

CollabAR – Investigating the Mediating Role of Mobile AR Interfaces on Co-Located Group Collaboration

Thomas Wells

School of Computing and Communications
Lancaster University, United Kingdom
t.wells@lancaster.ac.uk

Steven Houben

School of Computing and Communications
Lancaster University, United Kingdom
s.houben@lancaster.ac.uk

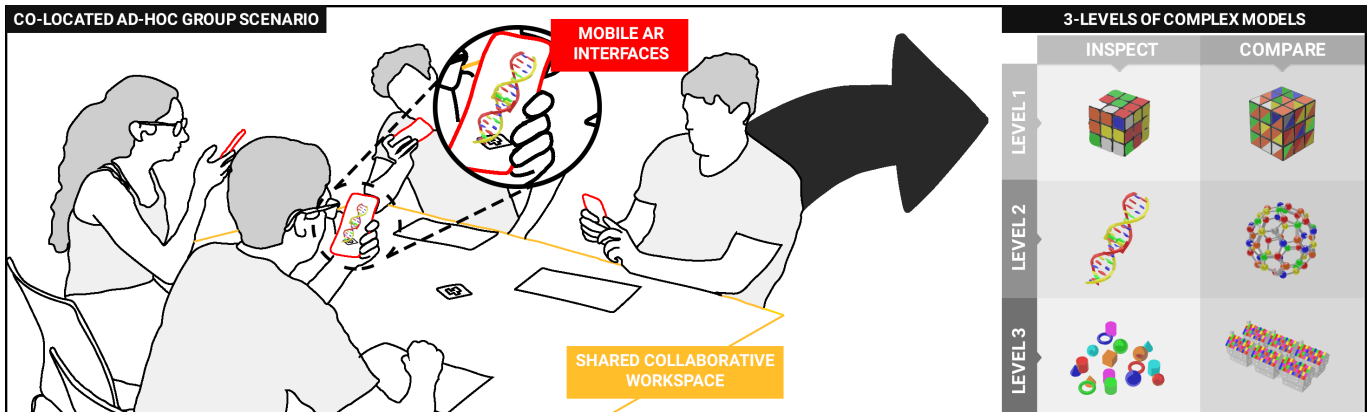


Figure 1. CollabAR is a Mobile Augmented Reality (M-AR) system that allows for co-located group interaction. Using CollabAR, we study collaborative group practices through a range of virtual models with different levels of complexity.

ABSTRACT

Mobile Augmented Reality (AR) technology is enabling new applications for different domains including architecture, education or medical work. As AR interfaces project digital data, information and models into the real world, it allows for new forms of collaborative work. However, despite the wide availability of AR applications, very little is known about how AR interfaces mediate and shape collaborative practices. This paper presents a study which examines how a mobile AR (M-AR) interface for inspecting and discovering AR models of varying complexity impacts co-located group practices. We contribute new insights into how current mobile AR interfaces impact co-located collaboration. Our results show that M-AR interfaces induce high mental load and frustration, cause a high number of context switches between devices and group discussion, and overall leads to a reduction in group interaction. We present design recommendations for future work focusing on collaborative AR interfaces.

Author Keywords

Mobile augmented reality; co-located collaboration.

CCS Concepts

•Human-centered computing → User studies in HCI;

INTRODUCTION

Handheld mobile devices have democratized access to Augmented Reality (AR) interfaces [21]. Mobile AR (M-AR) leverages cameras, sensors and processing power of mobile devices to enable new forms of blended reality which people can access, process and use on-the-go. A user can pull out their hand-held device and augment virtual objects onto the real world, e.g. to view furniture in their home without the need to purchase it. Mobile AR interfaces have opened up new opportunities for domains such as education, games, design, industry 4.0 or knowledge work [1].

While prior research suggests that AR interfaces can support collaborative work [6], most hand-held, through-the-screen AR (hereby referred to as mobile AR) interfaces are built predominantly with a focus on individuals and *not for collaboration*. Although some work explored co-located collaboration for mobile AR (such as [19, 34, 38]), there is generally a lack of understanding of how mobile-AR affects collaboration and group interaction. While the importance of co-located group work has been extensively demonstrated for touch-based interfaces [7, 8, 17, 32, 51], it is *unclear how mobile AR interfaces mediate, impact or shape group collaboration*. As concluded in the review on AR by Kim et al. [25], we need more insights on how collaborative systems can be created for AR.

As a first step towards a principled understanding of how co-located collaboration is mediated through AR, we conducted an observational study in which we explore how current mobile AR interfaces impact co-located collaboration. The overall goal is to provide insights into practices, problems and challenges encountered by groups, and identify a range of

recommendations for how to design and build effective mobile AR interfaces for collaborative work. In our study, we examine how groups of four people collaborate on basic tasks using standard mobile devices and a basic mobile AR interface (Figure 1). The tasks involve discovering, inspecting and comparing various virtual models within a meeting setting (where virtual objects are projected on a table). We designed 6 models that are exemplars of complex virtual data typically used in medical, educational, or knowledge work domains. We investigate *how we can build an understanding of existing group dynamics and practices in co-located ad-hoc situations*. Our empirical study looks at collaborative practices and codifies them into three sets of group dynamics: (i) collaboration styles; (ii) context switching; (iii) device handling and manipulation.

This paper contributes *insights into how current mobile AR interfaces affect co-located group collaboration*. We observed how groups utilise different collaboration styles and context-switching strategies to aid conversation and cooperation. We found that when the complexity of 3D objects increases, there is a decrease in the number of context-switches and a reduction of discussion and engagement in the groups. When collaborating on tasks with complex models, participants focus on their own devices with little interaction or discussion. Our study demonstrates that while AR can support collaborative work, the absence of clear collaboration mechanisms reduced collaboration and forced groups to invent workarounds.

RELATED WORK

As our study examines the mediating role of mobile AR for co-located group work, it builds on prior studies and systems that investigate (i) collaborative mobile augmented reality, and (ii) co-located group interaction with mobile devices.

AR Collaboration

While Augmented Reality (AR) has its origins in research in the 60s [2], the collaborative aspect is relatively under-explored [25]. As AR resides in the physical space between people, collaborative AR essentially attempts to “*augment the face-to-face collaborative experience*” [5, 30], thus providing “*a richness of interaction seemingly unmatched by any other means of communication*” [20]. Examples of AR have demonstrated that there is a value in these technologies to support co-located and distributed collaboration [31]. An early study using TransVision [42] explored how AR-based co-located collaboration supports collaborative design utilising handheld AR displays. A study using the Studierstube system [49] looked at how co-located users can visualise 3D objects using AR glasses. A further study by Reitmayr and Schmalstieg [41] utilising the same system investigated the possibilities of collaborative workspaces, but utilising M-AR devices. Butz et al. [10] use a head-worn display that allows co-located users to interact with shared virtual information in the form of virtual 3D widgets and physical objects. In more recent research, Benko et al. created a collaborative, projection-based AR system [3]. This overlays 3D objects into the real world using cameras and projectors. Due to the immersive image creation, users of the system do not require additional AR displays to view any virtual objects.

Ar Pad [33] is an example of hand-held, through-the-screen mobile AR that is designed for face-to-face collaboration. Similarly, Henrysson et al. explored the role of collaborative AR for game scenarios on early mobile phones [19]. Nilsson et al. studied how AR can play a role in crisis management by supporting cross-organisational collaboration [34]. Other work explored the role of eye contact as part of collaborative AR [38]. Despite these point examples, Billingham et al. concluded in a recent survey that there is a systemic lack of user studies that look at collaborative AR as there are “almost none that examined communication process measures” [4]. A more recent systematic exploration of the mixed-reality space by Ens et al. confirms that “collaborative MR technology is now mature enough to focus squarely on human needs” [12]. This includes looking at how hand-held AR can be used to look at complex collaboration in space as well as time and symmetry.

While commercial hand-held mobile AR applications often rely on multiple users or even entire communities, the collaborative aspect is implicit and secondary. For example, Ingress [43] is a team-based AR video game with 7 million players that allowed users to see and fight for control over a portal which existed in the real world. Similarly, Pokemon GO [36] was another hugely successful AR application of this kind which allowed users to battle at real world locations and in real time. A final example is the IKEA furniture app named IKEA Place [27], which allowed users to place and view virtual furniture over any real world location. All of these applications have some inherent collaborative scenarios, and generally thrive on the aspect of community such as battles in the real world. However, whilst users may be co-located, in reality they are interacting with their own version of the virtual content and not the same object. The user experience is quite insular and we conclude that there is a need to create a better understanding of the mediating role of AR for co-located collaborative scenarios.

Co-Located Collaboration on Mobile Devices

While there is little research on co-located collaboration for AR, co-located interactions and collaboration practices are better understood in touch-based interfaces for phones, tablets and tabletops. While there is a huge research space on Cross-Device Computing [8], we focus on work that explores the benefits, impact and performance of co-located collaboration. A considerable amount of work has explored how touch-based interfaces mediate collaboration. Early work looked at closely-coupled collaborations using shared stationary displays [11, 18]. Prior research also looked into how people collaborate and orient information on interactive tabletops which naturally afford group work [26]. This allowed for specific guidelines for how to construct co-located, collaborative work on a tabletop display [46] and studies on emerging collaboration concept, such as e.g., territoriality [45], group territory [35], collaborative coupling strategies [50], or collaboration [9].

With many people requiring mobility to accomplish their jobs, mobile devices appear to be the natural step to support collaboration and bring many advantages for users [16]. Lucero et al. [29, 28] were one of the first to show how mobile devices enable co-located collaboration for photo sharing or brain-

storming. Prior work also showed that in teams, collaboration and communication positively influence task performance, as well as their outcomes [24], and mobile devices were found to improve group collaboration [47]. In contrast, it is also suggested that in small group collaborations, the use of tablets negatively affect communication [17]. More recently there has been a push towards better understanding ‘co-located interaction’ as a concept [32, 13]. This is reflected in recent technologies aimed at collaborative mobile experiences (such e.g. HuddleLamp [39] or Connichiwa [44]) but studies into more detailed device-handling, collaboration practices or sensemaking [7, 22, 37, 40, 51]. We build on these insights and findings and use them as a framework to study the collaborative aspects of AR in co-located settings.

COLLABAR SYSTEM

The goal of this paper is to study the mediating role of mobile handheld AR interfaces on co-located collaborative work. To study this kind of collaboration, we constructed *CollabAR*, a basic AR interface for mobile devices that allows users to view and modify the same virtual object.

Interface

The interface shows a full-screen camera feed and overlays virtual objects on that feed when AR markers are detected (Figure 2). The application is specifically designed as a simplified exemplary interface with features that are commonly found in commercially available AR technology. CollabAR is a web-based interface that can be accessed through a standard browser on most modern mobile devices (including phones and tablets). Multiple users can take their phone, launch the web URL and point the device at the same AR marker. Through the interface, users can see the same 3D model and collaborate accordingly. Users can physically move towards or around the marker to zoom in and out of the model or see it from a different angle. As the AR marker is static, viewing the models from different sides is done by physically moving around.




Figure 2. The three views of the CollabAR Interface: nothing in view (left), model in view (middle) and the locking border (right).

Basic Gestures

To manipulate 3D models, we support a set of basic interactions similar to those available in off-the-shelf AR applications. Users interact with the 3D models through a tap-and-drag gesture to rotate the object. To rotate the object, the user touches the screen and simply moves their finger up or down on the

touchscreen to tilt the model upwards and downwards. Similarly moving their finger left and right will rotate the model accordingly. This is a common way to interact with 3D models in touch-based environments. When a user is modifying the model (rotating), a clutching mechanism locks the model (Figure 3) so only one user can *manipulate* the model at the same time. We thus implemented a *turn-based* system that only allows one user a time to modify the model. All other users are locked out of the manipulation mode, but can still see the model on their interface and walk around the AR marker to see it from different sides. Furthermore, non-modifying users can see the modifications made by the locked-user in real-time. To visualise this locking mechanism every user is given a specific colour. When a user modifies the model and receives the lock, all other connected devices visualise a coloured border around the camera feed to indicate which user is modifying the model (Figure 2 – right).

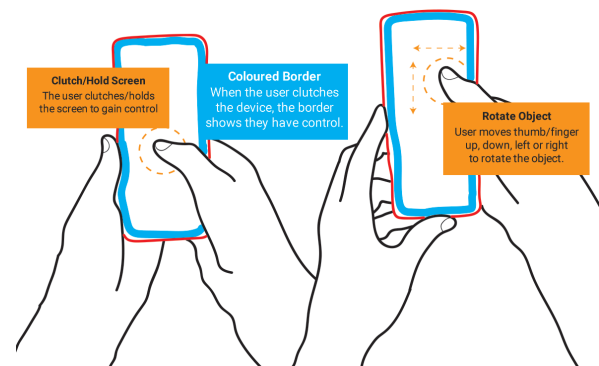


Figure 3. Basic interactions with CollabAR

Implementation

CollabAR is a web application which allows a user to point their camera at a marker to render a 3D object. Each 3D object is rendered locally in the browser to eradicate problems with latency when loading the object. We created a web application as it has fewer restrictions on which device can be used. To access this application users need a device and a web browser which has WebGL and WebRTC capabilities.

The system is built around three interaction phases. In the configuration phase, users approach the table, pull out their phone and load the web page. The interaction phase starts with the engagement with the object, pulling out a marker and placing it on the table to then simply point the camera at the marker and have a 3D model appear which can be manipulated. The disengagement phase occurs when users stop interacting with the model. The user simply closes the browser and puts away their phone and the marker. The collaboration part of the system was created by utilising Google’s real-time database Firebase [15], which facilitates the updating of data in the database and notifying clients of any new events within the database. This allowed us to send information such as the coordinates of the object to the database, which would then be updated in real-time on all clients. We recorded a device ID allowing us to detect when a user was interacting with the device and model, and lock the interactions to that user.

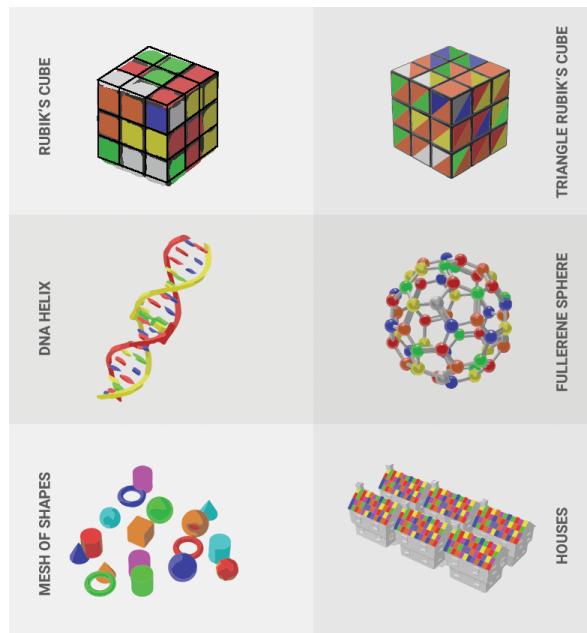


Figure 4. 6 models of varying complexity used for the tasks.

Virtual Models

We designed and created 6 models that were used during the study. These models are abstract exemplars of complex virtual data which are typically found in different AR domains such as education, healthcare, or other knowledge work settings. We designed the models to have different levels of complexity by considering the scale, geometric shape, density of colour and sub-shapes, and intricacies (Figure 4). The models include a *Rubik's Cube*, a *DNA helix*, a *mesh of shapes*, a *triangle Rubik's Cube*, a *Fullerene Sphere*, and a *block of houses*.

STUDY DESIGN

Using *CollabAR*, we studied collaborative activities in a co-located setting. Participants performed collaborative tasks with virtual models that have three different levels of complexity. The goal of the study was to understand how mobile AR facilitates and mediates collaborative group practices and dynamics during these tasks. In this section, we will describe our study design and setup.

Participants

We recruited 20 participants in total (11 identified as female, 9 as male) in groups of 4 using snowball sampling. There were no special requirements for this study, other than the participants were not colour blind (as there are colour specific tasks). The tasks were abstract and not domain-specific, so no information about the task content was given before the study.

All participants were aged between 18 and 24 and, whilst not a requirement, 7 participants had used AR before the study. Some example applications include a Dulux application, to augment virtual furniture into a physical space; Pokemon GO, a game that utilises AR to capture Pokemon; and the Google Pixel AR camera, which allows the user to overlay 3D objects into a room. All participants owned and frequently used at least one mobile device (such as a mobile phone or tablet).

Apparatus

Participants were invited to a circular meeting room setting which allowed for free movement around the table (Figure 5). Chairs were provided, participants were informed before the study began that they did not have to remain seated. After completing a consent form, participants completed a pre-study questionnaire which collected basic demographic data, including prior experiences with AR technologies and technology participants regularly used. All participants were provided with a Moto G4 (XT1642) mobile device which had a 5.5" display. Each participant was seated around the table, with the experimenter on the opposite end. Pens and NASA-TLX forms were provided for completing the questionnaire after every task and also in the case that participants wanted to take notes during the study.



Figure 5. The study setup in a meeting room

After an introduction to the task (see next section), participants were given a 3-minute training session on how to use *CollabAR*. Participants were then asked to complete tasks based on models that were given to them. There was no time limit given to the participants for each task: they simply had to complete the task by saying the final answer out loud. After each task was completed, the participants were asked to complete a post-task NASA-TLX form to capture the perceived workload. Following the completion of all the tasks, we concluded with a semi-structured group interview.

Tasks

We had a total of 6 models over 3 levels of varying difficulty for our study. Each level of difficulty came with three tasks, the first two were to inspect the model and the third task was to compare two models together. During the first inspect task, the participants would be asked a question such as '*count how many red tiles are on this cube*'. The participants would then inspect the cube as a group, counting the total number of red tiles. In the second inspect task, the participants would be asked something specific about the model, such as '*find how many shapes are touching each other*'. Similarly to the first task they would inspect the model, complete whatever task was given and come to a consensus before returning the answer. For the third task in each level we would introduce a second model and ask them to compare the original model along with one they had not seen. An example of a question asked is '*comparing the third and sixth model, list the colours that do not appear in both models*'.




Figure 6. The timeline visualises the coding of the video recording per group according to observed categories.

Tasks begun after the question had been read to the group. Once participants reached a consensus on their answer and said ‘finished’, they would progress on to the next task. We chose the questions for each task based on their simplicity as we wanted to ensure that the collaboration changes were based solely of the complexity of the model, rather than the complexity of the tasks.

Data Collection and Analysis

Each session was video recorded to analyse groups’ collaboration practices including their collaboration styles and context switches. We define a context switch as an instance in which the participant looks away from the virtual content. The videos were used to ensure that all interactions were captured for later analysis. Alongside the videos, the experimenter took notes which would coincide with the video. A NASA-TLX questionnaire was completed individually following every task to record the subjective workload. These questionnaires were used to initiate discussion points during a semi-structured group interview at the end of the session.

The data collected from video recordings were analysed using ChronoViz [14]. We coded the video data for collaboration practices, participants’ interactions with each other or with their devices. Analysis of these materials followed a grounded theory approach, utilising open coding of the video material along with notes that have been taken by the researcher. We utilised and adapted an existing coding framework which was previously developed by Tang et al. [50]. Similarly to a study

done by Brudy et al. [7], we adapted this framework to include groups of four participants instead of pairs and to also cater to multiple devices. However, we also added an additional style which was ‘Distributed View with Discussion’, as this happened a lot in each group. We focused particularly on individuals within the groups and where their attention was focused. We differentiated this from a context switch based on the duration. The collaboration styles we explored were:

- C1 Active Discussion** – Active discussion denoted any face-to-face discussions which included all participants. Due to this, there were limited interactions with their mobile devices.
- C2 Single-Shared View** – A single shared view denoted all participants focusing their attention to a single device.
- C3 Disjoint and Shared View** – Disjoint and shared view denoted the focus of 2-3 members on a single device, whilst others focused on their own device.
- C4 Disjoint and Distributed View** – Denoted 1-2 group members focusing on their device, whilst the other two were engaged in active discussion (not using their own devices).
- C5 Distributed View** – Denoted that participants were focused on their device with little to no discussion between each other. Complete focus on the task at hand.
- C6 Distributed View with Discussion** – Denoted that each participant was focused on their device, whilst continuing the conversation with the others in the group.

RESULTS

Overall, we observed how groups utilise different collaboration styles and context switches to aid conversational needs within groups. Results will be grouped by level of difficulty. There are a total of 3 levels of difficulty, level 1 (low) through to level 3 (high) and the difficulty is based on model complexity and scale. Our study indicates that there are significant difference between the different levels of difficulty between each model in terms of *mental demand* and *frustration*. Our findings also show that whilst the complexities of 3D objects increases, there is a decrease in the number of times participants interact with each other whilst observing the model. Instead, participants tended to focus on their own devices.

Use Patterns

Figure 6 provides a detailed breakdown of the analysis of the group collaboration. The entire timeline for each group is shown with a breakdown of the different categories including context switches, collaboration styles, device handling, and task start- and stop times. The timelines reveal that overall there was a high amount of context switches (all groups). More specifically, over all groups there was a total of 1,589 context switches (μ : 317.8, min: 183, max: 422, σ : 24.4) over the 9 tasks that each group was given. For some groups (1, 3, 5) there was a high degree of collaborative activity while for the other two groups there was less frequent collaboration. The timeline also shows how long the tasks took for each group. As detailed in Figure 7, tasks were relatively short and quick to complete (μ : 128.19s, min: 6s, max: 386s, σ : 83.15s). In general, the comparison tasks took longer than the inspection tasks, but there is a clear variability between groups.




Figure 7. Task completion time for each task.

For overall device mobility, we observed 44 occurrences of moving to landscape mode, 50 to portrait mode, 41 to place down the device, and 26 to pick it up. For overall collaboration styles, we observed 36 Active Discussion, 2 Single-Shared View, 57 Disjoint and Shared, 128 Disjoint and Distributed view, 104 Distributed View, 167 Distributed View with Discussion, and 27 Disjoint Shared and Disjointed. Figure 8 shows a break down of each group and the number of instances for each metric that was measured. Compared to other metrics, the amount of context switching was extremely high. A repeated measures ANOVA shows that on inspection tasks, there is no significant effect for the levels of difficulty ($F(2, 8) = 1.690, p > 0.1$) when it comes to context switches. However, further pairwise comparisons reveals that there is




Figure 8. Overall occurrences for each group.

significant difference between the third level of difficulty (μ : 17.8) and the other two levels (L2 μ : 34.8, L1 μ : 38.8).

When we examine the difference in task completion time (Figure 7), we can see that although the frequency of context switches has almost doubled between levels 2 and 3, the time taken to complete each task is similar. Statistical analysis shows that during the categories of device handling and device orientation, there was no significant difference between each measure on a difficulty level. The handle measure specifically showed that participants would rarely put down or pickup their devices ($F(2, 8) = 1.106, p > 0.1$). The maximum number of times a participant would put down their device was 1, which was toward the end of a task. Participants would generally hold their device for the entire time, even doing other tasks.

Overall, there was no significant difference between the collaboration style changes within the inspection tasks ($F(1.001, 8) = 0.659, p > 0.1$). There was a significant difference between the first (μ : 12.2) and third (μ : 7.8) level of difficulty, showing that the less difficult model had more collaboration style changes, particularly with participants looking through each other's devices and engaging in active discussion.

Similar results were garnered through the comparison tasks. We found a significant difference across difficulty levels ($F(2, 8) = 9.003, p < 0.05$). The pairwise comparisons revealed that there is significant difference between level 1 (μ : 51.8) and level 3 (μ : 26.4), which is similar to what we found in the inspection tasks. However, there is also a significant difference between level 2 (μ : 57.2) and level 3. These analyses clearly show that there is a lot more focus on the models during the more difficult tasks.

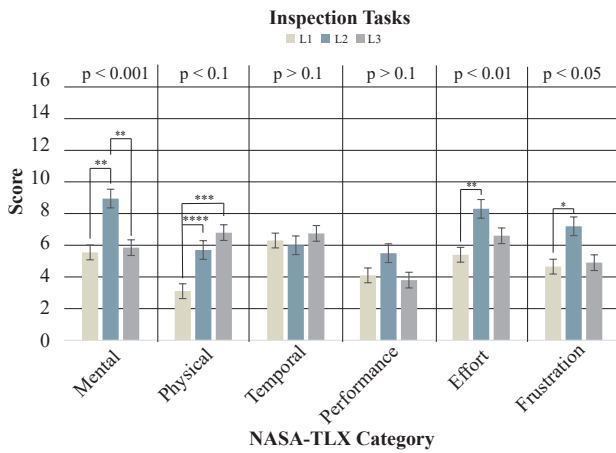
The collaboration style metric showed that there was a significant difference between the levels of difficulty ($F(2, 8) = 13.159, p < 0.05$). The main difference found within the styles was that on the highest difficulty level, there was significantly fewer collaboration style changes. We observed that during the comparison task with the highest level of difficulty, there was a lot more focusing on the task at hand with little discussion. Now and then there were pockets of distributed views with discussion, where participants would discuss what they could see without taking their focus away from the device (Figure 6).

Device handling and orientation also had no significant difference within-groups. During collaborative tasks we can see

that there was a slight increase of placing down and picking up a device ($F(1.019, 4.075) = 4.214, p > 0.1$) however, it still remained a rare occurrence. Participants would generally not put down their device too often, as is the same with inspection tasks. Similarly with orientation, there was a slight increase in frequencies across difficulty levels, however this was statistically insignificant ($F(2, 4.467) = 2.179, p > 0.1$). Observation of the tasks showed that instead of turning their device landscape to see a wider view of the objects, participants tended to either move further back to fit more into the portrait view, or to get closer to view a specific part. This could be in part due to the form factor of the display.

Usability

Within this section we report the statistics we garnered from running repeated measures ANOVA tests on each of our NASA-TLX metrics. Pairwise comparisons are done using Bonferroni correction. We split this into two categories: (i) the inspection tasks and (ii) the comparison tasks.

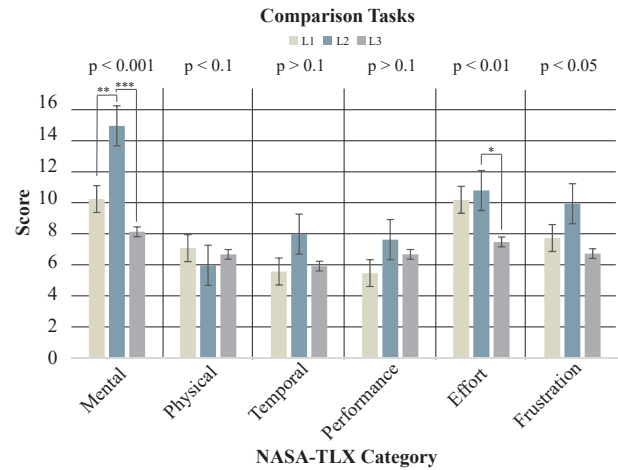


| NASA-TLX Metric | Difficulty Levels (μ) | | | F Statistic |
|-----------------|-----------------------------|------|------|---|
| | L1 | L2 | L3 | |
| Mental Demand | 5.55 | 8.95 | 5.85 | ($F(141.73, 233.6) = 11.528, p < 0.05$) |
| Physical Demand | 3.10 | 5.70 | 6.80 | ($F(144.40, 284.93) = 9.63, p < 0.05$) |
| Temporal Demand | 6.30 | 6.00 | 6.75 | ($F(2, 38) = 0.421, p > 0.1$) |
| Performance | 4.10 | 5.50 | 3.80 | ($F(2, 38) = 2.39, p > 0.1$) |
| Effort | 5.40 | 8.30 | 6.60 | ($F(2, 38) = 6.87, p < 0.05$) |
| Frustration | 4.65 | 7.20 | 4.90 | ($F(79.03, 394.97) = 3.802, p < 0.05$) |

Figure 9. Statistics of the comparison tasks shown as figures. The top figure is a graph showing the pairwise comparisons and significance between each level. The bottom figure is the table of the mean values and the F-statistic for each metric.

Upon first observation, it can be seen that – over both Inspect and Compare categories (Figure 9 and 10) – the perceived load over all metrics is quite high. This is because to complete a task, in the most part, there is a lot of multi-tasking involved. The participants would have to move around a lot, speak to their peers, and hold the device just to complete one part of the task. Altogether, this means there is a high overall perceived load. The analysis of NASA-TLX shows that there is significant differences in Mental Demand, Effort and Frustration between both categories (Inspect and Compare), over the three different levels of difficulty. This shows that the abstract exemplars we provided were different enough in difficulty that when this level increases, it becomes more mentally difficult

for the users to complete tasks. Within the Inspect category, there is also significant difference found within the Physical Demand section. This tells us that the different scales of the objects made participants more physical involved in order to view the models on their viewports.



| NASA-TLX Metric | Difficulty Levels (μ) | | | F Statistic |
|-----------------|-----------------------------|-------|------|-------------------------------------|
| | L1 | L2 | L3 | |
| Mental Demand | 10.20 | 14.90 | 8.10 | ($F(2, 38) = 15.77, p < 0.05$) |
| Physical Demand | 7.05 | 5.95 | 6.65 | ($F(2, 38) = 3.80, p > 0.1$) |
| Temporal Demand | 5.55 | 7.95 | 5.90 | ($F(2, 38) = 2.110, p > 0.1$) |
| Performance | 5.45 | 7.60 | 6.65 | ($F(2, 38) = 1.76, p > 0.1$) |
| Effort | 10.15 | 10.75 | 7.45 | ($F(1.52, 38) = 2.757, p > 0.05$) |
| Frustration | 7.70 | 9.90 | 6.70 | ($F(2, 38) = 3.01, p = 0.06$) |

Figure 10. Statistics of the comparison tasks shown as figures. The top figure is a graph showing the pairwise comparisons and significance between each level. The bottom figure is the table of the mean values and the F-statistic for each metric.

User Experience & Observations

We report findings from our observations and interviews in three themes: (i) context switches; (ii) device handling; and (iii) collaboration styles.

Theme 1 — Context Switches

AR allows collaborators to dynamically switch focus between a workspace and a communication space [48]. We observed that context switches occurred when referencing locations, for example when a participant would point out “the face closest to [the researcher]” – P7, or “the face closest to the camera” – P1. These context switches also happened when participants wanted confirmation of a particular statement they had said, or even just in natural conversation. If participants addressed a specific participant, they would generally glance in their direction, even if the participant did not look back.

Despite the high number of context switches (Figure 6 and 8), there was a lot less context switching during the inspection tasks in contrast to the comparison tasks. In particular, with the exception of group three (18%), the most difficult task had the lowest number of context switches (G1: 6%, G2: 8%, G4: 9%, G5: 12%). Even though the second level of difficulty was deemed the most mentally demanding in our quantitative findings, the third level of difficulty received the least amount of context switches, as participants were really concentrating on the device. The same is also true within the comparison




Figure 11. Disjoint and Distributed View was the second most common collaboration style.

tasks, where the least amount of context switches occurred on the most difficult task for all groups (G1: 11%, G2: 14%, G3: 5%, G4: 4%, G5: 7%). Some participants mentioned that the sixth virtual model was the most difficult due to how much they had to focus. Participants mentioned that “*if you got closer, you couldn’t see both sides [of the model] and it got more confusing*” – P4. This could explain why the context switching happened less, as they did not want to lose where the side they were interacting with.

Theme 2 — Collaboration Styles

We were interested in what collaboration patterns would emerge between participants within the different levels of difficulty. The most used collaboration style in groups 2-5 was Distributed View with Discussion (34%, 36%, 36% and 30%). This means that each group would actively discuss the task at hand with each other while looking at their own device rather than each other. Group 1 tended to use a little more Disjoint and Distributed View (29% see Figure 11) when compared to Distributed View with Discussion (26%) in which we observed that they would often check each other’s devices to confirm that they were seeing the same object. The video data shows that there is a common trend with people looking through their device onto the model but using their words and hands to place spatial signals to other users when discussing. For example, one participant said “*it is the side closest to the camera*” – P14 when asked to find which face had an L shape made of red tiles. While *spatial referencing* contributes towards coordination, our observations show that it is commonly used in AR-mediated collaboration as a work-around and replacement for direct face-to-face interaction.

We observed that every group had a high percentage of their collaboration styles as Disjoint and Distributed View (29%, 28%, 19%, 23%, 26%), with two users sharing one device to view the model. This was generally initialised by a trigger to do a context switch from their own personal device to the group. Such triggers include a brief conversation about content that the other participant could view; a brief look at the participants sitting next to them; an explicit invitation to look at another device; or generally just walking to the other side of the table to view and confirm. In some cases, participants would walk around the table to the other side to look at the same view. But rather than looking at their group member’s device, they would look through their own device instead. On

one or two occasions of this happening, a participant would glance at another group member’s device to confirm they are seeing the same thing. Participants would point out a certain part of the model or bring attention to it to ensure they were looking at the same thing through their own devices.

There was little to no Single-Shared View usage with only group 2 and 3 using this shared view *once*. Common practice was that participants would sit opposite to each other so they could view both sides of the model. They would only move around when they needed to be where all participants could view the same part of the model. Instead, they would trust each other when answers were given as part of the task. If participants had to confirm something and look at one part of the model, they would do so through their own devices instead of looking at the same device. This is likely caused by the form factor of the devices: because they were small, it was easier to focus on their own device and view.

Theme 3 — Device Handling

We expected that the orientation of the device would be changed significantly during the tasks to allow for more content to fit on the smaller form factor screens. However, our data shows that only a few participants turned their device to create a different field of view. Our observations show that most participants decided to keep their phones in portrait instead of landscape. When asked about this, one participant said “*I never really use my phone in landscape mode. It never crossed my mind to do that*” – P9. Another reason is that the device is easier to hold and more stable in portrait mode, thus increasing mobility allowing participants to more easily do other tasks or perform context switches.

Most 3D objects were scaled in such a way that if the marker was in the middle of the table, the model would fit completely in the viewport of CollabAR when standing next to the table. From there, the participant could get closer to inspect the object in more detail. If the models were scaled larger, more participants may have felt the need to turn their device landscape. This observation is supported by some of the groups, who turned and moved their device a lot more with the larger models, stating in the interview that “*sometimes it was difficult to concentrate on the larger objects because of the other things in the room*” – P7, such as a TV and cupboard. It is also supported by findings from the comparison tasks more specifically. Here, participants attempted to fit both models within their screen to be able to make the comparison. They would turn their phones to landscape - which meant we saw an increase in the orientation of the device during the comparison tasks when compared to the inspection tasks.

Summary of Results

Use Patterns – Video analysis shows diverse and changing collaboration styles (Figure 6), high degree of context-switching (Figure 8), and decrease in collaboration for complex models.

Usability – NASA-TLX data shows overall high load, and significant differences across models for mental and physical load as well as overall effort (Figure 9 and 10).

Observations – We observed a range of group practices and appropriations, including hand signalling to coordinate work,

approaches to switch context between the AR interface and the face-to-face collaboration, glancing through other peoples' devices to confirm tasks, or physically moving around the table (rather than using the interface) to inspect the model in order to ensure operation consistency for other users.

DISCUSSION

Our study shows how mobile-AR mediates collaboration. Although participants managed to complete the tasks, we identified recurring problems, observations and work-around strategies to mitigate the absence of *collaborative features*. Below we discuss the findings and present design recommendations.

Context-Switching and Device Handling

Our findings indicate extreme amounts of context switches for all the levels of difficulties for each model, especially considering the relatively short task duration times. Our data indicates that collaboration in M-AR *specifically* has a very high mental demand as users divide their attention between managing the device and collaborating with other users. Particularly, the combination of configuring the device correctly in space to see the AR marker, with actually looking 'through' the device to see the model produces a significant amount of *configuration work* [23]. The number of context switches decreases slightly as models get more complex, which – if we class a switch of context as an 'interruption' – shows that there is more focus on the models but less explicit collaboration. This was mitigated by workarounds such as spatial signalling, waving or conversing with others without face-to-face interaction. Overall, our study shows that *as intricacies and complexity of models increase, there is less collaboration and users report a high level of frustration with completing the tasks*.

Throughout the entirety of our study, we observed very little device handling - by which we mean the placing down and picking up of devices. Participants argued it was important for the system to maintain the viewport and the angle of the model so they could go back to the same point when analysing and discussing the tasks. As there are no tools available to help participants quickly recover from putting their device down and picking it back up again to continue where they left off, they would simply hold it in place while interacting with other participants or looking at other devices. While holding the device, participants were constantly adjusting, moving, and turning the device to ensure the AR models were properly visible. Especially in a collaborative setup, this constant re-alignment and fine-tuning of the device orientation and angle is a considerable distraction from the actual collaborative activities and imposes a high physical and cognitive demand on users. *The basic requirement to constantly hold up a device to be able to see the virtual model has a direct negative impact on collaboration*. Whilst this could be overcome by utilising HMDs to collaborate, it removes the portability and accessibility of this type of collaboration which is what makes this type of AR important.

Impact of Model Complexity on Collaboration

For complex models, the number of collaboration style changes decreased as participants focused more on their own devices. Overall, the more complex the model, the less dynamic collaboration occurs. Across all of the tasks, the most

used type of collaboration style is 'Distributed View with Discussion'. In this style, each participant focused on their own device while actively discussing the task at hand. This type of collaboration style occurred 167 times through the study and was the dominant way for participants to collaborate. Participants, thus, collaborate almost literally 'through' the device and the many observed context-switches are the shift from looking at the device to face-to-face discussions. The second most popular collaboration style occurred 128 times and was Disjoint and Distributed View. Here, 2 or 3 participants were focusing their attention on one single mobile device (sharing one viewport), while others in the group focused on their own device. This happened more often during complex models as participants wanted to confirm what they were seeing through another device even though what they were seeing was the same on both devices. Having a shared view allowed the participants to focus their attention on the same thing and they felt more confident in the knowledge that they 'could see the same thing on the same device' so there was no confusion.

The least used collaboration style is the Single-Shared View - which is when all participants focus their attention on a single device. Over all tasks and groups, this type of collaboration style only happened twice, when the whole group wanted to confirm one part of a model. Upon observation, it appeared to be difficult for all participants to huddle around one mobile device due to the angle and position of the device but also because there were 4 participants. We observed that no more than two participants would usually share a single viewport for a short amount of time. One of the report issues is that sharing one single viewport put constraints on the angle and position of the device, as multiple participants needed to see the AR model. This implies that the participants could not get closer to the models and view intricacies due to physical limitations without a considerable amount of configuration work.

Mobility in Tabletop AR

As AR relates digital models to a location, position and anchor in the real world, we observed a considerable amount of mobility work. Zooming and inspecting different sides of the model, or comparing multiple models all required participants to walk around the meeting table. Our results show that there is high physical demand throughout all tasks and that this increases when models are larger or more complex. Participants would view the model from different directions and, thus, move in all directions to obtain different vantage points over the model. Our observations even show that participants in some cases would have to stand on their tiptoes to view the top of the model. Furthermore, there was a lot of standing up and sitting down throughout the study. Participants also moved closer or away from the device to zoom, causing them to lean far over tables to see the model closer in the physical space. Although CollabAR allowed participants to rotate and scale the model, most groups simply walked or moved around the model to see different sides. We believe this is to ensure that the model would remain 'stable' for the entire group, and not make a modification while multiple participants are collaborating with the same model. This strong mobility requirement has implications for accessibility and general usability of AR.

Design Recommendations

We propose the following five concrete design principles to improve co-located collaborative AR experience. These are directly grounded in our empirical findings and oriented towards building applications and systems for collaborative work:

- D1 View-Clutching** – Our observations shows that a lot of effort went to aligning and configuring the AR view. A clutch mechanism could allow a user to hold the viewport in place whilst looking away from the device. This would mitigate the need for re-configuration and re-alignment steps when the user's attention shifts between looking at the device and looking at other users during the collaboration. Clutching the view would also allow users to move the device out of the way without losing the model on the screen.
- D2 Spatial Awareness** – In our study, context-switches were often caused by mobility (movement) around the table. This could be mitigated by supporting visual feedback on the location, distance and orientation of the other users in the room or by giving users accurate spatial references concerning the model. Particularly when multiple users are inspecting the virtual object, the system could provide indicators for where users are in relation to the physical marker or object representing the virtual model.
- D3 Visual Guidance** – To help users recover from context-switches or change to another collaboration style, a system can provide a visual guidance mechanism which would help users 'find' an appropriate view when they pick up the device. To avoid longitudinal and manual configuration of the view, the system could suggest where, when, and how to look at the model by providing visual indicators for the right direction and distance. This particular design principle could also help provide various views for users, which allows for different collaboration styles to be used.
- D4 Awareness Cues** – While the coloured border helped users see if interactions were locked, more visual feedback can be provided to the user on where other users are looking and who is interacting/modifying the model in view. Similarly to screen-based groupware, collaborative M-AR interfaces could provide awareness cues for what content is being used by each person. The 'turn-based' coloured frame visualisation in CollabAR is one example, but more advanced visual cues could be included in the interface.
- D5 View Filters** – To handle increasing complexity in virtual models and objects, the system can remove the background clutter and noise from the physical space and allow users to toggle between the full-screen camera view with the AR model in view and a more annotated, collaborative view that is optimised to see collaborative features, awareness cues, visual indicators, or even parts of the model which are visible to each user. These different filters could also support the articulation of certain viewing angles, which would enable complementary information when users observe the model from different angles.

LIMITATIONS AND FUTURE WORK

Understanding the impact of M-AR interfaces on collaboration is important to be able to push the field forward. There

is a need for new design principles and studies which will pave the way to new frameworks for M-AR to consolidate different AR applications. Our work focuses on how current M-AR applications mediate ad hoc collaboration. We provide insights through analysis of collaboration metrics such as context switching, collaboration styles and device handling. We present empirical findings of how users collaborate in a co-located setting and extrapolate design principles.

Our study was limited in size and scope with only 20 participants with limited diversity in demographics. Although we attempted to create an environment similar to that of the real world (standard meeting room), the study itself was controlled without specific focus on one application area. However, while limitations exist in the scope and size of the study, it sets a clear path for further research in this area. The results of this study provide general reflections and insights into how collaborative AR applications work in co-located settings using hand-held, through-the-screen AR. Our work encourages further development in this area of M-AR and we believe the design principles and concepts introduced in this paper can inform future studies and AR technologies.

There is a lot of room to break down the context switches in more depth and studying this further which would gather some interesting results in itself. Further to this, the mobility of the study (people walking around the table) is something that could also be taken into account and how this affects the collaboration. A limitation of our study was the limited types of tasks for users. Future research could look into how a similar collaboration framework lends itself to more advanced tasks which take a longer amount of time. Through these studies, there could be further conversational analysis, as well as looking at how different age groups could play a role. It could also be beneficial for comparative research to be undertaken, for example, how might context switching differ between hand-held vs. head-mounted AR displays? How could markerless AR solutions change the user experiences of re-locating and manipulating objects? And how might different form factor displays affect shared-view experiences?

CONCLUSION

While prior work argues that Augmented Reality (AR) facilitates collaboration, there is little empirical work into how AR affects co-located collaboration. We explored how current mobile AR interfaces mediate collaborative work. Our results show that while people can collaborate through a mobile AR interface, there is a high physical and cognitive load. We also observe extreme amounts of context switches between looking at the AR model through the device, and face-to-face collaboration. As virtual models become more complex, the amount of collaboration decreases. Based on our empirical findings, we present design principles for the implementation of mobile AR interfaces that support co-located collaboration.

ACKNOWLEDGEMENTS

We thank our participants and anonymous reviewers for their comments and feedback on this manuscript.

REFERENCES

- [1] Ronald Azuma, Yohan Baillet, Reinhold Behringer, Steven Feiner, Simon Julier, and Blair MacIntyre. 2001. Recent advances in augmented reality. *IEEE computer graphics and applications* 21, 6 (2001), 34–47.
- [2] Ronald T Azuma. 1997. A survey of augmented reality. *Presence: Teleoperators & Virtual Environments* 6, 4 (1997), 355–385.
- [3] H. Benko, E.W. Ishak, and S. Feiner. 2006. Cross-Dimensional Gestural Interaction Techniques for Hybrid Immersive Environments. DOI: <http://dx.doi.org/10.1109/vr.2005.27>
- [4] Mark Billinghurst, Adrian Clark, and Gun Lee. 2015. A Survey of Augmented Reality Augmented Reality. *Foundations and Trends in Human-Computer Interaction* (2015). DOI: <http://dx.doi.org/10.1561/11000000049>
- [5] Mark Billinghurst and Hirokazu Kato. 2002. Collaborative augmented reality. *Commun. ACM* (2002). DOI: <http://dx.doi.org/10.1145/514236.514265>
- [6] M. Billinghurst, S. Weghorst, and T. Furness. 1998. Shared space: An augmented reality approach for computer supported collaborative work. *Virtual Reality* (1998). DOI: <http://dx.doi.org/10.1007/BF01409795>
- [7] Frederik Brudy, Joshua Kevin Budiman, Steven Houben, and Nicolai Marquardt. 2018. Investigating the role of an overview device in multi-device collaboration. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 300.
- [8] Frederik Brudy, Christian Holz, Roman Rädle, Chi-Jui Wu, Steven Houben, Clemens Nylandsted Klokmoose, and Nicolai Marquardt. 2019. Cross-Device Taxonomy: Survey, Opportunities and Challenges of Interactions Spanning Across Multiple Devices. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, 562.
- [9] Stéphanie Buisine, Guillaume Besacier, Améziiane Aoussat, and Frédéric Vernier. 2012. How do interactive tabletop systems influence collaboration? *Computers in human behavior* 28, 1 (2012), 49–59.
- [10] A. Butz, T. Hollerer, S. Feiner, B. MacIntyre, and C. Beshers. 1999. Enveloping users and computers in a collaborative 3D augmented reality. In *Proceedings - 2nd IEEE and ACM International Workshop on Augmented Reality, IWAR 1999*. DOI: <http://dx.doi.org/10.1109/IWAR.1999.803804>
- [11] Joan Morris DiMicco, Anna Pandolfo, and Walter Bender. 2004. Influencing group participation with a shared display. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work, CSCW*.
- [12] Barrett Ens, Joel Lanir, Anthony Tang, Scott Bateman, Gun Lee, Thammathip Piumsomboon, and Mark Billinghurst. 2019. Revisiting collaboration through mixed reality: The evolution of groupware. *International Journal of Human Computer Studies* (nov 2019). DOI: <http://dx.doi.org/10.1016/j.ijhcs.2019.05.011>
- [13] Joel E Fischer, Stuart Reeves, Barry Brown, and Andrés Lucero. 2018. Beyond “Same Time, Same Place”: Introduction to the Special Issue on Collocated Interaction. *Human-Computer Interaction* 33, 5-6 (2018), 305–310.
- [14] Adam Fouse, Nadir Weibel, Edwin Hutchins, and James D Hollan. 2011. ChronoViz: a system for supporting navigation of time-coded data. In *CHI’11 Extended Abstracts on Human Factors in Computing Systems*. ACM, 299–304.
- [15] Google. 2019. Firebase Realtime Database. (2019). <https://firebase.google.com/docs/database/>
- [16] Luis A. Guerrero, Sergio F. Ochoa, José A. Pino, and César A. Collazos. 2006. Selecting computing devices to support mobile collaboration. In *Group Decision and Negotiation*. DOI: <http://dx.doi.org/10.1007/s10726-006-9020-3>
- [17] Jonathan Haber, Miguel A Nacenta, and Sheelagh Carpendale. 2014. Paper vs. tablets. 89–96. DOI: <http://dx.doi.org/10.1145/2598153.2598170>
- [18] Preben Hansen and Kalervo Järvelin. 2005. Collaborative Information Retrieval in an information-intensive domain. *Information Processing and Management* (2005). DOI: <http://dx.doi.org/10.1016/j.ipm.2004.04.016>
- [19] Anders Henrysson, Mark Billinghurst, and Mark Ollila. 2005. Face to face collaborative AR on mobile phones. In *Proceedings of the 4th IEEE/ACM international symposium on mixed and augmented reality*. IEEE Computer Society, 80–89.
- [20] Jim Hollan and Scott Stornetta. 1992. Beyond being there. In *Conference on Human Factors in Computing Systems - Proceedings*.
- [21] Tobias Höllerer and Steve Feiner. 2004. Mobile augmented reality. *Telegeoinformatics: Location-based computing and services* 21 (2004).
- [22] Leila Homaeian, Nippun Goyal, James R Wallace, and Stacey D Scott. 2018. Group vs individual: Impact of touch and tilt cross-device interactions on mixed-focus collaboration. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 73.
- [23] Steven Houben, Paolo Tell, and Jakob E Bardram. 2014. Activityspace: Managing device ecologies in an activity-centric configuration space. In *Proceedings of the Ninth ACM International Conference on Interactive Tabletops and Surfaces*. ACM, 119–128.
- [24] Petra Isenberg, Danyel Fisher, Sharoda A. Paul, Meredith Ringel Morris, Kori Inkpen, and Mary Czerwinski. 2012. Co-located collaborative visual analytics around a tabletop display. *IEEE Transactions on Visualization and Computer Graphics* (2012). DOI: <http://dx.doi.org/10.1109/TVCG.2011.287>

- [25] Kangsoo Kim, Mark Billinghurst, Gerd Bruder, Henry Been-Lirn Duh, and Gregory F Welch. 2018. Revisiting trends in augmented reality research: A review of the 2nd decade of ISMAR (2008–2017). *IEEE transactions on visualization and computer graphics* 24, 11 (2018), 2947–2962.
- [26] Russell Kruger, Sheelagh Carpendale, Stacey D Scott, and Saul Greenberg. 2003. How people use orientation on tables: comprehension, coordination and communication. In *Proceedings of the 2003 international ACM SIGGROUP conference on Supporting group work*. ACM, 369–378.
- [27] Dumi Lee. 2017. Ikea Place is an AR app that lets you put furniture on the street. (2017). <https://www.theverge.com/2017/9/20/16339006/apple-ios-11-arkit-ikea-place-ar-app>
- [28] Andrés Lucero, Jussi Holopainen, and Tero Jokela. 2011. Pass-them-around: collaborative use of mobile phones for photo sharing. In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 1787–1796.
- [29] Andrés Lucero, Jaakko Keränen, and Hannu Korhonen. 2010. Collaborative use of mobile phones for brainstorming. 337. DOI : <http://dx.doi.org/10.1145/1851600.1851659>
- [30] Stephan Lukosch, Mark Billinghurst, Leila Alem, and Kiyoshi Kiyokawa. 2015a. Collaboration in augmented reality. *Computer Supported Cooperative Work (CSCW)* 24, 6 (dec 2015), 515–525. DOI : <http://dx.doi.org/10.1007/s10606-015-9239-0>
- [31] Stephan Lukosch, Mark Billinghurst, Leila Alem, and Kiyoshi Kiyokawa. 2015b. Collaboration in Augmented Reality. *Computer Supported Cooperative Work (CSCW)* 24, 6 (dec 2015), 515–525. DOI : <http://dx.doi.org/10.1007/s10606-015-9239-0>
- [32] Sus Lundgren, Joel E Fischer, Stuart Reeves, and Olof Torgersson. 2015. Designing mobile experiences for collocated interaction. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*. ACM, 496–507.
- [33] D Mogilev, Kiyoshi Kiyokawa, Mark Billinghurst, and J Pair. 2002. Ar pad: An interface for face-to-face ar collaboration. In *Conference on Human Factors in Computing Systems: CHI'02 extended abstracts on Human factors in computing systems*, Vol. 20. 654–655.
- [34] Susanna Nilsson, Bjorn Johansson, and Arne Jonsson. 2009. Using AR to support cross-organisational collaboration in dynamic tasks. In *2009 8th IEEE International Symposium on Mixed and Augmented Reality*. IEEE, 3–12.
- [35] Shuo Niu, D Scott, Julia Nguyen, Derek Haqq, Lindah Kotut, Timothy Stelter, and Edward Fox. 2020. Investigating Paradigms of Group Territory in Multiple Display Environments. In *To appear in the ACM International Conference on Supporting Group Work (GROUP)*. DOI : <http://dx.doi.org/10.1145/3375193>
- [36] Janne Paavilainen, Hannu Korhonen, Kati Alha, Jaakko Stenros, Elina Koskinen, and Frans Mayra. 2017. The Pokémon GO experience: A location-based augmented reality mobile game goes mainstream. In *Proceedings of the 2017 CHI conference on human factors in computing systems*. ACM, 2493–2498.
- [37] Christopher Plau and John Stasko. 2009. Presence & placement: Exploring the Benefits of Multiple Shared Displays on an Intellectual Sensemaking Task. In *GROUP '09: Conference on Supporting Group Work*. 179. DOI : <http://dx.doi.org/10.1145/1531674.1531701>
- [38] Erik Prytz, Susanna Nilsson, and Arne Jönsson. 2010. The importance of eye-contact for collaboration in AR systems. In *2010 IEEE International Symposium on Mixed and Augmented Reality*. IEEE, 119–126.
- [39] Roman Rädle, Hans-Christian Jetter, Nicolai Marquardt, Harald Reiterer, and Yvonne Rogers. 2014. HuddleLamp. 45–54. DOI : <http://dx.doi.org/10.1145/2669485.2669500>
- [40] Roman Rädle, Thomas Luger, Clemens N Klokmoose, Thomas Plank, Harald Reiterer, and Hans-Christian Jetter. 2017. Is Two Enough? 4548–4560. DOI : <http://dx.doi.org/10.1145/3025453.3025537>
- [41] G. Reitmayr and D. Schmalstieg. 2001. Mobile collaborative augmented reality. In *Proceedings - IEEE and ACM International Symposium on Augmented Reality, ISAR 2001*. DOI : <http://dx.doi.org/10.1109/ISAR.2001.970521>
- [42] Jun Rekimoto. 1996. Transvision: A hand-held augmented reality system for collaborative design. *Proc. Virtual Systems and Multimedia* (1996).
- [43] S Rutherford. 2015. Meet Endgame: Proving Ground - Google's AR Mobile Game | Tom's Guide. (2015). <https://www.tomsguide.com/us/endgame-proving-grounds-beta,news-20650.html>
- [44] Mario Schreiner, Roman Rädle, Hans Christian Jetter, and Harald Reiterer. 2015. Connichiwa - A framework for cross-device web applications. In *Conference on Human Factors in Computing Systems - Proceedings*, Vol. 18. 2163–2168. DOI : <http://dx.doi.org/10.1145/2702613.2732909>
- [45] 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*. ACM, 294–303.
- [46] Stacey D Scott, Karen D Grant, and Regan L Mandryk. 2003. System guidelines for co-located, collaborative work on a tabletop display. In *ECSCW 2003*. Springer, 159–178.

- [47] Julian Seifert, Adalberto Simeone, Dominik Schmidt, Christian Reinartz, Paul Holleis, Matthias Wagner, Hans Gellersen, and Enrico Rukzio. 2012. MobiSurf: Improving co-located collaboration through integrating mobile devices and interactive surfaces. In *ITS 2012 - Proceedings of the ACM Conference on Interactive Tabletops and Surfaces*. DOI: <http://dx.doi.org/10.1145/2396636.2396644>
- [48] Verlag Springer, Mark Billingham, and Hirokazu Kato. 1999. *Collaborative Mixed Reality*. Technical Report. 261–284 pages. <https://pdfs.semanticscholar.org/35f8/82016ef051c6853dbd0f31ba9c3b09d0d914.pdf>
- [49] Z. Szalavári, D. Schmalstieg, A. Fuhrmann, and M. Gervautz. 1998. "Studierstube": An environment for collaboration in augmented reality. *Virtual Reality* (1998). DOI: <http://dx.doi.org/10.1007/BF01409796>
- [50] 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*. ACM, 1181–1190.
- [51] James R. Wallace, Stacey D. Scott, and Carolyn G. MacGregor. 2013. Collaborative sensemaking on a digital tabletop and personal tablets: Prioritization, comparisons, and tableaux. In *Conference on Human Factors in Computing Systems - Proceedings*. DOI: <http://dx.doi.org/10.1145/2470654.2466458>