

Evaluating the Extension of Wall Displays with AR for Collaborative Work

Raphaël James

Université Paris-Saclay, CNRS, Inria
Orsay, France

Anastasia Bezerianos

Université Paris-Saclay, CNRS, Inria
Orsay, France

Olivier Chapuis

Université Paris-Saclay, CNRS, Inria
Orsay, France




Figure 1: Photo-edited view of collaborators using our Wall+AR prototype displaying image cards, a setup similar to our study. The wall display is seen on the right, and shared virtual surfaces on the left and center of the image. The user on the right is looking at their personal space that moves with them, and is not visible to others. The virtual surfaces and personal space are shown only in Augmented Reality, while the wall is a physical display in the room.

ABSTRACT

Wall displays are well suited for collaborative work and are often placed in rooms with ample space in front of them that remains largely unused. Augmented Reality (AR) headsets can seamlessly extend the collaboration space around the Wall. Nevertheless, it is unclear if extending Walls with AR is effective and how it may affect collaboration. We first present a prototype combining a Wall and AR headsets to extend the Wall workspace. We then use this prototype to study how users utilize the virtual space created in AR. In an experiment with 24 participants, we compare how pairs solve collaborative tasks with the Wall alone and with Wall+AR. Our qualitative and quantitative results highlight that with Wall+AR, participants use the physical space in front and around the Wall extensively, and while this creates interaction overhead, it does not impact performance and improves the user experience.

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.

CHI '23, April 23–28, 2023, Hamburg, Germany

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9421-5/23/04...\$15.00

<https://doi.org/10.1145/3544548.3580752>

CCS CONCEPTS

• **Human-centered computing** → **Collaborative interaction**;
Mixed / augmented reality; Laboratory experiments.

KEYWORDS

Wall Display; Augmented Reality; Collaboration; Empirical Study

ACM Reference Format:

Raphaël James, Anastasia Bezerianos, and Olivier Chapuis. 2023. Evaluating the Extension of Wall Displays with AR for Collaborative Work. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*, April 23–28, 2023, Hamburg, Germany. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3544548.3580752>

1 INTRODUCTION

Wall displays, also referred to as Hiperwalls and Large High-Resolution Displays (LHRD), are well suited for collaborative work as they can accommodate multiple people simultaneously [8, 35, 39]. Due to their high resolution and size, they have been adopted by institutions that analyze or monitor large quantities of data, in research (e.g., biology [3]), operations (e.g., [2]), and industry (e.g., [1]). They are nonetheless heavy physical displays that are hard to move and expensive to reconfigure and extend. But, they are often placed in rooms with ample space in front of them to allow multiple

users to move, space that remains largely unused (see examples both in research and in practice [1, 3, 8, 35, 39]). So while wall displays are not easy to physically reconfigure or extend, the physical space available in front of them provides a unique opportunity to extend them *virtually*, for example, through augmented reality.

Traditional real-world multi-display environments, such as command-and-control rooms, combine various displays (wall displays, desktops, and digital tabletops) that research has shown each serves a particular purpose: wall displays and tabletops are commonly used for group awareness and collaboration, and tablets or desktops for personal work [20, 56, 64, 80]. Research on the topic has gone further, combining wall displays, in particular, with portable devices, such as smartwatches, mobile phones, or tablets, that can act as private displays or input devices when users are further away from the wall, and direct touch is not possible [19, 31, 52, 77].

Recent research work has started to combine existing wall displays with augmented reality (AR) headsets in the context of visual exploration. Here wall displays show publicly high-resolution renderings of core data visualizations. In these cases, the physical wall display acts as a public display seen with or without AR headsets, but also provides a rendering resolution and a field of view unmatched by AR headsets [17, 33]. Whereas AR headsets add personal information virtually, either on top or around the wall display content, in the form of additional visualizations [68], information [75], or highlights [36]. This work points to the potential benefits of adding AR head-mounted displays (HMD) to wall display environments, especially in the context of displaying private information. Some of this work, e.g., [42, 68], considers cases where virtual content is also publicly shared across all users with a headset. This opens a new research avenue: *using AR to increase the shared workspace available to collaborators by utilizing the physical space in front and around the wall display*.

This past work proposes but does not study the impact of using AR to extend the shared space around the wall display during collaboration. Is this extended AR space actually used when available, or do collaborators choose to work on the physical surface instead? Does the extended AR space affect how users collaborate and perform their tasks?

To answer these questions, we set out to empirically study the impact of extending wall displays with AR in collaborative contexts. In particular, we focus on fundamental questions regarding differences in the use of space (physical and virtual) and the impact on collaboration before and after adding AR. As a first step, we implement a prototype that allows users to use the virtual space around a wall display (see Figure 1). This Wall+AR system adds virtual space in the form of surfaces, and combines several techniques for users to organize, manipulate and move content between the wall display and the virtual space. We then use this system to run a comparative study with pairs of participants conducting collaborative tasks, using only the wall display or the wall display extended with AR headsets.

Our contribution is thus two-fold: a system that extends a wall display using AR in terms of visual space and interaction support; and the results of an empirical study that compares this extension with a wall display alone. Our comparison highlights that with the Wall+AR system, participants extensively used the physical space in front of the wall display. Virtual surfaces are used for

storing, discarding, and presenting data. Surprisingly, participants often use the virtual surfaces as their main interactive workspace, abandoning the wall display. We observed that adding AR to a wall display creates interaction overhead, such as physical and mental demand. Nevertheless, it also brings a real benefit over using the wall alone: the Wall+AR system is preferred and found more enjoyable and efficient than the wall alone, and we did not measure any loss in performance despite the interaction overhead.

2 RELATED WORK

Wall displays are extremely useful collaborative work environments. They have been found to improve performance [9], content organization [50], sensemaking activities [5], and have been shown to increase discovery and improve data analysis both in laboratory conditions [67] and real work settings [66]. These displays have a high resolution that allows them to render a very large amount of information [10, 21], as well as a field of view that surpasses what AR headsets alone can achieve today [17, 33]. And, of course, multiple people can see and use a wall display together [8, 35, 39] without requiring them to wear specialized equipment such as AR headsets. They are thus not going to be replaced any time soon. Nevertheless, they are heavy, hard to move, and expensive to reconfigure and extend. In our work, we focus on extending them using augmented reality and on studying the impact of such an extension. We thus cover related work on approaches that use AR to augment physical displays and on studies on collaborative interaction for data manipulation or sensemaking, focusing, in particular, on space use and vertical surfaces such as wall displays.

2.1 Physical Displays Combined with AR

Recent years have seen a plethora of work combining AR with other displays. For example, in the context of multi-display environments, AR has been used to create a continuous interface between different displays by using projection on a surface [69] or directly in the air [16]. Or to create new visualizations around and over physical tablets using AR headsets [41], to link different displays together for cross-device [70] or cross-display [82] interaction, to improve depth perception in 3D and personalized points of view [44], and to support file transfer [46]. Our work focuses instead on enlarging one single physical display in collaborative contexts.

For collaboration, AR has been used to add personal content over large horizontal and vertical displays. For instance, AR can add information related to the users on top of a map seen on a tabletop [59, 65], help users navigate a network seen on a wall display [36], add complementary text to the data points and graph nodes rendered on the wall [38, 74], or even show sign language subtitles on a TV [78]. Our work is complementary, we study the use of AR to extend the shared workspace area on wall displays rather than to add personal information.

Since the early work by Feiner and Shamash [23], AR has been used to extend working surfaces. AR can extend the workspace of desktop computers [23, 72, 81], smartphones [14, 61, 83], and even smartwatches [26]. The motivation in these projects is to augment the screen real-estate of "small" screens by simply enlarging the screen, offloading widgets, or providing additional information. Although the work on enlarging the screen size with AR generally

focuses on relatively small screens, some [60] enlarge the top of a circular wall display to address the physical vertical size limitation of the wall. Our motivation is similar, we use AR to augment a vertical wall display in order to extend the available real-estate, although our focus is on studying the effects of such an extension in collaboration. While there is existing work that uses AR to augment 2D visualizations displayed in a wall display, for example, with 3D visualization, (personal) links between visualizations, and additional contextual visualization presented just in the front of the wall [42, 53, 68], this work does not study collaboration.

There is also research on creating dedicated environments for conducting collaborative exploratory data analysis. For example, setups with a circular wall display and a tabletop at its center are used to render 2D visualizations, while AR headsets are used to complement these visualizations with 3D data [17]. This setup was later simulated in a fully Virtual Reality (VR) environment with several users using VR HMD [43]. In a similar vein, recent work [47] proposes using a fully VR environment to manipulate floating 2D documents to reproduce/simulate a sensemaking task executed in the past in a wall display environment [5]. Interestingly, the performance in the fully simulated VR environments [17] highlights that despite advancements in terms of interaction, resolution, and field of view, VR head-mounted displays (HMD) can still not fully reach the capabilities of physical displays. Always in the topic of sensemaking, there are also studies of how groups of users with AR HMDs place and organize floating AR 2D documents in a furnished office [51]. Most of this work does not empirically study the effect of these setups, the ones that do [43, 51] are discussed next.

2.2 Studies on Co-located Collaboration

There is a large amount of work on collaboration around tabletops. This includes work on collaborative strategies [32, 76], that go from tightly-coupled collaboration (e.g., using sequential strategies) to loose collaboration (e.g., using parallel strategies). And work on territoriality [71], where three main space territories have been identified during collaborative work: personal, group and storage. Nevertheless, this previous work on tabletops does not apply directly to wall displays, AR and VR. In such contexts, users move around to take advantage of the environment, which is not the case with tabletops (but see [25] for an exception). In particular, because users move in front of wall displays, there are not always clear territories [7, 12], the relative position and distance between the users (i.e., proxemics [27]) can impact collaboration style [55, 79], and such environments need to support fluid transitions between loose and tight collaboration [35, 49]. We review next in detail these findings from collaborative studies conducted either with wall displays or in AR and VR. These studies consist of manipulating "data" (images, documents, virtual post-it notes, etc.) in classification tasks, puzzle tasks, sensemaking tasks, and storytelling tasks.

Most studies have been conducted on a single wall display. Azad et al. [7] performed an observational field study of the behavior of groups on and around public wall displays. They combined it with a lab study over a puzzle task to investigate concurrent behavior between individuals and groups. They found that wall displays, like tabletops, should support public, personal, and storage territories, but that the location of personal space and proximity zone

(buffer zone between others) must be refined depending on the user's position. The results differ in Jakobsen and Hornbæk [35]'s study of how pairs of users collaborate, navigate and interact with a multitouch wall display during a problem-solving task. Their study suggests that "multitouch wall displays can support different collaboration styles and fluid transitions in group work". As a consequence, participants did not divide the wall into territories. Their study also suggests a correlation between the distance between the partners and collaborative coupling (smaller distance implying tight collaboration). Wallace et al. [79] found a similar result using collaborative puzzle tasks, and moreover suggest that the user's range in front of the display can characterize degrees of collaboration. Finally, Sigitov et al. [73] studied collaboration coupling and territoriality when pairs used a curved wall display. They suggest more types of territories than in previous works, on and in front of the wall display. They also observe participants dividing the task spatially among themselves, working in parallel. The above work suggests that interpersonal distance (proxemics) can indicate degrees of collaboration coupling, and that territories likely exist but are fluid in nature and location.

Nevertheless, there is likely a complex interplay between collaboration strategy, interpersonal distance (proxemics), task, and interaction. Liu et al. [48] considered a classification task on a wall display, where different strategies were enforced on pairs of participants. They found that the strategies, from tightly-coupled to parallel work, influence the space usage in front of the wall and the relative position of the partners. They also found that with appropriate interaction techniques, partners can collaborate closely, even at a distance. A similar result is obtained in a storytelling task [49], where participants can use cooperative gestures. Thus, the relative interpersonal distances predicted by proxemics [27] and tight collaboration coupling [35] may not apply when shared interaction techniques are provided to the users. Mayer et al. [55] also observed that in a cooperative condition, participants worked mainly side-by-side but that in a competitive condition, they crossed and physically blocked each other. This indicates that the distance between participants was smaller in the competitive than in the cooperative condition.

In VR, Lee et al. [43] studied how groups of 3 participants solve visual analytic problems in a fully VR simulated room. They particularly examined the role and use of surfaces in this environment, with a first task restricting the system to 2D visualizations and virtual walls acting as wall displays to pin the visualizations as support. Then in a second task, they introduced 3D visualization and a virtual table at the center of the room. They found that territories were defined by initial individual workspace placement around the room, were never negotiated, and that "participants never entered a territory of another unless for tightly-coupled work". Finally, in the context of Augmented Reality in a room that is either empty or contains furniture, Luo et al. [51] studied a collaborative task involving document layout for sensemaking. They found that users place virtual items around a room by grouping them on the physical walls or the surrounding furniture.

The above works suggest that the situation is complex and that the space usage on the (virtual) displays and in the room depends on the collaborative strategies, the task, and the setup (e.g., tabletop vs. wall). Our setup is unique, we go beyond physical walls [7, 35, 48, 73]



Figure 2: Different virtual surfaces. On the left, an image of a typical virtual Surface with CARDS. On the middle and right, views of the *Personal Space*: first, a view from the user's headset, with CARDS organized in a belt configuration that follows the user; and next, a rendering that highlights the relative position and size of the belt with respect to the user.

and consider situations where AR is added not only in physical rooms [43, 51], but in rooms where a large physical display is also present. Given the rich results seen in previous setups, we expect our configuration will provide additional insights to the use of space on and around wall displays. Nevertheless, we rely on this past work to motivate the tasks used in our study, a classification task that, a priori, can be solved using a parallel strategy (loose collaboration) and a storytelling task where participants have to collaborate closely.

3 PROTOTYPE

In this section, we present our prototype that combines a wall display with several synchronized AR Headsets. Our goal was to build a system that allows users to layout and organize different types of information (e.g., images, graphs, texts, maps) in order to make sense of, classify, order, or compare them. This type of activity is common in various contexts, such as organizing physical papers on a desk [54], arranging icons on a desktop, or moving post-it notes around on a whiteboard during a brainstorming session [15]. In particular, wall displays have been used in several such tasks: for scheduling the CHI 2013 conference [37, 50], navigating photos in a public city hall [62], and for various sensemaking tasks [5, 35] such as identifying anomalies in a set of documents [24].

For this purpose, our prototype includes several interaction techniques for users to organize, manipulate and move content between the wall display and the space around it. It also contains functionality to record and playback interactions to help us with our experimental analysis.

Our prototype renders content inside a 7 by 4.5 meters room. On one of the larger sides of the room is our physical wall display of 5.91×1.96 meters, with a resolution of $14,400 \times 4,800$ pixels (60 ppi), composed of 75 LCD displays (with 3 mm bezels) and driven by 10 workstations. The AR is rendered through HoloLenses (version 1). For the software, we used Unity 3D with identical scenes between the HoloLenses. A "master" program controls the HoloLenses and the wall, and we used the Unity UNet Multiplayer and Networking framework to synchronize the content between all the devices and to send input commands from the HoloLenses' to the rest of the system (other HoloLenses and wall display via the master). We calibrated the wall in the HoloLenses scenes using a Vuforia marker rendered on the wall. The marker is recognized by the HoloLens and used to calculate the position of the wearer in the scene relatively to the wall. The source code is available at <https://gitlab.inria.fr/ilda/arviz> and could be adapted to other setups.

We followed an iterative design process to develop our system, testing techniques among the authors and two other users before reaching the final prototype. We explain next the details and motivation behind the design of the displayed content, our interaction techniques, and replay functionality.

3.1 Virtual Elements

Our prototype contains three main types of virtual objects: BOUNDARIES, SURFACES, and CARDS (seen together in Figure 1). As we are studying collaboration, all objects are visible by every user with an AR headset, with the exception of *Personal Space* discussed later.

CARDS. In real-world usage, we expect the wall display and the extended virtual space to be able to render documents, images, visualizations, or more complex objects such as application windows. Motivated by previous work investigating space use on tabletops [71], physical walls [7], and furniture using AR [51], our prototype displays basic content in the form of CARDS. Our prototype can display CARDS of any shape and size, and their content can include images or text. However, for experimental purposes (see Sect. 4), we kept their size fixed and deactivated the possibility to add, remove, or resize CARDS. We initially allowed CARDS to be placed anywhere in the space within the AR environment. Nevertheless, early tests showed us that depth placement is not easy and makes CARD organization and layout challenging. We thus decided to restrict their layout in the virtual space on planes that we call SURFACES.

SURFACES. These are virtual workspaces where users can place and organize CARDS (see Figure 2-left). This allows users to group CARDS and perform operations on them (detailed in Sect. 3.2). We chose to render these virtual surfaces in a way that resembles the physical wall display to convey the impression that these surfaces can act as extensions to the wall display. Thus their height matches that of our physical wall (1.96 m). By default, their width is 2 m, smaller than the wall, to allow two of them to be placed side-by-side along the shorter side of our room. Nevertheless, their width can be increased if they contain many items.

Users can create as many SURFACES as they want and reposition them (and their content) in the environment. This choice is motivated by past work on wall displays that showed that colleagues tend to move around the space, and thus their interpersonal distance and location of their workspace territories may change [7]. Empty surfaces can also be deleted. At any time, users can rearrange the content inside a SURFACE using a re-layout function that cleanly organizes content in a grid and resizes the surface appropriately to contain all content. We first allowed users to position SURFACES

freely in the 3D space, but in our tests, we noticed it was difficult for users to position them accurately and lay them out in space. Due to their size, this led to a lot of clutter in the virtual space and some occlusion of other elements of the scene. Thus for SURFACES, similar to CARDS, we decided to constrain their position and movement on magnetic planes around the room, which we call BOUNDARIES.

We introduced one type of surface that is special, the *Personal Space*. By default, surfaces are visible to all users, but the *Personal Space* is a virtual surface only visible to the user who owns it. This *Personal Space* is supposed to be private and is thus placed as close to the user as possible, i.e., within their personal zone as defined by proxemics theory [27]. Inspired by previous work, e.g., [11, 22], this virtual zone resembles a belt made of the items stored inside (see Figure 2-middle and right). Elements in the *Personal Space* are placed in a circle around the user like a semi-cockpit [22], always facing the user and moving with them. This personal workspace allows users to bring CARDS around them for closer inspection. But, it also acts as a storage space, easily accessible, where the user knows they can quickly access stored documents [71] and move them around the space.

BOUNDARY. These are magnetic guides for constraining the placement of SURFACES around the physical space. We set up the BOUNDARIES as a rectangular area of 7 by 4.5 meters to match the size of our wall room (excluding the wall side). We initially considered BOUNDARIES on the floor and ceiling. However, tests showed that due to the headset's weight, it was tiring for users to tilt their heads for long periods to interact with content on the ceiling. We also do not allow placing surfaces on the floor as the *Personal Space* occludes it.

Objects in our prototype have a hierarchy: BOUNDARIES are static and defined before the start of the application, and SURFACES can be moved and must be placed on BOUNDARIES. CARDS can move and are placed on SURFACES. Every object is visible to every user by default, with the exception of *Personal Spaces* and the content contained within them.

3.2 Interaction

Our system was designed to create a visual and interaction continuum between the wall display and the augmented environment. We thus introduced a set of techniques to allow fluid content movements and content organization between the wall display and the virtual space represented by SURFACES. This section describes how input functions in our prototype, as well as our techniques for selecting and organizing content within the continuum between virtual space and physical wall display.

3.2.1 Input. Interaction in our prototype is carried out through the AR headset. We use a combination of head-cursor and clicker provided by default by the Hololense headset. Users can "point" at an item of interest by looking at them and use the clicker to select or manipulate it. Even though the Hololens hand gesture recognition is supported by our prototype, in our experiment, we chose to use a clicker because using hand-gestures in front of the head is tiring [13, 30]. To improve awareness of others' actions, we represent all user cursors as colored telepointers [29] in the shape of a cross. A unique color is assigned to each user's cursor and their selections (see Sect. 3.2.2).





Figure 3: Visual feedback on the CARDS and menus. On the left: a corner highlight added in AR when the user gazes at the CARD (not seen by others). In the middle, the CARD selection states, visible by all users: the CARD in its unselected state; the CARD selected by a user who is assigned the green color; and by the user who is assigned yellow. The last image shows our contextual menu, only visible to the user that invokes it. In this case, the contextual menu has been invoked on a CARD that exists on a virtual surface, so we see the available options to move all content of the surface towards another surface (top-right) or the *Personal Space* (bottom-right), or to move the CARD to the *Personal Space* (bottom-left).

3.2.2 CARD Selection & Movement. To support content organization, users can select and move content, which in our case is represented by CARDS (see Figure 3 for visual feedback provided by the prototype). A click on a CARD selects it, and a second de-selects it. A click on an empty space de-selects all selected CARDS. Selected CARDS are highlighted in the color of the user who selected them. To avoid continuous clicking, CARDS can also be selected using crossing selection [4, 6]: while holding the clicker button, every CARD that is crossed by the user's cursor will be selected and become part of the selected group. A CARD or a selection group can be dragged along with the head cursor, until the clicker button is released. If a drag is not released on a SURFACE (including the *Personal Space*) or on the wall, the selected content returns to its initial position. We rely on social protocol to deal with interaction conflicts [57], enforcing a simple coordination mechanism: if multiple users select the same CARD, the last person to select it has ownership.

3.2.3 SURFACE Creation & Movement. A user can create a virtual SURFACE by clicking on a virtual button that always follows each user, placed high up so as not to interfere with other virtual content. Once the button is clicked, a new SURFACE is created, following the user's cursor, until the user releases it. A SURFACE can be moved by dragging the bar at the top, similar to how application windows are moved on a desktop. The movement of SURFACES is constrained by the BOUNDARIES defined around the room, and when they are released, they snap to the closest boundary.

3.2.4 Advanced Content Management. Apart from single or multiple CARD selection and movement, we also provide advanced content management options to help users reorganize their virtual space more efficiently. They can be accessed with a long clicker press that brings up a contextual pie menu (Figure 3). We describe these options next.

Move the content. When the menu is invoked on a SURFACE (including the wall) or a group of CARDS, it allows, respectively, to start moving all the cards of the surface , even if they are not selected, or just the grouped selection , towards another surface. This allows users to quickly rearrange the content of shared virtual surfaces (and the content of the wall display).









Figure 4: View for the Replay tool. In the middle, we see a schematic of the entire scene, where the viewer can choose to adapt their view of the scene with a Camera placed in the scene (see red 1, added for annotation purposes). On the top right (Inspector), the viewer can change the parameters for the playback, pause/play the scene, and choose the type of progress bar to use. On the bottom right, the camera view, with a progress bar of the processed messages. The camera view can be changed through the Gameview target display (see red 2, added for annotation purposes). Here we see a third-person view from the camera positioned in the schematic, but we can adopt a first-person view that follows one of the cameras attached to the participants.

Move a CARD or groups of CARDS to the Personal Space. Depending on whether the menu is invoked on an unselected CARD or on CARDS that form a grouped selection, the menu provides the option to move the single CARD , or the entire group of selected CARDS , towards the *Personal Space*. If the menu is invoked on a surface, there is also the option to move the content of the entire surface onto the *Personal Space* . Once on, their *Personal Space* CARDS are only seen by the user and follow them around the space.

Expand the Personal Space. When the menu is invoked on the *Personal Space*, it activates the inverse operation. Users can drag individual CARDS out of the personal space or choose to extract all CARDS  through a menu option. These CARDS get attached as a group to their cursor and can then be placed on any virtual surface or the wall. Finally, the *Personal Space* can be expanded  to a new shared SURFACE that contains all the original content.

3.3 Replay Sessions

To study how users move and use space, we needed to keep a record of their interactions. A camera can record the physical room and wall display, but not the virtual space. We thus include a tool to record and replay user sessions, as seen in Figure 4. This tool, a script for Unity 3D, can record and replay the log of messages sent between the wall display and the AR HoloLens. By opening the dedicated scene and selecting the generated log file, the entire session will be replayed, showing changes in user position and cursor movements, as well as any surface creations, selections, and card movements. Finally, the replay tool can show either a bird's eye view or take the point of view of one of the users inside the replayed session following the camera attached to them.

4 USER STUDY

Our study aims to understand if and when it is helpful to extend a wall display with virtual spaces in the context of collaboration. And what is the impact, and potential cost, of this extension on the use of space and collaborative work. For example, we assume that the added virtual space in the form of surfaces will be appreciated when the wall display real-estate is too cluttered. However, it is unclear how this additional space will be used, under which tasks these virtual surfaces are needed, how many are helpful, and if some surface configurations are preferable. Furthermore, the additional virtual space may come at a cost in terms of interaction, as content needs to be moved across larger AR distances; or in terms of collaboration quality, if participants find it harder to coordinate across multiple virtual surfaces.

4.1 System Conditions

To answer these questions, we built a system that increases the display and interactive space available for users in front of the wall through the use of AR headsets (see Sect. 3). We study pairs of users working either on the wall display alone, our baseline condition (condition WALL); or on a setup using our prototype that combines the wall with AR surfaces (condition WALL+AR). An image of our setup can be seen in Figure 1.

We used the same basic input functionality, relying on the HoloLens head-cursor and clicker, for both conditions WALL and WALL+AR (see Sect. 3.2). In the WALL condition, the techniques related to AR surfaces are obviously disabled. We made this choice of consistent input to ensure we observe effects related to virtual workspace use and collaboration without introducing a bias that may stem from different input capabilities or discomfort in wearing the headset in some conditions only. This choice is to ensure experimental consistency, but it is also justified by research trends. First, in our context, one can imagine hand ray-casting as an interaction alternative, nonetheless using the head-cursor and the clicker is less tiring and is an efficient technique when no high precision pointing is needed [40, 58], which is the case in our tasks. Second, when it comes to wearing a headset to interact with a physical display, AR HMDs are becoming increasingly lighter and similar to vision glasses worn every day, see, e.g., [18, 34, 45].

4.2 Tasks

In each condition, we asked pair of participants to conduct two different tasks inspired by previous work (see Sect. 2.2): A classification task that could be performed using loose collaboration and a storytelling task (story task for short) that enforces tightly-coupled collaboration.

A Dixit image set was also used by [49] in their study on wall displays alone. Each Dixit card contains many colors and usually presents an abstract scene. We requested they group all cards in one of three groups of color (red, blue, and green). The CLASSIFICATION task operationalizes and simulates a collaborative situation where pairs need to coordinate and make decisions (as the cards contain many colors), while remaining a simple task that does not require domain knowledge. Moreover, it simulates loose collaboration tasks, as the work can be parallelized, since each participant could be working on a single group category.



Figure 5: Wall display at the start of the experiment, showing one of the image datasets.

In the classification task, participants have to make grouping decisions and select, move and sort picture cards, a task similar to past studies investigating workspace use and movement patterns in front of a display [49, 51, 71]. We requested pairs to group 54 image cards from the popular game Dixit¹, which has been used in the past in wall display studies [49]. Each Dixit card contains many colors and usually presents an abstract scene. We requested that they group all cards in three groups of colors (red, blue, and green). The classification task operationalizes and simulates a collaborative situation where pairs need to coordinate and make decisions because the cards contain many colors. However, it remains a simple task simulating loose collaboration, as the work can be parallelized, and each participant could work on a single group category.

Bradel et al. [12] discuss how users can engage in two different kinds of collaboration: independent workspace collaboration with large personal working spaces (territories), also referred to as loose collaboration; and integrated workspace collaboration with large shared territories, often referred to as close or tight collaboration. The classification task described above falls under the loose collaboration category, as it is highly parallelized, and participants could, if desired, divide the work and the workspace between them. To try and stimulate both types of collaboration, we thus introduced a second task, a story task. Here we ask participants to start from the collection of 54 cards, and create and tell a story using only 10 of them. In this more open-ended task, participants are required to make decisions about images and build a story together. This task encourages close collaboration, discussions between partners, and the use of shared workspaces.

Images and Layout. We selected 2 datasets of 54 cards from the Dixit card game (one per System condition). They are (i) colorful cards that prevent a straightforward classification based on color (in the classification task); and (ii) have abstract picture content that can promote discussion within the pair (in the story task). For both tasks, participants are presented with the wall display covered by cards (Figure 5). The number of cards (54) was chosen so that all cards were fully visible, but so that the wall was purposely crowded. This was to simulate situations where the available wall real-estate is at its limit. In other words, situations where we expect the use of the virtual space may be of interest. This will provide insights into if and how participants choose to use the virtual space (when it is available) and allow observing their strategies when dealing with a crowded display when only the wall display is available.

¹libellud.com/nos-jeux/dixit/

4.3 Hypotheses and Measures

As we set out to understand the impact of extending a wall display environment with AR and studying workspace use, our study is largely observational [49, 51, 71]. We, nevertheless, form some high-level research questions and our hypothesis about them. We next explain the measures we used to answer these questions.

RQ1 *Is the extension of a wall display environment useful? When is it used?* We hypothesize that participants will naturally move content in the AR surfaces as the wall display real-estate is cluttered. We hypothesize that the AR space and surfaces will serve secondary purposes (e.g., storage areas) and that the wall will remain the central working surface for two reasons: (i) the content is on the wall when the tasks start, so it is natural to continue working on it; and (ii) because the wall is such a central landmark in the physical room.

RQ2 *How is the AR space used?* We aim to observe more specific uses of the virtual space, for example, where surfaces are placed, how many, if they are moved, if the personal space is useful, etc., and identify differences in workspace use between the wall alone and the extended virtual environment.

RQ3 *Does the addition of AR affect collaboration strategy?* We hypothesize that the working strategies and practices will remain largely unchanged across the setups, given that the tasks are fairly simple in nature.

RQ4 *What is the cost of adding AR?* Extending the working area virtually around the wall display creates a bigger interaction area. We hypothesize that this will require more and longer interaction sequences, thus slowing down the pairs' performance and may fatigue participants.

We collect a variety of subjective and objective measures to access and compare the two setups (WALL and WALL+AR): Observed pair *strategy* in solving the task; *in-pair distance* between participants as a measure of tight/loose collaboration; Measured *virtual surface use* in terms of frequency and placement; Number of *interactions*, and interaction Distance traveled (e.g., card movement), as a measure of interaction cost; total *Distance* traveled by participants, as a measure that could possibly indicate fatigue but also engagement; *Time* to complete the task, as a measure of cost. Finally, we elicited *Subjective feedback* in the form of a Likert scale questionnaire, using (i) the four NASA-TLX questions on efficiency, ease of use, and mental and physical demand; (ii) two questions on partner awareness and communication from Harms and Biocca [28]; and (iii) six custom questions on space usage and collaboration relevant to our research questions.

4.4 Participants & Apparatus

Participants. We recruited 24 participants in 12 pairs: 10 women, 13 men, and 1 unspecified. Participants were aged 21 to 46 (average 25.8, median 24), with normal or corrected-to-normal vision. Sixteen participants had experience using an AR device, such as the HoloLens. Participants were HCI researchers, engineers, or graduate students in Computer Science. All pairs of participants were recruited together (volunteered in pairs), and were familiar with each other, being friends, colleagues or students in the same class.

Apparatus. We used the prototype described in the previous section with three HoloLenses, one per participant, and one for the experimenter.

4.5 Experiment Design & Procedure

Design. The experiment is a within-participants design with one factor, the system condition, with two values: WALL and WALL+AR. Their presentation order was counterbalanced across pairs. We fixed the task order as our primary goal was not to compare the tasks but the system conditions: pairs of participants always start with the classification task and then run the story task. We always start with the classification task because (i) it is simpler than the story task; and (ii) as our second story task requires users to study the content of the pictures, we could use the same datasets between the two tasks, allowing participants to become familiar with them from the start. Nevertheless, we ensured that, for each pair, the datasets were different across conditions (we counterbalanced the system condition and our two Dixit datasets across pairs).

Procedure. Participants work in pairs in two sessions (on different days), one session per system condition. When participants arrive for the first session, they sign a consent form and a demographics questionnaire. At the end of each session (system condition) participants fill out a questionnaire, and, at the end of the second session, they fill-in a global preference questionnaire. For the full duration of the experiment, the operator wears a headset too, and informs the participants prior to the study that the operator can also see the full Augmented Reality scene. This is to help in the training phase and ensure participants understand they can use any surface (virtual or physical) to display their work to the operator.

At the beginning of each system session, the pairs trained until both participants were comfortable using the system (this lasted 10 to 15 minutes). The operator explained the different interaction techniques and instructed each participant to try all the interaction techniques at least twice.

The system is restarted after the training and after each task, all virtual surfaces are removed and all cards are placed back to their original position. Each system session lasted about 1 hour, including the training and answering the questionnaires.

5 RESULTS

We first discuss the general collaboration and workspace use strategies adopted by pairs, as well as surface placement in the physical space. Then we report on the usage of different techniques, quantitative measures (e.g., traveled distance, distance between partners, time), and finally, the questionnaires. All statistical analyses reported are paired t-tests unless otherwise specified. Due to a technical issue with data for pair G8 in the WALL+AR condition, some analyses regarding WALL+AR use and comparisons between WALL and WALL+AR do not take into account G8.

5.1 Collaboration Strategies & Workspace Use

To analyze the collaborative strategies, we used a thematic analysis on the recorded sessions using the replay tool, notes taken by the operator, and the supplementary material (Section 1 of the PDF) showing the virtual screenshots of the final results for each task

and pair. One author coded the strategies used by each pair: collaboration coupling over time and steps in the task, placement and role of the virtual surfaces (if any), use of the workspace (wall and virtual surfaces), cards placement, displacement and layout, formation of territories and their use, etc. A second author double-checked the coding using the replay tool and independently summarized the coding as described below. Some codes were pre-determined (deductive) based on our hypotheses and related work, such as collaboration coupling (close vs. loose), type of territories, and surfaces used to show the story; and others came from the data (inductive), such as the five emerging strategies of 5.1.1. and how the cards were chosen in the story task (together vs. independently).

5.1.1 Wall - Classification. We observed four main strategies. In *parallel*, the pair worked with all the colors simultaneously (5 pairs: G0, G3, G5, G6, G11). In *divide*, the pair assigned a color to each of them and then handled the remaining color together (3 pairs: G1, G2, G4). In *mix*, pairs adopted a mix of the two previous strategies, they started out by assigning one color to each of them, and then after a few cards had been placed, the pair handled all the colors together (3 pairs: G8, G9, G10). And in *sequential*, the pair worked together on each color, proceeding color by color (one pair G7).

Interestingly, the pairs that used the parallel strategy placed the images into lines, one for each color. The other pairs used arbitrary-shaped blocks to organize images (see Figure 6-top). Indeed, the parallel pairs decided where to place each color before even starting to move the cards, and it seems that using the same linear organization as the original placement of the cards was a natural decision. In fact, these pairs started the task by swapping cards between different lines, and we even observed two pairs (G0, G11) exchanging cards synchronously between the lines. Another parallel pair (G5) fully divided the work by splitting the wall in half, and each partner then sorted the colors into lines on their side of the wall.

The non-parallel pairs split the wall into blocks, as each partner decided to group their color in front of them. We noted that one of these groups (G8), used a particular strategy, they built a heap of cards at the center of the wall to create space on the sides and then started building the color groups.

5.1.2 WALL+AR - Classification. All pairs created surfaces to put the cards of a given color. The pairs either used three surfaces (6 pairs: G2, G4, G6, G7, G8, G10), one for each color, or two surfaces (6 pairs: G0, G1, G3, G5, G9, G11) the wall being used for the remaining colors. See Figure 6-bottom.

All pairs (except G10, see below) started the task by creating two surfaces and putting them on each side of the wall. Then, each partner used his/her surface (the one closest to them) to classify a color. The pairs that created three surfaces either created this third surface at the beginning of the task (G2, G4, G6, G8), or later when the first two colors were classified (G7). In both cases, this third surface was used to classify the remaining color and was placed at the back of the room (opposite of the wall). For all those pairs, both partners used this third surface to handle the last color.

The partners of the G10 pair started by using their personal space: after choosing a color each, they put the cards of the corresponding color in their personal space. Then, they transformed their personal space into surfaces on each side of the wall and adjusted the content of these surfaces.



Figure 6: Examples of final results for the classification task. (top) WALL condition, using a strategy of placing cards in lines (G0) and in blocks (G7). (bottom) WALL+AR condition, with a strategy that uses both the wall and virtual spaces for classification (G0), and a pure virtual space classification strategy that does not use the wall at all (G4).



Figure 7: Examples of a final result in the story task for the WALL condition (G1, left), and the WALL+AR condition (G3, right)

5.1.3 Summary Classification. When comparing the results with the WALL condition, we can observe that adding AR affects the collaborative strategies. In AR, all pairs divided the work (at least at the beginning of the task), each using one surface close to their location, showing loose collaboration coupling. In contrast, only half of the pairs in the WALL condition divided the work, while others worked in a tightly coupled manner. Pairs created virtual surfaces to overcome the lack of free space on the wall, and even half of the pairs removed everything from the wall to organize the three colors in 3 virtual surfaces (typically on the left, right and facing the wall).

5.1.4 WALL - Story. As expected, in this task, all pairs worked in close collaboration. Ten pairs chose the cards together for their story. However, two pairs, G9 and G10, decided that they would choose five cards each independently. Nevertheless, all pairs build the story together (this was enforced by the task).

The chosen cards were then moved to a specific part of the wall (10 pairs, G1, G2, G3, G4, G5, G6, G8, G9, G10, G11) or selected using the selection feature of the prototype (2 pairs, G0 and G7). All pairs (except G5) created free space on the wall to make room to create and present the story (recall that the task starts with the wall covered by cards). Some pairs (G4, G9, G11) even started the task by freeing space before choosing their cards. The other seven pairs either made free space just after they chose the cards to build the story, or created the story and freed the space more or less at the same time. G5 created the story over the initial cards' layout (i.e., overlapped other cards) without taking care to make room to present their resulting story. Figure 7-left shows an example of a resulting story.

5.1.5 WALL+AR - Story. All pairs worked in close collaboration (same as in the WALL condition). Most pairs chose the cards for the story together (11 pairs), except G1, G4, and G9. For these three pairs, the partners selected the cards they preferred each to create a pool of cards for the story (in a "pool" surface, see below).

All pairs created one or more surfaces to create their story. Six pairs (G0, G3, G5, G6, G7, G11) created one surface and picked cards from the wall to place them on the surface to create the story and present it to the operator (see Figure 7-right). One pair, G10, used the same strategy, but after the story was completed, the partners moved all the (unused) cards that had remained on the wall towards a newly created "trashcan" surface and then moved their story onto the wall to present it. Note that all these pairs placed the surfaces on the left or right of the wall.

G2 used a somehow different strategy than the other pairs. After using a surface to select the cards for the story, they moved the cards remaining on the wall to a "trashcan" surface (using their personal space) and moved back the story-selected cards to the wall to create and present their story.

The three pairs that selected the cards independently created a first surface to place the cards they selected as a pool of cards. Then, two pairs (G1, G9) created a second surface on the side of the first one to create the story with the elements from the first surface (G1 discarded it when empty). G4 used a similar strategy but created three other surfaces to be able to present the story in a line.

5.1.6 Summary Story. Compared to the WALL alone, we observe that adding AR did not affect the collaborative strategies: all pairs worked closely together, in a tightly coupled manner, and adopted similar strategies to select and work on images. For example, in both

cases, the most common strategy is to select images together, and only 3 pairs selected candidate images individually. However, AR did affect workspace use. In the WALL condition, pairs had to adopt strategies to make space, moving items to the side, sometimes even before they started considering the story. Whereas in the WALL+AR condition, all pairs immediately created at least one virtual surface to create their story on the left or right of the wall. Surprisingly, in almost all cases, the AR surface(s) were used as both the working area and final presentation area of the story. We observed only a few instances where virtual surfaces were used only as storage of unused cards.

5.2 Interactions and AR Technique Use

From our interaction logs, we analyzed all elementary actions (move a card, select a card, move a selection) for both conditions, and for WALL+AR, we also counted surface and personal space related actions that the partners of each pair performed. We use these counts to analyze different aspects.

5.2.1 Number of Interactions. For WALL, we recorded an average of 52.5 ± 7.6 elemental actions for the classification task and an average of 60.9 ± 12.4 for the story task. For WALL+AR, we recorded an average of 52.37 ± 7.7 elemental actions for the classification task and 37.3 ± 11.0 for the story task. The number of actions is very similar between WALL and WALL+AR for the classification task. However, there is an important and significant difference between WALL and WALL+AR for the story task ($p = 0.003$, $d = 0.91$). This smaller number of actions in the WALL+AR can be explained by the fact that in the story task with WALL+AR, most pairs just interacted with the story's cards (10 cards or a little more). At the same time, with WALL, the pairs had to interact with the story's cards, but also many other cards to make room for laying out their story. This does not happen with the classification task because pairs had to move more or less all the cards in both conditions.

5.2.2 Interaction Types. Without surprise, the most used elemental actions were moving a card ($58.5\% \pm 7.2$ of the actions for WALL, $38.9\% \pm 7.0$ for WALL+AR), then selecting a card ($31.6\% \pm 5.4$ for WALL and $38.5\% \pm 6.6$ for WALL+AR) and moving a selection ($9.8\% \pm 2.2$ for WALL, $8.5\% \pm 2.2$ for WALL+AR). These three elemental actions represent, of course, 100% of the action for the WALL condition, and $85.9\% \pm 2.8$ of the actions for WALL+AR ($81.0\% \pm 3.5$ for the classification task and $90.1\% \pm 3.5$ for the story task). However, we can notice some disparities between pairs, and even between partners of a pair, in the usage of the above actions. Some participants mainly moved individual cards, while others tended to select cards and move these selections.

We now focus on the actions that specifically concern the WALL+AR condition.

As described in the previous section, all pairs created surfaces, with about 2 or 3 surfaces for the classification task and about 1 or 2 surfaces for the story task. Most operations consisted of moving cards from the wall to the surfaces or moving cards on the surfaces (and in a few cases moving cards from a surface to the wall). Pairs rarely moved surfaces after they positioned them at creation time, with an average of about one surface move by task. Surfaces deletion were used sparsely (9 surface deletions across

all pairs), and moving all the content of a surface was used only once. However, all pairs but one (G7) used the surface re-layout feature. In total, all the surface operations represent $14.9\% \pm 3.0$ of the actions for the classification and $8.9\% \pm 0.8$ for the story task.

The personal space was used by 8 pairs (9 participants for $3.3\% \pm 2.2$ of the actions for the classification task and 4 participants and $0.8\% \pm 0.9$ of the actions for the story task). Thus the personal space was used moderately, but some participants still found it helpful. The most interesting examples were described in the previous subsection, but it seems that the possibility to transform the content of the personal space into a surface was appreciated by some participants.

5.2.3 Summary of Interactions. Participants made, on average, the same number of actions in the WALL+AR and WALL conditions for the classification task, but surprisingly fewer actions in WALL+AR for the story task. All WALL actions, and the majority of the WALL+AR actions, involve card moves, either one-by-one or as a group. When considering WALL+AR, most actions were movements of cards from the wall towards one of the created virtual surfaces, followed by movements to rearrange content on the virtual surfaces, and a few actions to move content back to the wall. Virtual surfaces were generally placed in a position and rarely moved or deleted afterward, but their content was often reorganized. Only a few groups used the personal space to move content around.

5.3 Additional Objective Measures

We report next a set of objective measures: partners and cards traveling, position, and task time.

5.3.1 Participant Position and Distance Traveled. At the beginning of the tasks, the partners positioned themselves side-by-side in front of the wall close to the center, one slightly on the left and the other slightly on the right, at a distance of about 3 m from the wall. In the classification task, the pairs kept this position during all the task with minimal crossing and only little traveling, especially in the WALL condition. In the story task, they moved around more and occasionally inverted their relative position in front of the wall, especially in the WALL+AR condition. (Supplementary material Section 2 of the PDF).

Figure 8-(a) shows the traveled distance by pair by task and condition (we used the headset's position to compute this measure). Pairs traveled far more with WALL+AR than with WALL, and the differences are significant ($p = 0.007$, $d = 1.05$ for the classification task, and $p = 0.008$, $d = 0.68$ for the story task). This difference can be easily explained as the pairs interacted with a larger workspace with WALL+AR than with WALL.

5.3.2 Card / Interaction Distance Traveled. However, an interesting phenomenon occurs when we measure the total distance traveled by cards (the most common interaction). As expected, we can observe in Figure 8-(b), that the distance is far higher for WALL+AR than for WALL in the classification task ($p = 0.002$, $d = 1.18$), but for the story task, the difference is small, and not significant ($p = 0.850$, $d = 0.07$). These contrasting results can be explained by the number of interactions in the story task across conditions (discussed in the previous subsection). In the story task with WALL+AR, the pairs performed fewer actions, just interacting with the story cards, while

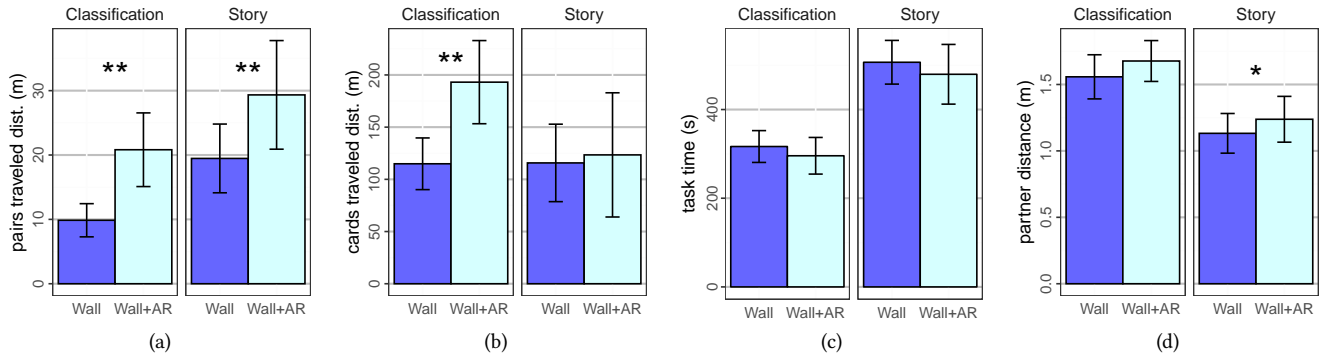


Figure 8: Average of (a) the traveled distance by both partners, (b) the traveled distance by cards, (c) the task time, and (d) the average distance between the partners (by condition and task). Error bars show the 95% CI.

with WALL the pairs had to interact with many more cards to make room for the story. Indeed, on average, with WALL+AR, the pairs performed less than 2/3 of the number of actions than with WALL for the story task, leading to smaller total distances. In contrast, this number of elementary actions was similar for both conditions in the classification task, which led to larger total traveled distances in WALL+AR.

5.3.3 Time. We hypothesized that interacting with a larger workspace that needs more traveling and additional operations, such as creating surfaces, has a cost on the task time, especially in a loose collaboration task such as the classification task. We were surprised to observe very similar task times (Figure 8-c) and no significant difference ($p = 0.516$, $d = 0.22$) between WALL and WALL+AR for the classification task. The difference for the story task is not significant either, but this task requires more analysis and reflection, which dominates the task time, so the lack of difference is less surprising. Overall, it seems that extending a wall display with AR does not necessarily impact performance.

5.3.4 Distance between Participants. As a measure for loose and close collaboration, we measured the average distance between the partners of a pair during the tasks, similarly to [35], for instance. Figure 8-d shows the results. We found no significant difference between WALL and WALL+AR for the classification task ($p = 0.320$, $d = 0.29$), and a significant difference with a small effect size for the story task ($p = 0.048$, $d = 0.32$, 13 cm difference). However, we found a significant difference with a large effect size when comparing the classification and the story tasks irrespective of condition ($p = 0.002$, $d = 1.11$, a difference of 42 cm). This suggests a correspondence between the distance between the partners and the proximity of the collaboration (proxemics [27]) expressed by our two tasks: personal distance and tight collaboration for story, and social distance and loose collaboration for classification.

5.3.5 Summary of Additional Objective Measures. Even though participants clearly moved more around the room in the WALL+AR condition, this did not affect their time as we found no evidence of a difference in time to complete the tasks between conditions. Due to the large virtual room, their total interaction distance (card

moving distance) was higher with WALL+AR in the classification task. However, this was not the case in the story task, where interaction distance was smaller in WALL+AR since the virtual surface allowed them to focus on the cards of interest (in WALL they had to constantly move cards around to make space). Finally, our findings suggest a correlation between the distance between partners and the degree of collaboration, in agreement with proxemics theory.

5.4 Subjective results: Questionnaires

At the end of each condition session, we asked the participants to rate on 7 points Likert scale: their mental demand; their physical demand; how successful they were in accomplishing the task; how hard the tasks were; how irritated they were when performing the task; how aware they were about what their partner did; the quality of the communication with their partner; and whether they had enough space to perform the task. Results are shown in Figure 9-(a).

Overall, participants were positive about both conditions. However, paired Wilcoxon signed-rank tests² show that the WALL+AR condition received better scores than WALL regarding success ($p = 0.002$, although the difference is small) and available space ($p < 0.001$). On the other hand, participants found WALL+AR more physically demanding than WALL ($p < 0.001$). This last result is consistent with participants' traveled distance in the tasks, which was clearly higher with WALL+AR than with WALL.

At the end of each condition session, we also asked questions related to space usage (*i.e.*, territory): did you use a specific area to present the results; did you use a specific area to store cards; did you use a specific area to discard cards; did you work on specific areas with your partner; did you use all the space available on the wall. Results are shown in Figure 9-(b). Results slightly suggest that some specific areas have been used for discarding and storing cards and co-working (having no clear results here is not surprising given the nature of the classification tasks). On the other hand, the results suggest that a specific area has been used for presenting the results of the story task (but the areas, indeed, differ among the pairs - as discussed in the strategy section). When comparing the two conditions, the only significant result concerns the wall space

²we comment on all the significant results but only them.

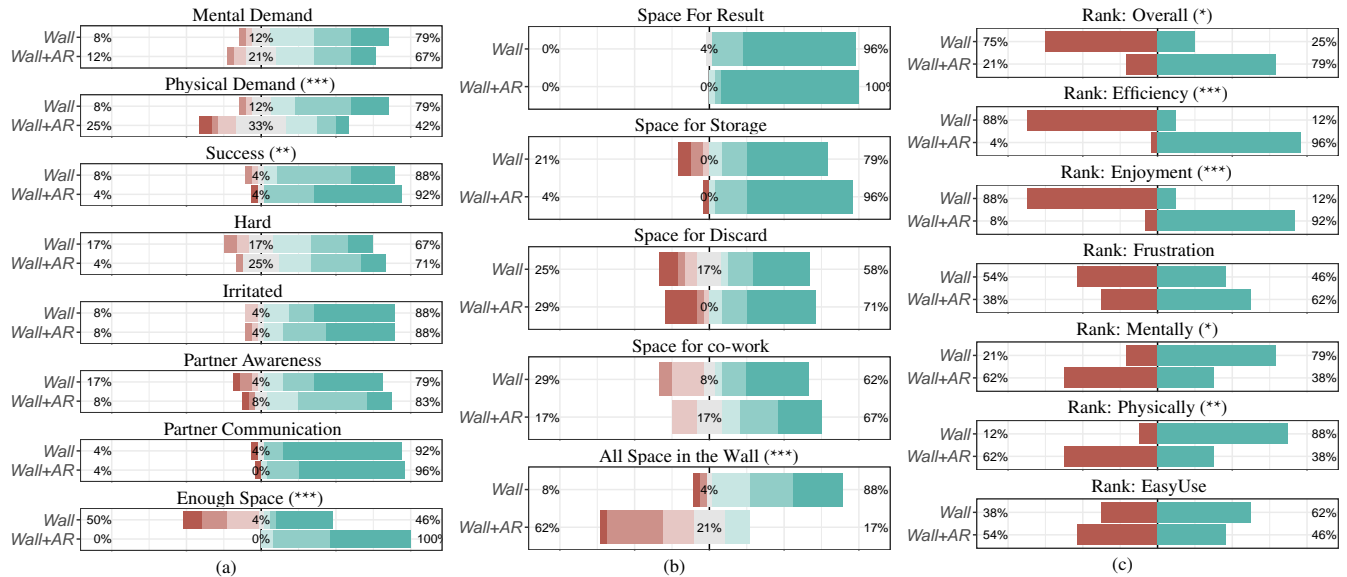


Figure 9: Results of (a) the standard questionnaire, (b) the space usage questionnaire, and (c) the ranking questionnaire. For easy reading, we put the "positive" answer on the right (in green).

usage ($p < 0.001$), where, as expected, pairs said they use all the space available on the wall with the WALL condition, but not for WALL+AR.

Finally, we asked participants to rank the two conditions (with possible ties) overall and relatively to: efficiency, enjoyment, frustration, mental and physical demand, and ease of use. Results are shown in Figure 9-(c). Overall, participants preferred WALL+AR ($p = 0.011$), and found WALL+AR more efficient ($p < 0.001$) and more enjoyable ($p < 0.001$). On the other hand, participants found WALL less physically ($p = 0.008$) and mentally demanding ($p = 0.041$). Results on physical demand align with our findings on movement around the wall that was higher for WALL+AR. However, they do not explain the result on mental demand. Here it is likely that the WALL+AR condition is more complex, *e.g.*, with many more interaction possibilities, and thus created more mental demand.

Summary of Subjective Results. Participants overall preferred the WALL+AR condition, and found it more enjoyable and efficient. They also found it provided them with more appropriate amount of space for their tasks. Nevertheless, as expected, it is more physically and mentally demanding than the WALL only condition.

6 DISCUSSION AND LIMITATIONS

We next revisit how our results answer our original research questions on combining physical wall environments with augmented reality. We highlight limitations of our work and discuss remaining open questions and future directions.

RQ1. Is the extension of a wall display environment useful? When is it used? We observed that, indeed, the additional workspace provided by the WALL+AR interface is beneficial when the wall display is cluttered and at its limit regarding available space. The subjective responses from participants confirmed this. They reported that

the wall display was not enough for their task, and they overall preferred the extended AR environment. Moreover, we measured that the additional virtual space, in some cases, can even reduce the number of elements users have to manipulate.

We had hypothesized that virtual surfaces would be used mainly for secondary purposes, such as storage, and that the wall would serve as the primary workspace surface. While we did observe virtual surfaces used as secondary storage (pool of images) and trashcans (discard piles and temporary storage to clear the main workspace), such usage was, in fact, marginal. In most cases, virtual surfaces took the central stage in the pair's work. For example, in the classification task, where participants created 3 groups, they were used systematically as the main grouping containers, probably because they have the advantage of explicitly separating the space. These containers often started as personal workspaces in the classification task. Sometimes pairs went as far as creating three virtual containers and leaving the wall empty. The third surface might not be optimal in terms of interactions, nevertheless, we believe that this strategy allowed the pairs to (i) make an explicit choice for every single card and validate their grouping; and (ii) visually organize and present all color groups consistently. This third space was created by one of the partners, but quickly transitioned to a group space and shared equally.

In the story task, all groups immediately moved the main cards they wanted to use off the wall and onto a surface and kept working there as a group. This indicates that for our participants, virtual surfaces acted as flexible containers that could be created on-the-fly, and easily took the role of the main *working area*. It would be interesting to investigate if these behaviors persists when the wall display is less crowded. We suspect that due to the grouping flexibility of virtual surfaces the findings related to content organization may also transfer to situations where the wall is not as crowded.

This also raises a question for future investigation: are the observed effects due to the nature of AR, or could a fully instrumented space (e.g., a room surrounded by wall displays) lead to similar findings? Even if we discard the cost of building and maintaining such rooms, we do feel some of our findings are unique to AR surfaces. Participants treated surfaces as containers to divide items and easily move them around. This easy division and movement cannot be accommodated by fixed physical walls. Past work where all content existed only in AR [43] also identified the movement of content in different locations around the room and the creation of distinct work areas (fluid territories). We thus confirmed it in a mixed environment that combines physical and AR displays. Other observations will likely hold in a purely physical setup only, e.g., in a room surrounded by wall displays. For example, participants tended to start working on a surface closest to their side, we suspect this will be the case if surfaces are replaced by physical walls.

Wall displays exist in multiple settings and remain today superior in resolution and field of view to AR HDMs [17, 33]. However, it is interesting to consider if our results could hold if all surfaces (including a perfect "wall") are rendered only in AR. We believe some of our findings would still hold. Such as the creation of transient territories that may start as personal but transition to group spaces, also seen in previous work in VR [43]. Or the need for more tight collaboration coupling strategies when a single (virtual) surface is available, as this is likely driven by the lack of space and the need for coordination. Other results may be influenced by the strong physical presence of the wall and may be unique to our setup. For example, past work on the placement of AR visualizations around a room [43] did not show any pattern on which side or area of the room to place information on. Whereas we saw strong patterns of putting surfaces first on the left and right of the wall, almost as direct extensions to it, influencing, in turn, participant movement and interaction distance. Thus the physical wall, even though it may not be the primary interaction surface, seems to "anchor" the placement of other surfaces.

RQ2. How is the AR space used? Participants placed most virtual surfaces directly on the left and right of the wall, enforcing the metaphor of an extended virtual space – surfaces on the back of the room were rarer. In the story creation task, participants largely interacted together on all surfaces indicating close collaboration coupling. Thus, in this case, virtual surfaces can be considered as group territories. In the classification task, users adopted loose collaboration coupling: they tended to create surfaces close to their location (e.g., participants starting on the left side of the wall created a surface on the left), and largely kept ownership of these surfaces. Thus, in this case, surfaces could be considered personal territories similar to previous work on VR content only [43]. Nevertheless, this was not always the case. For example, in several instances, they created a third surface that was then clearly shared by both participants. This is consistent with past work studying territoriality in wall displays alone [12]: the notion of territories is fluid and their nature is hard to predict as participants transitioned between periods where they created and worked on surfaces together and alone. We expect this is partly due to the flexible nature of surfaces that participants could easily appropriate for parallel or group work, and the fact that participants can move freely in the room.

Our classification task represents a common organizational task that requires frequent content movement and manipulation, but little analysis or reflection. It is thus possible our findings may differ under other collaborative tasks. More complex analysis tasks may require remaining stationary for longer periods, for example, to read documents or charts in a sensemaking task. This could create a feeling of ownership of specific physical locations around the room (and virtual surfaces placed there), leading to the creation of virtual territories of a more permanent or personal nature. How and where people form virtual territories likely also depends on the layout and whether it is divisible or not. For example, a planning activity around a single map requires both stationary interactions and movement to reach different map areas. In these instances, virtual territories may take new forms, such as personal copies of group territories that are transient in nature and easy to incorporate back into group territories.

RQ3. Does the addition of AR affect collaboration? In the story task, we did not observe differences in the collaboration strategies between WALL+AR and WALL, apart from the fact that all pairs used a virtual surface as their main working area. Nevertheless, in the classification task, we notice differences. The WALL+AR setup led to more uniform strategies between participants, mainly focused around creating discrete surfaces and working independently for parts of the task. While with the WALL, we observed various classification strategies, ranging from entirely parallel to tightly coordinated and sequential. In addition, we measured that the distance between participants tended to be larger in the WALL+AR condition, as they were able to interact with content around the room (beyond the wall display). This alignment between degree of collaboration strategy and distance between partners is in agreement with proxemics theory, as in previous work on wall displays alone [35, 79]. Our two results may indicate that the reduced available space in the WALL may encourage tighter collaboration and coordination, as space is at a premium and pairs need to carefully negotiate their actions and space use. A similar effect was observed in past work on wall displays during network analysis, where interactions that created clutter led to tighter collaboration and coordination [63].

RQ4. What is the cost of adding AR? Our results suggest the existence of a trade-off. On the one hand, the perceived efficiency of the WALL+AR interface and the importance of the space it provides; and on the other hand, the lower physical and mental demand of the WALL interface that could be interpreted as less fatiguing. Our participants also moved more in the WALL+AR condition, this could further indicate that the WALL+AR condition causes more fatigue: past literature, e.g. [9, 48, 55], has made the connection between movement distance and physical demand or fatigue. The combination of subjective comments on physical demand and distance traveled suggests this may be the case in our results. Nevertheless, it is also possible that participants were willing to walk more in the WALL+AR condition because they were more engaged in the task than with the WALL condition (discussed next).

Trade-offs were also seen in the case of measured interactions: fewer actions in WALL+AR when dealing with fewer objects, but smaller interaction distances with the WALL. In this last trade-off there is likely an interplay of setup and available workspace. In WALL the small interaction distances and more frequent interactions

in the story task, are likely a direct result of the limited available workspace. While the larger interaction distances in WALL+AR are likely due to a combination of factors: the large physical space taken up by the AR surfaces and the larger workspace they create.

We note that the WALL+AR setup was found more enjoyable, although we cannot exclude a novelty effect or an impact of our participants that come from a university and that are familiar with technology. Nevertheless, we feel our participants are representative of the target audiences of immersive technology for groupwork. More importantly, we found no measurable difference in time performance across the setups. Collectively, we deem that the cost of introducing AR to extend a wall display environment is not as high as we expected (interaction time, interaction cost), given the clear benefit in terms of available working space and user satisfaction.

Limitation and Future Work

For experimental purposes, participants used the AR headset to use the same basic input (default head-cursor and clicker) irrespective of condition. While this allows us to remove any bias related to technical differences across input modalities (*e.g.*, head pointing vs. ray-casting) and perceived fatigue due to wearing or not the headset, it is an artificial requirement. In real-world situations, we expect colleagues to use the wall display alone until additional surface space is needed, for example, to flexibly organize content (as we saw in our study) or just to make space. It is thus possible that such factors may affect performance and preference during real-world use. And it would be interesting to evaluate the cost of "putting on the headset", in other words studying when it is worth it for colleagues to decide to pass from a purely physical setup to one where headsets are required.

For this initial investigation, we created a setup that uses basic interaction techniques for content selection and movement. Nevertheless, more advanced interactions such as zoom, pan, and resize may affect the results reported here. For example, global view manipulations may prevent parallel work on virtual surfaces as colleagues may refrain from interacting without coordinating first; or may, on the contrary, encourage using more personal virtual surfaces to avoid disturbing their partners. This requires further investigation.

As most of the work on collaboration with wall displays, we only consider pairs of users in our investigation. Although we believe some observations might be generalized for collaborative work with more than two users (*e.g.*, one surface by user for a classification task), future work should study the case of groups of three or more users. Moreover, our work focuses on manipulating images, similar to the abstract family of tasks used in previous work on wall displays (*e.g.*, [35, 49]). It is difficult to generalize our results in terms of space use and collaborative strategy to interfaces that are (i) hard to "split" and reorganize, such as large visualization dashboards, maps, and more generally geolocated data; or that (ii) have a lot of visual details (that high-resolution walls can render) and may require stationary reading. It remains future work to consider an AR+Wall prototype for such contexts, for example, in the form of a focus + context display, given that AR headsets still cannot match the high resolution of wall displays.

7 CONCLUSION

Wall displays are extremely useful collaborative working environments that can be seen and used by multiple users and show a large amount of information. Nevertheless, they are hard to adapt or extend. It seems natural to use a readily available technology, AR headsets, to extend wall displays when their real-estate is no longer sufficient. However, the benefits and drawbacks of such an addition are not clear. To answer this question, we first introduce a set of techniques for extending the wall virtually in the form of additional surfaces and appropriate interactions to organize, manipulate and move content between the wall display and the AR virtual space. We next use this setup to study the differences in how pairs of users collaborate and use the available workspace in a wall display environment and in a wall display extended by AR headsets, with two collaborative tasks.

Our results highlight that such an extension is useful, and participants used the physical space in front and around the wall display extensively to place virtual content. Virtual surfaces were occasionally used as expected for storing and discarding data. More surprisingly, virtual surfaces were most often used as the primary interactive workspace, with participants abandoning the wall display. Adding AR to a wall display brings a real benefit over using the wall alone, and this extended setup was preferred, and found more enjoyable and efficient than the wall alone. But it does create interaction overhead, and increases physical and mental demand. We note, however, that we did not measure any loss in performance, despite this interaction overhead.

These findings provide empirically measured benefits of extending wall displays with AR, and insights into how they influence collaboration and space use. We discuss open questions that remain when it comes to applying such extensions in practice. However, our work demonstrates how such an extension is feasible and beneficial.

SUPPLEMENTAL MATERIAL

The supplemental material consists of a pdf file containing (i) virtual screenshots of the final results of the experiment for each task and pair; and (ii) heatmaps of participants and surfaces during the tasks of the experiment for each pair (but G8).

Additional material is available online at <https://ilda.gitlabpages.inria.fr/arviz/>. It includes the source code of the prototype, documentation, as well as a web application allowing to replay the sessions of the experiment.

ACKNOWLEDGMENTS

We wish to thank all the members of the Inria-LISN ILDA team for their feedback, and Olivier Gladin, Emmanuel Courtoux and Eugénie Brasier for their help.

This research was supported by French government funding managed by the National Research Agency (ANR) under: (i) the Investments for the Future program (PIA) grant ANR-21-ESRE-0030 (CONTINUUM); and (ii) the Collaborative Research Project (PRC) grant ANR-21-CE33-0022 (InterPlay).

REFERENCES

- [1] visited on Sep 2022. ADNOC's Panorama Digital Command Center. <https://www.adnoc.ae/en/News-and-Media/Press-Releases/2020/ADNOC-Panorama-Digital-Command-Center-Generates-Over-1-Billion-in-Value/>

- [2] visited on Sep 2022. Hiperwall in Global Satellite Services Provider command center. <https://www.sharpeddisplays.us/success-stories/intelsat/143>
- [3] visited on Sep 2022. IBPC institute Visualization Wall. <http://www.ibpc.fr/en/visualisation-wall-910.htm>
- [4] Johnny Accot and Shumin Zhai. 2002. More than Dotting the i's — Foundations for Crossing-Based Interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '02)*. ACM, 73–80. <https://doi.org/10.1145/503376.503390>
- [5] Christopher Andrews, Alex Endert, and Chris North. 2010. Space to Think: Large High-resolution Displays for Sensemaking. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. ACM, 55–64. <https://doi.org/10.1145/1753326.1753336>
- [6] Georg Apitz and François Guimbretière. 2005. CrossY: A Crossing-Based Drawing Application. *ACM Trans. Graph.* 24, 3 (jul 2005), 930. <https://doi.org/10.1145/1073204.1073286>
- [7] Alec Azad, Jaime Ruiz, Daniel Vogel, Mark Hancock, and Edward Lank. 2012. Territoriality and Behaviour on and around Large Vertical Publicly-Shared Displays. In *Proceedings of the Designing Interactive Systems Conference (DIS '12)*. ACM, 468–477. <https://doi.org/10.1145/2317956.2318025>
- [8] Sriram Karthik Badam, Fereshteh Amini, Niklas Elmquist, and Pourang Irani. 2016. Supporting visual exploration for multiple users in large display environments. In *Conference on Visual Analytics Science and Technology (VAST '16)*. IEEE, 1–10. <https://doi.org/10.1109/VAST.2016.7883506>
- [9] Robert Ball, Chris North, and Doug A. Bowman. 2007. Move to improve: promoting physical navigation to increase user performance with large displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*. ACM, 191–200. <https://doi.org/10.1145/1240624.1240656>
- [10] Michel Beaudouin-Lafon, Olivier Chapuis, James Eagan, Tony Gjerlufsen, Stéphane Huot, Clemens Klokmoose, Wendy Mackay, Mathieu Nancel, Emmanuel Pietriga, Clément Pillias, Romain Primet, and Julie Wagner. 2012. Multisurface Interaction in the WILD Room. *IEEE Computer* 45, 4 (April 2012), 48–56. <https://doi.org/10.1109/MC.2012.110>
- [11] M. Billinghurst, J. Bowskill, M. Jessop, and J. Morphet. 1998. A wearable spatial conferencing space. In *Second International Symposium on Wearable Computers (ISWC '98)*. IEEE, 76–83. <https://doi.org/10.1109/ISWC.1998.729532>
- [12] Lauren Bradel, Alex Endert, Kristen Koch, Christopher Andrews, and Chris North. 2013. Large high resolution displays for co-located collaborative sensemaking: Display usage and territoriality. *International Journal of Human-Computer Studies* 71, 11 (2013), 1078 – 1088. <https://doi.org/10.1016/j.ijhcs.2013.07.004>
- [13] Eugénie Brasier, Olivier Chapuis, Nicolas Ferey, Jeanne Vezien, and Caroline Appert. 2020. ARpads: Mid-air Indirect Input for Augmented Reality. In *Proceedings of the International Symposium on Mixed and Augmented Reality (ISMAR '20)*. IEEE, Porto de Galinhas, Brazil, 13 pages. <https://hal.archives-ouvertes.fr/hal-02915795>
- [14] Eugénie Brasier, Emmanuel Pietriga, and Caroline Appert. 2021. AR-Enhanced Widgets for Smartphone-Centric Interaction. In *Proceedings of the 23rd International Conference on Mobile Human-Computer Interaction*. ACM, Article 32, 12 pages. <https://doi.org/10.1145/3447526.3472019>
- [15] Catherine E. Brouwer and Jelle Van Dijk. 2011. Brainstorming: Talk and the representation of ideas and insights. In *Proceedings of the Participatory Innovation Conference (PIN-C '11)*. PIN-C, 15–20. <https://pin-c.sdu.dk/2011.html>
- [16] 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'99)*. IEEE, 35–44.
- [17] Marco Cavallo, Mishal Dolakia, Matous Havlena, Kenneth Oehlert, and Mark Podlasek. 2019. Immersive Insights: A Hybrid Analytics System For Collaborative Exploratory Data Analysis. In *25th ACM Symposium on Virtual Reality Software and Technology (VRST '19)*. ACM, Article 9, 12 pages. <https://doi.org/10.1145/3359996.3364242>
- [18] Chenliang Chang, Kiseung Bang, Gordon Wetzstein, Byoungcho Lee, and Liang Gao. 2020. Toward the next-generation VR/AR optics: a review of holographic near-eye displays from a human-centric perspective. *Optica* 7, 11 (Nov 2020), 1563–1578. <https://doi.org/10.1364/OPTICA.406004>
- [19] Olivier Chapuis, Anastasia Bezerianos, and Stelios Frantzeskakis. 2014. Smarties: An Input System for Wall Display Development. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. ACM, 2763–2772. <https://doi.org/10.1145/2556288.2556956>
- [20] Apoorve Chokshi, Teddy Seyed, Francisco Marinho Rodrigues, and Frank Maurer. 2014. Eplan Multi-Surface: A Multi-Surface Environment for Emergency Response Planning Exercises. In *Proceedings of the Ninth ACM International Conference on Interactive Tabletops and Surfaces (ITS '14)*. ACM, 219–228. <https://doi.org/10.1145/2669485.2669520>
- [21] Alex Endert, Christopher Andrews, Yueh Lee, and Chris North. 2011. Visual encodings that support physical navigation on large displays. In *Proceedings of Graphics Interface 2011 (GI '11)*. CHCCS, 103–110. <http://dl.acm.org/gate6.inist.fr/citation.cfm?id=1992917.1992935>
- [22] Barrett M. Ens, Rory Finnegan, and Pourang P. Irani. 2014. The Personal Cockpit: A Spatial Interface for Effective Task Switching on Head-Worn Displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. ACM, 3171–3180. <https://doi.org/10.1145/2556288.2557058>
- [23] Steven Feiner and Ari Shamash. 1991. Hybrid User Interfaces: Breeding Virtually Bigger Interfaces for Physically Smaller Computers. In *Proceedings of the 4th Annual ACM Symposium on User Interface Software and Technology (UIST '91)*. ACM, 9–17. <https://doi.org/10.1145/120782.120783>
- [24] Georges Grinstein, Theresa O'Connell, Sharon Laskowski, Catherine Plaisant, Jean Scholtz, and Mark Whiting. 2006. Vast 2006 contest-a tale of alderwood. In *Symposium On Visual Analytics Science And Technology*. IEEE, 215–216.
- [25] Jens Emil Grønbaek, Majken Kirkegaard Rasmussen, Kim Halskov, and Marianne Graves Petersen. 2020. KirigamiTable: Designing for Proxemic Transitions with a Shape-Changing Tabletop. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. ACM, 1–15. <https://doi.org/10.1145/3313831.3376834>
- [26] Jens Grubert, Matthias Heinisch, Aaron Quigley, and Dieter Schmalstieg. 2015. MultiFi: Multi Fidelity Interaction with Displays On and Around the Body. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, 3933–3942. <https://doi.org/10.1145/2702123.2702331>
- [27] Edward T. Hall. 1960. *The Hidden Dimension*. Anchor Books, USA.
- [28] Chad Harms and Frank Biocca. 2004. Internal Consistency and Reliability of the Networked Minds Measure of Social Presence. In *Seventh Annual International Workshop: Presence* (Universidad Politécnica de Valencia, Valencia).
- [29] S. Hayne, M. Pendergast, and S. Greenberg. 1993. Gesturing through cursors: implementing multiple pointers in group support systems. In *Proceedings of the Twenty-sixth Hawaii International Conference on System Sciences (HICSS '93, Vol. iv)*. IEEE, 4–12 vol.4. <https://doi.org/10.1109/HICSS.1993.284158>
- [30] Juan David Hincapié-Ramos, Xiang Guo, Paymahn Moghadasian, and Pourang Irani. 2014. Consumed Endurance: A Metric to Quantify Arm Fatigue of Mid-Air Interactions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. ACM, 1063–1072. <https://doi.org/10.1145/2556288.2557130>
- [31] Tom Horak, Sriram Karthik Badam, Niklas Elmquist, and Raimund Dachsel. 2018. When David Meets Goliath: Combining Smartwatches with a Large Vertical Display for Visual Data Exploration. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, 1–13. <https://doi.org/10.1145/3173574.3173593>
- [32] Petra Isenberger, 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* 18, 5 (2012), 689–702. <https://doi.org/10.1109/TVCG.2011.287>
- [33] Yuta Itoh, Tobias Langlotz, Jonathan Sutton, and Alexander Plopski. 2021. Towards Indistinguishable Augmented Reality: A Survey on Optical See-through Head-Mounted Displays. *ACM Comput. Surv.* 54, 6, Article 120 (jul 2021), 36 pages. <https://doi.org/10.1145/3453157>
- [34] Yuta Itoh, Tobias Langlotz, Jonathan Sutton, and Alexander Plopski. 2021. Towards Indistinguishable Augmented Reality: A Survey on Optical See-through Head-Mounted Displays. *ACM Comput. Surv.* 54, 6, Article 120 (jul 2021), 36 pages. <https://doi.org/10.1145/3453157>
- [35] Mikkel R. Jakobsen and Kasper Hornbæk. 2014. Up Close and Personal: Collaborative Work on a High-Resolution Multitouch Wall Display. *ACM Transactions on Computer-Human Interaction* 21, 2 (2014), 1–34. <https://doi.org/10.1145/2576099>
- [36] Raphaël James, Anastasia Bezerianos, Olivier Chapuis, Maxime Cordeil, Tim Dwyer, and Arnaud Prouzeau. 2020. Personal+Context navigation: combining AR and shared displays in Network Path-following. In *Proceedings of Graphics Interface (GI '20)*. CHCCS/SCDHM, 267 – 278. <https://doi.org/10.20380/GI2020.27>
- [37] Juho Kim, Haoqi Zhang, Paul André, Lydia B. Chilton, Wendy Mackay, Michel Beaudouin-Lafon, Robert C. Miller, and Steven P. Dow. 2013. Cobi: A Community-Informed Conference Scheduling Tool. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology (UIST '13)*. ACM, 173–182. <https://doi.org/10.1145/2501988.2502034>
- [38] Taeheon Kim, Bahador Saket, Alex Endert, and Blair MacIntyre. 2017. VisAR: Bringing Interactivity to Static Data Visualizations through Augmented Reality. *CoRR* abs/1708.01377 (2017), 4 pages. [arXiv:1708.01377](http://arxiv.org/abs/1708.01377)
- [39] Ulrike Kister, Patrick Reipschläger, Fabrice Matulic, and Raimund Dachsel. 2015. BodyLenses: Embodied Magic Lenses and Personal Territories for Wall Displays. In *Proceedings of the International Conference on Interactive Tabletops & Surfaces (ITS '15)*. ACM, 117–126. <https://doi.org/10.1145/2817721.2817726>
- [40] Anne Köpsel, Päivi Majaranta, Poika Isokoski, and Anke Huckauf. 2016. Effects of auditory, haptic and visual feedback on performing gestures by gaze or by hand. *Behaviour & Information Technology* 35, 12 (2016), 1044–1062. <https://doi.org/10.1080/0144929X.2016.1194477>
- [41] Ricardo Langner, Marc Satkowski, Wolfgang Büschel, and Raimund Dachsel. 2021. MARVIS: Combining Mobile Devices and Augmented Reality for Visual Data Analysis. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21)*. ACM, Article 468, 17 pages. <https://doi.org/10.1145/3411764.3445593>

- [42] Benjamin Lee, Maxime Cordeil, Arnaud Prouzeau, Bernhard Jenny, and Tim Dwyer. 2022. A Design Space For Data Visualisation Transformations Between 2D And 3D In Mixed-Reality Environments. In *CHI Conference on Human Factors in Computing Systems (CHI '22)*. ACM, Article 25, 14 pages. <https://doi.org/10.1145/3491102.3501859>
- [43] Benjamin Lee, Xiaoyun Hu, Maxime Cordeil, Arnaud Prouzeau, Bernhard Jenny, and Tim Dwyer. 2021. Shared Surfaces and Spaces: Collaborative Data Visualisation in a Co-located Immersive Environment. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2021), 1171–1181. <https://doi.org/10.1109/TVCG.2020.3030450>
- [44] Gun A. Lee, Ungyeon Yang, and Wookho Son. 2006. Layered multiple displays for immersive and interactive digital contents. In *International Conference on Entertainment Computing (IFIP ICEC '06)*. Springer, 123–134. https://doi.org/10.1007/11872320_15
- [45] Gun-Yeal Lee, Jong-Young Hong, SoonHyoung Hwang, Seokil Moon, Hyeokjung Kang, Sohee Jeon, Hwi Kim, Jun-Ho Jeong, and Byoungcho Lee. 2018. Metasurface eyepiece for augmented reality. *Nature Communications* 9, 1 (2018), 4562. <https://doi.org/10.1038/s41467-018-07011-5>
- [46] Sikun Lin, Hao Fei Cheng, Weikai Li, Zhanpeng Huang, Pan Hui, and Christoph Peylo. 2017. Ubii: Physical World Interaction Through Augmented Reality. *IEEE Transactions on Mobile Computing* 16, 3 (2017), 872–885. <https://doi.org/10.1109/TMC.2016.2567378>
- [47] Lee Lisle, Kylie Davidson, Edward J.K. Gitre, Chris North, and Doug A. Bowman. 2021. Sensemaking Strategies with Immersive Space to Think. In *Virtual Reality and 3D User Interfaces (VR '21)*. IEEE, 529–537. <https://doi.org/10.1109/VR50410.2021.00077>
- [48] Can Liu, Olivier Chapuis, Michel Beaudouin-Lafon, and Eric Lecolinet. 2016. Shared Interaction on a Wall-Sized Display in a Data Manipulation Task. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, 2075–2086. <https://doi.org/10.1145/2858036.2858039>
- [49] 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 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, 6730–6741. <https://doi.org/10.1145/3025453.3025594>
- [50] Can Liu, Olivier Chapuis, Michel Beaudouin-Lafon, Eric Lecolinet, and Wendy E. Mackay. 2014. Effects of Display Size and Navigation Type on a Classification Task. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. ACM, 4147–4156. <https://doi.org/10.1145/2556288.2557020>
- [51] Weizhou Luo, Anke Lehmann, Hjalmar Widengren, and Raimund Dachselt. 2022. Where Should We Put It? Layout and Placement Strategies of Documents in Augmented Reality for Collaborative Sensemaking. In *CHI Conference on Human Factors in Computing Systems (CHI '22)*. ACM, Article 627, 16 pages. <https://doi.org/10.1145/3491102.3501946>
- [52] Mayra Donaji Barrera Machuca, Winyu Chinthammit, Weidong Huang, Rainer Wasinger, and Henry Duh. 2018. Enabling symmetric collaboration in public spaces through 3D mobile interaction. *Symmetry* 10, 3 (2018), 69. <https://doi.org/10.3390/sym10030069>
- [53] Tahir Mahmood, Erik Butler, Nicholas Davis, Jian Huang, and Aidong Lu. 2018. Building Multiple Coordinated Spaces for Effective Immersive Analytics through Distributed Cognition. In *International Symposium on Big Data Visual and Immersive Analytics (BDVA '18)*. IEEE, 1–11. <https://doi.org/10.1109/BDVA.2018.8533893>
- [54] Thomas W. Malone. 1983. How Do People Organize Their Desks? Implications for the Design of Office Information Systems. *ACM Trans. Inf. Syst.* 1, 1 (jan 1983), 99–112. <https://doi.org/10.1145/357423.357430>
- [55] Sven Mayer, Lars Lischke, Jens Emil Grønbaek, Zhanna Sarsenbayeva, Jonas Vogelsang, Paweł W. Woźniak, 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, 1–10. <https://doi.org/10.1145/3173574.3174113>
- [56] Will McGrath, Brian Bowman, David McCallum, Juan David Hincapié-Ramos, Niklas Elmqvist, and Pourang Irani. 2012. Branch-Explore-Merge: Facilitating Real-Time Revision Control in Collaborative Visual Exploration. In *Proceedings of the 2012 ACM International Conference on Interactive Tabletops and Surfaces (ITS '12)*. ACM, 235–244. <https://doi.org/10.1145/2396636.2396673>
- [57] 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, 262–265. <https://doi.org/10.1145/1031607.1031648>
- [58] Mathieu Nancel, Olivier Chapuis, Emmanuel Pietriga, Xing-Dong Yang, Pourang P. Irani, and Michel Beaudouin-Lafon. 2013. High-Precision Pointing on Large Wall Displays Using Small Handheld Devices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, 831–840. <https://doi.org/10.1145/2470654.2470773>
- [59] Susanna Nilsson, Björn Johansson, and Arne Jonsson. 2009. Using AR to support cross-organisational collaboration in dynamic tasks. In *International Symposium on Mixed and Augmented Reality (ISMAR '09)*. IEEE, 3–12.
- [60] Arthur Nishimoto and Andrew E Johnson. 2019. Extending Virtual Reality Display Wall Environments Using Augmented Reality. In *Symposium on Spatial User Interaction (SUI '19)*. ACM, Article 7, 5 pages. <https://doi.org/10.1145/3357251.3357579>
- [61] Erwan Normand and Michael J. McGuffin. 2018. Enlarging a Smartphone with AR to Create a Handheld VESAD (Virtually Extended Screen-Aligned Display). In *IEEE International Symposium on Mixed and Augmented Reality (ISMAR '18)*. IEEE, 123–133. <https://doi.org/10.1109/ISMAR.2018.00043>
- [62] 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, 1285–1294. <https://doi.org/10.1145/1357054.1357255>
- [63] Arnaud Prouzeau, Anastasia Bezerianos, and Olivier Chapuis. 2017. Evaluating Multi-User Selection for Exploring Graph Topology on Wall-Displays. *IEEE Transactions on Visualization and Computer Graphics* 23, 8 (2017), 1936–1951. <https://doi.org/10.1109/TVCG.2016.2592906>
- [64] Arnaud Prouzeau, Anastasia Bezerianos, and Olivier Chapuis. 2018. Awareness Techniques to Aid Transitions Between Personal and Shared Workspaces in Multi-Display Environments. In *Proceedings of the ACM International Conference on Interactive Surfaces and Spaces (ISS '18)*. ACM, 291–304. <https://doi.org/10.1145/3279778.3279780>
- [65] Erik Prytz, Susanna Nilsson, and Arne Jonsson. 2010. The importance of eye-contact for collaboration in AR systems. In *International Symposium on Mixed and Augmented Reality (ISMAR '10)*. IEEE, 119–126. <https://doi.org/10.1109/ISMAR.2010.5643559>
- [66] Fateme Rajabiyaazdi, Jagoda Walny, Carrie Mah, John Brosz, and Sheelagh Carpendale. 2015. Understanding Researchers' Use of a Large, High-Resolution Display Across Disciplines. In *Proceedings of the International Conference on Interactive Tabletops & Surfaces (ITS '15)*. ACM, 107–116. <https://doi.org/10.1145/2817721.2817735>
- [67] Khairi Reda, Andrew E. Johnson, Michael E. Papka, and Jason Leigh. 2015. Effects of Display Size and Resolution on User Behavior and Insight Acquisition in Visual Exploration. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, 2759–2768. <https://doi.org/10.1145/2702123.2702406>
- [68] Patrick Reipschläger, Tamara Flemisch, and Raimund Dachselt. 2021. Personal Augmented Reality for Information Visualization on Large Interactive Displays. *IEEE Transactions on Visualization and Computer Graphics* 27 (2 2021), 11 pages. Issue 2.
- [69] Jun Rekimoto and Masanori Saitoh. 1999. Augmented Surfaces: A Spatially Continuous Work Space for Hybrid Computing Environments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '99)*. ACM, 378–385. <https://doi.org/10.1145/302979.303113>
- [70] Marc Satkowski, Weizhou Luo, and Raimund Dachselt. 2021. Towards In-situ Authoring of AR Visualizations with Mobile Devices. In *International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct '21)*. IEEE, 324–325. <https://doi.org/10.1109/ISMAR-Adjunct54149.2021.00073>
- [71] Stacey Scott, Sheelagh Carpendale, and Kori Inkpen. 2004. Territoriality in Collaborative Tabletop Workspaces. In *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW '04)*. ACM, 294–303. <https://doi.org/10.1145/1031607.1031655>
- [72] Marcos Serrano, Barrett Ens, Xing-Dong Yang, and Pourang Irani. 2015. Gluey: Developing a Head-Worn Display Interface to Unify the Interaction Experience in Distributed Display Environments. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '15)*. ACM, 161–171. <https://doi.org/10.1145/2785830.2785838>
- [73] Anton Sigitov, André Hinkenjann, Ernst Kruijff, and Oliver Staadt. 2019. Task Dependent Group Coupling and Territorial Behavior on Large Tiled Displays. *Frontiers in Robotics and AI* 6 (2019), 17 pages. <https://doi.org/10.3389/frobt.2019.00128>
- [74] Tianchen Sun, Yucong Ye, Issei Fujishiro, and Kwan-Liu Ma. 2019. Collaborative Visual Analysis with Multi-level Information Sharing Using a Wall-Size Display and See-Through HMDs. In *Pacific Visualization Symposium (PacificVis '19)*. IEEE, 11–20.
- [75] Tianchen Sun, Yucong Chris Ye, Issei Fujishiro, and Kwan-Liu Ma. 2019. Collaborative Visual Analysis with Multi-Level Information Sharing Using a Wall-size Display and See-Through HMDs. In *Proceedings of IEEE Pacific Visualization Symposium (PacificVis '19)*. IEEE, 11–20.
- [76] 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, 1181–1190. <https://doi.org/10.1145/1124772.1124950>
- [77] Theophanis Tsandilas, Anastasia Bezerianos, and Thibaut Jacob. 2015. SketchSliders: Sketching Widgets for Visual Exploration on Wall Displays. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, 3255–3264. <https://doi.org/10.1145/2702123.2702129>

- [78] Vinoba Vinayagamoorthy, Maxine Glancy, Christoph Ziegler, and Richard Schäfer. 2019. *Personalising the TV Experience Using Augmented Reality: An Exploratory Study on Delivering Synchronised Sign Language Interpretation*. ACM, 1–12. <https://doi.org/10.1145/3290605.3300762>
- [79] James R. Wallace, Nancy Iskander, and Edward Lank. 2016. Creating Your Bubble: Personal Space On and Around Large Public Displays. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, 2087–2092. <https://doi.org/10.1145/2858036.2858118>
- [80] James R. Wallace, Stacey D. Scott, Eugene Lai, and Deon Jajalla. 2011. Investigating the Role of a Large, Shared Display in Multi-Display Environments. *Computer Supported Cooperative Work* 20 (2011), 529–561. <https://doi.org/10.1007/s10606-011-9149-8>
- [81] Xiyao Wang, Lonni Besançon, David Rousseau, Mickael Sereno, Mehdi Ammi, and Tobias Isenberg. 2020. Towards an Understanding of Augmented Reality Extensions for Existing 3D Data Analysis Tools. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. ACM, 1–13. <https://doi.org/10.1145/3313831.3376657>
- [82] Andrew D. Wilson and Hrvoje Benko. 2010. Combining Multiple Depth Cameras and Projectors for Interactions on, above and between Surfaces. In *Proceedings of the 23rd Annual ACM Symposium on User Interface Software and Technology (UIST '10)*. ACM, 273–282. <https://doi.org/10.1145/1866029.1866073>
- [83] Fengyuan Zhu and Tovi Grossman. 2020. BISHARE: Exploring Bidirectional Interactions Between Smartphones and Head-Mounted Augmented Reality. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. ACM, 1–14. <https://doi.org/10.1145/3313831.3376233>