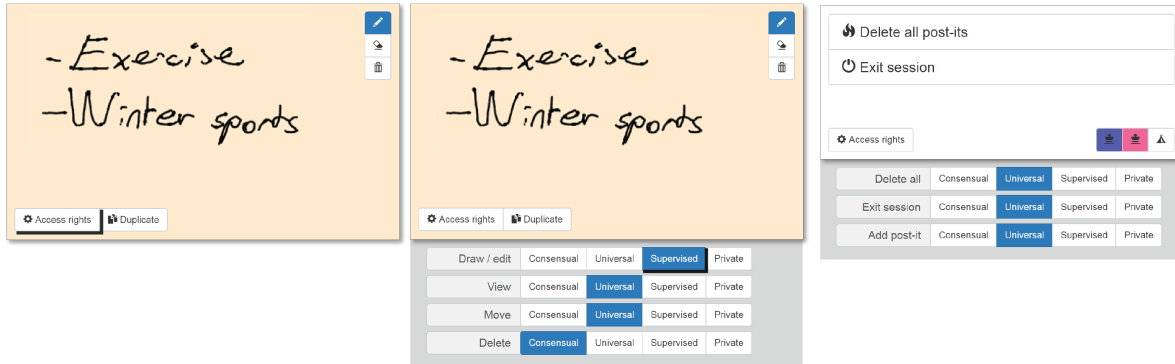


Fig. 4. Each element (such as the post-it element on the left) or the global action menu (right) had a collapsible “Access rights” menu that allowed participants to assign and modify their access preferences for different actions by pointing at one of the 4 access types using touch or a pen. The “duplicate” option allowed the creation of an item with the same access preferences. By default, the access configuration for all actions were set to universal. Other buttons allowed switching between drawing and erasing modes, as well as deleting the element.

4.4 Procedure

Each session started by introducing the purpose of the study, the interaction basics and the access type control panel to the participants. The participants were then given a demonstration of how different access types (universal, consensual, supervised and private) behave differently in each of the visual attention situations. After participants felt comfortable with interaction basics and understood access types (after 5–10 minutes), they proceeded to complete the tasks for each scenario in the order of 1) project planning, 2) brainstorming and 3) document sharing. The scenarios were respectively allocated 30, 15 and 15 minutes, based on their complexity and the amount of actions that need to be assigned different access types (Table 2). Apart from the task and time constraint, the sessions were unstructured. The participants could perform different actions and change their access type preferences anytime during the task, allowing them to experience and discuss different configurations before making a final decision. Additionally, each application featured two draggable information sheets as a reminder of different access types and the task. The participants were instructed to assign access types jointly (except for personally owned items) in order to make them take different considerations into account and observe their verbal reasoning.

Toward the end of each task, the participants were asked to finalize their preferences for actions and explain why they assigned the particular access types for different actions. After completing all of the tasks, the participants were interviewed about their general impression of visual attention-based access. The sessions approximately took one and a half hours.

4.5 Data Collection

The system logged participants’ access type assignment actions and their final preferences. Additionally, the sessions were video-recorded to observe participants’ interactions and discussions. The videos were later transcribed to link participants’ interactions and discussions with their access type assignment actions. The semi-structured

post-study interview inquired about participants' overall impression of visual attention-based access by drawing on their experiences and how they compare it with traditional access control.

4.6 Participants

20 participants (9 female), aged 18 to 35 ($m = 26.4$, $sd = 4.5$), were recruited for 10 sessions through university email channels and bulletin boards. The participants included 11 undergraduate and master level students, 7 researchers (PhD candidate and post-doc) and 2 designers (1 front-end developer and 1 interaction designer). Participants, on average, reported moderate previous experience with brainstorming ($m = 3.3$, in a scale of 1–5), project planning ($m = 3.1$) and sharing personal media ($m = 3.2$). They were compensated with two cinema ticket vouchers and their informed consent was collected for data logging and video recording.

5 RESULTS

We present the quantitative results based on log data and qualitative results derived from participant comments and actions. The qualitative results are denoted by the session number and the initials of the scenario during which the comment has been recorded (a participant remark recorded in session 4 during project planning scenario is indicated as "S4, PP" and only as "S4" for data from the final interview).

5.1 Access Type Assignment Process

In many cases, the assignment of a specific access type for an action was the result of participants' ongoing interactions with each other and the prototype. Although the sessions were unstructured, we observed that the assignment process was often influenced by 3 different factors:

Discussions among participants were observed in the form of references to prior experiences (e.g. workplace meetings or student committees) and verbal reasoning about why a specific action should be assigned a particular access type. Some of the arguments were presented as practical considerations. For example, in one session a participant questioned the other participant's idea of assigning a supervised access type: *"But this supervised still gives some possibility of misuse that one will change or two will change secretly... I think maybe it should be mostly consensual"* (S1, PP). In other cases, they were presented as personal preferences. For example, in the following exchange, participants (S2, PP) were expressing their priorities (in this case awareness vs. privacy) when deciding on whether viewing personal calendars should be private or supervised:

A: *It is good that you can look my calendar so you know.*

B: *No, I don't want anyone to look my calendar it is my own private stuff.*

Practical implications of assignments were another resource for deciding on a particular access type, as the prototype allowed participants to immediately experience the effects of their assignment. In some cases, the practical implications led to a revision of the initially assigned access type. In one session (S3, B), participants assigned consensual access to moving a post-it, but the assignment later proved to be impractical. This first led to a help request from the other participant and then to a revision of the original assignment:

B: *Can you look this way?*

A: *Ok* [A comes near B]

B: *The consensual moving is really...* [B reverts moving back to universal]

Progression of the task led to changes in participants' access preferences for some elements. For example, for to-do and post-it items it was common for participants to start editing the element with the default universal access type setting (that provides the least restriction) and then modify the access type once the element content was filled in.

5.2 Quantitative Results: Distribution of Final Access Preferences

We logged 1002 access type assignments for different actions from the final state of the application interfaces across the sessions. The number of assignments varied between sessions (min = 82, max = 166, sd = 24.3), as participants created different amounts of to-do (m = 5.8, sd = 2.3) and post-it (m = 6.3, sd = 3.4) elements. Thus, to calculate the overall access type distribution, we report the grand means that are aggregated within each session. Table 3 provides an overview of how access types were assigned to different actions.

We analyzed the distribution of access preferences for different kind of actions based on element-level–global distinctions and action typology. The data show salient differences between user preferences for these two dimensions across different applications. Universal access, which provides the least restriction, was rarely assigned to element-level delete (7.1%) and never to global delete actions, but it was overwhelmingly common for creating new items (95.0%). Universal access was also frequently assigned to view (72.3%) and element-level move actions (68.6%). On the other hand, the results show a high preference for consensual access for actions with global scope such as exiting the session (56.7%), global deletion (76.7%) and aligning elements (35%). For comparison, the preference for consensual access was lower for element-level delete (35.4%) and move (7.3%) actions. Instead, we recorded a higher incidence of supervised and private access for element-level actions (Table 3). These results are in line with previous design insights that propose joint control for actions with global scope [40] and our expectation that elements such as delete will be assigned more restrictive access types.

We also observed a degree of co-occurrence between different action types that belong to the same element, when we collectively analyzed to-do (n = 58), post-it (n = 63) and document (n = 80) elements that all supported editing, viewing, moving and deletion actions (Figure 5). For example, for elements that had editing actions set to private access (n = 35), a high proportion of them would also be assigned private access for deletion (n = 27), viewing (n = 23) and moving (n = 15). On the other end of the spectrum, for elements that had deletion action set to universal access (n = 13), a majority of them would also be assigned universal access for moving (n = 12), viewing (n = 9) and editing (n = 8). Thus, participants' access preferences for certain actions can help infer their preferences for other actions that belong to the same element.




Fig. 5. Two examples of co-occurrence: The distribution of access types for different actions of the same element when editing is set to private (left) and when deletion is set to universal (right). The data is across to-do, post-it and document elements. Each continuous line represents a group of elements that have the same access control settings for 4 action types.

The distribution data provide us with a summative view of overall patterns. However, it does not answer what accounts for the differences between user preferences for the same actions and interesting patterns in the data. For example, while the assignment of private access to view actions is foreseeable, we also recorded many instances in which private access has been assigned to manipulation actions and in rare cases even to actions with global scope such as global delete (3.3%). Below we report qualitative data that give insights about participants' reasoning when assigning different access types.

Table 3. Access right preferences for different actions in different scenarios across all sessions. The number of individual to-dos and post-its varied between different sessions. Thus, the results show grand means that are aggregated within each session. The bottom block shows the aggregated grand means based on action type.

				Access Preferences (%)				Univ.	Cons.
	Action	Level	Type	Univ.	Cons.	Supv.	Priv.	Supv.	Priv.
Project planning	Add to-do	Element	Create	90.0	10.0	0.0	0.0		
	Draw / edit to-do	Element	Edit	22.8	26.8	45.9	4.5		
	View to-do	Element	View	89.5	4.0	0.8	5.7		
	Move to-do	Element	Move	60.8	11.3	27.0	0.8		
	Change owner to-do	Element	Transfer	19.0	11.3	58.1	11.7		
	Delete to-do	Element	Delete	8.0	34.3	51.2	6.5		
	Edit calendar	Element	Edit	15.0	10.0	25.0	50.0		
	View calendar	Element	View	45.0	0.0	25.0	30.0		
	Align all to-dos	Global	Move	60.0	30.0	10.0	0.0		
	Delete all to-dos	Global	Delete	0.0	80.0	20.0	0.0		
Exit session	Global	Exit	30.0	60.0	10.0	0.0			
Brainstorming	Add post-it	Element	Create	100.0	0.0	0.0	0.0		
	Draw / edit post-it	Element	Edit	64.8	15.8	15.4	4.1		
	View post-it	Element	View	94.6	0.0	1.2	4.1		
	Move post-it	Element	Move	87.3	8.2	1.2	3.2		
	Delete post-it	Element	Delete	8.2	63.0	25.0	3.8		
	Delete all post-its	Global	Delete	0.0	70.0	30.0	0.0		
	Exit session	Global	Exit	20.0	50.0	30.0	0.0		
Document sharing	Annotate document	Element	Edit	30.0	2.5	32.5	35.0		
	View document	Element	View	60.0	0.0	16.2	23.8		
	Move document	Element	Move	57.5	2.5	10.0	30.0		
	Remove document	Element	Delete	5.0	8.8	27.5	58.8		
	Pile all media	Global	Move	30.0	40.0	30.0	0.0		
	Remove all media	Global	Delete	0.0	80.0	10.0	10.0		
	Exit session	Global	Exit	0.0	60.0	40.0	0.0		
	Element	Create	95.0	5.0	0.0	0.0			
	Element	View	72.3	1.0	10.8	15.9			
	(Element)	Move	68.6	7.3	12.8	11.4			
	(Global)	Move	45.0	35.0	20.0	0.0			
	Element	Edit	33.1	13.8	29.7	23.4			
	Element	Transfer	19.0	11.3	58.1	11.7			
	Global	Exit	16.7	56.7	26.7	0.0			
	(Element)	Delete	7.1	35.4	34.6	23.0			
	(Global)	Delete	0.0	76.7	20.0	3.3			

Table 4. Summary of participant comments for different motivations for assigning different access types.

	Motivation	Example comment
Universal	Unimp./Reversible	“...but this one the edit is universal just because it is not that as important as such and somebody can expand on it.” (S10, PP)
	Convenience	“...but while we are setting up this plan because we are alone in this room we can keep this as universal for usability.” (S9, PP)
	Trust / Cooperation	“This should be a list that everyone should be able to add to when they figure something out so I let it as universal. I trust my colleagues.” (S5,PP)
	Awareness (view)	“I also want everybody in this room to see what I have written here because it is some common data useful for project.” (S9, PP)
Consensual	Prevent accidents	“Consensual actions is really good because you would accidentally you know destroy stuff from the screen if you are not both looking.” (S6)
	Agreement	“I would say consensual makes sense in the sense that we all agree that we save and quit unless there is some hierarchy...” (S9, PP)
	Group awareness	“Consensual is pretty nice for something like this ... like we are planning to be there make sure that everybody is aware of what is happening.” (S3)
Supervised	Permission	“yes you can not remove it without my permission and yeah that is fine right.” (S2, DS)
	Owner’s awareness	“If you made a mistake in your work it can be supervised and you can see the person modifying your content.” (S7, DS)
	Scalable	“This is more like a topic which can be debated upon a lot so instead of making it consensual... I think it is up to the chair it should be.” (S4, B)
Private	Privacy	“If you have ideas that you are not certain of yet and you don’t want others to bother you about them...” (S5, B)
	Non-attention acc.	“For me editing calendar is private I want to do it myself.” (S4, PP)

5.3 Qualitative Results

We analyzed the video recordings and post-study interviews through an open coding process to classify observations and participants’ statements for each access type. This was followed by the grouping of participant statements to identify distinct motivations within each access type (Table 4). The analysis particularly focused on how participants referred to the specific affordances of visual-attention based access control.

5.3.1 Visual Attention-based Access for Preventing Conflicts. As expected, a common motivation for assigning consensual and supervised access types was conflict prevention in the absence of joint attention. Consensual access has been frequently assigned to actions with global scope (Table 3). Participants stressed that they “*have to both agree to end it [the session]*” (S7, PP) or stated that “*if I want to delete all to-dos it means that everyone agrees on something*” (S4, PP). Besides maintaining agreement, another motivation for consensual access was preventing accidents: “*I don’t think anyone should be able to delete all of them easily so we should all be there to see*” (S2, B). In a few instances, consensual access was assigned after the action of one participant interrupted the other (to prevent further interruptions): “*If it troubled you we should do it... so they can pile only if it is consensual*” (S2, DS).

Participants similarly used supervised access type to prevent unwanted actions and described it as a “permission” mechanism: *“I don’t want you to move this without my permission”* (S5, PP). Another participant stated that *“... it is better to have this supervised access types so the other could not delete it [personal task] or change it as he wants”* (S6, PP).

Conversely, lack of need for conflict prevention can partially explain the reasoning behind universal access. Universal access has frequently been preferred for creating, moving and viewing elements (Table 3). Actions such as creating new objects or moving are easily reversible, and this was reflected in participant comments: *“Anyone in principle can add stuff because we can always delete stuff”* (S2, PP). A related motivation was the cooperative setting of the task: *“If it is a competition, I would understand using them [private access] but this is just like brainstorming and creating to-do together”* (S8, PP). Another participant stated that both of them *“are on the same side it does not make sense that there is some guy with malicious intent”* (S2, B).

5.3.2 Visual Attention-based Access for Maintaining Awareness. In some cases, however, consensual and supervised access types have been motivated not through conflict prevention but as a means for facilitating awareness. In these cases, the participants utilized visual attention-based access restrictions as a way to ensure that the action performed will be attended by themselves or the group. For example, in one session participants set “change owner” to consensual so that *“we know who is actually in charge of this”* (S10, PP). During the task and later in the interview, consensual access was suggested as an explicit tool for enforcing the attention of the group, to help the *“group to focus on the single thing when needed”* (S10, PP). In another example, one participant assigned supervised access to moving: *“I would like to know where this goes... so it [moving] should be supervised by me.”* (S5, DS). Even though moving is easily reversible, supervised access makes tracking the changes easier by increasing the restrictions on the other user.

In these instances participants traded individual power with awareness. On the other hand, restrictions to individual power were also perceived as inconvenient. For example, one participant described consensual moving as a *“big hassle to have every one look at same place even if it is just two of us. I noticed that when I was not being able to move stuff”* (S3). Accordingly, a lack of restrictions to individual power can explain the preference for universal access, which was perceived as convenient: *“This is like nothing right now ... Because universal is the most easiest setting I guess ... this is not an idea that is necessary for this ... like random stuff”* (S8, PP). One participant stated that he would be willing to prioritize easy editing at the expense of awareness: *“Okay, I changed it to universal because at the moment I want to add ideas I am putting an input and if the people are not aware at the time of putting the input it is okay”* (S4, PP).

5.3.3 The Influence of Element Content. Our observations and participant remarks during the session revealed that the content of individual items can partly account for the variance of access types within element-level actions. In project planning and document sharing tasks, we observed that newly created items (with no content) were often left in their default universal access and were assigned another access type only after participants filled in some content. In some cases, access rights remained universal even after editing due to the unimportance of the content, such as such as for *“temporary ideas”* (S9, B) or *“random stuff”* (S8, PP). As the task progressed, participants set aside certain items as their *“general schedule”* (S1, PP), *“best ideas”* (S7, B) or things *“they both agreed on”* (S5, B) and assigned consensual access to the actions associated with the item.

On the other hand, supervised access type was often motivated by the personal content of the element: *“And those ones are more supervised because it is our actual personal tasks”* (S7, PP). In these cases, participants stressed the need for their visual attention for accessing the document: *“Because they are [pointing to his articles] articles if anyone wants to make change to this I have to be there to see what is happening”* (S4, DS). Supervised access has also been assigned to view actions (Table 3). In one instance, this was motivated by preventing other people from viewing personal information first: *“... cause sometimes in the office you come to the calendar and people have already seen how your days are going to look like...”* (S10, PP). The same consideration also explains the assignment

of private access. Participants referred to certain items as “*personal stuff*” or “*personal notes*” and made viewing these elements private: “*my bookmarks should also be private, no one can view it*” (S4, DS).

5.3.4 Private Access as a Proxy for Traditional Access. Even though personal content of an item was a reason for assigning private access, we also found counter-examples to this motivation, particularly when private access has been assigned to actions other than viewing. For example, one participant explicitly ruled out privacy as a concern when assigning private access type: “*It is not other person should not see it... because it is my bookmarks so I want to be the only person who can delete it from the screen.*” (S7, DS). We also recorded other statements that emphasize limiting the access to the owner: “*I don’t want anyone else to be able to edit them, because this is my own personal text*” (S5, DS). Unlike supervised actions, private actions are strictly restricted to a single person as they become unavailable if another person is visually attending. This also minimizes the risk of giving unwanted access. One participant highlighted this aspect of private access as a reason for selecting it over supervised access: “*talking about private stuff that you don’t want to have anyone else access to it seems like you can accidentally give someone access, but if it is private then it is really private*” (S3).

We classify this use of private access as attempts by participants to manage permissions in a way that is similar to traditional access control, namely based on who performs the action rather than who attends to the action. In other words, participants assigned private access not due to privacy concerns but to strictly restrict access to themselves. While private access can be used in this way, it comes with disadvantages that we observed during the session. When another user is visually attending to the same area, private access either requires him or her to look away or it unnecessarily restricts the access of the owner. This has been highlighted as a shortcoming of visual attention-based access during one post-task interview: “*... so there is private you can edit the document when no one is looking... in supervised everyone else can when you are looking... but there should be like that kind of private if someone is looking you can edit but he can’t*” (S7). We, therefore, interpret this use of private access as a shortcoming of visual attention-based access control.

5.3.5 Scalability. Even though the study involved two users, participants reflected on the scalability of joint attention for consensual access, particularly the difficulties that can occur for larger groups. One participant described the situation as “*hard to get everyone involved at the same time then it is hard to make any decisions*” (S5). Another participant stated “*If there is two of us, consensual is easy but if there is tens of like 20 coworkers it is hard everyone has to be present to do all of this and it is kind of unnecessary*” (S8, PP).

The concerns about scalability can also explain the assignment of supervised access to actions with global scope (such as global deletion or exiting the session) when there is a session master. Instead of requiring the visual attention of all users (as in consensual), participants proposed using supervised access for global actions: “*If it is consensual then when we exit session everyone should still be here but if supervised someone can leave little early*” (S1, PP). While supervised access can be used to address scalability issues, participants also suggested alternative solutions for larger groups during the post-task interviews, such as dividing the users into subgroups (“*I think there should be groups, you can include group 1, 2, 3 and put there some people.*”, S2) or employing degrees of consensus among the user group (“*if we are five people we choose who is consensual, it can be consensual between two of us or all five.*”, S10).

5.3.6 The Trade-off between Convenience and Certainty. During the final interview, participants were also asked to give their general impression of visual attention-based access and compare it to dedicated confirmation mechanisms (to make this more concrete participants were given the examples of hitting confirmation buttons on personal devices or a shared screen). Not having to provide manual input was described as convenient: “*It is very convenient because we don’t have to move our hands just look.*” (S6). This was related to avoiding another step: “*It may be more natural, I think, visual based [access]. When you have to push a button or receive something on your phone there is another step that adds to it and if you are doing it with visual based access right it can be quicker.*”

(S7). On these occasions visual attention-based access was compared favorably with explicit confirmations: “*You have freedom everyone looks at the screen and I am gonna delete this and do this... if everyone is confirming then it is a lot bigger process.*” (S6).

On the other hand, visual attention-based access was at times described as uncertain. This was related to the perceived uncertainty of measuring attention (“*It seems a bit less exact to be... measuring attention does not seem quite that exact way to manage the right...*”, S3) and potential cheating by users (“*or if you tried to spy on, is it based on head position right? You can kind of cheat.*”, S6). Participants also stated that “*explicit confirmation can be more legally suitable*” (S4) and for some situations “*there should be a second layer of authorization*” (S5). One participant’s comment directly pointed to a trade-off between convenience and certainty: “*When I am touching it, requires that I go and touch something. It is much more certain kind of manifestation of my attention so it requires effort from the user then the level of certainty increases also. Whereas this gaze-based attention inference it decreases this certainty thing so it could be that I am looking there, but I might just be thinking and not paying attention. But on the other side, it also helps fluidity of the interaction, and it does not enforce users to make explicit actions but is somehow like... could be much more blended in the interaction*” (S9).

6 DISCUSSION

6.1 Main Findings and Design Implications

In this section, we summarize the main findings from the study from which we derive various design implications that can be explored in future work.

6.1.1 Participants Employed Visual Attention-based Access Types to Maintain Agreement and Owner’s Control. Consensual and supervised access were chosen over universal and private access types when participants were willing to grant access as long as they were visually attending to the action. In these cases, visual attention-based access was perceived as a sufficient means for conflict prevention.

- This finding suggests that manual confirmation mechanisms proposed in earlier work such as collaborative gestures for global-level actions [40] or touch confirmations [26] can partly be substituted with visual attention-based access. On the other hand, designers should be aware of the uncertainties introduced by visual attention-based access when deciding on to what extent and how they can replace traditional access control.

6.1.2 Granting Access with Head-orientation Was Found Convenient but Uncertain. Not having to manually touch the screen for confirmations was highlighted as a convenient feature. At the same time, visual attention-based access introduced uncertainty that was attributed both to a mismatch between head orientation and participants’ actual locus of visual attention and also to situations when visual attention does not indicate awareness (i.e., when participants remarked that they may be looking but not paying cognitive attention). One strategy that participants employed to overcome uncertainty was the use of private access as a proxy for traditional access control. However, as we observed, this strategy comes with the disadvantage of unnecessary restrictions either to the owner’s access or to other users’ awareness. Thus, future work should consider alternative design options like those proposed below.

- One design approach is decreasing the level of uncertainty by employing more accurate measures of visual attention such as eye tracking [35]. Another approach would be to develop input handling techniques that account for the uncertainty when inferring visual attention from behavioral data. For example, the required level of certainty about users’ visual attention can vary depending on the action type; that is, deleting an item can require a higher threshold of certainty when compared to editing. Uncertainty can also be managed by letting the users know about how the system perceives their visual attention. In the current study, this was achieved by showing users their head orientation on the screen as a circle, but insights from

eye movement research warn that salient representations of visual attention can be intrusive [59]. A more acceptable solution would be providing subtle cues on items that indicate who is visually attending to an item and what actions are available. Finally, the interface can provide additional means for overriding or undoing an action (e.g. cancel buttons) to handle situations where awareness does not equate consent.

- An alternative approach could be using visual attention-based access complementary to traditional access models instead of a complete replacement. For example, private manipulation actions such as editing can be managed using traditional access control models if the system is able to determine who performs the action (e.g., through touch identification, proximity or personal input devices).

6.1.3 *Participants Assigned Access Types Not Only for Conflict Prevention but Also for Facilitating Awareness.*

In some situations, participants preferred supervised and consensual types not for preventing conflicts but as a means to ensure that the owner/supervisor is aware of the action or to direct other users' attention. This is an interesting finding as it provides a counter-example to our conceptualization of access management as an implicit effect of visual attention. Instead, *users can deliberately utilize access control mechanisms in order to control awareness.*

- When implementing visual attention-based access designers and researchers should be sensitive to user adaptations that can emerge once detailed measures of awareness becomes part of the interaction. Although not observed in this study, of interest is how actionable use of awareness information in groupware would affect users' privacy concerns, their methods for creating accountability and "plausible deniability" (the ability to perceive information without being held accountable [25, 43]).

6.1.4 *Both Action Type and Interaction History Were Relevant for Visual Attention-based Access.* The study featured fine-grained access controls for data-gathering purposes, but this level of fine grained control can be too demanding considering that most online collaboration tools make a basic edit and view access distinction. However, we observed that access type assignments were to a certain degree influenced by the action type and interaction history of the item, such as the assignment of consensual access once the item content is agreed by both users.

- A possible design implication of this is potential automation of access types based on interaction history. The system can keep track of the visual attention situations in which an item is created or edited and adapt the access type accordingly. For example, an item that is created under joint attention can automatically be assigned universal view access, or an item that is edited under joint attention can require the visual attention of the same users for further modifications. Similarly, the access rights for viewing an item can be extended to other users who had already viewed the item under the owner's supervision. Furthermore, access types can be clustered based on the user preference data to simplify user choice (e.g., if editing an object is set to private, then the system can automatically assign the same access to deletion).

6.1.5 *Strict Consensus Is Likely to Be Challenging in Larger Groups.*

- Another take-away is that although visual attention-based access can be more convenient when compared to manual confirmations, strictly requiring the joint attention of multiple users can be arduous. One possible solution is to configure consensual action to work with lower thresholds of joint attention such as the majority of users or with a minimum number of supervising users, such as when an action requires the attention of at least this/these person(s). The threshold can be configured dynamically based on the action type, since actions like deleting can be more critical than editing or moving. Another possible solution that is not explored in the current study is to employ alternative input handling techniques such as notifying other users instead of denying access.

6.2 Limitations and Directions for Future Work

Although our study was limited to two user interactions, we envision the use of visual attention-based access for larger groups and other physical settings. Previous studies suggest that increased number of individuals and multiple groups working in parallel can result in even more fragmented visual attention [31, 57]. Horizontal and vertical displays surfaces can come in various sizes and configurations that facilitate different levels of collaboration and visual monitoring [5]. The current study setting (single-wall display in a meeting room) ensured that the two participants were always in the proximity of each other, and could peripherally monitor the shared space and reorient their attention with relative ease. These conditions might not be valid for larger groups working on distributed or larger shared spaces. Even though the identified user motivations can help inform design work for other settings, the generalizability of user preferences observed in the current study remains an issue for future work, and partly depends on how flexibly joint attention is implemented for larger groups.

We also implemented visual attention-based access for collocated and synchronous interactions. Thus, our exploration of visual attention-based access only covers a single quadrant of the traditional groupware time/space matrix. However, visual attention situations and the access types framed in the paper equally apply to remote collaborations. For example, edit and view rights are often used in remote collaboration tools [1, 2], but we know of no tool that allows extending editing rights to other users based on whether the owner(s) is paying attention. Implementing similar mechanisms, however, would require taking additional factors into account. First, unlike collocated collaboration where the physical constraints of the space naturally limit the amount of users, many people can act on the same workspace area during remote collaboration. Keeping track of massively concurrent input that can be the case for remote collaboration is a challenge even when the user looks precisely at the edited workspace area. Second, as users' arm or body movements are not observable, remote collaboration gives fewer visual cues about users' actions, although past research has shown that this can be remedied to a certain extent [6, 22]. Third, remote groupware would require other means for sensing visual attention, which can range from crude measures of user presence to fine-grained eye tracking data. However, less straightforward is the extension to asynchronous collaboration, as the access types we described are based on concurrent visual attention situations.

In fact, the concepts of groupware and collaboration alone can imply sustained interactions within a defined set of users with pre-established roles, an assumption that might not hold valid in other settings such as spontaneous, brief and anonymous interactions with urban displays in public space (e.g., [19, 31, 42, 44]). These settings often involve larger displays, many transitory users and ephemeral content that might not be as personal or persistent as in groupware. At the same time, previous studies acknowledge the role of public visibility of actions in preventing inappropriate content from anonymous users [19] and observe that users can still feel ownership of the content they create for public displays [44]. Yet we are not aware of any prototype that flexibly handles user input based on the visual attention or the presence of others in the public space. Applications in this domain would require alternative technical solutions for identifying users and sensing visual attention as instrumented solutions (e.g., markers, eye trackers) can be impractical.

We also limited our scope to awareness as facilitated by visual perception. This was informed by the fact that the interactions with our applications relied on visual monitoring to be perceived; manual (touch and pen) input is visually observable through arm movements and the output provided by the system is visual. However, other interfaces can enable interaction using different modalities such as speech input or audio feedback. Input and output in other modalities would require different measures of awareness. For example, the system can utilize proximity information to infer whether a user hears another user's speech input.

Finally, as with every exploratory study, there are limitations to what we can claim. The user study we conducted was oriented towards gathering data about user preferences and motivations for visual attention-based access, rather than quantifying its performance. Further work will be required to assess visual attention-based

access in comparison to traditional access models in realistic use settings. Nonetheless, our work points to a number of user motivations and design considerations for future work to build on.

7 CONCLUSION

We proposed the concept of visual attention-based access as a form of a contextual access control model that manages access based on who visually attends to an action or an interface element. Based on insights from previous work, we highlighted that different approaches to awareness, namely relying on self-organization and access control, place different demands on users' visual attention. Visual attention-based access provides a contextual switch between these two approaches by relying on self-organization when users visually attend to an action and restricting access when they don't, in a way working around the trade-off between individual power and awareness [17, 22].

The user study provided us with data about access preferences for different types of actions and uncovered different motivations for utilizing different access types. We observed that in addition to preventing conflicts, visual attention-based access can be employed to ease keeping track of the workspace and direct others' attention. Visual attention-based access has been perceived as convenient but also uncertain, a finding that calls for various design measures to be explored in future work. Overall, our work contributes to a more detailed understanding of utilizing visual attention and awareness data as a real-time input in groupware.

ACKNOWLEDGMENTS

We would like to thank the anonymous reviewers for their feedback and Hendrik Schneider for his help with setting up the interactive screens. This work has been partly supported by the Academy of Finland project EMOEMP (Grant Agreement #305577).

REFERENCES

- [1] 2018. Google docs. <https://www.google.com/intl/en/docs/about/>. Accessed: 2018-02-10.
- [2] 2018. Office 365. <https://support.office.com/en-us/article/collaborate-in-word-2016-b3d7f2af-c6e9-46e7-96a7-dabda4423dd7?ui=en-US&rs=en-US&ad=US>. Accessed: 2018-02-10.
- [3] Maneesh Agrawala, Andrew C. Beers, Ian McDowall, Bernd Fröhlich, Mark Bolas, and Pat Hanrahan. 1997. The Two-user Responsive Workbench: Support for Collaboration Through Individual Views of a Shared Space. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '97)*. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 327–332. <https://doi.org/10.1145/258734.258875>
- [4] Michelle Annett, Tovi Grossman, Daniel Wigdor, and George Fitzmaurice. 2011. Medusa: A Proximity-aware Multi-touch Tabletop. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology (UIST '11)*. ACM, New York, NY, USA, 337–346. <https://doi.org/10.1145/2047196.2047240>
- [5] Carmelo Ardito, Paolo Buono, Maria Francesca Costabile, and Giuseppe Desolda. 2015. Interaction with Large Displays: A Survey. *ACM Comput. Surv.* 47, 3, Article 46 (Feb. 2015), 38 pages. <https://doi.org/10.1145/2682623>
- [6] Ignacio Avellino, Cédric Fleury, Wendy E. Mackay, and Michel Beaudouin-Lafon. 2017. CamRay: Camera Arrays Support Remote Collaboration on Wall-Sized Displays. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York, NY, USA, 6718–6729. <https://doi.org/10.1145/3025453.3025604>
- [7] Michel Beaudouin-Lafon and Alain Karsenty. 1992. Transparency and Awareness in a Real-time Groupware System. In *Proceedings of the 5th Annual ACM Symposium on User Interface Software and Technology (UIST '92)*. ACM, New York, NY, USA, 171–180. <https://doi.org/10.1145/142621.142646>
- [8] Jeremy P. Birnholtz, Tovi Grossman, Clarissa Mak, and Ravin Balakrishnan. 2007. An Exploratory Study of Input Configuration and Group Process in a Negotiation Task Using a Large Display. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*. ACM, New York, NY, USA, 91–100. <https://doi.org/10.1145/1240624.1240638>
- [9] Lauren J. Bricker, Marla J. Baker, and Steven L. Tanimoto. 1997. Support for Cooperatively Controlled Objects in Multimedia Applications. In *CHI '97 Extended Abstracts on Human Factors in Computing Systems (CHI EA '97)*. ACM, New York, NY, USA, 313–314. <https://doi.org/10.1145/1120212.1120410>
- [10] Xiang Cao, Clifton Forlines, and Ravin Balakrishnan. 2007. Multi-user Interaction Using Handheld Projectors. In *Proceedings of the 20th Annual ACM Symposium on User Interface Software and Technology (UIST '07)*. ACM, New York, NY, USA, 43–52. <https://doi.org/10.1145/1240624.1240638>

- [//doi.org/10.1145/1294211.1294220](https://doi.org/10.1145/1294211.1294220)
- [11] Jean Carletta, Robin L. Hill, Craig Nicol, Tim Taylor, Jan Peter de Ruiter, and Ellen Gurman Bard. 2010. Eyetracking for two-person tasks with manipulation of a virtual world. *Behavior Research Methods* 42, 1 (2010), 254–265. <https://doi.org/10.3758/BRM.42.1.254>
 - [12] Michael J. Covington, Wende Long, Srividhya Srinivasan, Anind K. Dev, Mustaque Ahamad, and Gregory D. Abowd. 2001. Securing Context-aware Applications Using Environment Roles. In *Proceedings of the Sixth ACM Symposium on Access Control Models and Technologies (SACMAT '01)*. ACM, New York, NY, USA, 10–20. <https://doi.org/10.1145/373256.373258>
 - [13] Sarah D'Angelo and Andrew Begel. 2017. Improving Communication Between Pair Programmers Using Shared Gaze Awareness. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York, NY, USA, 6245–6290. <https://doi.org/10.1145/3025453.3025573>
 - [14] Sarah D'Angelo and Darren Gergle. 2016. Gazed and Confused: Understanding and Designing Shared Gaze for Remote Collaboration. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 2492–2496. <https://doi.org/10.1145/2858036.2858499>
 - [15] Paul Dietz and Darren Leigh. 2001. DiamondTouch: A Multi-user Touch Technology. In *Proceedings of the 14th Annual ACM Symposium on User Interface Software and Technology (UIST '01)*. ACM, New York, NY, USA, 219–226. <https://doi.org/10.1145/502348.502389>
 - [16] Jakub Dostal, Uta Hinrichs, Per Ola Kristensson, and Aaron Quigley. 2014. SpiderEyes: Designing Attention- and Proximity-aware Collaborative Interfaces for Wall-sized Displays. In *Proceedings of the 19th International Conference on Intelligent User Interfaces (IUI '14)*. ACM, New York, NY, USA, 143–152. <https://doi.org/10.1145/2557500.2557541>
 - [17] Paul Dourish and Victoria Bellotti. 1992. Awareness and Coordination in Shared Workspaces. In *Proceedings of the 1992 ACM Conference on Computer-supported Cooperative Work (CSCW '92)*. ACM, New York, NY, USA, 107–114. <https://doi.org/10.1145/143457.143468>
 - [18] Morten Esbensen, Paolo Tell, Jacob B. Cholewa, Mathias K. Pedersen, and Jakob Bardram. 2015. The dBoard: A Digital Scrum Board for Distributed Software Development. In *Proceedings of the 2015 International Conference on Interactive Tabletops & Surfaces (ITS '15)*. ACM, New York, NY, USA, 161–170. <https://doi.org/10.1145/2817721.2817746>
 - [19] Patrick Tobias Fischer and Eva Hornecker. 2012. Urban HCI: Spatial Aspects in the Design of Shared Encounters for Media Facades. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, New York, NY, USA, 307–316. <https://doi.org/10.1145/2207676.2207719>
 - [20] Yaser Ghanam, Xin Wang, and Frank Maurer. 2008. Utilizing Digital Tabletops in Collocated Agile Planning Meetings. In *Agile 2008 Conference*. 51–62. <https://doi.org/10.1109/Agile.2008.13>
 - [21] Carl Gutwin, Scott Bateman, Gaurav Arora, and Ashley Coveney. 2017. Looking Away and Catching Up: Dealing with Brief Attentional Disconnection in Synchronous Groupware. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW '17)*. ACM, New York, NY, USA, 2221–2235. <https://doi.org/10.1145/2998181.2998226>
 - [22] Carl Gutwin and Saul Greenberg. 1998. Design for Individuals, Design for Groups: Tradeoffs Between Power and Workspace Awareness. In *Proceedings of the 1998 ACM Conference on Computer Supported Cooperative Work (CSCW '98)*. ACM, New York, NY, USA, 207–216. <https://doi.org/10.1145/289444.289495>
 - [23] Carl Gutwin and Saul Greenberg. 2002. A Descriptive Framework of Workspace Awareness for Real-Time Groupware. *Comput. Supported Coop. Work* 11, 3 (Nov. 2002), 411–446. <https://doi.org/10.1023/A:1021271517844>
 - [24] Carl Gutwin, Gwen Stark, and Saul Greenberg. 1995. Support for Workspace Awareness in Educational Groupware. In *The First International Conference on Computer Support for Collaborative Learning (CSCL '95)*. L. Erlbaum Associates Inc., Hillsdale, NJ, USA, 147–156. <https://doi.org/10.3115/222020.222126>
 - [25] Jeff Hancock, Jeremy Birnholtz, Natalya Bazarova, Jamie Guillory, Josh Perlin, and Barrett Amos. 2009. Butler Lies: Awareness, Deception and Design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09)*. ACM, New York, NY, USA, 517–526. <https://doi.org/10.1145/1518701.1518782>
 - [26] Christian Holz and Patrick Baudisch. 2013. Fieberio: A Touchscreen That Senses Fingerprints. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology (UIST '13)*. ACM, New York, NY, USA, 41–50. <https://doi.org/10.1145/2501988.2502021>
 - [27] Eva Hornecker, Paul Marshall, Nick Sheep Dalton, and Yvonne Rogers. 2008. Collaboration and Interference: Awareness with Mice or Touch Input. In *Proceedings of the 2008 ACM Conference on Computer Supported Cooperative Work (CSCW '08)*. ACM, New York, NY, USA, 167–176. <https://doi.org/10.1145/1460563.1460589>
 - [28] Scott E. Hudson and Ian Smith. 1996. Techniques for Addressing Fundamental Privacy and Disruption Tradeoffs in Awareness Support Systems. In *Proceedings of the 1996 ACM Conference on Computer Supported Cooperative Work (CSCW '96)*. ACM, New York, NY, USA, 248–257. <https://doi.org/10.1145/240080.240295>
 - [29] P. Isenberg, D. Fisher, S. A. Paul, M. R. Morris, K. Inkpen, and M. Czerwinski. 2012. Co-Located Collaborative Visual Analytics around a Tabletop Display. *IEEE Transactions on Visualization and Computer Graphics* 18, 5 (May 2012), 689–702. <https://doi.org/10.1109/TVCG.2011.287>
 - [30] Shahram Izadi, Harry Brignull, Tom Rodden, Yvonne Rogers, and Mia Underwood. 2003. Dynamo: A Public Interactive Surface Supporting the Cooperative Sharing and Exchange of Media. In *Proceedings of the 16th Annual ACM Symposium on User Interface Software and*

- Technology (UIST '03)*. ACM, New York, NY, USA, 159–168. <https://doi.org/10.1145/964696.964714>
- [31] Giulio Jacucci, Ann Morrison, Gabriela T. Richard, Jari Kleimola, Peter Peltonen, Lorenza Parisi, and Toni Laitinen. 2010. Worlds of Information: Designing for Engagement at a Public Multi-touch Display. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. ACM, New York, NY, USA, 2267–2276. <https://doi.org/10.1145/1753326.1753669>
- [32] Mikkel R. Jakobsen and Kasper Hornbæk. 2014. Up Close and Personal: Collaborative Work on a High-resolution Multitouch Wall Display. *ACM Trans. Comput.-Hum. Interact.* 21, 2, Article 11 (Feb. 2014), 34 pages. <https://doi.org/10.1145/2576099>
- [33] Ulrike Kister, Patrick Reipschläger, Fabrice Matulic, and Raimund Dachselt. 2015. BodyLenses: Embodied Magic Lenses and Personal Territories for Wall Displays. In *Proceedings of the 2015 International Conference on Interactive Tabletops & Surfaces (ITS '15)*. ACM, New York, NY, USA, 117–126. <https://doi.org/10.1145/2817721.2817726>
- [34] Devdatta Kulkarni and Anand Tripathi. 2008. Context-aware Role-based Access Control in Pervasive Computing Systems. In *Proceedings of the 13th ACM Symposium on Access Control Models and Technologies (SACMAT '08)*. ACM, New York, NY, USA, 113–122. <https://doi.org/10.1145/1377836.1377854>
- [35] Christian Lander, Sven Gehring, Antonio Krüger, Sebastian Boring, and Andreas Bulling. 2015. GazeProjector: Accurate Gaze Estimation and Seamless Gaze Interaction Across Multiple Displays. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology (UIST '15)*. ACM, New York, NY, USA, 395–404. <https://doi.org/10.1145/2807442.2807479>
- [36] Roman Lissermann, Jochen Huber, Martin Schmitz, Jürgen Steimle, and Max Mühlhäuser. 2014. Permulin: Mixed-focus Collaboration on Multi-view Tabletops. In *Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 3191–3200. <https://doi.org/10.1145/2556288.2557405>
- [37] Can Liu, Olivier Chapuis, Michel Beaudouin-Lafon, and Eric Lecolinet. 2017. CoReach: Cooperative Gestures for Data Manipulation on Wall-sized Displays. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York, NY, USA, 6730–6741. <https://doi.org/10.1145/3025453.3025594>
- [38] Mitsunori Matsushita, Makoto Iida, Takeshi Ohguro, Yoshinari Shirai, Yasuaki Kakehi, and Takeshi Naemura. 2004. Lumisight Table: A Face-to-face Collaboration Support System That Optimizes Direction of Projected Information to Each Stakeholder. In *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work (CSCW '04)*. ACM, New York, NY, USA, 274–283. <https://doi.org/10.1145/1031607.1031653>
- [39] Sven Mayer, Lars Lischke, Jens Emil Grønbaek, Zhanna Sarsenbayeva, Jonas Vogelsang, Pawel W. Wozniak, Niels Henze, and Giulio Jacucci. 2018. Pac-Many: Movement Behavior when Playing Collaborative and Competitive Games on Large Displays. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, USA, Article 539, 10 pages. <https://doi.org/10.1145/3173574.3174113>
- [40] 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 (CHI '06)*. ACM, New York, NY, USA, 1201–1210. <https://doi.org/10.1145/1124772.1124952>
- [41] Meredith Ringel Morris, Kathy Ryall, Chia Shen, Clifton Forlines, and Frederic Vernier. 2004. Beyond "Social Protocols": Multi-user Coordination Policies for Co-located Groupware. In *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work (CSCW '04)*. ACM, New York, NY, USA, 262–265. <https://doi.org/10.1145/1031607.1031648>
- [42] Jörg Müller, Florian Alt, Daniel Michelis, and Albrecht Schmidt. 2010. Requirements and Design Space for Interactive Public Displays. In *Proceedings of the 18th ACM International Conference on Multimedia (MM '10)*. ACM, New York, NY, USA, 1285–1294. <https://doi.org/10.1145/1873951.1874203>
- [43] Bonnie A. Nardi, Steve Whittaker, and Erin Bradner. 2000. Interaction and Outeraction: Instant Messaging in Action. In *Proceedings of the 2000 ACM Conference on Computer Supported Cooperative Work (CSCW '00)*. ACM, New York, NY, USA, 79–88. <https://doi.org/10.1145/358916.358975>
- [44] Peter Peltonen, Esko Kurvinen, Antti Salovaara, Giulio Jacucci, Tommi Ilmonen, John Evans, Antti Oulasvirta, and Petri Saarikko. 2008. It's Mine, Don't Touch!: Interactions at a Large Multi-touch Display in a City Centre. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*. ACM, New York, NY, USA, 1285–1294. <https://doi.org/10.1145/1357054.1357255>
- [45] Ken Pfeuffer, Jason Alexander, and Hans Gellersen. 2016. GazeArchers: Playing with Individual and Shared Attention in a Two-player Look&Shoot Tabletop Game. In *Proceedings of the 15th International Conference on Mobile and Ubiquitous Multimedia (MUM '16)*. ACM, New York, NY, USA, 213–216. <https://doi.org/10.1145/3012709.3012717>
- [46] Dominik Schmidt, Julian Seifert, Enrico Rukzio, and Hans Gellersen. 2012. A Cross-device Interaction Style for Mobiles and Surfaces. In *Proceedings of the Designing Interactive Systems Conference (DIS '12)*. ACM, New York, NY, USA, 318–327. <https://doi.org/10.1145/2317956.2318005>
- [47] Kjeld Schmidt. 2002. The Problem with 'Awareness': Introductory Remarks on 'Awareness in CSCW'. *Computer Supported Cooperative Work (CSCW): The Journal of Collaborative Computing* 11, 3 (2002), 285–298.
- [48] 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 (CSCW '04)*. ACM, New York, NY, USA, 294–303. <https://doi.org/10.1145/1031607.1031655>

- [49] Baris Serim and Giulio Jacucci. 2016. Pointing While Looking Elsewhere: Designing for Varying Degrees of Visual Guidance During Manual Input. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 5789–5800. <https://doi.org/10.1145/2858036.2858480>
- [50] Baris Serim, Khalil Klouche, and Giulio Jacucci. 2017. Gaze-Adaptive Above and On-Surface Interaction. In *Proceedings of the 2017 Conference on Designing Interactive Systems (DIS '17)*. ACM, New York, NY, USA, 115–127. <https://doi.org/10.1145/3064663.3064744>
- [51] Garth B. D. Shoemaker and Kori M. Inkpen. 2001. Single Display Privacyware: Augmenting Public Displays with Private Information. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '01)*. ACM, New York, NY, USA, 522–529. <https://doi.org/10.1145/365024.365349>
- [52] Desney S. Tan and Mary Czerwinski. 2003. Information Voyeurism: Social Impact of Physically Large Displays on Information Privacy. In *CHI '03 Extended Abstracts on Human Factors in Computing Systems (CHI EA '03)*. ACM, New York, NY, USA, 748–749. <https://doi.org/10.1145/765891.765967>
- [53] Anthony Tang, Melanie Tory, Barry Po, Petra Neumann, and Sheelagh Carpendale. 2006. Collaborative Coupling over Tabletop Displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '06)*. ACM, New York, NY, USA, 1181–1190. <https://doi.org/10.1145/1124772.1124950>
- [54] William Tolone, Gail-Joon Ahn, Tanusree Pai, and Seng-Phil Hong. 2005. Access Control in Collaborative Systems. *ACM Comput. Surv.* 37, 1 (March 2005), 29–41. <https://doi.org/10.1145/1057977.1057979>
- [55] Edward Tse, Jonathan Histon, Stacey D. Scott, and Saul Greenberg. 2004. Avoiding Interference: How People Use Spatial Separation and Partitioning in SDG Workspaces. In *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work (CSCW '04)*. ACM, New York, NY, USA, 252–261. <https://doi.org/10.1145/1031607.1031647>
- [56] Daniel Vogel and Ravin Balakrishnan. 2004. Interactive Public Ambient Displays: Transitioning from Implicit to Explicit, Public to Personal, Interaction with Multiple Users. In *Proceedings of the 17th Annual ACM Symposium on User Interface Software and Technology (UIST '04)*. ACM, New York, NY, USA, 137–146. <https://doi.org/10.1145/1029632.1029656>
- [57] Lauren Westendorf, Orit Shaer, Petra Varsanyi, Hidde van der Meulen, and Andrew L. Kun. 2017. Understanding Collaborative Decision Making Around a Large-Scale Interactive Tabletop. *Proc. ACM Hum.-Comput. Interact.* 1, CSCW, Article 110 (Dec. 2017), 21 pages. <https://doi.org/10.1145/3134745>
- [58] Mike Wu and Ravin Balakrishnan. 2003. Multi-finger and Whole Hand Gestural Interaction Techniques for Multi-user Tabletop Displays. In *Proceedings of the 16th Annual ACM Symposium on User Interface Software and Technology (UIST '03)*. ACM, New York, NY, USA, 193–202. <https://doi.org/10.1145/964696.964718>
- [59] Yanxia Zhang, Ken Pfeuffer, Ming Ki Chong, Jason Alexander, Andreas Bulling, and Hans Gellersen. 2017. Look together: using gaze for assisting co-located collaborative search. *Personal and Ubiquitous Computing* 21, 1 (2017), 173–186. <https://doi.org/10.1007/s00779-016-0969-x>

Received February 2018; revised May 2018; accepted September 2018