



HAL
open science

Augmented and Physical reality environments used in collaborative visual analysis

Raphaël James

► **To cite this version:**

Raphaël James. Augmented and Physical reality environments used in collaborative visual analysis. Human-Computer Interaction [cs.HC]. Université Paris-Saclay, 2023. English. NNT : 2023UP-ASG011 . tel-04031932

HAL Id: tel-04031932

<https://theses.hal.science/tel-04031932>

Submitted on 16 Mar 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Augmented and Physical reality environments used in collaborative visual analysis

*Environnements de Réalité Physique et Augmentée utilisés
pour l'analyse visuelle collaborative*

Thèse de doctorat de l'université Paris-Saclay

École doctorale n° 580, Sciences et Technologies de l'Information et de la
Communication (STIC)

Spécialité de doctorat: Informatique

Graduate School : Informatique et Sciences du Numérique. Référent : Faculté des
sciences d'Orsay

Thèse préparée dans l'unité de recherche **Laboratoire interdisciplinaire des sciences
du numérique (Université Paris-Saclay, CNRS, Inria)**, sous la direction
de **Anastasia BEZERIANOS**, Professeure et le co-encadrement
de **Olivier CHAPUIS**, Chargé de Recherche HC

Thèse soutenue à Paris-Saclay, le 16 Février 2023, par

Raphaël JAMES

Composition du jury

Petra Isenberg

Directrice de recherche, Inria Saclay

Gilles Bailly

Directeur de recherche, Sorbonne Université

Laurent Grisoni

Professeur des universités, Université de Lille

Christophe Hurter

Professeur des universités, ENAC

Valérie Gouranton

Maîtresse de conférences, Université de
Rennes

Présidente du Jury

Rapporteur et Examineur

Rapporteur et Examineur

Examineur

Examineur

Acknowledgement

I want to thank all the people who allowed me to do this thesis, from the moment I took the direction of the academic path, to the good preparation and execution of the thesis defense.

First, thanks to Anastasia and Olivier, my thesis supervisors. This thesis was not easy, but which thesis is? You have always been there to help me, to accompany me. When I had to work between France and Australia on a research project, when COVID hit at the worst moment, or when technical problems occurred and had to be solved. I had been warned that supervisors of a thesis become in a way my second parents, but I reaffirm that you have been "the parents" of this thesis. I thank you from the bottom of my heart for allowing me to do a thesis in such beautiful conditions.

A big thank you to the jury, who did not read these acknowledgments but have read this manuscript in its entirety. Thank you for having taken the time to read this work, and for asking me these so interesting questions during my defense, it is an exercise that I appreciated a lot, and for which I am grateful.

Then, many thanks to my parents. They pushed me to exploit my abilities and satisfy my thirst for knowledge. Thanks to my mother, I did a preparatory class instead of entering a computer science program. And these 2 years of preparatory classes were the beginning of my journey so far. These classes taught me rigor and failure when it comes to fields where I thought I would excel. Thanks too to my grandparents who supported me to get out of my comfort zone and do what I love.

I would also like to mention dance, a discipline I practiced for 13 years. This sport, its practitioners, and its teachers, taught me physical and mental discipline, taught me the fighting spirit, and the will to continue to do what we like despite the failures. In this experience, I would particularly like to thank Océane Lesoeur, who accompanied me for more than 10 years, from competitions in the 4 corners of France to learning this way of life which is the dance. I also want to thank Sébastien Labergère, who has given me a lot of human support, who believed in me and taught me to believe in myself, to trust myself, and to show this trust. He, with Laure Colmard, Frantz Ducroix, and Séverine Asnar Mourad, also taught me by combining exigence and benevolence, which inspired me during my thesis and still inspires me in my way of teaching others. Finally, I want to thank all the dancers I met and with whom I was able to practice with passion a passion, in the same way, that I embarked on this Ph.D. to work passionately on my passion.

But I would not have reached the end of these 4 years of work without moral support. For that, I thank my love, Maëva Bernard, whom I met during this thesis, who brought me a breath of fresh air during the worst period to live which was COVID. She supported me during the endless nights of work, the endless hours, and the incessant discussions about my thesis, and encouraged me to take it easy, to rest to be able to

work on my work afterward. And for that, for all the people who will read these thanks, and who are already engaged in a thesis: rest is crucial, don't feel guilty for not being able to work at times, it's probably a signal that you need a rest or a new perspective to start again. Likewise, it's good to be inspired and to work without seeing the clock ticking, but learning to set boundaries in your work is what will keep you going. Seeing your friends and talking with your lab mates is also important.

And for that, I thank my friends. First of all, my friends from engineering school. Always there to chat, to have a drink, to start a game of Phasmophobia or Valheim and to change the ideas, or to host (for 6 months) a Ph.D. student in need when things don't go exactly as planned... Alexandre Vincent, fellow traveller, thank you for having chosen to embark on this crazy adventure of doing a master's degree at the other end of Paris, driving every day of the week through the Parisian traffic jams to embark on the best courses we've ever had. Thanks to you and MC for sacrificing hours of sleep for crazy projects based on deodorant bottles and Augmented Reality markers. Mickaël, thank you for taking life as it comes and thank you for living extraordinary adventures, and for having the audacity to make choices that no one would make. Maïwenn, thank you for being so candid, and for reminding me that beyond work there are other things that are worth appreciating. Alexandre Terrien, thank you for your depth of thought, your discussions, and your sensitivity. And thank you to the 5 of you for always being there to have a drink, to cheer up, or to simply chat. Also, thanks to Valentin Famelart, for discussions about everything and nothing, and for coming back to get me even when I was head over heels in thesis work. I promise I will finally be able to play Horizon now. Thanks to my long-time friends, those who have been there since high school: Antoine, Emeline, Marie, Laure, Tristan, Marc, François, Valentin. That's it, I'm joining you, I finally finished my studies too!

I would also like to thank Jean-Paul Forest, a professor who during my years of engineering school led me to think about computer science not only as a field of engineering, but also as a field of research, and who led me to use my mathematical knowledge in my programming.

And if my supervisors have the role of parents, I must thank the family that was the ILDA team. Emmanuel and Caroline who have this gift to put newcomers at ease, whose advice and guidance were invaluable in directing, writing, and presenting my research work. Vanessa who is a real ray of sunshine from Chile. I thank Eugénie Brasier, thesis companion, team leader, board game connoisseur, and galley companion for developing projects on HoloLens! Thanks to Manu, Mehdi, Vincent, Camille, Dylan. And as he is everywhere, I mention Julien, because it is sure that you will make a thesis in this magnificent team at this rate!

Thanks to Mickaël Sereno, with whom I could discuss technique and research, and who has been a great companion of research, good luck with your new role of dad! Thanks to all the members of AVIZ and Ex-Situ with whom I could discuss, or who

passed my experiments.

I would also like to thank all the people I met during my internship at Monash: Tim Dwyer, Maxime Cordeil, Arnaud Prouzeau, Barrett Ens. They showed me another way of doing research and gave me the opportunity to do a research internship in Australia. I would also like to thank Benjamin Lee, with whom I shared superb research discussions on our respective subjects. And thanks to Cameron and Stephanie, who accompanied me during this Australian adventure!

Finally, thanks to all the interns I have met, and to the students, whom I have taught during these 4 years. Teaching has allowed me to affirm my knowledge, and to learn more about how to transmit the very particular knowledge that I have acquired during this thesis. It is for me an essential exercise of the thesis, to give back the knowledge accumulated in the preparation for becoming a researcher. If you are in the thesis, I recommend you try it.

Abstract

The amount of data produced every day is constantly growing. This mass of raw information requires finding solutions to analyze and visualize it. Wall displays are a good candidate to solve this issue, allowing to display a large amount of data. However, they are shared devices made to accommodate several people, which does not allow for the existence of a private workspace. Moreover, wall displays are still physical devices that can not be easily extended when the content to display becomes too large.

In this manuscript, I explore the combination of Augmented Reality (AR) and a wall display to: first, add a personalized information space onto a wall display; and second, extend the rendering space. We focus on using an augmented reality headset (AR-HMD) to display content complementary to the context offered by the wall. We study the synergy of these two devices through two research questions:

RQ1 - How can AR assist exploration and navigation of a network on a wall display, through personal view and navigation?

RQ2 - How does extending the screen space of a wall display using AR changes users' space exploitation during a collaborative session?

For RQ1, I study the case of using AR with a wall display for navigating networks. I present four techniques, proposing different visual aid mechanisms for navigation. I then evaluate these four techniques in two experiments, one for low accuracy path following in the network (path selection) and one for high accuracy path following (path tracing). We report that for the path selection task, a persistent connection between the cursor in AR and the network on the wall obtained better results. In the case of path tracing, we observe that a lighter connection offered better performance. Moreover, we show the feasibility of a system where the interaction is done privately in AR for network navigation on shared screens.

For RQ2, I focus on the space available in front of the wall display and consider using AR to extend the wall display. I present a system extending a wall display with AR, taking advantage of the space in front of the wall. We compare our system that combines a wall display with AR with a wall display alone using two collaborative tasks. We observe that users extensively use the available virtual space with our system. Although this creates an additional cost of interaction, we observe no performance difference and a real benefit of this extra space.

The complexity of setting up the previously studied systems led us to study a cheaper way to use a wall display:

RQ3 - Can we emulate a wall display inside a VR headset?

This question has many aspects, and I focus on the capacity for VR headsets to

reach the necessary resolution to replicate the user experience of a high-resolution wall display. For this, we study the optic model of VR headsets and compare it to the human vision model. Our analysis indicates that current headsets need a higher resolution to emulate a wall display. We confirm our analysis by running a pilot study comparing a wall display, two VR headsets, and an emulation of a perfect VR headset.

I conclude this manuscript by discussing the different ways for a wall display to be combined with AR headsets, the requirements for replacing wall displays by VR/AR headsets, and by elaborating on the future work this thesis opens.

Synthèse

Cette thèse de doctorat porte sur la combinaison de la Réalité Augmentée (RA) et des murs d'écrans pour créer des espaces d'information personnalisés et étendre l'espace de rendu. Elle part du constat que les murs d'écrans manquent d'espace privé ou personnel, mais aussi qu'il est difficile d'augmenter physiquement leur taille. Cette thèse explore donc le potentiel des casques de Réalité Augmentée pour afficher un contenu complémentaire au contexte offert par le mur, en relevant le défi de trouver des solutions pour analyser et visualiser la quantité massive de données produites chaque jour.

La première question de recherche (QR1) est de savoir comment la Réalité Augmentée peut aider à l'exploration et à la navigation d'un réseau sur un mur d'écran par le biais d'une vue et d'une navigation personnelles. Je présente quatre techniques proposant différents mécanismes d'aide visuelle à la navigation. L'évaluation de ces techniques par le biais de deux expériences, l'une pour un suivi de chemin de faible précision dans le réseau et l'autre pour un suivi de chemin de haute précision, montre qu'une connexion persistante entre le curseur dans la Réalité Augmentée et le réseau sur le mur a obtenu de meilleurs résultats pour la tâche de tracé de chemin. Dans le cas de la sélection du chemin, une connexion plus légère offre de meilleures performances. Ce chapitre montre également la faisabilité d'un système dans lequel l'interaction se fait en privé dans la Réalité Augmentée pour la navigation dans le réseau sur des écrans partagés.

La deuxième question de recherche (QR2) porte sur l'espace disponible devant le mur d'écran et envisage l'utilisation de la Réalité Augmentée pour étendre le mur d'écran. Ce chapitre présente un système qui étend un affichage mural avec la Réalité Augmentée, en tirant parti de l'espace devant le mur. Ce système est comparé à un mur d'écran seul dans le cadre de deux tâches collaboratives, et les résultats montrent que les utilisateurs utilisent largement l'espace virtuel disponible avec le système. Bien que cela crée un coût d'interaction supplémentaire, il n'y a pas de différence de performance. Les observations montrent une uniformisation des stratégies de collaboration avec le système mur d'écran et Réalité Augmentée lors de la réalisation d'une tâche de classification. De plus, les participants se sentent plus performants, et préfèrent utiliser le système apportant cet espace supplémentaire.

La troisième question de recherche (QR3) vise à déterminer s'il est possible d'émuler un mur d'écran à l'intérieur d'un casque de Réalité Virtuelle (RV). Ce chapitre analyse le modèle optique des casques de Réalité Virtuelle et le compare au modèle de vision humaine pour déterminer si les casques actuels peuvent reproduire l'expérience utilisateur d'un mur d'écran haute résolution. L'analyse indique que les casques actuels ont besoin d'une résolution plus élevée pour émuler un affichage mural. Une étude pilote, comparant un mur d'écran, deux casques de Réalité Virtuelle et une émulation

d'un casque de Réalité Virtuelle parfait, confirme cette analyse. Le questionnaire utilisateur suggère que même avec un casque de Réalité Virtuelle parfait, les utilisateurs préféreraient l'utilisation du mur d'écran.

Dans l'ensemble, cette thèse montre le potentiel de la combinaison de la Réalité Augmentée et des écrans muraux pour créer des espaces d'information personnalisés et étendre l'espace de rendu. A travers la création de plusieurs techniques proposant une aide personnelle à la navigation sur un réseau sur un mur, en Réalité Augmentée mais aussi l'étude d'un système augmentant la taille d'un mur d'écran avec de la Réalité Augmentée, cette thèse propose différents exemples de synergie entre un mur d'écran et un casque de Réalité Augmentée. Cette thèse ouvre de nouvelles voies pour explorer les possibilités de combiner la RA et les affichages muraux afin de fournir des espaces d'information plus immersifs et personnalisés pour les individus et les groupes. Mais elle montre aussi que de tels systèmes peuvent être bénéfiques, pour une utilisation dans des contextes multi-utilisateurs.

Contents

Acknowledgements	i
Abstract	v
Synthèse	vii
Contents	ix
List of Figures	xii
1 Introduction	1
1.1 Design dimensions	3
1.2 A wall shared by all	6
1.3 A wall limited in size	7
1.4 A wall in Virtual Reality	7
1.5 Thesis outline and contributions	8
2 Related Work	11
2.1 Working with large displays	11
2.1.1 Large display use-context & applications	11
2.1.2 A large vertical display or horizontal display?	12
2.1.3 Large displays advantages and drawbacks	14
2.1.4 Enhancing wall displays with devices	15
2.1.5 Summary	16
2.2 Screen combined with AR	17
2.2.1 Early work	17
2.2.2 Extending a display	18
2.2.3 Superimposing on a display	19
2.2.4 Alternative to a display	20
2.2.5 Summary	21
2.3 Placement of this thesis	22
3 Personal+Context navigation: combining AR and shared displays in Network Path-following	25
3.1 Introduction	25
3.2 Related Work	28
3.2.1 Navigating Paths in Graphs	28
3.2.2 Interacting with Graphs on shared large displays	29
3.3 Design Goals	29

3.3.1	Personal Navigation	31
3.4	Techniques	31
3.4.1	Persistent connection: <i>Sliding Metaphor</i>	32
3.4.2	Transient connection: <i>Magnet Metaphor</i>	35
3.5	Experimental Setup	38
3.6	Experiments	39
3.6.1	Path Selection Experiment	44
3.6.2	Path Tracing Experiment	46
3.7	Discussion and Perspectives	48
3.8	Conclusion	50
4	Extending Wall Displays with AR	53
4.1	Introduction	53
4.2	Related Work	55
4.2.1	Studies on Co-located Collaborative Interaction	56
4.3	Prototype	57
4.3.1	Virtual Elements	58
4.3.2	Interaction	60
4.3.3	Replay Sessions	62
4.4	User Study	63
4.4.1	System Conditions	63
4.4.2	Tasks	64
4.4.3	Hypotheses and Measures	65
4.4.4	Participants & Apparatus	66
4.4.5	Experiment Design & Procedure	66
4.5	Results	67
4.5.1	Collaboration Strategies and Workspace Usage	67
4.5.2	Interactions and AR Technique Use	71
4.5.3	Additional Objective Measures: Partners and Cards Traveling, Position, and Task Time	72
4.5.4	Subjective results: Questionnaires	74
4.6	Discussion and Limitations	76
4.7	Conclusion	80
5	Virtual Wall Display	81
5.1	Introduction	81
5.2	Related Work	83
5.2.1	Screen Size and Resolution	83
5.2.2	HMD versus Other Displays	84
5.2.3	HMD in Immersive Analytics	84
5.2.4	Emulating Walls in VR HMDs	84

5.3	Background and Geometrical Models of Vision	85
5.3.1	Principle of the monocular eye for a physical screen	85
5.3.2	Principle of the monocular eye for a VR headset	87
5.3.3	Emulating a Wall	89
5.4	The Classification Task	91
5.4.1	The classification Task in VR	92
5.5	Pilot User Study	93
5.5.1	Device conditions and apparatus	93
5.5.2	Task	93
5.5.3	Participants	94
5.5.4	Experiment Design & Procedure	94
5.5.5	Results	95
5.6	Discussion and Future Work	97
6	Discussion, Future Work and Conclusion	99
6.1	Summary	99
6.2	Discussion and Futur work	102
6.2.1	Personal AR Views on Shared Displays	102
6.2.2	Combining AR and Walls in Collaboration	103
6.2.3	Replacing wall displays with AR or VR?	103
6.2.4	Including wall displays in workflows	105
6.3	Conclusion	107
	Bibliography	109
A	Supplementary Material for Chapter 4	137
	View of the final results	139
	Heatmap of participants and Surfaces	145

List of Figures

1.1	Example of a wall display	2
1.2	Design Dimension of AR on context	4
1.3	Correlation between Design dimensions	5
2.1	Examples of AR screen extension	18
2.2	Examples of Superimposition of AR on physical screen	20
2.3	Examples of alternative screen in AR/VR	21
3.1	User looking at a shared display augmented with AR	26
3.2	Scenario for AR aids in the metro	27
3.3	Example of technique with rendering	33
3.4	Close-up of SlidingRing (combined view)	33
3.5	Close-up of SlidingElastic (combined view)	33
3.6	Link routing through Procrustes analysis	34
3.7	Close-up of MagneticArea	37
3.8	Close-up of MagneticElastic	37
3.9	Schematic of navigation techniques	38
3.10	Example of trial path from path-following experiment	40
3.11	Graphs used in path-following experiment	42
3.12	Completion time results	43
3.13	Questionnaire result of the selection experiment	46
3.14	Average time results	47
3.15	Questionnaire result for path-tracing experiment	48
4.1	Photo-edited view of our Wall+AR system for collaboration	54
4.2	Presentation of virtual surfaces types	59
4.3	Visual feedbacks for cards and menu	60
4.4	View for the Replay tool	62
4.5	Initial state of the Wall display for the experiment	65
4.6	Examples of final results for the classification task	68
4.7	Examples of final result in story task	69
4.8	Average for traveled distances of users, cards, task time and distance between partners	73
4.9	Results for the questionnaires	75
5.1	Classification task in VR	82
5.2	Simplified optic scheme for an eye looking at a screen	85
5.3	Different devices require different viewing distances and resolution	86
5.4	Different devices, their resolution and their optimal viewing distance	86
5.5	Simplified scheme for optics of a VR headset	88
5.6	Letter pixelisation on different resolution and upscaling	92
5.7	Result of the vision test	95

5.8	Task time by DEVICE×LETTERSIZES.	96
5.9	Distance to the wall and traveled distance	96
6.1	Scenario of transition between working alone to working in collaboration on wall display content	106

Introduction

Today more data is being generated than ever before. Over the last ten years, Big Data has become a critical field of research, with data scientists exploring different ways to transform terabytes of information into valuable insights. One of the popular methods to create understanding is to generate visual representations of the data [215]. In Visual Analytics, a wide variety of devices and visualization techniques are used to explore complex datasets.

Most recently, researchers have begun to explore how we can use novel immersive display technology, to enhance visual analytics solutions. This specific field of Visual Analytics: Immersive Analytics [41] "*investigates how new interaction and display technologies can be used to support analytical reasoning and decision-making. The aim is to provide multi-sensory interfaces for analytics approaches that support collaboration and allow users to immerse themselves in their data. Immersive Analytics builds on technologies such as large touch surfaces, immersive virtual and augmented reality environments, haptic and audio displays, and modern fabrication techniques.*"

These new display technologies range from room-scale environments, such as interactive wall displays, tables, and multi-display environments [212], to Virtual Reality and Augmented Reality Head-Mounted Displays (VR/AR HMDs) [59]. These immersive environments provide unique opportunities to gather multiple people and collaboratively analyze and understand complex visual information. The displays involved in those environments vary widely in size, resolution, and ways of interacting with them. Choosing which technology is more appropriate for the given tasks and needs of analysts, and designing for them, is a challenge.

Among the physical collaborative platforms, large displays, ranging from tabletops to vertical displays like whiteboards or wall displays, have been extensively studied in the context of information visualization analysis (e.g., [31, 166, 179]). Wall displays (such as Figure 1.1) with their high-resolution provide an ideal environment to display a large amount of data at once [239]. Their size offers a space that can be taken advantage of for sensemaking tasks [6], and can also be used to accommodate many users in front of them and allow them to work concurrently on the same content [12].



Figure 1.1: Example of a wall display with the Wilder system.

However, some of the wall displays' properties can become challenges when interacting with the display. For starters, when exploring content on a shared large display, different users may want to explore different views. However, as it is shared, the large display displays the same content to every user, so if someone does a modification, this affects everyone's view. This may disrupt other users when modifications create significant visual changes. Moreover, even though these displays are larger than traditional displays, they are still physical displays and they may reach their limits when the content to render exceeds their physical size. Like every physical device, they suffer from physical limitations, but expanding them is expansive and requires complex reconfiguration.

Another technology that is becoming increasingly popular and is central in the field of immersive visual analytics is Augmented Reality (AR). As per the definition given by Milgram and Kishino [147], Augmented Reality refers *"to any case in which an otherwise real environment is "augmented" by means of virtual (computer graphic) objects"*. In a later definition, Azuma [13] states 3 criteria to qualify a system of AR: (i) Combines real and virtual, (ii) Interactive in real-time, and (iii) Registered in 3D.

This last characteristic, its capacity to anchor virtual content to the physical space, is of interest to us as it allows the placement of visualization right where the content

of interest lies [15]. Thus it allows us to expand the area where we display content beyond the limits of the physical display. A variation of AR that is of great interest is using AR Head-Mounted Displays (AR HMD). This technology offers a mobile solution to have personalized content only visible to the headset's wearer.

Thus, the characteristics of both the wall display technology and AR HMD technology indicate a potential synergy between the two mediums by combining the high-resolution shared view of the physical wall display, with the easily extendable and personalized view offered by the AR HMD. Milgraam and Kishino [147] describe different categories for virtual technologies, forming milestones in a continuum ranging from a completely real environment, to a fully virtual environment. While we use Augmented Reality in our work, we use it to "Augment Virtuality", AR acting as a support for a context on a Large Display. However, our work let users freely use the real environment available to them, in addition to the Large Display and the headset. Thus I place this work closer to an in-between place Milgram and Kishino name "Hybrid Reality".

In our work, we focus on this combination of wall displays and AR HMD. Other work has considered combining AR HMD with physical displays in order to address their physical size limits. For example, recent work in Immersive Analytics has shown the use of AR with a physical display to extend the screen space of a smartphone [160], offer a complementary view to a desktop screen [230], or add information to a tabletop screen [158]. Those previous examples of augmenting physical screens with AR focus on small physical displays, while our work combines wall displays with AR.

In this thesis, I explore different designs mixing AR HMDs to wall displays to solve some limitations of large displays. In particular, I empirically study the impact of combining a wall display and AR to aid content navigation and mitigate the physical constraints of the physical wall display. I first explore the use of AR to provide a personal view and interaction for navigation on a shared display, without disturbing others. Then, I design a system that expands a physical wall display using virtual AR surfaces; and I study how extending the wall display's real-estate affects collaboration. I continue by discussing the impact of resolution when attempting to recreate a wall display entirely in Virtual Reality (VR) using VR Head-Mounted Displays (VR HMD). In particular, I study the performance of such a setup against an actual physical wall display, and discuss the cost of replacing a wall display with a virtual one.

1.1 Design dimensions

The work presented in this manuscript started during my master's internship, which explored the designed space of augmentation of a wall display by Augmented Reality. During this initial work, I introduced several different visualization designs that combine wall displays with AR. Using these designs and related work examples, I identified four characteristics (dimensions) that describe how AR content can be combined with a




Figure 1.2: Design Dimension of AR content on screen context

wall display. These dimensions were: *Dimensionality*, *Visibility*, *Personalization*, and *Coupling* (*Semantic Coupling* and *Display Space*, Figure 1.2).

Dimensionality describes whether the AR content has a 2D or 3D aspect. When AR content is superimposed on a physical display, it can be placed flat on the same plane than the display or extruded in 3D space. *Visibility* refers to how well we can see an AR item, in other words if the AR content is easily visible to every user or hardly visible because of transparency or not shown because of privacy rights. *Personalization* represents the amount of customization available on an AR component. We first defined *Coupling* as the connection between the AR content and the content on the wall display. However, further exploration of augmentation examples made us precise the *Coupling* dimension by dividing it into 2: the *Semantic Coupling* dimension and the *Display Space* dimension. These two aspects are the ones we focus on in more detail in this thesis.

We define *Semantic Coupling* as the semantic connection between the AR content and the content on the wall display. A close semantic coupling means the information of the content added in AR will be highly dependant of the context it is linked to. Taking the example of a scatterplot, a label displayed on top of a data point has a close semantic coupling because the content of the label describes directly the data point. On the other hand, a loose semantic coupling means the AR content will have little to no relation to the displayed content. For example, a notification that a new dataset is available may have little to no connection to the context on the wall display.

The *Display Space* corresponds to the balance between the augmented space relative to the physical screen space. It is a continuum that goes from having only a physical screen to replacing the screen by showing the content only in AR. Of course, there are many in-betweens with augmentations placed on top of the screen, bound to its physical boundaries like, for example, having an AR notification popping on top of a screen. Another possibility is AR which extends the limits of the physical screen like [160], who extends a smartphone space with AR.

Both dimensions take importance in the spatial positioning and relationship of the augmentation. While *Display Space* concentrates on the physical display, *Semantic Coupling* revolves around the relationship with the content on display and how we augment this content. We observe that those two dimensions are correlated. Most




Figure 1.3: Relationship between *Semantic Coupling* and *Display Space*.

of the time, when added content is closely semantically coupled with its context, it is also located in the same space as its content. If we return to the example of labels in a scatterplot, the labels are semantically closely coupled and close in display space. Moreover, when the augmentation is loosely coupled to the context, we see it placed farther away from the content. Indeed when we augment a screen to increase its workspace and manipulate more content, this is semantically loosely coupled and far in display space.

While examples other than the extremes are scarce, they do exist. For example, the excentric labeling technique [71] shows labels containing data points' values at a distance from the selected data point to reduce clutter and relationship confusion from having too many labels at one site. Let us consider an implementation where the data points are displayed on the wall display and the cursor in AR. We have an example of close semantic coupling: the data points' labels, with far display space as the labels are placed far from the points they represent.

In this thesis, we decided to study the extreme cases of *Semantic Coupling* and *Display Space*. In our [Chapter 3](#) we explore an example of tight *Semantic Coupling* and close *Display Space* by studying AR cursors on a map. In our [Chapter 4](#), we explore the other extreme, loose *Semantic Coupling* and far *Display Space* by studying the augmentation of a wall screen work space with AR.

1.2 A wall shared by all

The size of large displays often leads to multiple people using them at the same time, and can accommodate individuals working in parallel or groups collaborating together [12]. In such a situation, users risk disturbing others working on the shared display whenever they manipulate the view. Moreover, we cannot display information that requires privacy even if the wall display can have space to accommodate it.

One possibility to solve this issue is to display the necessary information on a personal device instead of the public display. While smartwatches, smartphones and tablets, already widespread personal devices, could be used, this would require a lot of back and forth between the context (large display) and the additional information (personal device). Using our previously defined design dimensions, we talk of information decoupled in display space from the context. This is where Augmented Reality, provides a possibility to display content closely coupled with the context.

In the case of close collaboration, the addition of AR on top of a large display has already shown interest. Nilsson et al. [158] showed we could display symbols on a tabletop for command and control situation. More recently, work has considered AR with wall displays, focusing on augmenting networks for sensitive information [205] and adding personal aids to visualizations [181].

Studying coupling variation of AR content with a wall display is a broad question. For this thesis, I decided to focus on a specific use-context. As wall displays are frequently used to display networks in control rooms like train networks, maritime traffic, or roads, I decided to focus on the specific use-context of network navigation. This decision leads to our first research question:

RQ1: How can AR assist exploration and navigation of a network on a wall display, through personal view and navigation?

To explore this personal view and navigation, we created two cursor techniques with two variations each to help navigation. Those techniques are displayed exclusively in the AR headset of the wearer, and both techniques are semantically tightly coupled with their wall display context (the network). One cursor technique displays a constant and close display coupling, and the other a more transient display coupling. We studied them during two path-following experiments to observe how the coupling variations impact navigation.

1.3 A wall limited in size

The previous example studied a use-context where the AR and wall display content are tightly coupled semantically and closely coupled in display space. For my next exploration, I chose to explore the other end of the spectrum, with loose semantic coupling, and far display coupling between AR and wall display content.

Wall display, as a technology, promote physical navigation [21] and spatial recall [104]. But large displays require space: they need to be placed in rooms big enough to fit their size. This space can be used by viewers to get different perspective of the content, for example moving farther away to get overviews and close to see details [94]. Or, it can be used as a mean to interact with the wall display [75]. However it is less exploited to display information [169].

On the other hand, AR has the specific ability to add virtual content to the physical space around us. In both of these technologies, the notion of space is key, so it is important to consider the use of space around the user. This aspect is even more important when considering the use-context of collaboration.

This is also true when considering the display space coupling. Additional information can be placed only in the space defined in AR. It does not serve the same purpose if the space augmented is on the physical display, around it, or far from the display. After considering closely coupled content added on top of the screen, I explore the space around the screen.

For this, I wanted to study the impact of extending the physical screen space using AR. More precisely, I wanted to see if this addition modifies user's behaviour when using the space available in front of the wall display, especially during collaborative sessions where the space needs to be shared. This leads to our second research question:

RQ2 : How does extending the screen space of a wall display using AR changes users' space exploitation during a collaborative session?

To investigate this question, we study the use of space and its impact on collaboration when the wall display is used alone and when it is used with AR. We created a prototype system to visually extend a wall display using AR, and to manipulate this AR content. We then conducted a comparative study on pairs of participants performing collaborative tasks.

1.4 A wall in Virtual Reality

Our two previous works explored the extremes of *Display Space* and *Semantic Coupling* dimensions. In both solutions, we combined AR, in the form of AR HMD, with existing

large vertical displays. This setup still has some challenges as it is technically complex: it requires alignment between the AR content seen in the HMD and the content shown on the physical wall display. As we are adding virtual information (through AR) to virtual content (through the physical display), a way to achieve a similar setup for a lower cost would be to use a *completely virtual* environment, simulating not only the AR content but also the entire wall display inside a complete VR environment.

In its visual form, VR technology shows virtual content instead of natural vision and immerses its users in a virtual world. During the past decade, VR HMDs have become more powerful and affordable. Consumer-grade headsets exist, from the cheapest Google Cardboard¹ which relies on your smartphone, to more performant solutions like the Meta Quest Pro² or the Pico 4³. Thus these technologies could simulate a wall display in VR, using interactions similar to the ones proposed previously for AR. This wall display simulation would be cheaper and less complex than building and running a physical wall display. Thus came the question:

RQ3 : Can we emulate a wall display inside a VR headset?

This is a complex question, as we will discuss later in the thesis. But to start addressing it, we focus the question by comparing large displays and VR HMDs when it comes to resolution in particular. We study the question through two angles: (i) theoretical, by studying the differences using standard vision models, and (ii) empirical, to validate through a pilot experiment the theoretical observations. In particular, we compare two types of VR HMDs and a high-resolution wall display in a pilot study, using a classification task previously used to compare a desktop and a wall display [133].

1.5 Thesis outline and contributions

In this thesis, we present in [Chapter 2](#) previous work on collaboration using wall display's alone, followed with a state of the art on how devices have been combined with a wall display. We then discuss how AR has been combined with physical devices more generally and then specifically with wall displays, the focus of this thesis.

[Chapter 3](#) explores different interaction techniques for network navigation in the use-case of a shared large public display. We developed two hands-free techniques with two visual variations for personal navigation of a network on a wall display. These techniques provide visual aids rendered in AR to interact and keep connection with the context on a large display without altering its content. We studied these techniques,

¹<https://arvr.google.com/cardboard/>

²<https://www.meta.com/us/quest/quest-pro/>

³<https://www.picoxr.com/fr/products/pico4>

presenting different levels of close display coupling, during two experiments of path-following of different difficulties. Our results show that the techniques providing the closest coupling performed better for high-precision path-following. For low-precision path-following, our technique with looser coupling is preferred. Moreover, we show that a personal view to provide navigation aid in AR to a shared public display is feasible.

In [Chapter 4](#), we study the extension of a wall display using AR. We created a system, Wall+AR, to interact with the virtual space around the wall display by adding virtual surfaces. We implemented techniques to allow content manipulation and grouping. The system proposes far display space coupling and loose semantic coupling with a seamless transition between the physical and virtual spaces. Then, we compare our system with a wall display alone by conducting an experiment with pairs of participants on collaborative tasks. We observed that participants used the virtual space as a primary surface instead of the wall display. Overall, our results highlight that while there is an interaction overhead at using Wall+AR, there was no performance impact, and users preferred it over using only the wall display.

In [Chapter 5](#), we investigate the possibility to emulate a high-resolution wall display in Virtual Reality. We focus on the resolution aspect, with both a theoretical analysis and a pilot study. The results suggest that current VR HMD technology is not yet ready to emulate a high-resolution wall display.

In [Chapter 6](#), I come back on the contributions of this thesis, the insight for combining wall displays and AR headsets it brought us, and the future work this thesis opens.

Related Work

My work explores the addition of AR on and around large displays to compensate for their weaknesses and support their strengths. I organize this related work around two themes: the use and benefits of large displays and how previous works use AR to augment physical displays in general and large displays in particular. Related work specific to each chapter is available in the respective related work sections within the chapter.

2.1 Working with large displays

Large displays are devices with a screen space large enough for multiple people to be in front of, see the content, and view it at different distances and positions. In the large display category, we can differentiate between horizontal displays and vertical displays. Many forms of large displays exist, from low-resolution projectors and information boards to high-resolution wall displays and whiteboards. They have many applications both in the public space and work contexts.

2.1.1 Large display use-context & applications

We commonly come across large displays in public spaces: from time boards in train stations and airports to recent advertisement screens flourishing in the streets. All these examples show large vertical displays used as non-interactive devices in public. In all those examples, large displays are used to present information in public spaces. Many works explore the addition of interaction to these displays to increase engagement. In the case of vertical displays, past research has either adapted the content to passersby to ease the reading [193] or has enhanced these displays to enable interaction through touch [164] or gestures [151]. Tabletops are also used in public spaces. We see them used for information search and planning in tourist information centers [143] or in museums to accompany the content and add interactivity [50, 113]. In this public context, large displays, vertical or horizontal, are public devices viewed and potentially manipulated by everyone.

We observe two different kinds of large displays with different purposes. First, low-resolution displays, in the form of either projectors or TVs, are used commonly for slide presentations in teaching or seminars. Beyond simply presenting slides, these somewhat low resolution displays have been used in the past as awareness displays in different contexts. For example, in teaching, Kreitmayer et al. [116] used them in a classroom context to help students track their activities. And Gong et al. [76] used a translucent projected screen as a support for teaching and added the possibility for students to interact with it. And such awareness displays are often also present in command and control centers providing awareness of current operations to monitor air traffic [219], road-traffic [166], crisis management [120], etc.

The second type of large displays are high-resolution displays (close to a regular monitor). They allow for detail and precision work when viewers are close to them. High-res TVs, tabletops, or wall displays are used to show a large amount of data at once, to show content closer to its scale, or to have multiple people interact in detail in front of them. For example, in automotive design, Buxton et al. [37] explains how the industry explores various applications of large displays, from digitalizing design methods like wall tape drawing to presenting virtual car models. More recently, Okuya et al. [161] have shown the potential to reduce the back and forth between model designers and 3D creators with browsable small multiple visualizations on the wall display. High-resolution wall displays are used to analyze data in the humanities [174], and in visual analytics for example for text analysis [6] or biomedical visualization [10]. They are also used to simulate the environment and add immersion, either for therapeutic [190] or training in simulations [52].

2.1.2 A large vertical display or horizontal display?

We can separate large displays between vertical and horizontal displays. While both types of displays can accommodate multiple users at once, some differences exist in how multiple people collaborate and use space. For example, Inkpen et al. [92] reported that users would work in a more efficient way on vertical display than on a horizontal display, by presenting less preparatory actions and pointing actions to the display.

Territoriality. Scott et al. [196] studied how, on a tabletop, users create different territories with different purposes, like storing items or sharing them. Azad et al. [12] noted that while we find similar interactions, during collaboration sessions, between tabletops and vertical displays, there are differences in where the different territories (Personal, Shared, and Public) appear. For example, in walls the territories differ in height/size, may move, and the perceived limits of these territories are more fuzzy. Rogers and Lindley [188] also noticed a difference in strategies between groups working on tabletops and groups working on vertical displays. They found that groups working

on a tabletop were more cohesive than groups working on the vertical display. However, later research pointed out that the limitation to one person interacting with the display and its size might cause this difference [99]. Jakobsen and Hornbæk [99] also pointed to frequent interferences in the personal space occurring during sessions which could indicate that territories were less spatially separated in the case of wall displays than in tabletops.

Collaboration. Like territories, the way people collaborate together changes between vertical and horizontal orientation. Collaboration coupling or as defined by Tang et al. [209] as “*the manner in which collaborators are involved and occupied with each other’s work*” shows changes when people collaborate in front of a display instead of around the side of display. Jakobsen and Hornbæk [99] showed that work on a multi-touch wall display and a tabletop have similarities in the collaborative coupling styles of the users during collaboration. Even though they did notice that close distance equals close collaboration coupling in wall displays, later research showed that in wall displays we could have close collaboration coupling with far distance [134].

Perception. Vertical and horizontal large displays also have differences at the perceptual level. Indeed a large vertical display allows users to stand at different distances from the display, either to view a particular section of the display or to have an overview of the content. Visual distortion can occur when the viewer is looking at a remote item or an item at an angle, both with tabletops [234] and wall displays [27]. However, when for tabletops the only solution to decrease this distortion is to come closer to the remote item, users of large vertical displays can take a step back to reduce the distortion due to view angles [27]. Also, in tabletops, content orientation is important: items read from a specific orientation, like text, need to be positioned facing the user. This requires either creating copies of content or rotating it during collaboration [117]. For large vertical displays, content orientation is not an issue as users are always facing the same way.

All these differences point out that while we can find inspiration in research results for tabletops, we cannot directly transfer the knowledge from horizontal displays to vertical displays. In our thesis, we focus on large vertical displays, these can both accommodate multiple users, and can be seen at very different distances. From now on, Large displays will refer specifically to Large *vertical* displays, and related work on tabletops will be explicitly mentioned.

2.1.3 Large displays advantages and drawbacks

Large and wall displays are used for different purposes than traditional "all-purpose" desktop computer, as they have different characteristics and thus different advantages and disadvantages. We discuss here past studies that have explored the benefits and drawbacks of using such large displays.

Advantages. Due to their high resolution and large size, large displays can render in one go an amount of data yet unmatched by other devices, beyond what traditional desktops can visualize [240]. The ability to display large amounts of information in full also impacts performance. As Ni et al. [156] reports, a large screen makes the user "less reliant on wayfinding aids to form spatial knowledge.". In the same vein, Andrews et al. [6] suggests that during sense-making tasks, that extra space allows analysts to better externalize their memory, organizing content spatially in a way that expresses different semantic layers. This large space also allows users to multitask by laying out multiple application windows on the same screen [28].

Another implication of their significant size over traditional desktops is that these large displays encourage users' physical navigation, i.e. moving in front of the display. Ball and North [20] found that physical navigation leads to better performance than having just a wider field of view. Wall displays have also been found to improve performance against virtual pan and zoom navigation during visual analysis tasks [21] and classification tasks [135]. When compared to a desktop during productivity tasks, Czerwinski et al. [55] reports that using a large display improves performance and is also preferred by users. And Reda et al. [178] even suggests that using large high-resolution displays leads to finding more insights during visual exploration sessions and improves engagement and exploration behaviour.

Large displays can accommodate multiple persons and groups in front of them in the case of vertical displays [12] or around them in the case of tabletops [196]. Users can work collaboratively or in parallel, with a fluid transition between different group dynamics and collaboration coupling [99].

Drawbacks. However, large displays also have limitations. Endert et al. [66] highlight that wall displays: "are seldom used to complete everyday work-related tasks.". The main reason they invoke is that they are challenging to start and set up and require custom applications. To be more versatile, they would need some physical flexibility to change between different configurations for different use, between a curved horseshoe disposition more suited to a single user, to a flat disposition that can accommodate many people. They also mention that future work is needed to see how these spaces can transition from personal workspaces to large-scale collaborative spaces. This lack of flexibility also impacts the amount of data displayed. Some advantages previously mentioned implying that no virtual navigation is needed because the physical boundaries

of the device can contain the whole quantity of data. Thus the rigidity and complexity of a wall display makes it challenging to increase their size and configuration when the visual content to render surpasses their rendering capabilities.

Shared large displays are often used by multiple people simultaneously, collaborating or not, which raises privacy concerns [208]. All the content is visually accessible to everyone, which restricts the use of personal/private information. Although new technological approaches like Parallel Reality Displays [57] are a promising possibility to offer personal views depending on the user's position, they lack interactivity. Moreover, as the content is shared, everyone can interact with the content and alter the shared view, potentially disrupting other users.

Overall, this past work has clearly proven the benefits of using wall displays over desktops, in particular in collaborative settings. Nevertheless, it has also identified two challenges that are tackled in this thesis: the lack of flexibility in extending/altering their configuration, and lack of privacy.

2.1.4 Enhancing wall displays with devices

The main means for interacting with large displays is touch. But other devices, for example mobile phones or desktops, can be used to (i) interact with large displays, and even to (ii) display additional content.

Input. The main interaction method for large displays can be considered to be touch. We find touch interaction for tabletops [86, 237], wall displays [134, 164] and whiteboards [47, 118]. Another way to interact with large displays is by using tangibles. A broad part of interaction with tangibles implies tabletops, as it is easier to have tangibles on a horizontal surface. There are examples of using tangibles on tablets for vertical displays [103] but at a distance, and recently, for on-display interaction, Courtoux et al. [51] proposed tangibles specifically for vertical displays.

However, a user needs to step back from the display to have an overview of the content. And using touch interaction, they can suffer from attention blindness, working on a specific part of the display without paying attention to the global picture. Also, touch interaction or tokens used over long periods of time can cause fatigue and pain, especially on large vertical displays [3]. This impact particularly wall displays, which reach the largest size. Thus several researchers have studied ways to interact at a distance. To interact at a distance, pointers used by presenters during slide shows are the main metaphor. Mice and trackpads can be used as pointers to perform basic interactions [73, 154]. As more elaborate trackpads, Smartwatches [106] can also be used like pointers, offering a "hand-free" alternative. Tablets [43, 103, 217], or smartphones [9] also allow different interactions and more expressiveness, relying less

on pointing directly to the wall display and more on indirect manipulation through the touch surface.

Output. Another way to take advantage of additional devices is to display content related to the wall display on these available screens. As Rashid et al. [176] states, there are three ways to connect semantically display outputs independently from their relative positions: clone, extend, or coordinate displays. When applying this with a large display, the main application of cloning the wall display to other devices is to support interaction on this large display either synchronously [18] or asynchronously [46].

Most examples of combining large displays with other screen devices are around coordinating displays. One way is to have one screen complement the large display to provide a focus + context system. Prouzeau et al. [169] developed a system where the desktop screen provides a zoomable view of a section of the wall display. Another possibility of such a combination is to show an alternative view of the wall display. Kister et al. [110] propose, in addition to a focus+context system, different alternative views of the graph displayed and rendered on a tablet. Horak et al. [87] use the smartwatch personal aspect to use it as a personal space to store and filter visualizations taken from the wall display in a visual data exploration setting. On tabletops, tablets can be used to display a private view of all or part of the content on the shared device [146, 241]. Apart from being mobile and light, and thus easy to carry when moving around a large display, tablets and smartphones can also be used for spatially-aware interactions. For example, Langner and Dachsel [121] use an external tracking system to calculate the distance of the tablet from the wall display, and they visualize on the tablet different data layers depending on this distance. Recent work even explores AR-HMD devices to offer new options to what the wall display can provide. As this is relevant to our work, we discuss it in more detail in [section 2.2](#).

The above work uses external devices mainly for interaction or for private views of information. Few cases consider external devices explicitly for extending the real-estate or configuration of large physical displays. While some frameworks exist to allow such configurations [18, 221], to our knowledge, no studies investigate the impact of this combination. A way to add more space around physical large displays that are hard to reconfigure, is to use HMDs, as we will discuss next.

2.1.5 Summary

Large displays and tabletops are devices that can accommodate many users, to interact with or view their content. They have many applications in the public space [86, 113, 143, 151, 164, 193] and the workspace [37, 52, 76, 116, 161, 190].

However, a rift exists between vertical and horizontal large displays. Tabletops and large vertical displays have different properties, such as orientation, the position

and distance of the user in front of the display. While studies' findings and research questions from one display type can inspire studies in the other, their results are not necessarily transferable. For example, in studies done on individuals, there are clear differences in perception between tabletops [234] and wall displays [27]. Similarly, in collaboration settings, groups show more fluid transitions in coupling styles on wall displays than on tabletops. They also show less clear territories [12, 99], which translate in different work strategies. In our work, we are studying collaboration when augmenting the wall display with AR, and focus on large vertical displays.

Large displays mainly use touch to interact but can allow distant interaction through a (mid-air) mice, laser pointer, mobile touchpads, mobile devices, mid-air gestures, or even body position. Additional device screens can be used to complement the large displays' output. Rashid et al. [176] states that this can be done by cloning, extending, or coordinating the display output. We see examples of studies about cloning or coordinating large displays with other screens but not of extending them.

While large displays are convenient ways to display a large amount of data, incite physical movement, and allow fluid coupling transition during collaboration, some significant drawbacks remain. Large displays are physically restrained, are expensive and complex to setup, and lack privacy.

2.2 Screen combined with AR

Physical screens are devices limited by their size: they can render a large amount of visual content but will never be able to show more than what their size allows. As AR places content in the space, it allows extending or adding information to those spatially constrained devices. We identify three ways to add AR to displays, depending on the coupling between the physical display and the AR content. A close display space coupling corresponds to superimposing information on the physical screen, while extending the physical screen with AR is considered far display space coupling. A decoupled relationship is when virtual displays replace the physical screen.

2.2.1 Early work

The early work combining physical displays and AR started by extending the physical screens. For a single display, Feiner and Shamash [69] proposed to increase the space available on a desktop and use the low-resolution AR as an information space while keeping the high-resolution display inside this zone as a workspace. For multiple displays, Rekimoto and Saitoh [182] use AR to create a continuous interface between different displays, represented by projection on top of the physical space. With another approach, using see-through technology instead of projection, Butz et al. [36] uses AR as a *"3D multimedia ether" that envelops all users, displays, and devices.* They

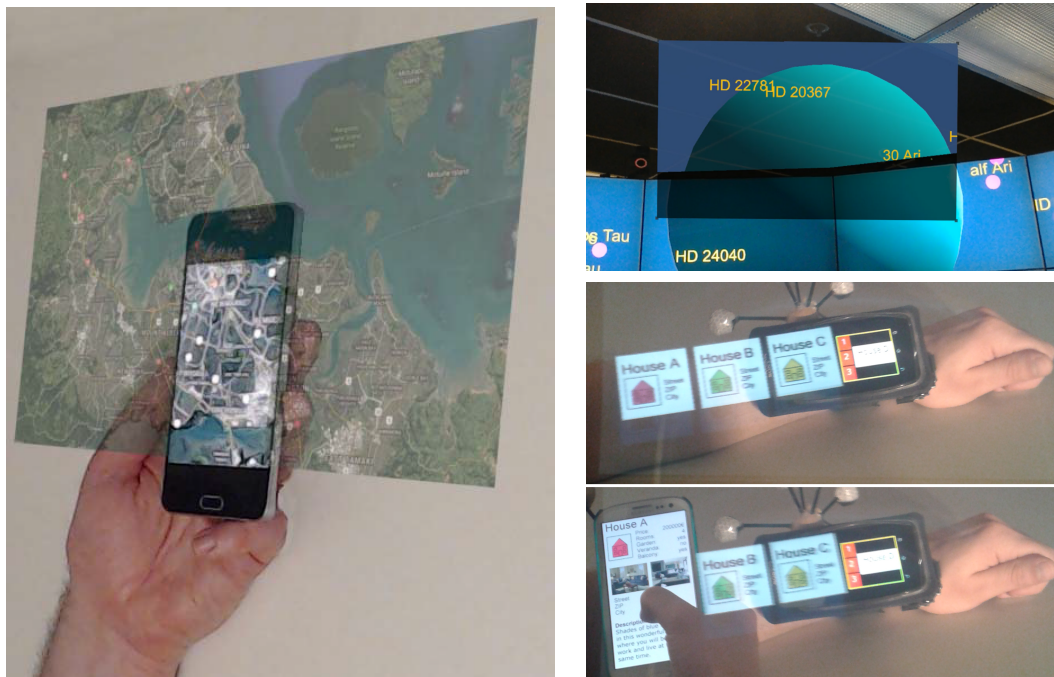


Figure 2.1: On the left, smartphone screen extension mockup for the VESAD project [160], where the AR could show other opened applications. On top right, extension of a CAVE2 system to show out of the screen view in AR by Nishimoto and Johnson [159]. On bottom right, the MultiFi prototype[78] where the AR extension displays more information about the visualization on the smartwatch.

extend the typical mechanism of standard GUI WIMP interfaces by enabling icons, menus, and pointers to be represented and interactive in a 3D environment. Their systems answer a need to bridge collaboration and personal devices.

2.2.2 Extending a display

Current works mainly divide themselves between augmenting a single display versus augmenting multiple displays. An extension can be used on a single screen to increase the size of small displays like smartphones [160] or/and smartwatches [78]. For example, a way to use the additional AR space is to offload controllers from the screen. Brasier et al. [33] show that 2D selection and, to some extent, sliders can be brought off the screen to solve occlusion issues when the interaction happens on a small smartphone screen. Along the same lines, we can transfer content from a smartphone to the AR space, and manipulate and control this AR content, and oppositely bring content from the AR space to the smartphone [191, 242].

When used with multiple displays, AR can create new visualization dimensions. Langner et al. [123] use AR to create visualizations on, around, between, and in-

between tablets. It can also be used in large display environments [159] to address the physical vertical limitation of a stereoscopic wall display.

Other works link different displays together like the virtual ether described by Butz et al. Benko et al. [25] provides a 3D view of an item on a surface by grabbing said item with your hand. Wilson and Benko [235] present a system where Spatial Augmented Reality (SAR) is used in conjunction with projected screens to transfer files across displays. Another possibility is to take advantage of the physical world by projecting a slice of a 3D visualization both on a screen and on a tangible to show its original shape [214].

AR can also be used to place virtual functions impacting physical objects near them [131, 198]. The most common use seen in single and multiple display environments is to create new views for the same data. Wang et al. [230] use AR as a supplementary desktop screen to an existing one to show a complementary 3D visualization. With the same purpose, Cavallo et al. [39, 40] integrate AR-HMD in room-sized multi-display environment to provide more ways to interact with the content (gaze, air-tap, voice) and to have 3D interactable visualization in addition to the 2D content from the other displays.

2.2.3 Superimposing on a display

The second possibility, superimposing AR on a screen, proposes either more information on the context or a different view for the same data. Spindler and Dachsel [201] use a sheet of paper where content is projected on top as a magic lens for the content displayed on a tabletop. Another use of AR on top of a display is to display specific content to specific users. Nilsson et al. [158] and Prytz et al. [172] use AR to provide personalized data visualization on top of a map used as the primary visualization. They aim to reduce clutter and improve collaboration in command and control contexts with operators of different cultures. For large vertical displays, Vinayagamoorthy et al. [222] offload sign language subtitles from the TV that would usually take space over the content. Kim et al. [109] use AR to add an interactive layer to static visualizations by displaying additional information to the data points. For wall displays, Sun et al. [205] have the same approach for knowledge graphs by having additional information attached to the nodes on the wall display, but also by giving access to private information only to the concerned headset wearer in AR. Our work also focuses on graph visualization. But it differs in two ways. First, we do not add graph information, our AR additions are rather assistance aids to help navigate the network displayed on the wall display. And second, we present a controlled experiment to empirically validate these AR additions.

Some systems propose physical screen extension and superimposition. One good example by Reipschläger and Dachsel [180] extends a surface screen by offloading

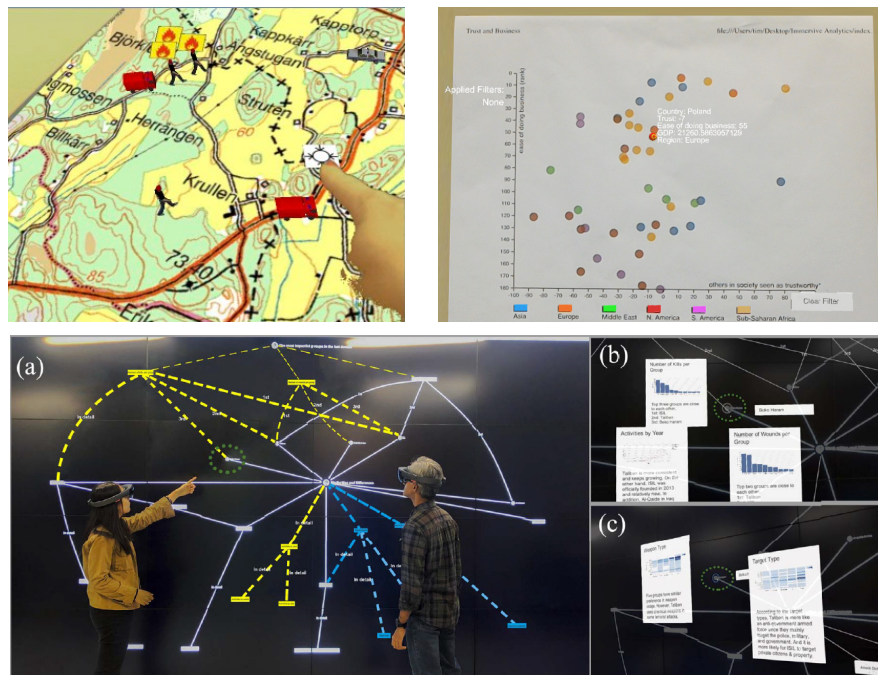


Figure 2.2: On top left, addition of domain-personalized glyphs on top of a map during a crisis management scenario by Nilsson et al. [158]. On top left, use of AR on a piece of paper to display an adjustable scatterplot. On bottom, Sun et al. [205] use of AR to deliver further information on an insight map shown by a wall display.

controllers on the side of the surface to keep the display as a pure drawing surface and adding a side view of the 3D result of the designed item. They also add the 3D rendering of the drawn object directly on the surface superimposed on the display content. The display size is ideal for superimposing information on wall displays, like adding unique data points on a map [139] or creating annotation and magic lenses [181]. However, both add the extension by showing alternative visualization of the same data, placed on the sides or in front in a separate AR layer.

2.2.4 Alternative to a display

The last possibility to add AR to displays is to replace the screen. With an interactive display, like a tablet [8, 88, 197], tabletop [35], or surface [180], the display role is taken over by AR, while the screen remains to interact with the environment. López et al. [137]’s work use the screen to provide a stereoscopic environment, with the support of the tablet as a controller. For desktops, Pavanatto et al. [163] discuss replacing the physical desktop screen with an AR one. Another way to replace physical screens, closely related to those works and worth mentioning, is to use Virtual Reality. Lee



Figure 2.3: Alternative screen missing figures: on top replacing a wall display by a VR HMD application [132], on bottom left, replacing additional desktop screen in AR [163], on bottom right, replacing large displays by VR HMDs [124]

et al. [126] provide a system with movable vertical surfaces, which could have been large displays. Lisle et al. [132] do document manipulation for sensemaking in VR using a HMD, with an initial setup comparable to a Large display setup [6] or wall displays [134].

2.2.5 Summary

Combining screens with AR is not a new concept [36, 69, 182]. In early works, we identified three ways to add AR to a screen: *Extend*, *Superimpose*, and/or *Replace*

the screen. To *extend* the screen, we can use AR to offload controllers [33, 180] or content [242] from the main display, to create content in between devices [123, 235], or to directly extend the screen space [25, 78, 160]. We can *superimpose* AR on a physical screen, to show an alternative version of the content [109, 181, 201] or add more information [139, 158, 172]. We can also use this proximity to the content to display information privately in context [205]. To *replace* a display with AR involves taking over the screen's role as a display. It comes either as a way to keep its interactive role [8, 35, 88, 180, 197], or to mimic existing displays [126, 132, 163]. These different kinds of addition hint at ways to compensate for large displays' shortcomings, such as the need to extend them and to deal with their privacy.

This work mainly focuses on technological contributions and very few actually study the impact of these additions to collaboration. Notable exceptions are studies on tabletops. Nilsson and Johansson [157] who showed that their tabletop and AR system was preferred over a paper map and helped reduce the information clutter coming from collaboration of different work roles. And Nilsson et al. [158] that showed a lower eye contact during collaboration sessions when wearing a HMD. Other studies on collaboration around physical displays augmented with AR do not study the impact of adding AR, in other words they do not compare the system they present with a baseline system such as the physical display alone.

2.3 Placement of this thesis

Large vertical displays are used for diverse application fields to study a large amount of data, and to accommodate multiple users working in parallel or collaborating. However, they have some limitations: their public nature can cause privacy issue and conflict between users, and they are limited in display space and configuration. Work on combining AR and physical displays more generally, has showed different ways to add information to a physical screen, information that can be private and placed in space without constraints. However, few studies exist on using AR with a wall display in particular.

The work of this thesis aims to explore how to combine AR and a wall display system to resolve some of the previously mentioned limitations of large vertical displays. More specifically, it studies different extremes of coupling the AR content with the content of the wall, both when it comes to display space coupling (i.e., placement of the AR content with regard to the content on the physical wall display) and semantic coupling (i.e., the semantic connection between AR and physical display content).

Large displays are shared spaces used in public and work spaces, with content viewed by every person looking at it. This can cause issues when using personal or private content or conflict between users interacting with the display's content. In our [Chapter 3](#), we study how to answer the lack of privacy of large displays. We propose

a solution, using the superimposition of AR on a large display to offer a private view closely coupled semantically to the public context. Other work [181, 205] also adds personal information on graph visualizations using AR, such as controllers, additional properties, or additional visualizations. But our work differs in two ways. First, we do not add graph information per se, our AR additions are rather assistance aids to help navigate the network displayed on the wall display. And second, we present a controlled experiment to empirically validate these AR additions. We discuss this related work more in [Chapter 3](#).

We have also seen earlier that large displays are physical screens limited in size, which can be complex to reconfigure spatially in the case of wall displays. In our [Chapter 4](#), we explore ways to extend large displays using AR HMD, an approach that has attracted little interest so far. Many works we presented before that extend screens using AR HMDs do it on smartphones [160] or tablets [123], but do not consider larger displays like wall displays. We chose to study the augmentation of the working space horizontally as opposed to [159], where the AR extends the display vertically. This represents a loose semantic and display coupling. We also use the large display room to add multiple surfaces to extend space and study its impact on collaboration.

Finally, there is an emerging body of work that emulates inside VR headsets work environments that resemble a collection of large and wall displays. This past work does not compare these VR headsets alternatives to physical wall displays, but it does raise an important question. Is it possible to replace wall displays by emulating them in VR HMDs? We study this question in [Chapter 5](#), focusing on the issue of resolution.

Personal+Context navigation: combining AR and shared displays in Network Path-following

This chapter is based on our paper "Personal+Context navigation: combining AR and shared displays in Network Path-following" [102] published in GI '20: Proceedings of Graphics Interface. This work was done in collaboration with Monash university. The co-authors for this paper are Anastasia Bezerianos, Olivier Chapuis, Maxime Cordeil, Tim Dwyer and Arnaud Prouzeau. Supplementary material for this project are available at https://gitlab.inria.fr/ilda/PersonalContext_Data and the video at <https://hal.archives-ouvertes.fr/hal-02860063/file/ARVizGIVideo-truefinal-1.mp4>.

Large displays accommodate multiple users at the same time, working either in parallel or collaboratively. This may cause conflict between users, as a change to the display view made by one user may impact another user's viewing experience. An AR HMDs can display information exclusively to its wearer, thus it could provide the personal view needed on a large display. In this chapter, we study how to provide a personal view, closely coupled in display space to the large display, and tightly coupled semantically to the display's context. In this chapter, we will refer to "display space coupling" as "display coupling" for brevity.

3.1 Introduction

Shared displays are well suited for viewing large amounts of data [6, 135, 174, 178] and for accommodating multiple users simultaneously [99, 134, 215]. Shared displays exist around us in different environments, such as dedicated analysis environments [122], command and control centers [150, 166], and public spaces such as metro stations or airports.

While a shared display provides a common view, users often require a private view to




Figure 3.1: On the left a user with a HoloLens navigating a network shown on a shared display, moving their head from left to right. On the right their personal view in the HoloLens at the start (top) and end (bottom) of their movement. The augmented content consists only of the white visuals connecting the headset center of view (cursor) to a link of the network on the shared display (*SlidingRing* shown, red marks added for illustration).

work independently, access sensitive data, or preview information before sharing it with others. To this end, existing work combines shared visualization displays with private views using dedicated devices such as desktops [72, 169, 195, 229] mobile phones [224], tablets [110] and watches [87]. However, these separate views are typically decoupled in display, making it difficult for users to maintain a connection between private and shared views [195].

To avoid such divided attention, recent work in immersive visual analytics [142] has combined shared displays with Augmented Reality (AR) headsets. These approaches consider the shared display and the personal AR view as having a close display coupling [109, 205]: the shared screen provides the context visualization, while the AR display superimposes private information. We investigate this combination in a specific setting.

We focus on publicly shared node-link *network* visualizations with close display coupling, with personal views and interaction in AR, in order to support navigation (Figure 3.1). Application scenarios abound: from control rooms where individual operators [169] reroute in AR their resources using a shared traffic network map as context; to private AR route finding on a public transport network map. Figure 3.2 illustrates such a scenario: multiple travelers focus on the same public information display, but may be interested in different aspects of the transportation network and do not want their personal route interests (e.g., their way home) advertised to onlookers. In these situations, AR can provide a personal view with navigational support that is tailored to the user's route priorities and preferences.

We explore two navigation techniques (with two visual variations for each), rendered only in the personal AR view, that aid the wearer in following their chosen route



Figure 3.2: Personal AR route on a public map at the metro station providing geographic context. While the public map is visible to all (including people without headsets), each traveller with a headset sees their own navigational aides. The personal augmentations do not disrupt other users of the map and they give the traveller privacy.

in the network visualization shown on a shared display. The techniques use only AR view-tracking as a means of user input, as gesture recognition and hand-held devices are not supported by all AR technologies and may be awkward to use in public settings [185]. Our hands-free interaction techniques help the viewer maintain a visual connection between their personal AR view, that may shift (for example due to small head movements), and their preferred route on the network on the shared display.

In two experiments on path following of different precision, we studied how different types of display coupling mechanisms between the personal AR view and the shared network display alter user performance. Our results show that persistent display coupling works well for high-precision path-following tasks, where controlling the view through the AR headset is hard; while more flexible transient display coupling works best for low-precision tasks, in particular when following paths of (personal) high weight. More globally, they demonstrate the feasibility of having the context visualization on an external display and a personal navigation view of the AR headset.

3.2 Related Work

Network and shared/collaborative visualization are well-studied topics. We focus next on the most relevant work, namely navigation and path following on network node-link diagrams (which we informally refer to as “graphs” in the sequel), the exploration of such graphs using shared displays, and the combination of AR and 2D displays.

3.2.1 Navigating Paths in Graphs

Navigating large graphs with dense links is difficult without assistance. Some solutions focus on a specific node and allow the user to explore neighboring nodes, gradually expanding their inspection of paths. Van Ham and Perer [220] use the concept of Degree of Interest to select a subgraph to present to the user and then extend it gradually. And May *et al.* [145] and Ghani *et al.* [74] use marks to highlight off-screen nodes of interest. Moscovich *et al.* [149] introduce techniques that use the topology of the graph to aid navigation, for example Bring-and-Go brings neighbors of a focus node close to it for easy selection.

All these techniques cannot work as-is in AR: although the personal view may include less elements, it still cannot override reality (the network on the shared display), adding clutter. More relevant to our work is the second technique from Moscovich *et al.* [149], Link-Sliding: after one of the neighbors is selected, the user can slide their view along a link to reach the original location of the neighbor. One of our techniques adapts this metaphor for AR settings. In our case the view of the user depends on their head orientation and cannot be forced, so instead of sliding the entire view we help the user maintain a connection to the link they are sliding on.

In a related approach, RouteLens [5] takes advantage of the topology of the graph and assists users in following links by snapping on them. Our second technique is inspired by the transient snapping [5], nevertheless we augment it with appropriate feedback to guide the users' navigation in AR.

Another family of neighborhood inspection techniques involve “lenses”. EdgeLens [236], curves links inside the lens without moving the nodes in order to disambiguate node and link relationships. Local Edge Lens [216], shows in the lens only the links connected to nodes inside the lens. PushLens [194], pushes links that would transit through the lens away instead of hiding them. In MoleView [89] the lens hides links depending on specific attributes, or bundles and unbundles them. Most of these lenses combined are seen in MultiLens [111]. Lens techniques cannot be easily applied to our case, as the video-see-through AR needs to overwrite reality, a technology not yet ready for real-world use. As with lenses, our techniques are visible in the constraint area seen through the AR display, but visuals are closely matched with the real-world content.

3.2.2 Interacting with Graphs on shared large displays

The use of shared displays for multi-user graph exploration is first demonstrated in CoCoNutrix by Isenberg *et al.* [93], where users interact using mice and keyboards. Lehman *et al.* [129] use proxemics (i.e., the distance between the user and the display) to indicate the level of zoom requested. Prouzeau *et al.* [168] compare two techniques to select elements in a graph using multi-touch. Both techniques impacted the shared workspace, potentially disturbing co-workers visually.

Lenses could limit the impact area of such interactions. For road traffic control, Schwartz [195] proposes a lens that gives additional information regarding the route on a map. Kister *et al.* [112] use the user's body position to control a lens for graph exploration. However, these lenses do not work well if two users work on the same area of the shared display, and do not address privacy issues raised in all-public workspaces. A solution is to provide users with an additional private display. Handheld devices can show detailed views [44], display labels [186], or used for interaction [110]. While such devices avoid disturbance of others and allow for privacy, they also divide the user's attention between displays [195] - the user has to match the content of their handheld device with the context on the large display, which can be cognitively demanding. Eaddy *et al.* [61] proposed voice instructions coming from an earpiece to guide the visual navigation of the user on a map. This solution monopolizes hearing, which may block sounds from the outside which can be dangerous in a public environment. On the other hand, AR overlays can seamlessly combine private information with that on the shared display, avoiding divided attention. We leverage this setup to assist users in graph navigation.

Most current AR headsets provide head-tracking capabilities (e.g., HoloLens-1¹ or Nreal²), which is what we use in our study. Previous work has used eye-tracking as a means to analyze graph visualization and navigation [155], but not as the input mechanism for navigation. Eye-tracking is a promising alternative to hand pointing for interacting with AR content [210] and for brief interactions on public displays [107]. While some recent AR headsets do support eye-gaze tracking (like HoloLens-2³ and MagicLeap⁴, this is not yet available across the board.

3.3 Design Goals

For the remaining of this chapter, we use the term "shared context" to talk about the visualization displayed on the shared external display and "personal view" for the AR

¹<https://docs.microsoft.com/en-us/hololens/hololens1-hardware>

²<https://www.nreal.ai/>

³<https://www.microsoft.com/fr-fr/hololens/hardware>

⁴<https://www.magicleap.com/en-us/magic-leap-1>

view shown to the user.

A key aspect of network navigation [84, 203, 223] is following paths of interest, while exploring their neighborhoods. Our techniques focus on aiding path following using an AR headset, while maintaining a visual link to their context (neighboring nodes and links) that exists in the network on a shared display. This gives rise to our first design goal:

G1. Tying tightly personal view to context. Contrary to situations where the AR view is related loosely to other displays or is a stand-alone visualization [14, 24], our techniques need to be tightly coupled semantically to the network on the shared display. This is similar to the need to tie labels to scatterplot points in previous work [109] or to visually link virtual elements to real surfaces [171]. For graph navigation, this means the techniques need to match closely the visual structure of the underlying network on the shared display. For example any additional information on a specific node or link rendered in AR, needs to be visually connected to the shared display representation of that link or node. In our designs we vary this connection, with both *persistent* and *transient* connection variations.

We opted for hands-free techniques, as in public settings hand gestures can be awkward [185] or reveal private information (e.g., final stop on the metro map). Or they may be headset-specific or even not possible with some technologies, e.g., [68]. On the contrary, any setup that assumes a display coupling of a public and private view requires view-tracking technology, that we can leverage for interaction. Nevertheless, view-tracking relies on head movement that can be inaccurate and hard to control, a situation that can be exasperated by the limited field of view of some headsets. This inspired our second design goal:

G2. Handling limited field of view and accidental head movement. The limited field of view of current AR headsets compounds the challenge of tying the external context display, that is often large, with the personal AR view that can cover only part of the shared display (that shows the network). This can hinder navigation, for example it is easy to get lost when following a link that starts and ends outside the field of view (especially in dense graphs). Unintended head movements aggravate this issue, as they can easily change the field of view. Thus techniques need to be robust to head movements and changes in the field of view, reinforcing the visual connection between shared and personal view. We vary how this connection is reinforced, either providing *simple visual links* to the shared view, or a *deformed re-rendering* of a small part of the shared view to render visible information that may be outside the field of view.

We note that existing AR technology superimposes a semi-transparent overlay on the real world. As such, it is still technically hard to "overwrite" the view of the real world – in our case the shared display. Thus, the additional visual aids we bring to the personal view cannot consist of independent renderings, nor of a large amount of additional information, as the underlying visualization will continue to be visible. We

thus avoided variations that require rendering large quantities of content in the field of view (e.g., remote nodes [216] or off-screen content [74, 145]), or that contradict the shared display (e.g., curving links [236] or removing them [89, 194, 216] since they will still be visible on the shared display).

3.3.1 Personal Navigation

When considering the motivation of having personal views on a network, we notice that path following priorities may differ depending on users' expertise and preferences. For example when looking at a public metro map, different travelers have personal route preferences (different destinations, speed, scenic routes, etc.). Thus one aspect that can be personal in network navigation is the notion of *weights* that different paths can have, depending on viewer preferences. These weights can be indicated *visually* in the AR headset (in our case as link widths), and can be taken into account in the *interaction*. Our techniques to aid path following and navigation, take into account personal path weights.

We expect that users have added their general preferences (scenic routes, high bandwidth, close to hospitals) in an initial setup phase of the system (e.g., when using it for the first time). Our techniques do not require the user to input specific/intended paths, only general preferences that many possible paths in the network can fulfill. It is then up to the user to choose and follow a specific/intended path, including ones that do not follow their general preferences. Thus our techniques do not assign weights, but make use of them.

3.4 Techniques

Our designs work without external devices, other than the head mounted display. In our setup the center of the personal view (referred to as *gaze-cursor*) acts as a virtual cursor. To navigate the graph the viewer moves their head, and thus their personal field of view, the way they would normally do when standing in front of a large shared display. Our designs do not alter the context graph visualization seen on the shared screen.

To facilitate the connection between the personal view and shared context view (*G1*) we explore two variations to aid path following: a *persistent* connection that snaps onto paths of the network permanently (Sliding), and a *transient* variation that attaches to the network temporarily (Magnetic). From now on, a subpath is a path made of links (or a link) connecting two nodes of degree 1 or > 2 , where all the intermediate nodes, if any, are of degree 2, i.e., have a single ingoing and outgoing link (the degree of a node is the number of links that are incident to the node).

3.4.1 Persistent connection: *Sliding Metaphor*

In these variations, when a viewer is following a path on the network (i.e., one or more links), a ring gets attached to the subpath. When they change their field of view by moving their head, the ring remains attached to the subpath (that exists in the shared display), but slides along it, thus maintaining a persistent connection with the selected subpath ($G1$). Sliding is inspired by Link Sliding [149] developed for desktops, that moves the viewport of the user to follow the selected path. In our case the viewport is controlled by the viewer's head (that in turn controls their gaze-cursor), while the graph is anchored on the shared display. So we cannot force the headset viewport to follow the path. Instead, we use the metaphor of the sliding ring that gets attached to the selected subpath, and add a visual cue in the viewer's personal field of view ($G2$). This visual connection can be a simple visual link (*SlidingRing*), or a deformed re-rendering of a small part of the shared context visualization (*SlidingElastic*).

Sliding Ring can be seen in Fig. 3.3 and 3.4. We draw a dashed trail from the center of the personal view, to the ring that slides on a path. This ring is the projection of the gaze-cursor on the path and it slides along it as the user moves their gaze. The trail guides the viewer's eye back to their selected subpath, even if that subpath is outside their AR field of view ($G2$).

Sliding Elastic can be seen in Figure 3.5. In this variation, the subpath in the personal view turns into a curve, pulled like an elastic towards the center of view of the user. The ring now slides on this curve. The curve can incorporate more than one node (and links). With this variation, the ring is always on the viewer's gaze-cursor (not simply attached to it via a trail), due to the deformed path curve that remains within the personal view. The end-points of the elastic curve remain tethered to the context graph ($G1$) when the viewer moves their head ($G2$), while the curve is a deformed copy of the subpath from that graph that follows the user's view. All nodes on the deformed (copied) subpath on the personal view are connected visually with dashed trails to the original nodes on the graph on the shared display. This sliding variation is more complex visually, nevertheless, as it copies the subpath, it maintains information about the local structure of the subpath in the personal view.

3.4.1.1 Interactive Behavior

For the sliding techniques we use the notion of subpaths, groups of one or more links where there are no branches and thus it is clear where the ring needs to slide to next. The ring slides on the selected subpath until it reaches the node at the end. At this point it gets attached to a new link (that may be the beginning of a longer path), depending on the direction of movement of the viewer's gaze-cursor.

At first we considered simply selecting the link closest to the center of the user's field of view. However, we found it often challenging to select between two links that




Figure 3.3: Example of a technique (SlidingRing) and its rendering. (Top-left) part of the shared wall display. (Bottom-left) Holographic content added by the HoloLens. (Right) Combined view seen by the user with the HoloLens on. Red border indicates the HoloLens field of view.



Figure 3.4: Close-up of SlidingRing (combined view). From left to right: as the user moves its field of view away from the link the ring is attached to, the trail (dashed line) guides the viewer back to the link, even when the ring is no longer in the AR view (right).




Figure 3.5: Close-up of SlidingElastic (combined view). From left to right: when moving away from the selected subpath, a curved copy of the path follows the user, even if the path is no longer in the HoloLens view. Again the view is the combination of the shared wall display view (graph) and the head-mounted display (elastic path and its connections only).




Figure 3.6: Choosing the link to slide into when the gaze-cursor is inside a node (cyan circle). From left to right: proxies (colored dots) are created for each link; proxies spaced equally; proxies rotated to minimize distances between proxies and original links (Procrustes analysis); slices of influence assigned to each link based on their proxies. The slice that the gaze-cursor exits on determines the link that the ring will slide to.

are close together (like top and left link in Figure 3.6). We thus provided wider zones of influence around the links, inspired by the interactive link-fanning approach of Henry Riche *et al.* [184] seen in Figure 3.6. We consider proxies of all the links around the node, that are then equally spaced around a circle. The circle is then rotated in order to minimize the distance between the proxies to the original link positions (Procrustes analysis⁵). The final proxy positions define zones (slices) of influence around them. These slices follow the relative direction of the original links, but are as wide as possible to ease gaze selection. When a link has higher personal weight for the viewer, it is assigned more proxies around the circle (equal to their weight), and thus wider slices.

While this algorithm determines the choice of the next link to follow, it does not alter the visual representation of links or nodes. We observed that this approach works well even when multiple links exit towards the same direction, as viewers tend to exaggerate their head movement in a way that differentiates the links (e.g., go higher for the top link). The performance of this approach will degrade for nodes of high degree (many links) – but in these situations any technique based on head tracking, that is hard to control precisely, will be challenged.

In the case of the *SlidingRing*, the viewer’s gaze-cursor can be either inside the node at the end of the subpath (as above), or further away maintaining the connection through the trail⁶. If the gaze-cursor is outside the node at the end of the subpath, we assume the viewer wants to quickly slide to the next link by roughly following with their gaze the direction of the path. We thus attach the ring to the link closer to the gaze-cursor, ignoring the fanning calculations mentioned above. This approach works

⁵https://en.wikipedia.org/wiki/Procrustes_analysis

⁶This does not happen in *SlidingElastic* because in the personal view the subpath curves to follow the gaze-cursor, thus the ring arrives to the node at the end of the subpath only when the gaze-cursor does.

well when the graph is homogeneous in terms of weights. Nevertheless, when the personal weights of the viewer are not uniform, we want to give priority to high-weight paths. These two goals (preferring the link with minimum distance and preferring one of high weight) can be conflicting, for example if a low-weight link is closer to the gaze-cursor. We assign each potential link around the node a value that thus combines their distance from the gaze-cursor and their weight. More specifically, we define two distances for each link p_i with weight w_i . The first distance $d_{p_i}^{gc}$ is the (Euclidean) distance from the gaze-cursor (gc) to link p_i . The second distance $d_{w_i}^{w_{max}}$ is defined as $d_{w_i}^{w_{max}} = (w_{max} + 1 - w_i)$. In this term, w_{max} is the highest weight in the graph, thus $d_{w_i}^{w_{max}}$ becomes smaller the higher the weight w_i of the link. We refer to this distance as *distance from max weight*. The value 1 is added to our calculation of distance from max weight in order to avoid the ring always sliding to the highest weight link (where $w_i = w_{max}$), disregarding the distance of the gaze-cursor. To make our final choice of path, we combine both distances in $v_i = d_{p_i}^{gc} \times d_{w_i}^{w_{max}}$ (to be seen as the inverse of link attraction/influence) and chose the link with smallest v_i .

3.4.1.2 Sliding Metaphor Summary

The sliding techniques anchor the personal view to a subpath on the shared display, maintaining a consistent visual link with that path ($G1$). They are thus well suited for following a path closely during navigation. Due to their persistent display coupling to paths, they are robust to changes in the field of view due to head movements ($G2$). *SlidingRing* consists of a simple dashed trail connecting the path and the viewer's field of view. *SlidingElastic* is more visually complex, providing a deformed copy of the path and its local structure, that is stretched to remains in the user's view, while the nodes at the ends remain anchored to the context graph.

3.4.2 Transient connection: *Magnet Metaphor*

In the magnet metaphor, the gaze-cursor gets "magnetically" attached to links and maintains a connection in case the user accidentally moves their head ($G2$). Only one link is attached to the magnet at a time. When the gaze-cursor moves, the attached link can change if a better candidate is found within the magnet's area of influence, thus this connection to the links on the graph of the context display is transient ($G1$). The visual connection can be again simple visual links (*MagneticArea*) or a deformed re-rendering of a small part of the shared context visualization (*MagneticElastic*). The magnet metaphor is inspired by area cursors [105] and bubble cursors [77]. As the area of influence moves, candidate links inside the area attract the magnet. Contrary to the bubble cursor, the magnet resists detachment from the selected link to deal with small accidental head movements ($G2$), but eventually detaches in order to attach

to a new candidate. The area of influence can be customized.⁷ The final aspect of the magnetic techniques is providing a feed-forward mechanism to indicate risk of detachment from the current link due to nearby links. We consider attraction in an area twice the size of the area of influence to identify detachment candidates.

Magnetic Area is seen in Figure 3.7. Similarly to the Rope Cursor [81], we draw lines from the gaze-cursor to the link that the magnet is attracted to. This simple link connects visually the gaze-cursor to the original graph on the shared display. If no links are in the area of influence, no line is drawn. The notion of a subpath is not used here, since this technique can detach from the graph. We show the attraction of other links using semi-transparent rays that fade as they move away from the gaze-cursor. The bigger the attraction the more visible the attraction rays become, acting as a feed-forward mechanism to warn the viewer that they risk detaching from the current link. To avoid visual clutter, we limit the attraction rays to the top 5 candidates.

Magnetic Elastic is seen in Figure 3.8. Similarly to *SlidingElastic*, an elastic copy of the selected subpath is pulled into the personal view. Contrary to the sliding variation it is not permanently attached to the path and can be detached if other candidate link attract it. We communicate this attraction from other links by fading out the elastic copy when it risks getting detached from the current path and attached to another. We experimented with adding gradient attraction lines in this variation as well, but decided against them due to the visual clutter of the design. We note that for the *MagneticElastic* the notion of subpath exists, nevertheless all value calculations are done on the link of the path the viewer is on. By link we refer to the original link on the graph (and not the elastic copy).

3.4.2.1 Interactive Behavior

When more than one link enters the area of influence of the magnet techniques, the closest to the gaze-cursor is selected. As with the sliding techniques, when the personal weights of the viewer are not uniform, we want to give priority to high weight links when considering candidates. To do so, we assign a value v_i to each link l_i inside the area of influence. This value takes into account both the distance of the gaze-cursor from the link $d_{l_i}^{gc}$ and the distance from max weight $d_{w_i}^{wmax}$ (see Sliding behaviour). Nevertheless, the simple product $v_i = d_{l_i}^{gc} \times d_{w_i}^{wmax}$ is not enough for techniques that are not permanently attached to links. We want a small resistance when the magnet is already attached to a link, preventing accidental detachments, especially from high-weight links.

First, we reduce detachments due to crossing. As the viewer follows a link l_a , their gaze may cross over another link l_j resulting in a distance of zero for the crossed link. This would result in the cursor detaching from l_a and attaching to l_j . While

⁷We set it to 5 cm on our shared display, as we empirically found this distance to be a good compromise to avoid accidental detachments when crossing over other links.



Figure 3.7: Close-up of MagneticArea. From left to right: the magnet is attached to a link (on the wall display); as the gaze-cursor moves other link candidates are drawn to it and attraction lines become more visible; until the magnet detaches and reattaches to another link.



Figure 3.8: Close-up of MagneticElastic. From left to right: an elastic copy of the subpath is created; that follows the user even if original subpath is out of HoloLens view, nevertheless it fades out as the other links start attracting it; until it detaches and snaps to another subpath.

this detachment may be a desired behavior, we want to avoid it triggering too easily, especially for high-weight link. We thus add a term c_1 to our calculations of v_i that prevents the distance $d_{l_i}^{gc}$ from causing immediate detachment. The new $v_i = d_{l_i}^{gc} \times d_{w_i}^{wmax} + c_1 \times d_{w_i}^{wmax} = (d_{l_i}^{gc} + c_1) \times d_{w_i}^{wmax}$ makes high-weight links more attractive when searching for the smallest value v_i . This is how values are calculated for candidate links inside the magnetic cursor's area of influence.

Second, to give priority to any currently attached link (making it resistant to detachments), we make a special calculation for the attached link l_a . We introduce the term $c_a < 1$ that further reduces the value of the currently attached link. The final calculation for the attached link is $v_a = (d_{l_a}^{gc} + c_1 * c_a) \times d_{w_a}^{wmax}$. The term $c_1 * c_a$ is always smaller than c_1 ($c_a < 1$), thus reducing v_a with respect to the values of other candidate links inside the area of influence. We found $c_1 = 0.1$ and $c_a = 0.75$ worked well for our setup.

3.4.2.2 Magnet Metaphor Summary

The magnetic variations temporarily attach the gaze-cursor to a subpath of the graph on the shared display. The attachment prevents the user from accidentally losing the path because of small changes in the field of view (G2). AR visuals are again closely tied to the context graph on the shared display (G1). *MagneticArea* consists




Figure 3.9: Schematic representations of the two interactive metaphors and their visual variations tested in our experiments (the BaseLine techniques is not shown). Visuals in **blue** indicate content rendered inside the AR headset and **black** visuals indicate content on the shared display.

of a line connecting the gaze-cursor to the selected link and rays that fade out to provide feed-forward when the viewer risks detachment. *MagneticElastic* provides an elastic copy of the path that fades when there is risk of detachment. Contrary to the sliding techniques, the link between the gaze-cursor and subpath can be broken if the gaze-cursor moves away from the path.

3.5 Experimental Setup

As the external large display we used a 5.91×1.96 m wide wall display with a resolution of $14,400 \times 4,800$ pixels (60 ppi), composed of LCD displays (with 3 mm bezels) and driven by 10 workstations. To avoid participant movement, we only used the central part of the wall display, about 2×1.96 m. For the AR part we used a HoloLens, an optical see-through head mounted display. On the wall display we used a simple program to display images on demand and on the HoloLens side we used Unity [211]. Both are controlled by a “master” program that sends UDP messages to the wall display (e.g., change the graph to display), and to the HoloLens (e.g., change the technique and specify the path to be followed).

We calibrated the HoloLens with the wall display, using 3 Vuforia [91] markers rendered on the wall display. The markers are recognized by the HoloLens and used to calculate an internal place-holder for the graph (position, orientation and scale). Once calibrated, the gaze-cursor of the HoloLens is projected on the wall display to calculate the distances used by the techniques.

3.6 Experiments

Our four techniques *MagneticArea*, *MagneticElastic*, *SlidingRing*, and *SlidingElastic*, Figure 3.9, provide personal path navigation on a graph on an external display (*G1*), in particular under situations where the personal field of view may change due to head movement (*G2*). We chose to evaluate them under *path following tasks* [127], that are central in graph navigation and motivated their design. As the techniques aid path following (rather than path choosing), we focus on their *motor* differences, i.e., how differences in persistence and linking to the shared context visualization affect path following performance. We thus removed noise related to path choosing, by explicitly indicating to users the ideal path to follow. We also consider as a *BaseLine* the simplest path following technique: the gaze-cursor displayed in the center of the field of view of the headset. We note that the *BaseLine* is indirectly affected by path weights since higher-weight paths are rendered as wider in the headset.

Path Following Tasks. Our techniques are also designed to support different precision – the sliding techniques are very precise once on a subpath (as it is impossible to lose it), while magnetic ones are flexible and allow for quick corrections through detachments. We thus tested our techniques on path following tasks of various precision. To better understand the weak and strong points of our techniques, we decided to consider two extremes in path following w.r.t. task precision:

1. *Path Selection*: participants had to go through all the nodes and links of the path in a given order. A simple attachment (or touch/hover with the gaze-cursor for *BaseLine*) of nodes and links is enough to consider that part of the path selected. This is a low precision task that consists of a sequence of simple selections that a user may want to perform when identifying a path of interest, without necessarily being interested in all the details of the path. An example use-case could be to quickly plan a trip on a metro map or a bus network.

2. *Path Tracing*: participants traverse the entire path in detail, tracing nodes and the entire length of the links. If the user leaves a link, they have to return to the link and resume where they left off, until they trace the entire link. This is a high precision task that a user may perform if they are interested in details along the path they are following. An example use-case may include tracing a metro path on a geographical map to look at the specific areas it goes through, or following a road to check which parts have emergency stops or bike lanes. In real life, examples of path tracing tasks are common when networks are part of a geographic map (e.g., roads, infrastructure maps), where users are interested both in the details of the points of interest and in their context, as discussed by Alvina et al. [5] (work that partially motivated our magnetic behavior). Beyond maps, path tracing tasks are used in network evaluations when there are concerns that the visual continuity of paths in the network may be




Figure 3.10: Example of one path from the experiment, on the Small World graph. Inside the red rectangle is the user's view through the HoloLens, that includes the red path they need to follow. The rendering outside the HoloLens view (dashed lines) are added for illustration only.

affected, for example when they cross bezels [56, 62] or curved screens [199]. In our case, visual continuity of the path may be affected by the limited field of view of the AR display (design goal $G2$).

Together, these two tasks represent extreme cases in terms of the need for interaction precision when going through a route: simple selection vs. tracing/steering along the path.

Task Operationalization. In real use, our techniques do not know the exact path a user would like to follow, they only know *global preferences* of the user such as preference for scenic routes or speed (that are represented as weights, see Section "Personal Navigation"). To test hypotheses in a controlled experiment, we need to isolate/control aspects relevant to the hypotheses (operationalization of the task). In our case, we have formed working hypotheses around navigation performance under different precision tasks, which we will formalize later in this section. We thus chose to

remove factors that make trials incomparable, such as personal biases, preferences and decision making. In real-life, navigation requires decisions influenced by participant preferences, and thinking delays. To remove these confounds, we decided to show participants a specific route to follow in a given color (red). This ensured that only interaction time was measured (not the time to think while choosing); and avoided participants from selecting routes of different lengths based on personal preferences that could make trials incomparable across participants.

In our experiments, all experimental indications were shown in the personal AR view. As mentioned above, for experimental purposes, the path to follow was shown in red, with the starting node highlighted with a green halo (Figure 3.10). In the Path Selection task, whenever the participant selected a node or link in the correct order, they would turn green. In the Path Tracing task, the links were progressively filled in green up to the point the participant had reached, tracing their progress.

We separated our evaluation into two experiments, one for each of the two tasks (Path-Following, Path-Tracing).

Paths and Weights. As we discussed in the Personal Navigation section, link weights are an example of personal information that may be specific to different users. As such, weights are a fundamental aspect of the techniques we introduced. The techniques favor high-weight links, and thus should perform better when following paths with high weights. To evaluate this aspect we considered two types of paths (factor PATH). We note that in real use, multiple paths can fulfill user preferences/constraints and could thus have high weights. For operationalization purposes we test two cases. (i) We tested *Weighted* paths where the links of the path to follow have a constant weight of 3 and all the other links have weight of 1. These simulate cases where the path that the system has identified as having high weight (based on a-priori preferences), matches the path that the user wants to follow in practice. (ii) We also tested *Homogeneous* paths where all the links of the graph have a weight of 1. With these paths of equal weight, we simulate cases where the system has assigned multiple paths with the same weight based on a-priori preferences. If the weights selected by the system do not match the needs of the user, we assume the user will deactivate them (leading them to the equal weight situation). Thus we did not consider *Homogeneous* paths that cross *Weighted* paths (an unfavorable situation for the techniques), as in real applications we expect that it would be possible to simply disable the weights.

In the experiment, links are visually scaled by their weights, hence a high weight link is 3 times wider than other links.

Graphs. To be able to generalize our results, we considered two graphs with different characteristics (Figure 3.11). The first type is a “5-quasi-planar” graph⁸ (the Paris

⁸A topological graph is k -quasi-planar if no k of its edges are pairwise crossing [204].




Figure 3.11: Graphs used in the experiments: (Left) The Quasi-Planar graph of the Paris metro map, 302 nodes and 369 links. (Right) the small world graph, 180 nodes and 360 links.

metro map). None of the paths participants had to follow contained crossings for this type of graph. The second type is a small world graph generated using NetworkX⁹. We generated the small world graph to have on purpose a similar number of links to the quasi-planar graph (369 and 360 respectively), but a higher link density. Link density is defined as $numberOfLinks/numberOfNodes$. In the small world graph this ratio was 2, almost double that of the quasi-planar graph (1.2). Consequently, the quasi-planar graph has more nodes than the small world graph (302 and 180 respectively). We tested graphs of different density, as this may affect our techniques (e.g., cause accidental detachment in the magnetic techniques). The quasi-planar graph is a typical example of networks such as roads, electricity, or water networks. Small world graphs are typical in real phenomena, e.g., social networks.

The small world graph is more complex than the quasi-planar graph, because it has more link crossings, more links attached to a node on average (higher degree), and also has longer links (e.g., links between communities) that are more challenging to trace. Our aim is to see whether the structure of the graphs impacts the differences between the techniques.

In our experimental trials, all paths that participants had to follow consisted of 7 links (and 8 nodes) of similar difficulty for each graph and did not contain cycles. All the paths we used for the small world graph had one long link. An example path is

⁹NetworkX (<https://networkx.github.io/>): we used the `connected_watts_strogatz` generator for the network structure, and the `yEd Organic Layout`.




Figure 3.12: (Top) Mean time to complete the path selection task per TECH for each GRAPH \times PATH condition (the number in the bars is the mean, and the error bars show the 95% CI for the mean over all data points). The lower the bar, the faster the technique. (Bottom) The 95% CI for mean differences between all the TECH, for the respective GRAPH \times PATH condition. *Note: A CI that does not cross 0 shows evidence of a difference between the two techniques - the further away from 0 and the tighter the CI, the stronger the evidence. Not crossing 0 indicates a “significant” difference in the corresponding paired t -test, i.e., $p < 0.05$. To compare two TECHS X and Y for a given GRAPH \times PATH, see first the two corresponding bars in the top figure for this condition, and then the CI that corresponds to the pair $X - Y$ at the bottom.*

seen in Figure 3.10.

Working hypotheses. Given the design of the techniques and the nature of the tasks, we made the following hypotheses:

- (H1) For the selection task, magnetic variations are more efficient for *Homogeneous* tasks, since they do not force users to follow the entire path when trying to make simple selections.
- (H2) For the tracing task sliding variations are more efficient than magnetic since the user needs to trace the path without detaching. Magnetic ones are likely more efficient than *BaseLine* as they can help keep the connection to the path, especially for *Weighted* paths.
- (H3) All the techniques are faster with the *Weighted* paths than with the *Homogeneous* paths. The differences are more pronounced for the small world graph than for the

quasi-planar graph, since the larger density may cause detachments in magnetic variations and distraction in *BaseLine*.

3.6.1 Path Selection Experiment

In this experiment participants conducted a low precision path following task, selecting in order links and nodes in a path.

Participants and Apparatus. We recruited 10 participants (8 women, 2 men), aged 25 to 42 (average 29.6, median 27.5), with normal or corrected-to-normal vision¹⁰. Five participants had experience using an AR device, such as the HoloLens. Participants were HCI researchers, engineers, or graduate students. As apparatus, we used the prototype described above.

Procedure and Design. The experiment is a $[5 \times 2 \times 2]$ within-participants design with factors: (i) TECH: 5 techniques: *BaseLine*, *MagneticArea*, *MagneticElastic*, *SlidingRing*, *SlidingElastic*; (ii) PATH, 2 types of paths: *Weighted*, *Homogeneous*; (iii) GRAPH, 2 types of graphs: *Q-Planar*, *SmallWorld*.

We blocked trials by TECH, and then, by PATH. We counter-balanced TECH order using a Latin square. We also counter balanced PATH, for each order one participant started with *Weighted* and another with *Homogeneous*. We fixed the graph presentation order, showing the simpler *Q-Planar* first.

For each TECH \times PATH, participants started with 6 training trials, followed by 6 measured trials. After each TECH block they rested while the operator checked the HoloLens calibration.

Participants were positioned 2m from the wall display and were instructed to avoid walking, to maintain a consistent distance from the wall display and personal field of view, across techniques and participants. They were asked to perform the task as fast as possible. Our main measure is the time to complete the task, since paths need to be completed for the trial to end (i.e., all trials are successful). Time started when participants selected the starting node, and ended when all nodes and links were selected in the correct order. The experiment lasted around 1 hour, concluding with participants ranking the techniques and answered questions on fatigue and perceived performance.

3.6.1.1 Results

We analyze, report, and interpret all inferential statistics using point and interval estimates [54]. We report sample means for task completion time and 95% confidence

¹⁰None of our participants had red-green colorblindness, but if replicating this work other experimental colors can be considered [231].

intervals (CIs), indicating the range of plausible values for the population mean. For our inferential analysis we use means of differences and their 95% confidence intervals (CIs). No corrections for multiple comparisons were performed [53, 165]. We also report subjective questionnaire responses.

Completion time. We removed one obvious outlier (a trial with a standardized residual of 18, while all others < 4). We did not find evidence for non-normal data. Figure 3.12 shows the mean completion time for each TECH grouped by GRAPH \times PATH (top) and the mean differences between TECH (bottom).

None of our techniques outperformed *BaseLine*. We see that *BaseLine* and *MagneticArea* exhibit, overall, the best performances with very similar mean task completion times across conditions (no evidence of difference). Both magnetic variations perform better than sliding ones when it comes to *Homogeneous* graphs, with strong evidence of this effect for the more complex *SmallWorld* graphs (partially confirming [H1]). Nevertheless, for *Weighted* the simple sliding variation *SlidingRing* performs well and even outperforms the complex elastic magnetic variation *MagneticElastic*.

We have evidence that *MagneticArea* and *SlidingRing* are always faster than their elastic versions *MagneticElastic* and *SlidingElastic* respectively. This is particularly clear for magnetic variations across the board, and for all cases except *Homogeneous SmallWorld* for sliding variations.

We observe mean times for each technique tend to be faster for the *Weighted* paths. However, we only have conclusive evidence they are indeed faster for *SlidingRing* (CI [1.4, 3.9]) and *SlidingElastic* (CI [2.0, 6.8]) for the *Q-Planar* graph. And for all the techniques except *MagneticElastic* for the *SmallWorld* graph (e.g., *BaseLine* CI [0, 1.7], *MagneticArea* CI [0.7, 1.9], *MagneticElastic* CI [-1.3, 2.7], *SlidingRing* CI [8.2, 15.5], *SlidingElastic* CI [3.2, 11.2]). This partially confirms [H3].

Questionnaire. Although in terms of time both *BaseLine* and *MagneticArea* performed similarly, the best rated technique is *MagneticArea* (average rank: 1.6). *BaseLine* is next (ar: 2.4), followed by *SlidingRing* (ar: 3.0), *MagneticElastic* (ar: 3.2), and *SlidingElastic* (ar: 4.3). Participants preferred *MagneticArea* and *SlidingRing* over their elastic counterparts.

We have very similar results when comparing responses for their perceived speed, accuracy and the easiness of use of the techniques (Figure 3.13). *MagneticArea* was always better rated than the other techniques, with *BaseLine* coming second.

Summary. The most preferred technique is *MagneticArea*, even though it objectively performs similarly to *BaseLine* (not confirming [H1]). *SlidingRing* also exhibits good time performance for weighted paths. There is evidence that with few exceptions, techniques performed better in *Weighted* paths (partially confirming [H3]). Finally, the




Figure 3.13: Results of the questionnaire of the selection experiment for perceived speed, accuracy and easiness of use of the techniques.

elastic versions performed worse than their non-elastic counterparts, possibly because the elastic variations have more visual clutter, and require users to go through the longer elastic subpath.

3.6.2 Path Tracing Experiment

In this experiment participants conducted a high precision task, tracing over each link and node in an indicated path.

Given the results of the first experiment we excluded the elastic version of the techniques. We thus consider only 3 techniques in this second experiment: *BaseLine*, *MagneticArea*, *SlidingRing*. We follow closely the design of the first experiment and consider the same two graphs and two types of PATH.

Participants. We recruited 12 new participants that did not participate in the first experiment (3 women, 9 men), aged 22 to 43 (median 26), with normal or corrected-to-normal vision. They were graduate students and engineers. Most had already used an AR or VR headset (five had used a HoloLens).

Procedure and Design. We used the same procedure and design as the previous experiment, but with 3 techniques. The experiment lasted about 45 minutes.

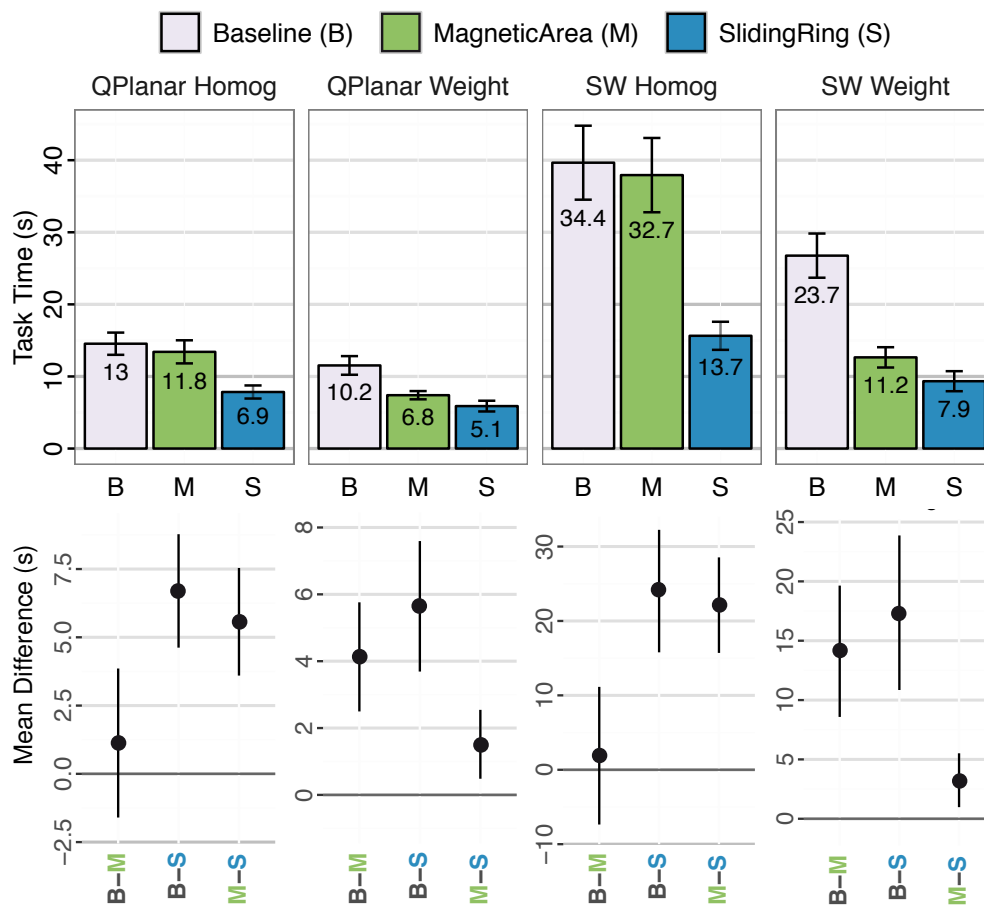


Figure 3.14: (Top) Average time to complete the path tracing task per TECH for each GRAPH \times PATH condition. (Bottom) The 95% CI for mean differences (12 points, one by participant) between all the TECH for the respective GRAPH \times PATH condition. See Figure 3.12 for reading CIs.

3.6.2.1 Results

Completion time. We removed two outliers (standardized residual > 4), and we could not observe any strong evidence for non-normal data. Figure 3.14 shows the task completion times for the task (top) and the 95%–CI for the difference in mean for the TECH by GRAPH \times PATH conditions (bottom).

We have strong evidence that *SlidingRing* is always faster than *MagneticArea* and *BaseLine*, with large differences with *BaseLine* overall, and with *MagneticArea* for *Homogeneous* paths. Moreover *MagneticArea* is clearly faster than *BaseLine* for *Weighted* paths, but this is not the case for *Homogeneous* paths. Thus, hypothesis (H2) is supported for the most part. The poorer performance of *MagneticArea* for *Homogeneous*




Figure 3.15: Results of the questionnaire of the tracing experiment for perceived speed, accuracy and easiness of use of the techniques.

paths is probably caused by accidental detachment when the gaze-cursor came close to a link not in the path (e.g., accidentally attaching to links that cross the path).

When comparing the performance of each TECH over the type of path, we observe that they are always faster (clear evidence) for the *Weighted* paths than for the *Homogeneous* paths (*BaseLine* for *Q-Planar* CI [1.8,4.1] and for *SmallWorld* CI [7.9,17.8], *SlidingRing* for *Q-Planar* CI [0.5,3.3] and for *SmallWorld* CI [3.5,8.9], *MagneticArea* for *Q-Planar* CI [4.0,7.9] and for *SmallWorld* CI [17.8,32.4]). This effect is particularly strong for *MagneticArea*. These results support (H3).

Questionnaire. Overall the results of the questionnaire are consistent with the results on time. When asked to rank the techniques based on preference, 9/12 participants ranked *SlidingRing* first, one ranked *MagneticArea* first, and one ranked *BaseLine* first. *MagneticArea* was generally ranked as the second choice (for 9/12 participants). While *BaseLine* was mostly ranked last (8/12).

We observe similar results for perceived speed, accuracy and ease of use of the techniques (see Figure 3.15). *SlidingRing* was judged faster, and easier to use than the other techniques. And responses for *MagneticArea* tended to be more positive than *BaseLine*.

Summary. *SlidingRing* exhibits the best results (objective and subjective) for path tracing tasks, partially confirming [H2]. The results for *MagneticArea* depend on the nature of the path, and might be a good choice for *Weighted* paths. As expected *BaseLine* exhibits poor performance. Finally, *Weighted* tasks were faster [H3].

3.7 Discussion and Perspectives

Our results show that depending on the goal of the users, different techniques are appropriate. In an explanatory path-following task where precision may be less impor-

tant, e.g., searching for a specific node or searching for the end of a link, a-priori no specific technique is needed (*BaseLine*). But viewers tend to prefer a snapping mechanism (*MagneticArea*) that is more forgiving to small view changes. When precision is important, for example when a well-determined path needs to be followed closely, a technique that is closely display coupled to the graph like *SlidingRing*, and to a lesser degree a snapping technique like *MagneticArea*, are clearly better. This is particularly true for personal paths that have higher weight, as sliding variations can follow them without risk of detaching (for instance, *SlidingRing* can be three times faster than *BaseLine*).

The elastic technique variations performed poorly. However, they can potentially be interesting when the viewer has identified a path of interest and needs to keep its local structure in their field of view. For example, in a situation where they may want to see parts of that path together with other locations on the graph. In the future we will perform a longitudinal study to observe the long-term situational use of our techniques. As new generation AR headsets are equipped with eye-tracking capabilities, we can utilize eye-tracking analysis as a means to better understand the true focus of participants when using our techniques in practice (as has been done in the past for understanding differences in graph visualization techniques [155]).

One implication of our work is that in a real system users will need to fluidly switch between techniques depending on their goal. More generally, we expect detailed tasks will require close display coupling between private and shared views, whereas in more coarse exploratory tasks this display coupling can be transient. We plan to investigate this technique transition in the future.

We expect that since our techniques address accidental head-movements (*G2*), they can also be applicable in newer AR headsets that use eye-gaze instead of head-tracking (eye-tracking being considerably more noisy and prone to small movements [70, 97]). Nevertheless, this requires empirical validation.

Although path following is a common task in graphs [127], in the future we plan to explore how the techniques, and their transition, fair in complex and high-level exploration tasks in networks. For example, they may be combined with additional interactions for graph navigation and exploration, such as filtering, relayout, etc. To support a larger set of interactions we may need to consider combinations of alternative input devices (e.g., clickers or smartphones) that can provide additional input. Moreover, we expect our findings to hold for other applications of steering-type tasks that are common in HCI literature [1], but this needs further study.

In our work we consider that viewers may have their own preferences for traversing certain paths. We model these personal link preferences with a simple weighting of each link. As the AR headset used is stereoscopic, it is tempting to use 3D to show such link weights (e.g., height above the display surface [238]). Indeed, graph visualization (such as node-link diagrams) in stereoscopic immersive environments has been proven

useful under certain conditions [4, 16, 119, 232]. Nevertheless, using depth or other 3D cues to indicate weight is not straightforward in our case, where the personal view is very closely coupled semantically to the context graph in the external display. In our early attempts we saw that the 3D copies of subpaths with high weight may get rendered further away from their original paths on the shared display, or artificially overlap other paths. This creates a discontinuity between the personal and shared view and requires further consideration.

In light of our promising results, as technology improves more techniques could be adapted for Personal+Context navigation in node-link diagrams. Obvious candidates are magic lenses [29], fanning [184], and Bring-and-Go [149]. For instance, when over a node, the user can trigger an adapted bring-and-go that brings neighbours into the AR field of view (rather than on screen [149]). Then the user can select a neighbour, resulting in an indication in the HoloLens that points towards the direction where the selected node can be found (since the viewport cannot be forced).

Our results show that having personal views tied to a shared external visualization can aid graph navigation. The fact that these personal views are always tied to the shared display, means they can be directly applied to a multi-user environment [205]. The shared graph remains visible to all users, and their personal preferences and views are private to their headset, not impeding the view of others. We next plan to investigate collaborative analysis settings, where colleagues may manipulate the graph on the shared display from their private view (e.g., scale it, move nodes, etc.). This may have several implications, such as change blindness [26] when one's AR view overlaps, but is out-of-sync, with these changes.

This raises the more general question of mismatch between reality and augmentation. Our techniques have been designed to add simple visuals, that are closely display coupled to the graph on the large display. Our decision is in part based on technological limitations of AR headsets (that can still not completely overwrite reality). Nevertheless, we feel this limited augmentation of reality, and limited movement for interactions, is appropriate in public settings (for example use in navigating public metro maps), given the recent discussions on isolation, acceptability and ethical concerns of using head-mounted displays in social settings [80].

3.8 Conclusion

We present two empirical studies on using personal views in augmented reality, that are tightly coupled semantically to a shared visualization on an external display. We consider a node-link network as the shared context visualization, and use the AR view to display personal weights and to provide feedback to aid navigation. This approach could allow several users to navigate the shared graph, receiving personalized feedback, without visually cluttering the shared visualization.

Our hands-free techniques are designed to help viewers maintain their connection to the shared network visualization, even when their personal field of view changes due to small head movements. Results show that our adaptation of the link sliding technique, that is closely display coupled to the shared graph, can bring a substantial performance improvement when precisely following a path. And that a technique using a magnetic metaphor performs well and is preferred over a standard gaze-cursor for a simpler path selection task.

More globally, our work shows that Personal+Context navigation, that ties personal AR views with external shared network visualizations, is feasible; but that the nature of interaction and visual support needed to maintain this connection depends on task precision.

Evaluating the Extension of Wall Displays with AR for Collaborative Work

This chapter is based on a paper done with Anastasia Bezerianos and Olivier Chapuis, and accepted at CHI 2023: "Evaluating the Extension of Wall Displays with AR for Collaborative Work" [101]. Additional material is available online at <https://ilda.gitlabpages.inria.fr/arviz/>. It includes the companion video, the source code of the prototype, documentation, as well as a web application allowing to replay the sessions of the experiment.

In the previous chapter, we saw we can add information to a large display, closely coupled in display space to it and tightly semantically coupled to its content. Large displays, and particularly wall displays due to their size, need space around them to be constructed and setup; a space that is not used to display information. In this chapter, we investigate a virtual display space far from its primary display, by adding display space in AR around a wall display. We study the impact of this far coupled display space and loose semantic coupling of content on the users behaviour in front of a wall display.

4.1 Introduction

Wall displays are well suited for collaborative work as they can accommodate multiple people simultaneously [17, 99, 112]. 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. 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.

Existing work on multi-display environments has looked at extending wall displays in




Figure 4.1: 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 display is a physical display in the room.

terms of input and output. Traditional multi-display environments, such as command-and-control rooms, combine various displays (wall displays, desktops, and digital tabletops) that each serves a particular purpose: wall displays and tabletops are commonly used for group awareness and collaboration, and tablets or desktops for personal work [45, 146, 169, 228]. Other work combines wall displays 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 display, and direct touch is not possible [43, 87, 218].

Recent work has started to combine wall displays with augmented reality (AR) headsets for visual exploration. Here wall displays show publicly core data visualizations, whereas AR headsets add personal information virtually, either on top or around the wall display content, in the form of additional visualizations [181], information [205], or highlights [102]. This work points to the potential benefits of adding AR head-mounted displays (HMD) to wall display environments, especially in the use-case of displaying private information. Some of this work, e.g., [125, 181], 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 space truly

used in practice, 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 4.1](#)). 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 an empirical study that compares this extension with a wall display alone.

Results from our comparison highlight 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 display alone, and we did not measure any loss in performance despite the interaction overhead.

4.2 Related Work

Wall displays are extremely useful collaborative work environments. They have been found to improve performance [21], content organization [135], sensemaking activities [6], and have been shown to increase discovery and improve data analysis both in laboratory conditions [178] and real work settings [174]. These displays have a high resolution that allows them to render a very large amount of information [23, 65], as well as a field of view that surpasses what AR headsets alone can achieve today [40]. And, of course, multiple people can see and use a wall display together [17, 99, 112] 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 already covered in [Chapter 2](#) approaches that use AR to augment physical displays. In this section, we examine studies on collaborative interaction for data manipulation or

sensemaking, focusing, in particular, on space use and vertical surfaces such as wall displays.

4.2.1 Studies on Co-located Collaborative Interaction

There is a large amount of work on collaborative interaction with tabletops. This work includes work on collaborative strategies [95, 209] from tightly-coupled collaboration (e.g., using sequential strategies) to loose collaboration (e.g., using parallel strategies); and territoriality [196], 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 use-contexts, users move around to take advantage of the environment, which is not the case with tabletops. In particular, because users move, there are not always clear territories [12, 31], and such environments need to support fluid transitions between loose and tight collaboration [99, 134].

In the last few years, there has been work studying wall displays, in particular, and more recently, AR and VR environments. We review next relevant work that considers collaborative studies close to our study, conducted either with wall displays or with AR or VR headsets. 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. [12] 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, focusing on interaction on wall displays. They found that wall displays, like tabletops, should support public, personal, and storage territories, but that personal space must be refined. The results differ in Jakobsen and Hornbæk [99]'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 display into territories. Liu et al. [133] 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 [134], where participants can use cooperative gestures. Finally, Sigitov et al. [200] studied pairs in two tasks, observing collaboration coupling and territoriality. Their system uses a curved wall display with smartphone-supported interaction. They suggest more types of territories than in pre-

vious works, on and in front of the wall display. They also observe participants dividing the task spatially among themselves, working in parallel.

With VR HMDs, Lee et al. [126] studied how groups of 3 participants solve visual analytic problems in a room fully simulated in a VR headset. 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, were never negotiated, and that "participants never entered a territory of another unless for tightly-coupled work".

Finally, in the case of Augmented Reality in a room that is either empty or contains furniture, Luo et al. [138] 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 [12, 99, 133, 200] and consider situations where AR is added not only in physical rooms [126, 138], 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.

4.3 Prototype

In this section, we present our prototype that combines a wall display with several synchronized AR Headsets. The 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 play-back 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 metres, 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 for our system, testing techniques both 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.

4.3.1 Virtual Elements

Our prototype contains three main types of virtual objects: BOUNDARIES, SURFACES and CARDS (seen together in Figure 4.1). 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 would 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 [196], physical walls [12], and furniture using AR [138], 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. 4.4), 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 4.2-left). This allows users to group CARDS and perform operations on them (detailed in Sect. 4.3.2). We chose to render these virtual surfaces in a way that resembles the physical wall display to convey to users 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. Empty surfaces can also be deleted. At any time, users can re-arrange 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.



Figure 4.2: 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 (created for illustration purposes).

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 virtual zone resembles a belt made of the items stored inside (see Figure 4.2-middle and right). Elements in the *Personal Space* are placed in a circle around the user like a semi-cockpit [67], 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 [196] 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, SURFACES can move 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.



Figure 4.3: Visual feedback on the cards and menus, all rendered either on the wall or in AR. On the left: the focus highlight that is added in AR to the card when the user gazes at it (personal content not seen by others). In the middle, the Card selection states, visible by all users: the card in its initial 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 the card to the *Personal Space*, to move all content of the surface towards another surface, or all content towards the *Personal Space*.

4.3.2 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.

4.3.2.1 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 HoloLens I 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 I hand-gesture recognition is supported by our prototype, in our experiment, we choose to use a clicker because using hand-gestures in front of the head is particularly tiring [32, 85]. To improve awareness of others' actions, we represent all user cursors as telepointers [83] in the shape of a cross, color-coded to be unique for each user. The color assigned to the user is used both for their cursor and their selections (see Sect. 4.3.2.2).

4.3.2.2 Card Selection & Movement.



To support content organization, users can select and move content, which in our case is represented by Cards (see [Figure 4.3](#) 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 *CARDS* selected by the user. 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 [2, 7]: 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 goes back to its initial position. We rely on social protocol to deal with interaction conflicts [148], enforcing a simple coordination mechanism: if multiple users select the same card, the last person to select it has ownership.



4.3.2.3 Surface Creation & Movement.

A user can create a virtual *SURFACE* by clicking on a virtual button that is always following 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.

4.3.2.4 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 4.3](#)). 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 re-arrange the content of shared virtual surfaces (and the content of the wall display).

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*

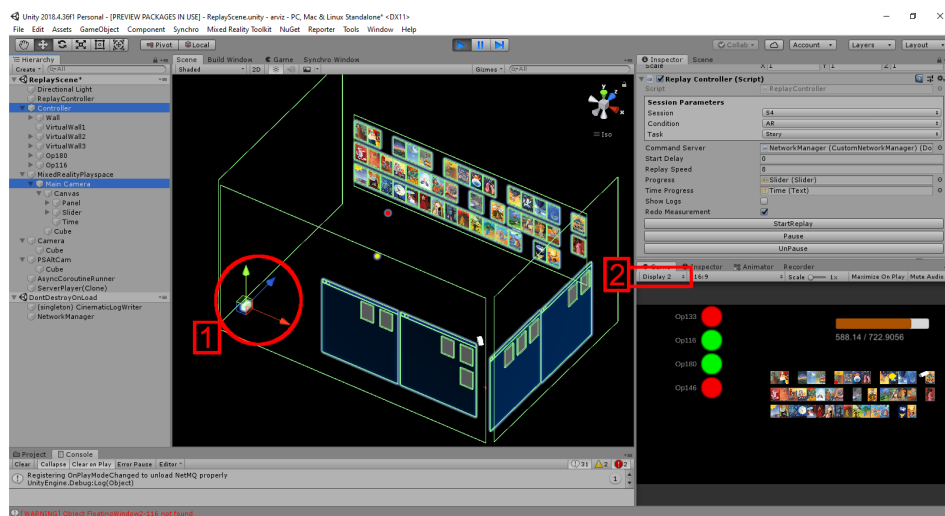





Figure 4.4: View for the Replay tool. In the middle, we see a schematic of the entire scene, where the viewer can chose to adapt their view of the scene with a Camera placed in the scene (see red 1, added for annotation purposes). On the top right (Inspector), the viewer can change parameters for the playback, pause/play the scene, and chose 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.

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 chose 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 content from the *Personal Space*.

4.3.3 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 4.4. This tool, a script for Unity 3D, can record and replay the log of messages sent between the wall display and the AR HoloLenses. By opening the dedicated scene and

selecting the generated log file, the entire session will be replayed, showing changes in user position and their 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.

4.4 User Study

Our study aims to understand if and when it is helpful to extend a wall display with virtual spaces in the collaboration context. 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 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.

4.4.1 System Conditions

To answer these questions, we built a system that increases the space available for users in front of the wall through the use of AR headsets (see Sect. 4.3). 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 4.1](#).

We used the same basic input functionality, relying on the HoloLens head-cursor and clicker, for both conditions WALL and WALL+AR (see Sect. 4.3.2). 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 space use and collaboration, without introducing a bias that may stem from different input capabilities and/or discomfort in wearing the headset in some conditions only. This choice is driven by our desire to ensure experimental consistency, but it is also justified by research trends. First, in our situation, one can imagine hand ray-casting as an interaction alternative, nonetheless using the head-cursor and the clicker is a less tiring and is an efficient technique when no high precision pointing is needed [114, 153], which is the case in our tasks. Second, when it comes to wearing a headset to interact with a physical display, AR HMD are becoming increasingly lighter and similar to vision glasses worn everyday, see, e.g., [42, 96, 128].

4.4.2 Tasks

In each condition, we asked pair of participants to conduct two different tasks inspired by previous work (see Sect. 4.2.1): A classification task that could be performed using loose collaboration and a storytelling task that enforce tightly-coupled collaboration.

In the classification task, participants have to make grouping decisions and select, move and sort picture cards, a task similar to past studies investigating space use [134, 138, 196]. We requested pairs to group 54 image cards from the popular game Dixit¹, that has been used in the past in wall display studies [134]. 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. [31] 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 space between them. To try and stimulate both types of collaboration, we thus introduced a second task, a storytelling 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 the 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 storytelling task). For both tasks, participants are presented with the wall display covered by cards (Figure 4.5). The number of cards (54) was chosen so that all cards were fully visible, but so that the wall display was purposely crowded. This was to simulate situations where the available wall display 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

¹libellud.com/nos-jeux/dixit/



Figure 4.5: Wall display at the start of the experiment, showing one of the image datasets.

observing their strategies when dealing with a crowded display when only the wall display is available.

4.4.3 Hypotheses and Measures

As we set out to understand the impact of extending a wall display environment with AR and studying the use of space, our study is largely observational [134, 138, 196]. 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 space 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 display will remain the central working surface for two reasons: (i) the content is on the wall display when the tasks start, so it is natural to continue working on it; and (ii) because the wall display 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 space use between the wall display alone and the extended virtual environment.

RQ3 *Does the addition of AR affect collaboration?* 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 when solving a task (in meters) 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 of fatigue; *Time* to complete the task, as a measure of cost. Finally, we elicited *Subjective feedback* in the form of a Likert scale questionnaire on: efficiency, ease of use, mental and physical fatigue, etc., that were in part inspired by NASA-TLX and adapted to our research questions.

4.4.4 Participants & Apparatus

Participants *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, either 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.

4.4.5 Experiment Design & Procedure

Design: The experiment is a within participants design with one factor, the system condition, with two values: WALL and WALL+AR. The presentation order for the system conditions was counterbalanced across pairs. We fixed the task order: pairs of participants always start with the classification task and then run the storytelling task. As our second task requires users to study the content of the pictures, we used the same datasets between the two tasks. 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 (in 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.

4.5 Results

We first discuss the general collaboration and space use strategies adopted by pairs, 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.

4.5.1 Collaboration Strategies and Workspace Usage

To analyze the collaborative strategies, we used the replay tool and notes taken by the operator during the experimental sessions. The supplementary material (Section 1 of Appendix A) shows the virtual screenshots of the final results for each task and pair.

4.5.1.1 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




Figure 4.6: 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 display and virtual spaces for classification (G0), and a pure virtual space classification strategy that does not use the wall display at all (G4).

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, and the other pairs used arbitrary-shaped blocks to organize images (see Figure 4.6-top). Indeed, the parallel pairs decided where to place each color before even starting moving 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 display in half, and each partner then sorted the colors into lines on their side of the wall.

The non-parallel pairs split the wall display in 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 display to create space on the sides and then started building the color groups.

4.5.1.2 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 display being used for the remaining colors. See Figure 4.6-bottom.

All pairs (except G10, see below) started the task by creating two surfaces and



Figure 4.7: Examples of a final result in story task for the WALL condition (G1, left), and the WALL+AR condition (G3, right)

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 display). 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.

4.5.1.3 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 display, 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). This third surface might not be optimal in terms of interactions, but 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.

4.5.1.4 Wall - Story

As expected, in this task, all pairs worked in close collaboration. Ten pairs chose together the cards 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 display (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 display to make room to create and present the story (recall that the task starts with the wall display 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 4.7-left](#) shows an example of a resulting story.

4.5.1.5 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 display to place them on the surface to create the story and present it to the operator (see [Figure 4.7-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 display towards a newly created "trashcan" surface and then moved their story onto the wall display 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 display to a "trashcan" surface (using their personal space) and moved back the story-selected cards to the wall display 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.

4.5.1.6 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 the use of space. 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.

4.5.2 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.

4.5.2.1 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.

4.5.2.2 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 display 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. Surface deletion was 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 task 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.

4.5.2.3 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 display 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.

4.5.3 Additional Objective Measures: Partners and Cards Traveling, Position, and Task Time

We report next a set of objective measures: distance, position, and time.

4.5.3.1 Participant Position and Distance Traveled.

At the beginning of the tasks, the partners positioned themselves side-by-side in front of the wall display close to the center, one slightly on the left and the other slightly on the right, at a distance of about $3m$ of 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

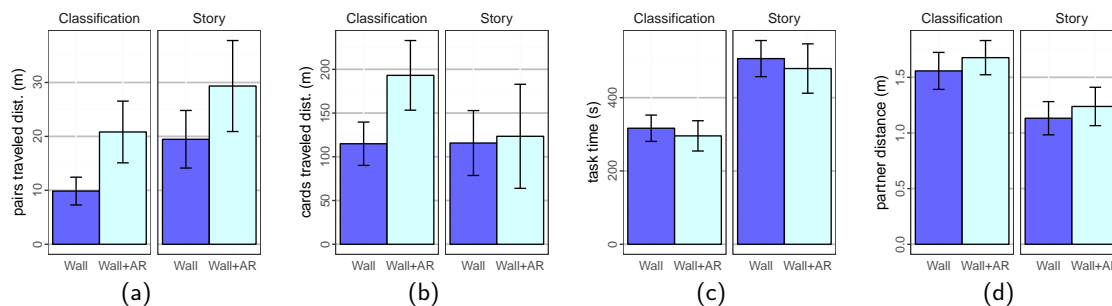


Figure 4.8: Average of (a) the traveled distance by both partners, (b) the traveled distance by cards, (c) the task time, and (d) the average distance between the partners (by condition and task). Error bars show the 95% CI.

and occasionally inverted their relative position in front of the wall, especially in the WALL+AR condition. (Section 2 of Appendix A).

Figure 4.8-(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.

4.5.3.2 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 4.8-(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 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.

4.5.3.3 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 4.8-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.

4.5.3.4 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 [99], for instance. Figure 4.8-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 expressed by our two tasks (tight for story and loose for classification).

4.5.3.5 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, the distance between partners was similar for the classification task across conditions, but a bit higher in the story task for WALL+AR.

4.5.4 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




Figure 4.9: 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).

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 4.9-\(a\)](#).

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 tiring than WALL ($p < 0.001$). This last result is consistent with participants' traveled distance in the tasks, which were 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 4.9-\(b\)](#). 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

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

section). When comparing the two conditions, the only significant result concerns the wall display space usage ($p < 0.001$), where, as expected, pairs said they use all the space available on the wall display 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 load, and ease of use. Results are shown in Figure 4.9-(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 tiring ($p = 0.008$) and less mentally demanding ($p = 0.041$). Results on physical demand align with our findings on movement around the wall display 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 tiring than the WALL only condition.

4.6 Discussion and Limitations

We next revisit how our results answer our original research questions on combining physical wall display 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 space 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 display would serve as the primary workspace surface. While we did observe virtual surfaces used as secondary storage (pool of images) and trashcans, such usage was in fact marginal. In fact, 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. Sometimes pairs went as far as creating three virtual containers and leaving

the wall display empty. Furthermore, in the story task, all groups immediately moved the main cards they wanted to use off the wall display and onto a surface and kept working there. – very few decided to bring the cards back to the wall after their work was done. 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*. Our pairs promoted virtual surfaces to first-class interactive surfaces. It would be interesting to investigate if this behavior persists when the wall display is less crowded – we suspect that due to the grouping flexibility of virtual surfaces these findings may also transfer to situations where the wall display 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 such rooms, we do feel some of our findings are unique to AR surfaces, because 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. Such actions in physical walls would be more challenging as they create an interaction overhead for re-selecting groups of objects, and they lack visual feedback of the items moving around the room. Of course, other observations will likely hold in a purely physical setup. For example, participants tending to start their work in the surface closest to their side. Future work could test these assumptions.

RQ2. How is the AR space used? Most virtual surfaces were placed just 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 more rare. 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 parallelizable task (classification), 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 display tended to create a surface on the left), and largely kept ownership of these surfaces. Thus, in this case, surfaces could be considered personal territories. 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 [31, 99]: the notion of the territory is fairly fluid and their nature is hard to predict, as participants transitioned from 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. Thus there is no one optimal or fixed position (and thus surface) they have easy access to. We note, however, that our parallelizable task was fairly simple, requiring frequent content movement and manipulation. It is possible

that in situations where pairs need to spend more time on specific types of information (document reading, visualization analysis), they may remain stationary for longer periods of time, thus creating a feeling of ownership of specific physical locations around the room, and virtual surfaces placed there.

The personal virtual space (a private surface that moves with the user) was only used by a few participants. They did not use them to store "personal" data (as the tasks did not include them). In our tasks, it was mainly used as a carrier bag to collect and move elements. We expect these surfaces may be of interest in situations requiring detailed personal work, particularly if the user needs to move a lot around the space (thus carrying their personal workspace).

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 parallelizing 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). These 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 used for network analysis, where interactions that created clutter led to tighter collaboration and coordination [167]. More work is required to understand the impact of available space on coordination and collaboration.

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. Trade-offs were also seen in the case of measured interactions: fewer actions in WALL+AR when dealing with few objects, but smaller interaction distances with the WALL. We note, however, 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. Nevertheless, our study presents evidence that such an addition is not only feasible but also has a low cost given its benefits. And it would be interesting to evaluate in depth 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 input (gaze-pointer + clicker) and 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., [99, 134]). It is difficult to generalize our results in terms of space use and collaborative strategy to interfaces that are (i) hard to "split", 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). It remains future work to consider an AR+Wall prototype for such use-case, for example, in the form of a focus + context display, given that AR headsets still cannot match the high resolution of wall displays.

4.7 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 space in a wall display environment alone 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 display alone, and this extended setup was preferred, and found more enjoyable and efficient than the wall display 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 in our work 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.

Can a Head Mounted Display Emulate an Ultra-wall?

In the previous chapter, we could observe that users do not hesitate to leave the wall and promote virtual surfaces to “first-class” interactive surfaces. Thus, why not abandon the wall, an expensive and heavy infrastructure, and work only in Augmented or Virtual Reality? That said, a physical wall has, a priori, several advantages. It is a *physical* shared surface and can have a very high pixel density. In this chapter, we, thus, compare an ultra-high resolution wall-sized display (ultra-wall) and a virtual wall of the same size rendered in virtual reality in an HMD from the point of view of their resolution.

5.1 Introduction

Ultra-high resolution wall-sized displays (ultra-walls) are large wall displays with a pixel density similar to the screen of a workstation (typically 100dpi). They allow to display a considerable amount of pixels and are appropriate platforms for visualizing and manipulating a large amount of data [21, 98, 135, 174, 178].

However, ultra-walls are expensive, need large rooms, and are complex devices. For instance, the cost of the wall used in Liu et al. [135] to compare a desktop and an ultra-wall can be estimated to at least 100k€ (we used an upgraded version of this wall for the pilot experiment of this chapter Figure 5.3-WILD), and the wall used in the two previous chapters had a cost of 300k€ (with its touch-sensitive system and mid-air tracking Figure 5.3-WILDER. In addition, ultra-walls need large rooms, and given the number of pixels to be rendered, they should be driven by a cluster of high-end workstations (at least for now). For these reasons, these platforms, although powerful, can be very power-consuming and complex to maintain and program.

In comparison, head-mounted displays (HMDs) are not too expensive, do not need large rooms, are easy to maintain, and are relatively easy to program as several environments exist for such platforms (e.g., Unity3D). Moreover, work in the recent Immersive Analytics trend [41, 183] suggests using HMD for visualization and data



Figure 5.1: Virtual wall: A user moving a disc while performing the classification task of Liu *et al.* [135] on a $5.5\text{ m} \times 1.8\text{ m}$ virtual wall display that is rendered in a HMD (Large letters: low information density condition).

analytics [49, 115, 225]. And although 3D visualization for abstract data has some drawbacks [64, 152], it is nevertheless possible to visualize abstract data in 2D in an HMD, by displaying them on 2D surfaces.

Thus, a natural research problem is to compare the performance of an HMD with an ultra-wall when visualizing and manipulating a large amount of data. This chapter explores this problem in a particular case where we emulate a wall in an HMD Figure 5.1.

The size and the pixel density of an ultra-wall determine the quantity of information that can be displayed. Both characteristics have been shown to be fundamental in the study of ultra-walls [65, 135, 239]. The big size of an ultra-wall allows to display a large number of elements of a dataset concurrently, as its high pixel density allows to display these elements (e.g., images, plots, or text descriptions) in high enough resolution to be easily distinguishable and readable.

Because of the qualities mentioned earlier, HMDs can potentially emulate a wall display of arbitrary size. In this chapter, we start by analyzing, in a theoretical way, the pixel density of an emulated wall for one of the current customer-ready HMD (the HTC Vive Pro) and one of the highest-resolution commercially available HMD (the Varjo XR-1). As we will see, both headsets' pixel density is sub-optimal and does not allow us to emulate the pixel density of an ultra-wall.

Then, to empirically validate the above computational approach, we use the classification task that has been proposed by Liu *et al.* [135] to compare an ultra-wall and a desktop. This task has been used in several other works to compare interaction techniques or paradigms [11, 100, 133, 160]. The advantage of this task is that it involves both interaction (pick-and-drop) and navigation (virtual or physical pan & zoom).

More importantly, the task operationalizes information density (and thus the quantity of information the device can render). Liu *et al.* [135] operationalizes this information density by simply considering classes represented by letters of different sizes.

Large letters operationalize low information density (see [Figure 5.1](#)), and small letters operationalize high information density.

The Liu et al. [135] task is abstract and formal and can be used as a benchmark. This allowed us to run the same classification tasks with an ultra-wall, a standard HMD, the Vive Pro, a commercial grade HMD, the Varjo XR-1, and a tentative emulation of a perfect HMD (using a Vive Pro).

5.2 Related Work

To better understand the differences between HMDs and wall displays, our related work focuses on the effects of the size and resolution of displays and previous findings of comparisons between HMDs and large displays.

5.2.1 Screen Size and Resolution

There is a long history in HCI about comparing screens and in particular their size and resolution. Pioneer work mainly compares standard monitor with larger screens (e.g., multi-monitor setting) in high level “desktop tasks” [28, 55, 90, 187]. Most of the results suggest that users are more efficient with larger screens as they have more space (pixels) to manage information.

Other works, under similar setup, suggest that large screens can bring a benefit in spatial orientation tasks just because of their size [207] or their wider field of view [206] and their larger number of pixels [156, 239].

All the above works consider users seated in front of the display, and thus physically static (beside eyes and head movement). However, with very large screens of high resolution, ultra-walls, it is not possible for the users to see (distinguish) all the pixels at a fixed position, and thus users should physically travel the wall room to take advantage of both the size and the high resolution, or should use virtual navigation (e.g., pan & zoom) to bring information closer to them.

Interestingly, it has been shown that, for an ultra-wall (2.7m wide, 96 dpi), physical navigation is more efficient than virtual navigation for simple search and pattern finding tasks [21]. Several subsequent works studied the cause of this effect. Ball and North [20] suggests that the field of view is not a key factor, Ragan et al. [173] suggests that the constant positioning of information is important, Jansen et al. [104] studies spatial memory, and Liu et al. [135] advocates that the main cause is the power of head movements to reach information (*i.e.* immersion when close to the wall).

Moreover, Liu et al. [135] insist that such results depend on the “density of information” and the full use of the high resolution of ultra-walls, a simple desktop being faster otherwise (see “The Classification Task” section below for more details).

5.2.2 HMD versus Other Displays

There are a lot of works that compare HMDs with desktops and room-scale immersive environments. Santos et al. [189] provides a good overview of such works done before 2009. Most of this work concerns tasks in a 3D virtual world (e.g. navigation in a maze or buildings, volume, visualization, etc). Results are mixed regarding performance, but HMDs are often preferred and some results suggest a better mental model of the environment with HMDs. Early on, Patrick et al. [162] suggests to replace a HMD by a "large" projected display, because at that time (2002), HMDs were very expensive (while there is no significant results in the comparison). This remark shows that fast change in technology can have a strong impact when comparing displays.

However, subsequent works are in the same lines: 3D tasks where resolution is secondary, have very few results on performance differences, but find some perceptual advantage for HMDs. For instance, Lhemedu-Steinke et al. [130] suggests a better concentration, involvement and enjoyment with a HMD than with a large display. Also, [213] suggests a better perceived spatial presence of the virtual environment with an HMD than with a desktop. On the other hand, some works report cybersickness issues with HMDs that do not occur with standard displays [34, 202].

5.2.3 HMD in Immersive Analytics

The interest around immersive visual analytics [41, 183] in the visualization and HCI community has increased in the recent years. In particular, HMDs have become very popular for visualizing abstract data with no obvious 3D representation. Although the use of 3D in this context is disputable [64, 152], some works have been conducted to take advantage of the immersion and the 3rd dimension to visualize multivariate data [22, 48], scatterplots [115, 170, 225], networks [49, 232], small multiples [136], etc.

Surprisingly, there are very few works that compare HMD with other types of displays in the context of immersive analytics. Cordeil et al. [49] compares an HMD and CAVE2, a $9\text{ m} \times 1.7\text{ m}$ horseshoe shaped wall display at 34 dpi with stereo 3D rendering for collaborative analysis of connectivity of simple networks. Results suggest that the HMD is more efficient than the room-scale environment, but collaboration is out of the scope of this chapter. To evaluate the possible advantage of immersive space-time cube geovisualization, Wagner Filho et al. [226] compares an HMD and a desktop, but there is no decisive results in term of performance.

5.2.4 Emulating Walls in VR HMDs

Recent work uses virtual reality HMDs to simulate various types of "physical" environments, in particular walls, to evaluate them without the need to build or deploy the




Figure 5.2: Simplified scheme for the optic of an eye looking at a physical screen. Angular size α needed to view a circle of width p_v on a screen at a distance d from the eye.

full physical environment. A pioneer work in this area is the work of Ville *et al.* [140] where they suggest conducting field studies on public displays in virtual reality. They observed similar phenomena in real-world and virtual public display setups, such as engagement with the displays. Similar types of work have been conducted to emulate AR [82], real-world authentication prototypes [144], cross-reality systems [79], *etc.*

However, to our knowledge, the only work that focuses on emulating a wall display is the work by Ville *et al.* [140]. Moreover, our goal is different, Ville *et al.* focus on field studies and high-level user engagement, as we want to compare an ultra-wall and HMDs with a focus on their resolution.

5.3 Background and Geometrical Models of Vision

The goal of this section is to describe the well-known geometrical model of vision. We examine through it the possibility of emulating a wall display in an HMD through the spectrum of pixel density.

5.3.1 Principle of the monocular eye for a physical screen

The relation between the diameter w of a circle (C) on a flat surface in front of a human at a distance d from this surface and the angular size μ of (C) in the “cyclopean” eye of a human (Figure 5.2) is given by the following formula:

$$w = 2 \cdot d \cdot \tan\left(\frac{\mu}{2}\right) \quad (5.1)$$




Figure 5.3: Different devices, because of their screen size, are used at different distances. A smartwatch's screen needs to be viewed closer than a smartphone, while a laptop can be placed farther than this smartphone. And wall displays allow a distance range on again another scale.

Device	<i>dpi</i>	pixel size	optimal visual distance
Smartwatch (Apple watch series 8)	330	0.0769mm	26.4cm
Smartphone (OnePlus 10T)	394	0.0640mm	22.0cm
Tablet (iPad)	264	0.0960mm	33.0cm
Laptop (15", FullHD)	145	0.1750mm	60.2cm
Wilder	62.85	0.4040mm	138.9cm
Wild(8K desktop screens)	280	0.0900mm	30.9cm

Figure 5.4: Different devices: resolution, corresponding pixel size and optimal viewing distance. A closer distance will result in pixelization; a farther distance will result in details not being perceived.

A human that does not need a visual correction through glasses is commonly defined as having 20/20 normal vision. Under this circumstance, a human can distinguish a circle among several side-by-side circles¹ (e.g., a series of 3 pixels black-white-black) if this circle is no smaller than $\alpha = 1'$ (i.e. one minute of arc) $= \frac{1}{60}^\circ$. This means that if this circle was a screen's pixel, the human vision optimal resolution would be about 60 pixels per degree (ppd). From this, we can determine with Equation 5.1 the smallest size of such a circle:

$$2 \cdot d \cdot \tan\left(\frac{1'}{2}\right) \sim 0.00029 \cdot d \quad (5.2)$$

By abuse of language, we call this smallest entity *a visual pixel*.

Conversely, we can compute the “ideal” visual distance from a computer screen given the size p of its pixels, i.e., the distance where the pixel size is equal to the size of a visual pixel, by just inverting Equation 5.2:

$$\frac{p}{2 \cdot \tan\left(\frac{1'}{2}\right)} \quad (5.3)$$

Within this distance, we have a pixelization effect, *i.e.* users can see the pixels, and beyond it, users might lose information rendered on screen.

Figure 5.3 and Figure 5.4 provide examples for standard devices, from a smartwatch to high-resolution wall displays. We can observe that, overall, **current screen technology for smartwatches, smartphones, laptops, and ultra-walls has reached human visual capacity and beyond.**

5.3.2 Principle of the monocular eye for a VR headset

We now turn our attention to HMDs. One particularity of HMDs is that the screens are (and must be) very close to the eyes. They are so close to the eyes that the human eyes cannot accommodate an image rendered on the screen. Common HMDs (Google cardboard, HCT Vive, Oculus Rift, Varjo RX-1 and RX-3, etc.) use convex lenses to create a virtual image of the screens at a far distance from the eyes (e.g. 1.5 m), avoiding any need for the eyes to perform an accommodation. Because of this, we can say that the VR HMDs' optical systems act as “*simple magnifiers*”.

Concretely, let us consider an HMD using convex lenses with focal f (and width and height w_l and h_l) located between the eyes and the screen (of width and height w_s and h_s) at a distance d_l of the screen (Figure 5.5). As stated by Wetzstein et al. [233] during their SIGGRAPH class, the Gaussian thin lens formulas tell us that the distance from the lens to the virtual image equals to:

$$d_v = \left| \frac{1}{\frac{1}{f} - \frac{1}{d_l}} \right| \quad (5.4)$$

¹this is bigger than the smallest entity that a human with normal vision can see




Figure 5.5: Simplified scheme for optics of a VR headset, f represents the focal distance of the lens, d_e the distance to the screen, d_i the distance between the lens and the virtual image.

with d_i the distance from the lens to the screen, and f the focal distance of the lens (Figure 5.5). From this, we can infer that the distance D from the eye to the virtual image is:

$$D = \left| \frac{1}{\frac{1}{f} - \frac{1}{d_i}} + d_e \right| \quad (5.5)$$

Moreover, the size in the virtual image of an object in the display is magnified by the following factor:

$$M = \frac{f}{f - d_i} \quad (5.6)$$

Note that if $f = d_i$, no virtual image is created and that if $f < d_i$, the virtual image is closer to the eyes than the screen (or even behind the eyes). Thus, f should be greater than d_i so that the virtual image is further away than the screen from the eyes. Indeed, we should have $f > d_i$, and d_i should be close enough to f so that the virtual image is at a distance where the eyes do not need accommodation.

To be concrete, let us take $f = 40 \text{ mm}$, $d_i = 39 \text{ mm}$, and $d_e = 18 \text{ mm}$ (expected value for a google cardboard for a smartphone). Then, the virtual image is at $D = 1578 \text{ mm}$ from the eyes and the magnification factor M is equal to 40 (e.g., an object of 1 mm in the display has a size of 40 mm in the virtual image).

If p is the pixel size of the screen, the size of this pixel in the virtual image is of $M \cdot p$ and its angular size is of:

$$A = 2 \cdot \arctan \frac{M \cdot p}{2 \cdot D}. \quad (5.7)$$

With this same formula, we can compute the (horizontal or vertical) field of view for one eye by taking for p the (horizontal or vertical) display size.

Conversely, we can compute the ideal size of the pixel on the display, *i.e.* the size so that the pixel in the virtual image has an angular size of 1 arc minute:

$$P = \frac{2 \cdot D \cdot \tan \frac{1'}{2}}{M} \quad (5.8)$$

If we continue with our example and use a pixel size of 0.0695 mm (400 dpi, pixel size of a typical high-end smartphone), we obtain an angular size of 0.10° , or a pixel density of 10 ppd. However, the ideal pixel size (to obtain 60 ppd) is of 0.0115 mm (2205 dpi).

Table 5.1 provides additional examples. Note that AR headsets such as HoloLens use a different technology than convex lenses, but the optic of such headsets has the same goal: build using the projected images on the glasses, a virtual image at a large distance from the eyes. Unfortunately, this table shows our lack of information, because manufacturers do not disclose the technical characteristics of their HMDs. For most commercial HMDs, we have no information on the focal distance and the distance from the lens to the screens (we provide our own estimations for the Vive Pro).

However, we can observe that **current VR headsets are far from being optimal regarding human visual capacity in their full field of view**. Indeed, the bests AR headsets reach 30 ppd (Varjo XR-3, Vive Pro 2), half of what would be needed. However, the Varjo's use an interesting technology using an additional 1920×1920 pixels display (one per eye and a semi-transparent mirror system) that is "projected" and concentrated in the center of the field of view (we call this display a "focus display", FD in Table 5.1). Varjo claims that this technology allows reaching the optimal resolution of 60 ppd, and even more, on a part of the field of view (*e.g.*, 27×27 degree).

However, covering the entire field of view with a resolution of 60 ppd would require screens with a resolution of about 5600×5600 pixels (more or less, depending on the field of view) or focus displays that can move to follow the users' gaze (and in particular eyes saccades, which seems unlikely, see Kim et al. [108], however).

5.3.3 Emulating a Wall

The emulation of a wall display in VR is obvious: we use a rectangle of the same size as the wall (the virtual wall) where we render the same scene that is rendered in the

model	f	d_l	M/D	display (pixels)	resolution	PPD	ideal resolution	FoV
phone vr	40	39	40/1.55 m	75×75mm	0.0635mm 16 ppm 400 dpi	15	0.0115mm 87 ppm 2213 dpi	80°
Vive Pro	33	32	33/1.06 m	1440×1600	0.0395mm 25 ppm 640 dpi	17	0.0095mm 105 ppm 2668 dpi	107° (both eyes)*
Vive Pro2	33 ^x	32 ^x	33/1.06 m ^x	2448×2448	0.0286 mm 35 ppm 890 dpi	25	0.0095mm 105 ppm 2668 dpi	116°H, 96°V (both eyes)
Varjo-XR1	?	?	?	1440×1600	?	18 (estimated)	?	87° (both eyes)
Varjo-XR1 FD	?	?	?	1920×1080	?	>60	?	28°H, 16°V
Varjo-XR3	?	?	?	2880×2720	?	30	?	115°H, 90°V (both eyes)
Varjo-XR3 FD	?	?	?	1920×1920	?	70	?	27°
HoloLens	?	?	?	1268×720	?	41 [§]	?	30°H, 17°V*
HoloLens 2	?	?	?	1440×936	?	20 [§]	?	43°H, 29°V*

^x Values inferred from the Vive Pro values.

* Values obtained from website [VR Compare](#), accessed on 17 Nov 2022.

[§] Values demonstrated by the blog post "[Hololens 2 Display Evaluation](#)" by Karl Gutttag accessed on 10 Nov 2022

Table 5.1: Specification for the focal distance, lens-to-screen distance, FoV (H: horizontal, V: vertical), and resolution of various VR HMDs (ppmm: pixels per millimeter; dpi (or ppi): pixels per inch; ppd: pixels by degree). All specifications given are for one eye only. Technical characteristics are most of the time not disclosed by the manufacturer, and are either retrieved from 3rd party analysis (see table notes) or marked in the table by a question mark when the values are not known.

physical wall using physical dimensions in the VR space. Then, we are interested under which circumstances the users can see details rendered on this virtual wall.

For this purpose, we can use the ppd of the HMD (that is, in general, provided by the manufacturers) to compute the size of a pixel when the virtual wall is at a distance d_v of the user camera (UC):

$$\frac{2 \cdot d_v \cdot \tan\left(\frac{1^\circ}{2}\right)}{ppd} \quad (5.9)$$

and in turn, compute the dpi of the virtual wall at a given distance of the (UC). It should be noted that this size and dpi depend on the distance from the (UC) to the emulated wall.

For instance, the dpi of a virtual wall at 1 m of the (UC), a comfortable distance to work with a wall, is of 24 dpi with the Vive Pro, 26 dpi with the Varjo XR-1 (standard

display), of 36 dpi with the Vive Pro 2, of 44 dpi with the Vajo XR-3 (standard display), and should have, a priori, an “optimal” dpi of 87 with the Varjo XR-1 and Varjo XR-3 focus displays.

5.4 The Classification Task

Liu *et al.* [135] compared an ultra-wall ($5.5\text{ m} \times 1.8\text{ m}$ at 100 dpi) with a desktop (30" display at 100 dpi) in a classification task that abstracts various concrete classification and scheduling tasks (such as classifying brain scans and schedule the CHI 2013 conference). The scene has the size of the ultra-wall and is constituted by 32 containers that can contain (at most 6) discs with a letter, see Figure 5.1. The task consists of moving the discs so that all the containers contain discs with the same letter. The task starts with a partially classified scene where, in each container, there is a majority of discs with the same letter, that are shown in green, and the participants should move miss-classified discs, shown in red, so that at the end all discs are green (correctly classified).

The difficulty of the task is controlled by the number of different letters (the classes). Liu *et al.* [135] consider an *Easy* task with two letters (two classes: “C” and “D”) and a *Hard* task with four letters (four classes: “H”, “K”, “N” and “R”). However, in the experiment below we will consider only the *Hard* tasks because it forces, a priori, stronger effects.

The most important factor in Liu *et al.* [135] is the size of the letters (factor `LETTERSIZESIZE`) that represents information density. Small letters correspond to a large amount of information (e.g., a lot of text can be displayed), and large letters represent a low amount of information. Liu *et al.* use 3 sizes.

- A large size, *LargeLetter*, $15.5 \times 20\text{ mm}$ letters, that has been chosen so that (i) no virtual navigation is need with the desktop: when the scene is scaled to be fully contained in the desktop screen the letters are easily readable; and (ii) no physical navigation is needed for tha wall either: all the letters can be easily scene when the user stand in front of the wall at a distance of about 2 m .
- A small size, *SmallLetter*, $1.8 \times 2.3\text{ mm}$ letters (usual desktop 12pts font). In this case, to perform the task, the participants had to navigate in the scene either virtually using panning and zooming with the desktop, or physically by moving in front of the ultra-wall.
- A medium size, *Medium*, twice the small size, that we will not consider in our experiment.

The main result of Liu *et al.* is that for large letters, where almost no navigation is needed, the desktop is faster than the ultra-wall, and that for small letters, where a




Figure 5.6: Pixelisation of the letter N for a screen A of 5x5 pixels on top row, and a screen B 3x3 pixels on the bottom row, both of size 1. We see that as the bottom row does not have enough resolution, we cannot tell apart the letter N for the letter H. By increasing the size of the rendered letter, we can display it on more pixels, and reach a point where the shape of the letter is as readable in screen B as in screen A.

lot of virtual or physical navigation is needed, the ultra-wall is faster than the desktop (35% faster for the hard task).

5.4.1 The classification Task in VR

To “replicate” the classification task in an HMD, we emulate a wall in virtual reality: we use a rectangle of the same physical size as the wall (the virtual wall) where we render the same scene using physical dimensions (see [Figure 5.1](#)).

For the *LargeLetter* letter size condition, letters are large enough to be easily seen when the user camera is at 2 m of the virtual wall, and, a priori, a participant can perform the task without moving from their starting position. We consider this condition for testing whether the basic interactions could impact the results: we expected no differences between the different *DEVICE*: Vive Pro, Varjo XR-1, and the wall (and the emulated condition, see below).

The *SmallLetter* letter size condition is more interesting. Indeed, a letter starts to be readable when it is rendered in a box with a width of at least 4 pixels and clearly readable when rendered in a 5 pixels wide box ([Figure 5.6](#)). The width of a letter is 1.8 mm, and thus we need to have pixels of at least $\frac{1.8}{4} = 0.45$ mm, which correspond to a wall at 56 dpi and an optimal distance of 1.55 m.

With the Vive Pro, [Equation 5.9](#) shows that the user camera (*UC*) should be at a distance of about 44 cm of the virtual wall so that a user can see the letters (with a Vive Pro 2 we estimate this distance to 64 cm and a little bit more with the standard display of the Vajo XR-3).

With the Varjo XR-1 and its focus display, this distance can be estimated to 1.55 m (60 ppd). However, this is true only on a small part of the field of view, and for the rest of the field of view, a letter will be seen when the user camera is about 46 cm of the virtual wall.

Thus, we considered a special condition *Emulated* where we try to emulate an optimal VR HMD at 60 ppd. We scale the size of the letters by a factor of 3.5, allowing a user to see the letters at a distance of 1.55 m in the Vive Pro. **Unfortunately, we made a mistake, and we scaled the letters by a factor of 4.2, making the classification task too easy.** However, as we will see, this led to an interesting result.

So with *Emulated*, we need to have pixels of at least $4.2 \cdot 0.45 = 1.89$ mm, which corresponds to a 13 dpi wall. With the Vive Pro, the users should be able to see the letters at a (virtual) distance of about 2 mm (and not 1.55 m as was our intention).

5.5 Pilot User Study

The goal of our study is to evaluate if a difference in resolution has an impact on the performance of users. We explained previously the simplified optical model for the vision of a screen and for the vision inside a VR headset. Now that we know the difference in the models, we want to verify if the differences in the devices can hinder or advantage users during a work session.

5.5.1 Device conditions and apparatus

For the wall display (*Wall*), we used a 5.50×1.80 m display with a resolution of 61440×17280 , composed of 32 8K LCD displays and driven by 8 workstations (280 dpi, see Figure 5.3-WILD). For the VR headsets, we used two different headsets: the Vive Pro ² (*Vive*) and the Varjo XR-1 ³ (*Varjo*). The Varjo XR-1 was, when we did the experiment, the best quality headset available, promising a resolution reaching human eye perceivable resolution with its focus display. We chose the Vive Pro as it is a standard, consumer-grade headset, with a good resolution. As explained in the previous section we also added an emulated condition by scaling the letters of the *SmallLetter* condition (*Emulated*, but we indeed scaled by 4.2 and not by 3.5 that is a priori the good scaling).

Thus, our primary factor is `DEVICE` with four values: *Wall*, *Vive*, *Varjo*, and *Emulated*.

5.5.2 Task

We used the classification task described in the previous section with four classes: “H”, “K”, “N” and “R”. For the letter size (factor `LETTERSIZESIZE`) we consider only two sizes, the *LargeLetter* and the *SmallLetter*.

²<https://www.vive.com/us/product/vive-pro/>

³<https://web.archive.org/web/20200929212842/https://varjo.com/products/xr-1/>

We used physical navigation for all conditions, with the cable for the VR headsets suspended so the participants would not trip on it. For the interaction, every condition uses a Vive controller. Participants would point at an item then click on the trigger to pick it, then point at the target position and drop it with the same action.

Our secondary factor is `LETTERSIZESIZE` with two values *LargeLetter* and *SmallLetter* (see previous the section). In fact, as we are studying the effect of resolution we are mainly interested by the *SmallLetter* size condition. The *LargeLetter* size is useful to see whether the different `DEVICES` perform equally well for the basic interactions (e.g., pick-and-drop): we expected very similar performance for the different `DEVICES` in the *LargeLetter* condition.

5.5.3 Participants

We recruited 6 participants with normal or corrected-to-normal vision over the 12 participants originally planned. Participants were HCI researchers, engineers or graduate students in Computer Science. We stopped the experiment after 6 participants once we noticed our error in the Emulated condition. Nevertheless, we report the results of the 6 participants here, as they present interesting tendencies.

5.5.4 Experiment Design & Procedure

Design. The experiment is a full factorial `DEVICE`×`LETTERSIZESIZE` within participants design. We blocked by `DEVICE`, and fixed the `LETTERSIZESIZE` presentation order to *LargeLetter* first, *SmallLetter* then. In the initial 12 participants design, we counter-balanced the device condition using a latin square. For each device, participants had 6 trials: 2 training trials with each `LETTERSIZESIZE` condition in the order defined above, then 2×2 experiment trials in the same order. We used the same initial wall disposition as Liu et al. [135].

Procedure. Each participant worked in two sessions, on two separate days, with one session holding 2 device conditions. A break was enforced in-between conditions inside one session. At the end of each condition, participants filled a questionnaire about the condition they completed.

At the beginning of each VR condition, we did a calibration phase with the participant to align the virtual world with the room space.

Then, a **vision test was performed**: the operator asked the participant to look at a letter on the wall and to move back until they hit the limit where they cannot tell the letter anymore (at each step we changed the letter, among “H”, “K”, “N” and “R”). The resulting distance was recorded. The operator then explained the interactions and the goal of the task, and ensured the participant felt comfortable.




Figure 5.7: Result of the vision test: maximal distance to the (virtual) wall to see the small letters of the experiment (in meters).

Before each trial, the participant were placed at a 2 m distance from either the physical wall in the *Wall* condition, or the virtual wall in the VR conditions. And then moved to the classification task.

5.5.5 Results

Vision test. Figure 5.7 shows the result of the vision test. Participants were able to see the *SmallLetter* at 50 cm with the *Vive*, which is slightly more from what we predicted (45 cm). For the *Varjo*, the participants needed to be at 1.20 m from the letters to see them, which is 35 cm less than predicted. This suggests that the focus display of the *Varjo* XR-1 does not reach the claimed 60 ppd (it is closer to 50 ppd).

The result for *Wall* is of about 1.9 m distance, 35 cm more than the expected value of 1.55 m. We have no explanation, but we will see that in practice the participants performed the classification task for the small letters at about 1 m from the screen.

For *Emulated* (condition run in *Vive* Pro), the distance is of about 2.2 m, more than our expectation of 2 m. This result and the result for *Vive* suggest that we may have underestimated the ppd of the *Vive* Pro.

Task Time. Figure 5.8 shows the task time by $DEVICE \times LETTERSIZE$. For *LargeLetter* condition we can observe, as expected, that the task times are very similar. This suggests that the input system worked similarly well for all *DEVICE*. For the *SmallLetter* condition we can observe that:

- The task time is similar for *Vive* and *Varjo*. This suggests that for a task where users have to move in front of a (virtual) wall and move their heads in all directions to see small details, the focus display of the *Varjo* is not really useful.
- As expected, the *Wall* was faster than the *Vive*, but also faster than the *Varjo* (this was the case uniformly for the 6 participants with clear differences). This suggests again than the focus display of the *Varjo* is not enough.

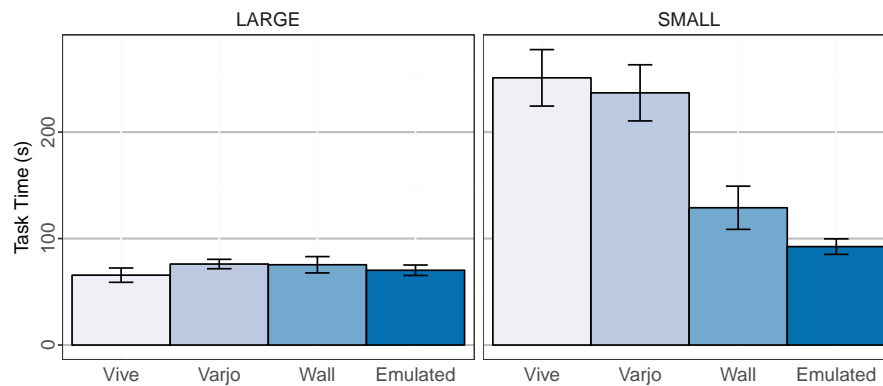


Figure 5.8: Task time by $DEVICE \times LETTERSIZE$.




Figure 5.9: (a) Average distance of the participant to the (virtual) wall in the small letters condition (in meters). (b) Mean of the total distance traveled by the participants for the small letters trial (in meters).

- *Emulated* was faster than the *Wall* (clearly for 5 participants, but one participant was 2 s faster with *Wall*). This confirms that we choose letters that were too big for *Emulated*.

We compared the task time for the *Wall* with the one from Liu *et al.* and we found very similar results (for the corresponding condition). Moreover, interestingly, the task time for the *Vive* and *Varjo* are very close to the time for the desktop condition of Liu *et al.*

Distance to the Wall and Traveled Distance. Figure 5.9 shows, for the *Small-Letter* condition: (a) the average distance of the participant to the (virtual) wall during a trial; and (b) the traveled distance by the participant during a trial. (In *LargeLetter* condition, participants did not need to come close to the wall and did not need to travel too much).

The results confirm the results on the task time: task time is “positively correlated” with the distance to the wall to perform the task (far away is better), and the traveled distance during the task (shorter is better).

Note that (i) the distance to the wall for *Wall* was of about 1 m (far smaller than the 1.9 m of the vision test and very close to the results of Liu *et al.*); (ii) this is also true for *Emulated*, but the difference is not as strong; (iii) in contrast, this distance is close to the vision test distance for the *Vive*; and (iv) despite its focus display this distance for the *Varjo* is close to the one of the *Vive*.

Regarding the traveled distance, we could note that this distance is half for *Emulated* than for *Wall*, and that this holds for all the participants.

Subjective. Regarding preferences, the six participants ranked the Wall condition 1st and the Emulated condition 2nd (4 participants ranked the Vive 3rd, and 3 participants ranked the Varjo 3rd, ex-aequo were allowed). Given the quantitative results above, it is not surprising that the participants preferred the Wall and the Emulated conditions. However, it is surprising, and interesting, that the six participants uniformly preferred the Wall condition over the Emulated condition, as ex-aequo were allowed and Emulated exhibits better quantitative performance than Wall (uniformly: for the six participants). This suggests that the participants appreciated the physical aspect of the wall display relatively to the virtuality of a headset.

5.6 Discussion and Future Work

Our results suggest that, as expected, current “standard” HMDs do not have the needed resolution to compete with an ultra-wall in a 2D task consisting of manipulation of “data of high information density” (although they might be in par with a standard desktop computer). They also suggest that the Varjo technology, which consists of adding a focus display with a very high resolution in the center of the field of view, is of limited utility in a search task involving small details on a large surface.

Our results also suggest that even with a possible VR HMD with perfect resolution (that will likely, allow reaching the performance of an ultra-wall), it is possible that the users would prefer the ultra-wall against the HMD. Indeed, the participants were faster and traveled less distance with the *Emulated* condition than with the *Wall* condition, but *all* participants preferred the *Wall* condition to the *Emulated* condition.

Of course, the above results should be confirmed with a 12 participants experiment. We plan to rerun our experiment in the near future, but with more recent headsets: the Vive Pro 2 and the Varjo XR-3 (and a correct emulated condition). However, our results suggest that an ultra-wall will still outperform these VR HMDs. Indeed, the Vive Pro 2 and the standard screens of the Varjo XR-3 are still far from being optimal (about 30 ppd), and our results suggest that the focus display of the Varjo XR-3 is not

enough for a task where many head movements are needed to observe small details on the wall.

It would also be interesting to test other types of headsets, in particular AR headsets (Hololens or VR HMDs with video see-through), to see whether having a view on the reality in addition to the emulated virtual wall is preferred by the participants to a purely virtual environment. Moreover, in a collaborative context, we hypothesize that AR HMDs will be better suited than VR HMDs with avatars.

Another aspect that we want to investigate is virtual navigation in front of a virtual wall (or several virtual walls, as in the previous chapter). Indeed, we used physical navigation with VR headsets in the experiment; thus, to emulate a wall display, we needed a large room. However, with VR headsets, we can use virtual navigation and then emulate a wall in a small space. We would like to know whether we can design virtual navigation techniques that reach walking performance in front of a wall. Of course, there is a considerable amount of work on virtual navigation techniques for VR headsets, but we did not find a simple answer to this question.

Discussion, Future Work and Conclusion

In this manuscript, I presented how adding Augmented Reality to support interaction on wall displays is feasible and bears advantages to compensate for wall displays' weaknesses. I also investigated the possibility of replacing a wall display by a VR HMD. The goal of this chapter will be to summarize the work done in the previous chapters and discuss the directions our work opened for future research.

6.1 Summary

We started this thesis by noting that wall displays are public displays, where all information is accessible to the whole audience. Thus, changes to the wall's content will impact every user, from cluttering the view to changing the entire content. To remove this potential clutter, we need personal views to display information so that users can work independently or preview changes to apply to the wall.

Networks and, more specifically, transportation networks are a frequent visualization on wall displays, either in the public space to inform users of buses or metro routes or in work spaces like control rooms to monitor and reroute road traffic. In both cases, users need to browse the content on the wall privately, in parallel with other users who do not share the same goal. Thus they would benefit from a private view. This observation led us to the first question:

RQ1: How can AR assist exploration and navigation of a network on a wall display, through personal view and navigation?

Through [Chapter 3](#), I answer this question by creating four techniques to aid the visual navigation of networks presented on large vertical displays. We show these techniques exclusively in the user's headset. By doing this, we provide a personal view to the user, which preserves the public view shared with all the users. The techniques

displayed in AR add visual aids to the gaze cursor to provide personalized feedback. As we are dealing with networks, this AR view shows personalized weights on the links of the node-link network. The four AR techniques present different variations of close display coupling with the context visualization seen on the large display. We empirically studied those techniques on two path-following tasks and found that techniques with very close and persistent display coupling to the content on the wall are preferred and more performant for tracing a path. In contrast, techniques that present more distant display coupling, ensure that the connection to the large display is visible even if the user moves away from a selected link; these techniques were preferred for following a path. Overall, participants prefer using visual AR cursors to assist navigation over using a simple gaze cursor.

More importantly, this study showed how an AR personal view semantically coupled with a public context on a large display could render personal preferences (in the form of weights) without cluttering the shared view on the public display. This work showed the possibility of adding an interactive personal layer to a non-interactive public visualization without altering the context visualization seen by external viewers. The proposed solution displays its content inside a headset. This can offer more privacy compared to alternative ways to provide a personal view, for example, by using a mobile phone [63], a smartwatch, or mobile AR [19] which can be subject to shoulder surfing. Furthermore, these techniques can be displayed on any screen with dynamic or static content and can even be shown on content laid on paper.

After studying the addition of AR to offer a personal view and navigation on a shared physical display, we looked at AR to add shared content around such a display. The limitation of a physical screen and the high cost of physically extending a wall display lead us to extend virtually a wall display using AR. However, we do not know how this virtual extension will be considered or exploited by users. This led to our second question:

RQ2 : How does extending the screen space of a wall display using AR changes users' space exploitation during a collaborative session?

For this question, we created in chapter 4 a system to consider the surrounding area around a wall display as an extension of its workspace. This extension is made using AR around the physical wall. This setup allows collaborators to work on and outside a wall display seamlessly, virtually breaking down the physical limits of the big screen. We studied what this new space (wall display + AR) changes for collaborative work and space usage compared to a wall display alone. More specifically, we studied pairs of participants over two tasks chosen to induce two different collaboration strategies: one that promotes close work and the other that allows parallel work. We observed

first if AR extension is beneficial when working with the wall display, then in which space users interacted, and finally whether it impacts their collaboration strategies.

To our surprise, we observed participants transitioning from the physical to the virtual space and back with no noticeable interaction cost. Even though the content was initially on the wall display, most pairs transferred the content to the virtual surfaces. They used the virtual space exclusively as a place to interact with the content and display their final result. While we did not see noticeable changes in strategies for the close collaboration task, participants synchronized their efforts when using the wall display alone, worked more independently, and used fewer actions when the AR was present. The addition of AR was found beneficial, as users reported more comfort using Wall+AR when the wall display was not enough for them to complete their task.

In the end, we showed that we could use a Wall+AR system with no performance loss over a wall display system. Nevertheless, the Wall+AR system increased physical load, which the weight of the headset can explain, and also increased the mental demand and interaction overhead. Regardless, it did not affect performance compared to the wall display system and improved the user experience.

The good reception of the virtual space in AR from our participants in the previous system made us consider how a VR headset could work instead of a wall display and an AR headset. VR headsets are cheaper than a standard wall display and allow navigating a virtual wall of the same size as a physical ones. However, emulated virtual walls cannot offer the physicality of wall displays and there is no guarantee that they can display as many elements on screen as a high-resolution wall display. That is why we ask:

RQ3 : Can we emulate a wall display inside a VR headset?

The capacity to emulate a wall display relies on multiple factors, such as the chosen interaction medium in VR or the way to navigate in this virtual space. We decided to investigate the visual aspect of the emulation by studying whether the current VR HMDs' resolution is enough to emulate a physical wall display or not. First, we analyzed this problem from a theoretical point of view. This analyze shows that the resolution of current HMDs is not yet good enough to match the pixel density of a high-resolution wall display.

Then we conducted a pilot experiment with 6 participants to validate the above analysis empirically. We compared a high-resolution wall display with a customer-grade VR HMD: the Vive Pro, a high-end commercial-grade VR HMD: the Varjo XR-1, and a tentative emulation of a perfect VR HMD inside a Vive Pro. We used a classification task with varying information density combined with a vision test to evaluate our different conditions. The results of our pilot study also suggest that the

resolution of current VR HMDs is not enough to emulate a high-resolution wall display. Participants performed the task closer to the virtual wall than the resolution lowered. Also, the farther participants were during the task, the faster they were to complete this task. Moreover, the questionnaire suggests that even with a possibly perfect VR HMD, participants prefer to use an ultra-wall display.

Of course, as this was a pilot study run with 6 participants, a complete study based on 12 participants is needed. We plan to run such a study using modern HMD and a corrected emulation of a perfect VR HMD to verify our insights.

6.2 Discussion and Future work

Our research studies low-level questions concerning hybrid Wall+ AR systems, and we must now continue to elaborate on this work. That is, to continue to study how augmented reality and wall displays can complement each other, support each other's strengths and compensate for their weaknesses. Subsequently, our results open the door to adding AR to the workflow of existing wall display systems.

To continue the work we presented in the previous chapter, we consider different topics we can directly work on after this thesis.

6.2.1 Personal AR Views on Shared Displays

Our [Chapter 3](#) shows the possibility of adding AR on public displays to add interactivity to them. This case of using AR privately can be extended outside of the node-link diagram case. It is already used to interact inside museums like Ding et al. [58] point out with the ArtLens2.0 application¹, or the Skin and Bones application² which displays on a mobile device a private view of the original content. Our research explores the technological aspect of adding AR on top of a public display and shows how it can be beneficial for performance and user preference. To go further, we should study the acceptance of using such AR techniques in public and how these techniques perform in situ. Meaning studying if acceptability differs in different situations, for example, in control room situations against in front of a public display in a train station.

We suspect that context is key. AR HMDs have already started being used and accepted in work situations [30, 60, 157, 175]. But acceptance for public use is still far from clear [177]. Nevertheless, we believe that in our work, we are looking at a particular public use-context: brief interactions in front of a shared public screen (e.g., to plot a metro trip). We think that as our scenario of use focuses on interacting for a short period and in a designated area, AR could be more accepted in this case than in other public spaces.

¹<https://www.clevelandart.org/artlens-gallery/artlens-app>

²<https://naturalhistory.si.edu/exhibits/bone-hall>

6.2.2 Combining AR and Walls in Collaboration

In our [Chapter 4](#) we compare the wall display alone and the Wall+AR system in our study. The results of this study can inspire further work to study the effect of a hybrid Wall+AR platform. During each task, the initial wall display was crowded to observe the impact of the extended AR space available to the participants. As we saw, the participants used the AR space as the primary interaction space. However, what would happen if the initial wall display contained fewer elements? Another experiment with an initial condition containing less content on the wall display would be interesting to observe if the virtual space would be used as much. We could also ask another question about where content is originally placed: if the content is placed initially in the extended AR space, would participants prefer to stay in the virtual space, or would they use the physical screen? Those experiments would help further understand if users left the wall display for the AR space only because it was the closest extra space available for the task or if another factor is at play. Some of our findings will likely change, for example, how often people move content in the AR space or if they come back to the physical wall. From the findings on the groups' positions during the classification task, we could expect users to group items close to them again, on the border of the initial surface, irrespective of setup.

Also, our study provided pairs with collaborative tasks to induce different collaboration strategies. We would be interested in extending our territoriality study on the Wall+AR system and a wall display system to explore the impact of the AR extension on more complex scenarios. We would like to test our system against a wall display alone on a competitive task and see how it impacts the pair dynamics and the use of public/personal spaces. Also, observing a composite task could help us understand the AR's impact more. Using a layout task, for example, the composition of a magazine cover, the creation of a website, or the design of a car, could lead to more diverse use of the space and maybe more creative use of the extended space. One other aspect that may have influenced our results is the data participants had to work with, images that are common in territoriality studies[12, 227]. Nevertheless, it is possible that other types of data, such as long text documents, will require users to stay stationary for more extended periods of time. This may result in them claiming "ownership" of areas in the room and forming personal territories often. Studying these aspects remains future work.

6.2.3 Replacing wall displays with AR or VR?

In their paper, Endert et al. [66] state *"Powerwalls are seldom used to complete everyday work-related tasks. They are difficult to "boot" and require custom applications."* Indeed wall displays are expensive devices that require heavy hardware and software setup and a large space to be placed in. All those constraints legitimate the question

of replacing this complex device with a much simpler one, like a VR HMD. VR HMD technology is now democratized and ranges from highly performant headsets that need to be plugged into a computer to less powerful standalone versions. They do not need a vast space to be used, are way cheaper than wall displays, and can display simulated VR environments that resemble physical wall displays [132, 192].

In Chapter 5, we compare a wall display against VR HMD to explore if VR HMDs can be an alternative to high-resolution wall displays. More precisely, we observe if the difference between two systems, a screen fixed in space (the wall display) and a screen relative to the user (the VR HMD), makes a difference in the users' performance in a classification task. Our early insight points out that replacing a wall display with a VR HMD significantly impacts the classification task, mainly because of a difference in resolution between both devices. This impact is enough to say that VR headsets do not come to the level of wall displays, possibly because of resolution.

However, the simple technological argument is not the only one to consider in the question of whether we can replace the screen with HMDs. For clarity, I'll refer to AR and VR HMDs as only HMDs. First, we can not force everyone to wear a headset. Despite advances in the field to reduce potential cybersickness, some users still suffer from it when using a VR HMD [38]. Second, wall displays are powerful for collaboration, and a VR headset setup will need to work on multiple issues to emulate a wall display on that part. For example, aside from requiring a robust network infrastructure to support collaboration, one big issue is seeing collaborators' social cues. Indeed HMDs cover a big part of the face, hiding facial expressions, which provide essential non-verbal cues during communication for information transmission. If we consider AR HMDs like the HoloLens2 that is transparent, we can see the other people's faces more, but we trade this against a less opaque screen that results in a less crisp and inferior rendering quality of the virtual content.

Nevertheless, there is apparent interest in emulating a wall display with HMDs. They allow for greater mobility, less setup cost, and need less real estate than a wall display. Indeed, recent advancements in hardware allowed for affordable consumer-grade standalone devices, which do not require a wired connection to a computing unit. This implies scenarios where users can work on a virtually rendered wall display outside their usual office environment, such as researchers refining their CHI presentations on their way to the conference. HMDs could also allow experts to join a collaboration session on a physical wall display when they can not be on-site, from their office or home. In this case, we could think on the physical wall display side of using either an AR headset or telepresence robots to show the remote user's presence. We have seen during the COVID pandemic that it is essential to have a remote alternative to our usual work practice. Even if this has yet to replace a co-located collaboration session, this grants the possibility to use the wall display more frequently by reducing both the financial cost of a wall display through emulation and the organization cost

of gathering a team around a device in a room.

6.2.4 Including wall displays in workflows

As we mentioned before, Endert et al. [66] says there is a setup issue with wall displays. I believe that to solve the setup issue, instead of replacing walls with HMDs, we should consider the integration of wall displays inside the workplace workflow. Indeed those systems often require engineers to create the content that will be displayed beforehand. And while tools like SAGE [141] can display content made for desktop screens, this content is not necessarily suited to be on a wall display. In our work, we explore a way to address the physical setup limitation of the wall display using AR; now, what needs to be addressed is a way to create content for the wall display without the need to go back and forth between our desktop and the physical wall. AR could help in that regard by offering a low-resolution view of what the result could be by either extending the desktop screen or projecting the wall display directly (see Figure 6.1). For better fidelity, we can also consider a view of the result inside a VR HMD. This solution can offer a way for an average user to edit content for the wall display while seeing the result with higher fidelity than on the small desktop screen.

When we have explored this possibility, we can think of the following scenario shown in Figure 6.1 to integrate the wall display as a simple tool in an analyst's workflow. A researcher prepares an analysis session on their laptop, processing visualizations from the data of their last experiment. They switch from laptop configuration to wall display configuration and put an AR headset on. The prepared visualizations appear all over a curved representation of a wall, and our analyst starts organizing the workspace grouping visualizations. They switch to a focus+context mode where the curved wall display now surrounds the laptop, and our analyst annotates precise data points. Now is the time for the collaborative session. They take off their headset and bring their laptop to the physical wall display where their colleague is. The prepared environment appears on the wall display by connecting their laptop to the wall display, and the analysts start sharing their interpretation of the data. A third analyst comes in with their data related to our case study. However, the space provided by the wall is not enough to contain and organize all the visualization. They grab a headset, which extends the physical wall display with new virtual spaces, for example, a second virtual wall, where the additional visualizations can now appear. After this, multiple possibilities are available. We can envision an expert in another field related to the study joining in the session from abroad, wearing a VR headset and having their avatar rendered in AR in the room. Alternatively, the analysts could manipulate the groups of visualizations with a tangible token that they bring back to their office once the session ends, containing the information from this collaboration session.




Figure 6.1: Scenario of a user working with wall display setup workflow. All physical devices have a purple border. On top, the user works on his own and can switch between a curved wall display in AR set up at a distance to have the most comfortable view of the content (A), a focus+context setup to access detail with high precision on the laptop surrounded by the context in AR (B), or a VR headset setup to see the result with the highest fidelity (C). Then they transition to the wall display to continue working, but this time on the physical wall display for a collaborative session (E). Another user comes and brings their data with them. The group put on their headset to extend the space on the wall(F), and maybe access other functionalities available from the use of AR, such as personal notifications

6.3 Conclusion

This thesis explored different ways to augment a wall display through Augmented Reality with Headsets. First, we created techniques to offer private views and interaction on a public wall display. Then we built a system to study the impact of extending with AR a wall display compared to using a wall display alone. After, we started investigating the possibility of replacing a wall display with a VR headset to study the trade-off between platforms.

We often think of using one device at a time, and this thesis goes in the direction of studying the synergy between physical screens and AR. I hope my work inspires researchers to further study the different configurations to combine AR with Large Displays or other physical displays.

Bibliography

- [1] Johnny Accot and Shumin Zhai. 1997. Beyond Fitts' Law: Models for Trajectory-based HCI Tasks. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI '97)*. ACM, 295–302. <https://doi.org/10.1145/258549.258760> 49
- [2] 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> 61
- [3] Shiroq Al-Megren, Ahmed Kharrufa, Jonathan Hook, Amey Holden, Selina Sutton, and Patrick Olivier. 2015. Comparing Fatigue When Using Large Horizontal and Vertical Multi-touch Interaction Displays. In *Human-Computer Interaction – INTERACT 2015*, Julio Abascal, Simone Barbosa, Mirko Fetter, Tom Gross, Philippe Palanque, and Marco Winckler (Eds.). Springer International Publishing, 156–164. 15
- [4] Basak Alper, Tobias Hollerer, JoAnn Kuchera-Morin, and Angus Forbes. 2011. Stereoscopic Highlighting: 2D Graph Visualization on Stereo Displays. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (Dec 2011), 2325–2333. <https://doi.org/10.1109/TVCG.2011.234> 50
- [5] Jessalyn Alvina, Caroline Appert, Olivier Chapuis, and Emmanuel Pietriga. 2014. RouteLens: Easy Route Following for Map Applications. In *Proceedings of the International Working Conference on Advanced Visual Interfaces (AVI '14)*. ACM, 125–128. <https://doi.org/10.1145/2598153.2598200> 28, 39
- [6] 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> 1, 12, 14, 21, 25, 55
- [7] Georg Aplitz 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> 61
- [8] Rahul Arora, Rubaiat Habib Kazi, Tovi Grossman, George Fitzmaurice, and Karan Singh. 2018. *SymbiosisSketch: Combining 2D & 3D Sketching for Designing Detailed 3D Objects in Situ*. ACM, 1–15. <https://doi.org/10.1145/3173574.3173759> 20, 22
- [9] Cagan Arslan, Yosra Rekik, and Laurent Grisoni. 2019. E-Pad: Large Display Pointing in a Continuous Interaction Space around a Mobile Device. In *Proceedings of the 2019 on Designing Interactive Systems Conference (DIS '19)*. ACM, 1101–1108. <https://doi.org/10.1145/3322276.3322284> 15

- [10] Jillian Aurisano, Khairi Reda, Andrew Johnson, Elisabeta G. Marai, and Jason Leigh. 2015. BactoGeNIE: a large-scale comparative genome visualization for big displays. *BMC Bioinformatics* 16, 11 (13 Aug 2015), S6. <https://doi.org/10.1186/1471-2105-16-S11-S6> 12
- [11] Ignacio Avellino, Cédric Fleury, Wendy E. Mackay, and Michel Beaudouin-Lafon. 2017. CamRay: Camera Arrays Support Remote Collaboration on Wall-Sized Displays. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, 6718–6729. <https://doi.org/10.1145/3025453.3025604> 82
- [12] 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> 1, 6, 12, 14, 17, 56, 57, 58, 103
- [13] Ronald T. Azuma. 1997. A Survey of Augmented Reality. *Presence: Teleoperators and Virtual Environments* 6, 4 (08 1997), 355–385. <https://doi.org/10.1162/pres.1997.6.4.355> arXiv:<https://direct.mit.edu/pvar/article-pdf/6/4/355/1623026/pres.1997.6.4.355.pdf> 2
- [14] Benjamin Bach, Ronell Sicat, Johanna Beyer, Maxime Cordeil, and Hanspeter Pfister. 2018. The Hologram in My Hand: How Effective is Interactive Exploration of 3D Visualizations in Immersive Tangible Augmented Reality? *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (Jan 2018), 457–467. <https://doi.org/10.1109/TVCG.2017.2745941> 30
- [15] Benjamin Bach, Ronell Sicat, Hanspeter Pfister, and Aaron Quigley. 2017. Drawing into the AR-Canvas: Designing Embedded Visualizations for Augmented Reality. *Workshop on Immersive Analytics: Exploring Future Interaction and Visualization Technologies for Data Analytics* (2017). 3
- [16] Felipe Bacim, Eric Ragan, Siroberto Scerbo, Nicholas F. Polys, Mehdi Setareh, and Brett D. Jones. 2013. The Effects of Display Fidelity, Visual Complexity, and Task Scope on Spatial Understanding of 3D Graphs. In *Proceedings of Graphics Interface 2013 (GI '13)*. CIPS, 25–32. <https://dl.acm.org/doi/abs/10.5555/2532129.2532135> 50
- [17] Sriram Karthik Badam, Fereshteh Amini, Niklas Elmqvist, and Pourang Irani. 2016. Supporting visual exploration for multiple users in large display environments. In *2016 IEEE Conference on Visual Analytics Science and Technology (VAST)*. 1–10. <https://doi.org/10.1109/VAST.2016.7883506> 53, 55
- [18] Sriram Karthik Badam and Niklas Elmqvist. 2014. PolyChrome: A Cross-Device Framework for Collaborative Web Visualization. In *Proceedings of the Ninth ACM International Conference on Interactive Tabletops and Surfaces (ITS '14)*. ACM, 109–118. <https://doi.org/10.1145/2669485.2669518> 16

- [19] Matthias Baldauf, Katrin Lasinger, and Peter Fröhlich. 2012. Private Public Screens: Detached Multi-User Interaction with Large Displays through Mobile Augmented Reality. In *Proceedings of the 11th International Conference on Mobile and Ubiquitous Multimedia (MUM '12)*. ACM, Article 27, 4 pages. <https://doi.org/10.1145/2406367.2406401> 100
- [20] Robert Ball and Chris North. 2008. The Effects of Peripheral Vision and Physical Navigation on Large Scale Visualization. In *Proceedings of Graphics Interface (GI '08)*. CIPS, 9–16. <http://dl.acm.org/citation.cfm?id=1375714.1375717> 14, 83
- [21] 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> 7, 14, 55, 81, 83
- [22] Andrea Batch, Andrew Cunningham, Maxime Cordeil, Niklas Elmqvist, Tim Dwyer, Bruce H. Thomas, and Kim Marriott. 2020. There Is No Spoon: Evaluating Performance, Space Use, and Presence with Expert Domain Users in Immersive Analytics. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (2020), 536–546. <https://doi.org/10.1109/TVCG.2019.2934803> 84
- [23] 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> 55
- [24] Daniel Belcher, Mark Billingham, S. E. Hayes, and Randy Stiles. 2003. Using augmented reality for visualizing complex graphs in three dimensions. *Proceedings - 2nd IEEE and ACM International Symposium on Mixed and Augmented Reality, ISMAR 2003* (2003), 84–93. <https://doi.org/10.1109/ISMAR.2003.1240691> 30
- [25] Hrvoje Benko, Eddie W. Ishak, and Steven Feiner. 2005. Cross-dimensional gestural interaction techniques for hybrid immersive environments. In *Proceedings of Virtual Reality (VR '05)*. IEEE, 209–216. <https://doi.org/10.1109/VR.2005.1492776> 19, 22
- [26] Anastasia Bezerianos, Pierre Dragicevic, and Ravin Balakrishnan. 2006. Mnemonic Rendering: An Image-based Approach for Exposing Hidden Changes in Dynamic Displays. In *Proceedings of the 19th Annual ACM Symposium on User Interface Software and Technology (UIST '06)*. ACM, 159–168. <https://doi.org/10.1145/1166253.1166279> 50
- [27] Anastasia Bezerianos and Petra Isenberg. 2012. Perception of Visual Variables on Tiled Wall-Sized Displays for Information Visualization Applications. *IEEE Transactions on*

- Visualization and Computer Graphics* 18, 12 (2012), 2516–2525. <https://doi.org/10.1109/TVCG.2012.251> 13, 17
- [28] Xiaojun Bi and Ravin Balakrishnan. 2009. Comparing Usage of a Large High-Resolution Display to Single or Dual Desktop Displays for Daily Work. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09)*. ACM, 1005–1014. <https://doi.org/10.1145/1518701.1518855> 14, 83
- [29] Eric A. Bier, Maureen C. Stone, Ken Pier, William Buxton, and Tony D. DeRose. 1993. Toolglass and Magic Lenses: The See-through Interface. In *Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '93)*. ACM, 73–80. <https://doi.org/10.1145/166117.166126> 50
- [30] Eleonora Bottani and Giuseppe Vignali. 2019. Augmented reality technology in the manufacturing industry: A review of the last decade. *IISE Transactions* 51, 3 (2019), 284–310. <https://doi.org/10.1080/24725854.2018.1493244> arXiv:<https://doi.org/10.1080/24725854.2018.1493244> 102
- [31] 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> 1, 56, 64, 77
- [32] Eugenie 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, 13 pages. <https://hal.archives-ouvertes.fr/hal-02915795> 60
- [33] Eugenie 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> 18, 22
- [34] James Brooks, Riley Lodge, and Daniel White. 2017. Comparison of a head-mounted display and flat screen display during a micro-UAV target detection task. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, Vol. 61. SAGE, 1514–1518. <https://doi.org/10.1177/1541931213601863> 84
- [35] Simon Butscher, Sebastian Hubenschmid, Jens Müller, Johannes Fuchs, and Harald Reiterer. 2018. Clusters, Trends, and Outliers: How Immersive Technologies Can Facilitate the Collaborative Analysis of Multidimensional Data. , Article 90 (2018), 12 pages. <https://doi.org/10.1145/3173574.3173664> 20, 22
- [36] 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. <https://doi.org/10.1109/IWAR.1999.803804> 17, 21

- [37] W. Buxton, G. Fitzmaurice, R. Balakrishnan, and G. Kurtenbach. 2000. Large displays in automotive design. *IEEE Computer Graphics and Applications* 20, 4 (2000), 68–75. <https://doi.org/10.1109/38.851753> 12, 16
- [38] Polona Caserman, Augusto Garcia-Agundez, Alvar Gámez Zerban, and Stefan Göbel. 2021. Cybersickness in current-generation virtual reality head-mounted displays: systematic review and outlook. *Virtual Reality* 25, 4 (01 Dec 2021), 1153–1170. <https://doi.org/10.1007/s10055-021-00513-6> 104
- [39] Marco Cavallo, Mishal Dholakia, Matous Havlena, Kenneth Ocheltree, and Mark Podlaseck. 2019. Dataspace: A Reconfigurable Hybrid Reality Environment for Collaborative Information Analysis. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. 145–153. <https://doi.org/10.1109/VR.2019.8797733> 19
- [40] Marco Cavallo, Mishal Dolakia, Matous Havlena, Kenneth Ocheltree, and Mark Podlaseck. 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> 19, 55
- [41] T. Chandler, M. Cordeil, T. Czauderna, T. Dwyer, J. Glowacki, C. Goncu, M. Klapperstueck, K. Klein, K. Marriott, F. Schreiber, and E. Wilson. 2015. Immersive Analytics. In *2015 Big Data Visual Analytics (BDVA)*. IEEE Computer Society, 1–8. <https://doi.org/10.1109/BDVA.2015.7314296> 1, 81, 84
- [42] Chenliang Chang, Kiseung Bang, Gordon Wetzstein, ByoungHo 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> 63
- [43] 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> 15, 54
- [44] Kelvin Cheng, Jane Li, and Christian Müller-Tomfelde. 2012. Supporting Interaction and Collaboration on Large Displays Using Tablet Devices. In *Proceedings of the International Working Conference on Advanced Visual Interfaces (AVI '12)*. ACM, 774–775. <https://doi.org/10.1145/2254556.2254708> 29
- [45] 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> 54

- [46] Haeyong Chung, Chris North, Jessica Zeitz Self, Sharon Chu, and Francis Quek. 2014. VisPorter: facilitating information sharing for collaborative sensemaking on multiple displays. *Personal and Ubiquitous Computing* 18, 5 (01 Jun 2014), 1169–1186. <https://doi.org/10.1007/s00779-013-0727-2> 16
- [47] Andrew Clayphan, Anthony Collins, Judy Kay, Nathan Slawitschka, and Jenny Horder. 2018. Comparing a Single-Touch Whiteboard and a Multi-Touch Tabletop for Collaboration in School Museum Visits. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 2, 1, Article 6 (mar 2018), 23 pages. <https://doi.org/10.1145/3191738> 15
- [48] Maxime Cordeil, Andrew Cunningham, Tim Dwyer, Bruce H. Thomas, and Kim Marriott. 2017. ImAxes: Immersive Axes As Embodied Affordances for Interactive Multivariate Data Visualisation. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology - UIST '17 (UIST '17)*. ACM, 71–83. <https://doi.org/10.1145/3126594.3126613> 84
- [49] M Cordeil, T Dwyer, K Klein, B Laha, K Marriott, and B H Thomas. 2017. Immersive Collaborative Analysis of Network Connectivity: CAVE-style or Head-Mounted Display? *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (Jan 2017), 441–450. <https://doi.org/10.1109/TVCG.2016.2599107> 82, 84
- [50] Nuno Correia, Tarquínio Mota, Rui Nóbrega, Luís Silva, and Andreia Almeida. 2010. A Multi-Touch Tabletop for Robust Multimedia Interaction in Museums. In *ACM International Conference on Interactive Tabletops and Surfaces (ITS '10)*. ACM, 117–120. <https://doi.org/10.1145/1936652.1936674> 11
- [51] Emmanuel Courtoux, Caroline Appert, and Olivier Chapuis. 2021. WallTokens: Surface Tangibles for Vertical Displays. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21)*. ACM, Article 421, 13 pages. <https://doi.org/10.1145/3411764.3445404> 15
- [52] Elizabeth Crundall, David R. Large, and Gary Burnett. 2016. A driving simulator study to explore the effects of text size on the visual demand of in-vehicle displays. *Displays* 43 (2016), 23–29. <https://doi.org/10.1016/j.displa.2016.05.003> 12, 16
- [53] Geoff Cumming. 2014. The New Statistics: Why and How. *Psychological Science* 25, 1 (2014), 7–29. <https://doi.org/10.1177/0956797613504966> arXiv:<https://doi.org/10.1177/0956797613504966> 45
- [54] Geoff Cumming and Sue Finch. 2005. Inference by eye: Confidence intervals and how to read pictures of data. *American Psychologist* 60, 2 (2005), 170–180. <https://doi.org/10.1037/0003-066X.60.2.170> 44
- [55] Mary Czerwinski, Greg Smith, Tim Regan, Brian Meyers, George G Robertson, and Gary K Starkweather. 2003. Toward characterizing the productivity benefits of very large displays.. In *Interact*, Vol. 3. 9–16. 14, 83

- [56] Rodrigo A. De Almeida, Clément Pillias, Emmanuel Pietriga, and Pierre Cubaud. 2012. Looking behind Bezels: French Windows for Wall Displays. In *AVI - 11th working conference on Advanced visual interfaces - 2012 (Proceedings of the 11th working conference on Advanced visual interfaces)*, ACM (Ed.). 124–131. <https://doi.org/10.1145/2254556.2254581> 40
- [57] Paul H. Dietz and Matt Lathrop. 2019. Adaptive Environments with Parallel Reality™ Displays. In *ACM SIGGRAPH 2019 Talks (SIGGRAPH '19)*. ACM, Article 34, 2 pages. <https://doi.org/10.1145/3306307.3328153> 15
- [58] Mandy Ding et al. 2017. Augmented reality in museums. *Museums & augmented reality—A collection of essays from the arts management and technology laboratory* (2017), 1–15. 102
- [59] Ciro Donalek, S. G. Djorgovski, Alex Cioc, Anwell Wang, Jerry Zhang, Elizabeth Lawler, Stacy Yeh, Ashish Mahabal, Matthew Graham, Andrew Drake, Scott Davidoff, Jeffrey S. Norris, and Giuseppe Longo. 2014. Immersive and collaborative data visualization using virtual reality platforms. In *2014 IEEE International Conference on Big Data (Big Data)*. 609–614. <https://doi.org/10.1109/BigData.2014.7004282> 1
- [60] Mathilde Drouot, Nathalie Le Bigot, Emmanuel Bricard, Jean-Louis de Bougrenet, and Vincent Nourrit. 2022. Augmented reality on industrial assembly line: Impact on effectiveness and mental workload. *Applied Ergonomics* 103 (2022), 103793. <https://doi.org/10.1016/j.apergo.2022.103793> 102
- [61] M. Eaddy, G. Blasko, J. Babcock, and S. Feiner. 2004. My own private kiosk: privacy-preserving public displays. In *Eighth International Symposium on Wearable Computers*, Vol. 1. 132–135. <https://doi.org/10.1109/ISWC.2004.32> 29
- [62] A. Ebert, S. Thelen, P. Olech, J. Meyer, and H. Hagen. 2010. Tiled++: An Enhanced Tiled Hi-Res Display Wall. *IEEE Transactions on Visualization and Computer Graphics* 16, 1 (2010), 120–132. <https://doi.org/10.1109/TVCG.2009.57> 40
- [63] Malin Eiband, Mohamed Khamis, Emanuel von Zezschwitz, Heinrich Hussmann, and Florian Alt. 2017. Understanding Shoulder Surfing in the Wild: Stories from Users and Observers. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, 4254–4265. <https://doi.org/10.1145/3025453.3025636> 100
- [64] Niklas Elmquist. 2017. 3D Visualization for Nonspatial Data: Guidelines and Challenges. <https://sites.umiacs.umd.edu/elm/2017/10/03/3d-visualization-for-nonspatial-data-guidelines-and-challenges/> Accessed May 6th, 2020. 82, 84
- [65] 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> 55, 82
- [66] Alex Endert, Lauren Bradel, Jessica Zeitz, Christopher Andrews, and Chris North. 2012. Designing Large High-Resolution Display Workspaces. In *Proceedings of the International Working Conference on Advanced Visual Interfaces (AVI '12)*. ACM, 58–65. <https://doi.org/10.1145/2254556.2254570> 14, 103, 105
- [67] 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> 59
- [68] Epson.com. 2019. *MOVERIO Augmented Reality Smart Glasses*. <https://epson.com/moverio-augmented-reality> Accessed 21 January 2020. 30
- [69] Steven Feiner and A Shamash. 1991. Hybrid user interfaces: Breeding virtually bigger interfaces for physically smaller computers. *Acm Uist* (1991), 9–17. <https://doi.org/10.1145/120782.120783> 17, 21
- [70] Anna Maria Feit, Shane Williams, Arturo Toledo, Ann Paradiso, Harish Kulkarni, Shaun Kane, and Meredith Morris Ringel. 2017. Toward Everyday Gaze Input: Accuracy and Precision of Eye Tracking and Implications for Design. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, 1118–1130. <https://doi.org/10.1145/3025453.3025599> 49
- [71] Jean-Daniel Fekete and Catherine Plaisant. 1999. Excentric Labeling: Dynamic Neighborhood Labeling for Data Visualization. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '99)*. ACM, 512–519. <https://doi.org/10.1145/302979.303148> 5
- [72] Andreas Rene Fender, Hrvoje Benko, and Andy Wilson. 2017. MeetAlive: Room-Scale Omni-Directional Display System for Multi-User Content and Control Sharing. In *Proceedings of the ACM International Conference on Interactive Surfaces and Spaces (ISS '17)*. ACM, 106–115. <https://doi.org/10.1145/3132272.3134117> 26
- [73] Clifton Forlines, Daniel Wigdor, Chia Shen, and Ravin Balakrishnan. 2007. Direct-Touch vs. Mouse Input for Tabletop Displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*. ACM, 647–656. <https://doi.org/10.1145/1240624.1240726> 15
- [74] S. Ghani, N. Henry Riche, and N. Elmqvist. 2011. Dynamic Insets for Context-Aware Graph Navigation. *Computer Graphics Forum* 30, 3 (6 2011), 861–870. <https://doi.org/10.1111/j.1467-8659.2011.01935.x> 28, 31

- [75] Mojgan Ghare, Marvin Pafla, Caroline Wong, James R. Wallace, and Stacey D. Scott. 2018. Increasing Passersby Engagement with Public Large Interactive Displays: A Study of Proxemics and Conation. In *Proceedings of the 2018 ACM International Conference on Interactive Surfaces and Spaces (ISS '18)*. ACM, 19–32. <https://doi.org/10.1145/3279778.3279789> 7
- [76] Jiangtao Gong, Teng Han, Siling Guo, Jiannan Li, Siyu Zha, Liuxin Zhang, Feng Tian, Qianying Wang, and Yong Rui. 2021. HoloBoard: A Large-Format Immersive Teaching Board Based on Pseudo HoloGraphics. In *The 34th Annual ACM Symposium on User Interface Software and Technology (UIST '21)*. ACM, 441–456. <https://doi.org/10.1145/3472749.3474761> 12, 16
- [77] Tovi Grossman and Ravin Balakrishnan. 2005. The Bubble Cursor: Enhancing Target Acquisition by Dynamic Resizing of the Cursor's Activation Area. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '05)*. ACM, 281–290. <https://doi.org/10.1145/1054972.1055012> 35
- [78] 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> 18, 22
- [79] Uwe Gruenefeld, Jonas Auda, Florian Mathis, Stefan Schneegass, Mohamed Khamis, Jan Gugenheimer, and Sven Mayer. 2022. VRception: Rapid Prototyping of Cross-Reality Systems in Virtual Reality. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI '22)*. ACM, Article 611, 15 pages. <https://doi.org/10.1145/3491102.3501821> 85
- [80] Jan Gugenheimer, Christian Mai, Mark McGill, Julie Williamson, Frank Steinicke, and Ken Perlin. 2019. Challenges Using Head-Mounted Displays in Shared and Social Spaces. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems (CHI EA '19)*. ACM, Article W19, 8 pages. <https://doi.org/10.1145/3290607.3299028> 50
- [81] Maxime Guillon, François Leitner, and Laurence Nigay. 2015. Investigating Visual Feedforward for Target Expansion Techniques. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, 2777–2786. <https://doi.org/10.1145/2702123.2702375> 36
- [82] Kota Gushima and Tatsuo Nakajima. 2021. Virtual Fieldwork: Designing Augmented Reality Applications Using Virtual Reality Worlds. In *Virtual, Augmented and Mixed Reality*. Springer, 417–430. https://doi.org/10.1007/978-3-030-77599-5_29 85
- [83] S. Hayne, M. Pendergast, and S. Greenberg. 1993. Gesturing through cursors: implementing multiple pointers in group support systems. In *[1993] Proceedings of the*

- Twenty-sixth Hawaii International Conference on System Sciences*, Vol. iv. 4–12 vol.4. <https://doi.org/10.1109/HICSS.1993.284158> 60
- [84] I. Herman, G. Melancon, and M. S. Marshall. 2000. Graph visualization and navigation in information visualization: A survey. *IEEE Transactions on Visualization and Computer Graphics* 6, 1 (Jan 2000), 24–43. <https://doi.org/10.1109/2945.841119> 30
- [85] 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> 60
- [86] Uta Hinrichs and Sheelagh Carpendale. 2011. Gestures in the Wild: Studying Multi-Touch Gesture Sequences on Interactive Tabletop Exhibits. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. ACM, 3023–3032. <https://doi.org/10.1145/1978942.1979391> 15, 16
- [87] Tom Horak, Sriram Karthik Badam, Niklas Elmqvist, and Raimund Dachzelt. 2018. When David Meets Goliath: Combining Smartwatches with a Large Vertical Display for Visual Data Exploration. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, Article 19, 13 pages. <https://doi.org/10.1145/3173574.3173593> 16, 26, 54
- [88] Sebastian Hubenschmid, Johannes Zagermann, Simon Butscher, and Harald Reiterer. 2018. Employing Tangible Visualisations in Augmented Reality with Mobile Devices. In *AVI 2018 Workshop on Multimodal Interaction for Data Visualization*. 20, 22
- [89] C. Hurter, A. Telea, and O. Ersoy. 2011. MoleView: An Attribute and Structure-Based Semantic Lens for Large Element-Based Plots. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (Dec 2011), 2600–2609. <https://doi.org/10.1109/TVCG.2011.223> 28, 31
- [90] Dugald Ralph Hutchings, Greg Smith, Brian Meyers, Mary Czerwinski, and George Robertson. 2004. Display Space Usage and Window Management Operation Comparisons between Single Monitor and Multiple Monitor Users. In *Proceedings of the Working Conference on Advanced Visual Interfaces (AVI '04)*. ACM, 32–39. <https://doi.org/10.1145/989863.989867> 83
- [91] PTC Inc. 2020. *Vuforia Engine Library [online]*. <https://library.vuforia.com/> Accessed 21 January 2020. 38
- [92] Kori Inkpen, Kirstie Hawkey, Melanie Kellar, Regan M, Karen Parker, Derek Reilly, Stacey Scott, and Tara Whalen. 2005. Exploring display factors that influence co-located collaboration: angle, size, number, and user arrangement. In *In Proc. HCI International*. 12

- [93] P. Isenberg, S. Carpendale, A. Bezerianos, N. Henry, and J. Fekete. 2009. CoCoNutTrix: Collaborative Retrofitting for Information Visualization. *IEEE Computer Graphics and Applications* 29, 5 (Sep. 2009), 44–57. <https://doi.org/10.1109/MCG.2009.7829>
- [94] P. Isenberg, P. Dragicevic, W. Willett, A. Bezerianos, and J.-D. Fekete. 2013. Hybrid-Image Visualization for Large Viewing Environments. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (Dec 2013), 2346–2355. <https://doi.org/10.1109/TVCG.2013.1637>
- [95] Petra Isenberg, Danyel Fisher, Sharoda A. Paul, Meredith Ringel Morris, Kori Inkpen, and Mary Czerwinski. 2012. Co-located collaborative visual analytics around a tabletop display. *IEEE Transactions on Visualization and Computer Graphics* 18, 5 (2012), 689–702. <https://doi.org/10.1109/TVCG.2011.28756>
- [96] 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/10.1145/345315763>
- [97] Robert J. K. Jacob. 1990. What You Look at is What You Get: Eye Movement-Based Interaction Techniques. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '90)*. ACM, 11–18. <https://doi.org/10.1145/97243.9724649>
- [98] Mikkel R. Jakobsen and Kasper Hornbæk. 2013. Interactive Visualizations on Large and Small Displays: The Interrelation of Display Size, Information Space, and Scale. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2336–2345. <https://doi.org/10.1109/TVCG.2013.17081>
- [99] 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, Article 11 (Feb. 2014), 34 pages. <https://doi.org/10.1145/257609913,14,17,25,53,55,56,57,74,77,79>
- [100] Mikkel R. Jakobsen and Kasper Hornbæk. 2015. Is Moving Improving? Some Effects of Locomotion in Wall-Display Interaction. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, 4169–4178. <https://doi.org/10.1145/2702123.270231282>
- [101] Raphaël James, Anastasia Bezerianos, and Olivier Chapuis. 2023. Evaluating the Extension of Wall Displays with AR for Collaborative Work. (2023). <https://doi.org/10.1145/3544548.358075253>
- [102] 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> 25, 54
- [103] Yvonne Jansen, Pierre Dragicevic, and Jean-Daniel Fekete. 2012. Tangible Remote Controllers for Wall-Size Displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, 2865–2874. <https://doi.org/10.1145/2207676.2208691> 15
- [104] Yvonne Jansen, Jonas Schjerlund, and Kasper Hornbæk. 2019. Effects of Locomotion and Visual Overview on Spatial Memory When Interacting with Wall Displays. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. ACM, Article 291, 12 pages. <https://doi.org/10.1145/3290605.3300521> 7, 83
- [105] Paul Kabbash and William AS Buxton. 1995. The “prince” technique: Fitts’ law and selection using area cursors. In *Proceedings of the SIGCHI conference on Human factors in computing systems (CHI '95)*. ACM, 273–279. <https://doi.org/10.1145/223904.223939> 35
- [106] Keiko Katsuragawa, Krzysztof Pietroszek, James R. Wallace, and Edward Lank. 2016. Watchpoint: Freehand Pointing with a Smartwatch in a Ubiquitous Display Environment. In *Proceedings of the International Working Conference on Advanced Visual Interfaces (AVI '16)*. ACM, 128–135. <https://doi.org/10.1145/2909132.2909263> 15
- [107] Mohamed Khamis, Axel Hoesl, Alexander Klimczak, Martin Reiss, Florian Alt, and Andreas Bulling. 2017. EyeScout: Active Eye Tracking for Position and Movement Independent Gaze Interaction with Large Public Displays. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology (UIST '17)*. ACM, 155–166. <https://doi.org/10.1145/3126594.3126630> 29
- [108] Jonghyun Kim, Youngmo Jeong, Michael Stengel, Kaan Akşit, Rachel Albert, Ben Boudaoud, Trey Greer, Joohwan Kim, Ward Lopes, Zander Majercik, Peter Shirley, Josef Spjut, Morgan McGuire, and David Luebke. 2019. Foveated AR: Dynamically-Foveated Augmented Reality Display. *ACM Trans. Graph.* 38, 4, Article 99 (jul 2019), 15 pages. <https://doi.org/10.1145/3306346.3322987> 89
- [109] 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). arXiv:1708.01377 <http://arxiv.org/abs/1708.01377> 19, 22, 26, 30
- [110] U. Kister, K. Klamka, C. Tominski, and R. Dachsel. 2017. GraSp: Combining Spatially-aware Mobile Devices and a Display Wall for Graph Visualization and Interaction. *Computer Graphics Forum* 36, 3 (June 2017), 503–514. <https://doi.org/10.1111/cgf.13206> 16, 26, 29

- [111] Ulrike Kister, Patrick Reipschläger, and Raimund Dachsel. 2016. MultiLens: fluent interaction with multi-functional multi-touch lenses for information visualization. *Proceedings of the 2016 ACM on Interactive Surfaces and Spaces - ISS '16* (2016), 139–148. <https://doi.org/10.1145/2992154.2992168> 28
- [112] 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> 29, 53, 55
- [113] Daniel Klinkhammer, Markus Nitsche, Marcus Specht, and Harald Reiterer. 2011. Adaptive Personal Territories for Co-Located Tabletop Interaction in a Museum Setting. In *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces (ITS '11)*. ACM, 107–110. <https://doi.org/10.1145/2076354.2076375> 11, 16
- [114] 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> 63
- [115] M. Kraus, N. Weiler, D. Oelke, J. Kehrer, D. A. Keim, and J. Fuchs. 2020. The Impact of Immersion on Cluster Identification Tasks. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (2020), 525–535. <https://doi.org/10.1109/TVCG.2019.2934395> 82, 84
- [116] Stefan Kreitmayer, Yvonne Rogers, Robin Laney, and Stephen Peake. 2013. UniPad: Orchestrating Collaborative Activities through Shared Tablets and an Integrated Wall Display. In *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '13)*. ACM, 801–810. <https://doi.org/10.1145/2493432.2493506> 12, 16
- [117] Russell Kruger, Sheelagh Carpendale, Stacey D. Scott, and Saul Greenberg. 2004. Roles of Orientation in Tabletop Collaboration: Comprehension, Coordination and Communication. *Computer Supported Cooperative Work (CSCW)* 13, 5 (01 Dec 2004), 501–537. <https://doi.org/10.1007/s10606-004-5062-8> 13
- [118] Kazutaka Kurihara, Naoshi Nagano, Yuta Watanabe, Yuichi Fujimura, Akinori Minaduki, Hidehiko Hayashi, and Yohei Tutiya. 2011. Toward Localizing Audiences' Gaze Using a Multi-Touch Electronic Whiteboard with SPieMenu. In *Proceedings of the 16th International Conference on Intelligent User Interfaces (IUI '11)*. ACM, 379–382. <https://doi.org/10.1145/1943403.1943470> 15
- [119] Oh-Hyun Kwon, Chris Muelder, Kyungwon Lee, and Kwan-Liu Ma. 2016. A study of layout, rendering, and interaction methods for immersive graph visualization. *IEEE*

- Transactions on Visualization and Computer Graphics* 22, 7 (2016), 1802–1815. <https://doi.org/10.1109/TVCG.2016.2520921> 50
- [120] Jonas Landgren and Fredrik Bergstrand. 2016. Work Practice in Situation Rooms – An Ethnographic Study of Emergency Response Work in Governmental Organizations. In *Information Systems for Crisis Response and Management in Mediterranean Countries*, Paloma Díaz, Narjès Bellamine Ben Saoud, Julie Dugdale, and Chihab Hanachi (Eds.). Springer International Publishing, Cham, 157–171. 12
- [121] Ricardo Langner and Raimund Dachzelt. 2018. Towards Visual Data Exploration at Wall-Sized Displays by Combining Physical Navigation with Spatially-Aware Devices. In *Poster Program of the 2018 IEEE VIS Conference*. 2 pages. 16
- [122] Ricardo Langner, Ulrike Kister, and Raimund Dachzelt. 2019. Multiple Coordinated Views at Large Displays for Multiple Users: Empirical Findings on User Behavior, Movements, and Distances. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (Jan 2019), 608–618. <https://doi.org/10.1109/TVCG.2018.2865235> 25
- [123] Ricardo Langner, Marc Satkowski, Wolfgang Büschel, and Raimund Dachzelt. 2021. *MARVIS: Combining Mobile Devices and Augmented Reality for Visual Data Analysis*. ACM. <https://doi.org/10.1145/3411764.3445593> 18, 22, 23
- [124] Benjamin Lee, Maxime Cordeil, Arnaud Prouzeau, and Tim Dwyer. 2019. FIESTA: A Free Roaming Collaborative Immersive Analytics System. In *Proceedings of the 2019 ACM International Conference on Interactive Surfaces and Spaces (ISS '19)*. ACM, 335–338. <https://doi.org/10.1145/3343055.3360746> 21
- [125] 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> 54
- [126] 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> 21, 22, 57
- [127] Bongshin Lee, Catherine Plaisant, Cynthia Sims Parr, Jean-Daniel Fekete, and Nathalie Henry. 2006. Task Taxonomy for Graph Visualization. In *Proceedings of the 2006 AVI Workshop on BEyond Time and Errors: Novel Evaluation Methods for Information Visualization (BELIV '06)*. ACM, 1–5. <https://doi.org/10.1145/1168149.1168168> 39, 49

- [128] Gun-Yeal Lee, Jong-Young Hong, SoonHyoung Hwang, Seokil Moon, Hyeokjung Kang, Sohee Jeon, Hwi Kim, Jun-Ho Jeong, and ByoungHo Lee. 2018. Metasurface eyepiece for augmented reality. *Nature Communications* 9, 1 (2018), 4562. <https://doi.org/10.1038/s41467-018-07011-5> 63
- [129] Anke Lehmann, Heidrun Schumann, Oliver Staadt, and Christian Tominski. 2011. Physical Navigation to Support Graph Exploration on a Large High-resolution Display. In *Proceedings of the 7th International Conference on Advances in Visual Computing - Volume Part I (ISVC '11)*. Springer-Verlag, 496–507. <http://dl.acm.org/citation.cfm?id=2045110.2045165> 29
- [130] Quinatei Lhemedu-Steinke, Gerrit Meixner, and Michael Weber. 2018. Comparing VR Display with Conventional Displays for User Evaluation Experiences. In *Conference on Virtual Reality and 3D User Interfaces (VR '18)*. IEEE, 583–584. <https://doi.org/10.1109/VR.2018.8446076> 84
- [131] 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> 19
- [132] Lee Lisle, Kylie Davidson, Edward J.K. Gitre, Chris North, and Doug A. Bowman. 2021. Sensemaking Strategies with Immersive Space to Think. In *2021 IEEE Virtual Reality and 3D User Interfaces (VR)*. 529–537. <https://doi.org/10.1109/VR50410.2021.00077> 21, 22, 104
- [133] 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> 8, 56, 57, 82
- [134] Can Liu, Olivier Chapuis, Michel Beaudouin-Lafon, and Eric Lecolinet. 2017. CoReach: Cooperative Gestures for Data Manipulation on Wall-sized Displays. *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (2017), 6730–6741. <https://doi.org/10.1145/3025453.3025594> 13, 15, 21, 25, 56, 64, 65, 79
- [135] 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. *Proceedings of the 32nd annual ACM conference on Human factors in computing systems - CHI '14* (2014), 4147–4156. <https://doi.org/10.1145/2556288.2557020> 14, 25, 55, 81, 82, 83, 91, 94
- [136] Jiazhou Liu, Arnaud Prouzeau, Barrett Ens, and Tim Dwyer. 2020. Design and Evaluation of Interactive Small Multiples Data Visualisation in Immersive Spaces. In *Conference on Virtual Reality and 3D User Interfaces (VR '20)*. IEEE, 588–597. <https://doi.org/10.1109/VR46266.2020.00-23> 84

- [137] David López, Lora Oehlberg, Candemir Doger, and Tobias Isenberg. 2016. Towards An Understanding of Mobile Touch Navigation in a Stereoscopic Viewing Environment for 3D Data Exploration. *IEEE Transactions on Visualization and Computer Graphics* 22, 5 (2016), 1616–1629. <https://doi.org/10.1109/TVCG.2015.2440233> 20
- [138] Weizhou Luo, Anke Lehmann, Hjalmar Widengren, and Raimund Dachsel. 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> 57, 58, 64, 65
- [139] Tahir Mahmood, Erik Butler, Nicholas Davis, Jian Huang, and Aidong Lu. 2018. Building Multiple Coordinated Spaces for Effective Immersive Analytics through Distributed Cognition. In *2018 International Symposium on Big Data Visual and Immersive Analytics (BDVA)*. 1–11. <https://doi.org/10.1109/BDVA.2018.8533893> 20, 22
- [140] Ville Mäkelä, Rivu Radiah, Saleh Alsherif, Mohamed Khamis, Chong Xiao, Lisa Borchert, Albrecht Schmidt, and Florian Alt. 2020. Virtual Field Studies: Conducting Studies on Public Displays in Virtual Reality. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. ACM, 1–15. <https://doi.org/10.1145/3313831.3376796> 85
- [141] Thomas Marrinan, Jillian Aurisano, Arthur Nishimoto, Krishna Bharadwaj, Victor Mateevitsi, Luc Renambot, Lance Long, Andrew Johnson, and Jason Leigh. 2014. SAGE2: A new approach for data intensive collaboration using Scalable Resolution Shared Displays. In *10th IEEE International Conference on Collaborative Computing: Networking, Applications and Worksharing*. 177–186. <https://doi.org/10.4108/icst.collaboratecom.2014.257337> 105
- [142] K. Marriott, F. Schreiber, T. Dwyer, K. Klein, N. Henry Riche, T. Itoh, W. Stuerzlinger, and B.H Thomas (Eds.). 2018. *Immersive Analytics*. Springer. 26
- [143] Paul Marshall, Richard Morris, Yvonne Rogers, Stefan Kreitmayer, and Matt Davies. 2011. Rethinking 'Multi-User': An in-the-Wild Study of How Groups Approach a Walk-up-and-Use Tabletop Interface. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. ACM, 3033–3042. <https://doi.org/10.1145/1978942.1979392> 11, 16
- [144] Florian Mathis, Joseph O'Hagan, Kami Vaniea, and Mohamed Khamis. 2022. Stay Home! Conducting Remote Usability Evaluations of Novel Real-World Authentication Systems Using Virtual Reality. In *Proceedings of the 2022 International Conference on Advanced Visual Interfaces (Frascati, Rome, Italy) (AVI 2022)*. Association for Computing Machinery, New York, NY, USA, Article 14, 9 pages. <https://doi.org/10.1145/3531073.3531087> 85

- [145] T. May, M. Steiger, J. Davey, and J. Kohlhammer. 2012. Using Signposts for Navigation in Large Graphs. *Computer Graphics Forum* 31, 3pt2 (6 2012), 985–994. <https://doi.org/10.1111/j.1467-8659.2012.03091.x> 28, 31
- [146] 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> 16, 54
- [147] Paul Milgram and Fumio Kishino. 1994. A taxonomy of mixed reality visual displays. *IEICE TRANSACTIONS on Information and Systems* 77, 12 (1994), 1321–1329. 2, 3
- [148] 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> 61
- [149] Tomer Moscovich, Fanny Chevalier, Nathalie Henry, Emmanuel Pietriga, and Jean-Daniel Fekete. 2009. Topology-aware navigation in large networks. *Proceedings of the 27th international conference on Human factors in computing systems - CHI 09* (2009), 2319. <https://doi.org/10.1145/1518701.1519056> 28, 32, 50
- [150] Jens Müller, Tobias Schwarz, Simon Butscher, and Harald Reiterer. 2014. Back to Tangibility: A post-WIMP Perspective on Control Room Design. In *Proceedings of the International Working Conference on Advanced Visual Interfaces (AVI '14)*. ACM, 57–64. <https://doi.org/10.1145/2598153.2598161> 25
- [151] Jörg Müller, Robert Walter, Gilles Bailly, Michael Nischt, and Florian Alt. 2012. Looking Glass: A Field Study on Noticing Interactivity of a Shop Window. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, 297–306. <https://doi.org/10.1145/2207676.2207718> 11, 16
- [152] Tamara Munzner. 2014. *Visualization Analysis and Design*. A. K. Peters Visualization Series, CRC Press. 82, 84
- [153] 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> 63
- [154] Mathieu Nancel, Julie Wagner, Emmanuel Pietriga, Olivier Chapuis, and Wendy Mackay. 2011. Mid-Air Pan-and-Zoom on Wall-Sized Displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. ACM, 177–186. <https://doi.org/10.1145/1978942.1978969> 15

- [155] R. Netzel, M. Burch, and D. Weiskopf. 2014. Comparative Eye Tracking Study on Node-Link Visualizations of Trajectories. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014), 2221–2230. <https://doi.org/10.1109/TVCG.2014.2346420> 29, 49
- [156] Tao Ni, Doug Bowman, and Jian Chen. 2006. Increased display size and resolution improve task performance in Information-Rich Virtual Environments. In *Proceedings of Graphics Interface 2006 (GI 2006)*. Canadian Human-Computer Communications Society, 139–146. <https://graphicsinterface.org/proceedings/gi2006/gi2006-18/> 14, 83
- [157] Susanna Nilsson and Björn Johansson. 2008. Acceptance of Augmented Reality Instructions in a Real Work Setting. In *CHI '08 Extended Abstracts on Human Factors in Computing Systems (CHI EA '08)*. ACM, 2025–2032. <https://doi.org/10.1145/1358628.1358633> 22, 102
- [158] Susanna Nilsson, Bjorn Johansson, and Arne Jonsson. 2009. Using AR to support cross-organisational collaboration in dynamic tasks. In *2009 8th IEEE International Symposium on Mixed and Augmented Reality*. 3–12. <https://doi.org/10.1109/ISMAR.2009.5336522> 3, 6, 19, 20, 22
- [159] 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> 18, 19, 23
- [160] 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> 3, 4, 18, 22, 23, 82
- [161] Yujiro Okuya, Olivier Gladin, Nicolas Ladevèze, Cédric Fleury, and Patrick Bourdot. 2020. Investigating Collaborative Exploration of Design Alternatives on a Wall-Sized Display. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. ACM, 1–12. <https://doi.org/10.1145/3313831.3376736> 12, 16
- [162] Emilee Patrick, Dennis Cosgrove, Aleksandra Slavkovic, Jennifer A. Rode, Thom Ver-ratti, and Greg Chiselko. 2000. Using a Large Projection Screen as an Alternative to Head-Mounted Displays for Virtual Environments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '00)*. ACM, 478–485. <https://doi.org/10.1145/332040.332479> 84
- [163] Leonardo Pavanatto, Chris North, Doug A. Bowman, Carmen Badea, and Richard Stoakley. 2021. Do we still need physical monitors? An evaluation of the usability of AR virtual monitors for productivity work. In *2021 IEEE Virtual Reality and 3D User*

- Interfaces (VR)*. 759–767. <https://doi.org/10.1109/VR50410.2021.00103> 20, 21, 22
- [164] 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. *Proceeding of the twenty-sixth annual CHI conference on Human factors in computing systems - CHI '08* (2008), 1285. <https://doi.org/10.1145/1357054.1357255> 11, 15, 16
- [165] Thomas V. Perneger. 1998. What's wrong with Bonferroni adjustments. *BMJ* 316, 7139 (1998), 1236–1238. <https://doi.org/10.1136/bmj.316.7139.1236> arXiv:<https://www.bmj.com/content> 45
- [166] Arnaud Prouzeau, Anastasia Bezerianos, and Olivier Chapuis. 2016. Towards Road Traffic Management with Forecasting on Wall Displays. In *Proceedings of the ACM International Conference on Interactive Surfaces and Spaces (ISS '16)*. ACM, 119–128. <https://doi.org/10.1145/2992154.2992158> 1, 12, 25
- [167] 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 (August 2017), 1936–1951. <https://doi.org/10.1109/TVCG.2016.2592906> 78
- [168] Arnaud Prouzeau, Anastasia Bezerianos, and Olivier Chapuis. 2017. Trade-offs Between a Vertical Shared Display and Two Desktops in a Collaborative Path-Finding Task. (2017), 214–219. <https://hal.archives-ouvertes.fr/hal-01506678> 29
- [169] 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> 7, 16, 26, 54
- [170] Arnaud Prouzeau, Maxime Cordeil, Clement Robin, Barrett Ens, Bruce H. Thomas, and Tim Dwyer. 2019. Scaptics and Highlight-Planes: Immersive Interaction Techniques for Finding Occluded Features in 3D Scatterplots. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI '19)*. ACM, Article 325, 12 pages. <https://doi.org/10.1145/3290605.3300555> 84
- [171] Arnaud Prouzeau, Antoine Lhuillier, Barrett Ens, Daniel Weiskopf, and Tim Dwyer. 2019. Visual Link Routing in Immersive Visualisations. In *Proceedings of the 2019 ACM International Conference on Interactive Surfaces and Spaces (ISS '19)*. ACM, 241–253. <https://doi.org/10.1145/3343055.3359709> 30
- [172] Erik Prytz, Susanna Nilsson, and Arne Jönsson. 2010. The importance of eye-contact for collaboration in AR systems. In *2010 IEEE International Symposium on Mixed and*

- Augmented Reality*. 119–126. <https://doi.org/10.1109/ISMAR.2010.5643559>
19, 22
- [173] Eric D. Ragan, Alex Endert, Doug A. Bowman, and Francis Quek. 2012. How Spatial Layout, Interactivity, and Persistent Visibility Affect Learning with Large Displays. In *Proceedings of the International Working Conference on Advanced Visual Interfaces (AVI '12)*. ACM, 91–98. <https://doi.org/10.1145/2254556.2254576> 83
- [174] Fateme Rajabiyazdi, 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>
12, 25, 55, 81
- [175] Sara Rankohi and Lloyd Waugh. 2013. Review and analysis of augmented reality literature for construction industry. *Visualization in Engineering* 1, 1 (29 Aug 2013), 9. <https://doi.org/10.1186/2213-7459-1-9> 102
- [176] Umar Rashid, Miguel A. Nacenta, and Aaron Quigley. 2012. Factors Influencing Visual Attention Switch in Multi-Display User Interfaces: A Survey. In *Proceedings of the 2012 International Symposium on Pervasive Displays (PerDis '12)*. ACM, Article 1, 6 pages. <https://doi.org/10.1145/2307798.2307799> 16, 17
- [177] Philipp A. Rauschnabel, Daniel W.E. Hein, Jun He, Young K. Ro, Samir Rawashdeh, and Bryan Krulikowski. 2016. Fashion or Technology? A Fashnology Perspective on the Perception and Adoption of Augmented Reality Smart Glasses. *i-com* 15, 2 (2016), 179–194. <https://doi.org/doi:10.1515/icom-2016-0021> 102
- [178] 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> 14, 25, 55, 81
- [179] Joshua Reibert, Patrick Riehmann, and Bernd Froehlich. 2020. Multitouch Interaction with Parallel Coordinates on Large Vertical Displays. *Proc. ACM Hum.-Comput. Interact.* 4, ISS, Article 199 (nov 2020), 22 pages. <https://doi.org/10.1145/3427327>
1
- [180] Patrick Reipschläger and Raimund Dachsel. 2019. DesignAR: Immersive 3D-Modeling Combining Augmented Reality with Interactive Displays. In *Proceedings of the 2019 ACM International Conference on Interactive Surfaces and Spaces (ISS '19)*. ACM, 29–41. <https://doi.org/10.1145/3343055.3359718> 19, 20, 22

- [181] Patrick Reipschläger, Tamara Flemisch, and Raimund Dachsel. 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. <https://doi.org/10.1109/TVCG.2020.3030460> 6, 20, 22, 23, 54
- [182] 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> 17, 21
- [183] William Ribarsky, Jay Bolter, Augusto Op den Bosch, and Ron van Teylingen. 1994. Visualization and Analysis Using Virtual Reality. *IEEE Comput. Graph. Appl.* 14, 1 (Jan. 1994), 10–12. <https://doi.org/10.1109/38.250911> 81, 84
- [184] Nathalie Henry Riche, Tim Dwyer, Bongshin Lee, and Sheelagh Carpendale. 2012. Exploring the design space of interactive link curvature in network diagrams. *Proceedings of the International Working Conference on Advanced Visual Interfaces* (2012), 506–513. <https://doi.org/10.1145/2254556.2254652> 34, 50
- [185] Julie Rico and Stephen Brewster. 2009. Gestures All around Us: User Differences in Social Acceptability Perceptions of Gesture Based Interfaces. In *Proceedings of the 11th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '09)*. ACM, Article 64, 2 pages. <https://doi.org/10.1145/1613858.1613936> 27, 30
- [186] Charlie Roberts, Basak Alper, JoAnn Kuchera-Morin, and Tobias Höllerer. 2012. Augmented textual data viewing in 3D visualizations using tablets. In *Proceedings of the Symposium on 3D User Interfaces*. IEEE, 101–104. <https://doi.org/10.1109/3DUI.2012.6184192> 29
- [187] George Robertson, Mary Czerwinski, Patrick Baudisch, Brian Meyers, Daniel Robbins, Greg Smith, and Desney Tan. 2005. The Large-Display User Experience. *IEEE Comput. Graph. Appl.* 25, 4 (July 2005), 44–51. <https://doi.org/10.1109/MCG.2005.8883>
- [188] Yvonne Rogers and Siân Lindley. 2004. Collaborating around vertical and horizontal large interactive displays: which way is best? *Interacting with Computers* 16, 6 (09 2004), 1133–1152. <https://doi.org/10.1016/j.intcom.2004.07.008> arXiv:<https://academic.oup.com/iwc/article-pdf/16/6/1133/2305668/iwc16-1133.pdf> 12
- [189] Beatriz Sousa Santos, Paulo Dias, Angela Pimentel, Jan-Willem Baggerman, Carlos Ferreira, Samuel Silva, and Joaquim Madeira. 2009. Head-mounted display versus desktop for 3D navigation in virtual reality: a user study. *Multimedia Tools and Applications* 41, 1 (2009), 161. <https://doi.org/10.1007/s11042-008-0223-2> 84

- [190] Corina Sas, Nigel Davies, Sarah Clinch, Peter Shaw, Mateusz Mikusz, Madeleine Steeds, and Lukas Nohrer. 2020. Supporting Stimulation Needs in Dementia Care through Wall-Sized Displays. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. ACM, 1–16. <https://doi.org/10.1145/3313831.3376361> 12, 16
- [191] Marc Satkowski, Weizhou Luo, and Raimund Dachselt. 2021. Towards In-situ Authoring of AR Visualizations with Mobile Devices. In *2021 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*. 324–325. <https://doi.org/10.1109/ISMAR-Adjunct54149.2021.00073> 18
- [192] Kadek Ananta Satriadi, Barrett Ens, Maxime Cordeil, Tobias Czauderna, and Bernhard Jenny. 2020. Maps Around Me: 3D Multiview Layouts in Immersive Spaces. *Proc. ACM Hum.-Comput. Interact.* 4, ISS, Article 201 (nov 2020), 20 pages. <https://doi.org/10.1145/3427329> 104
- [193] Constantin Schmidt, Jörg Müller, and Gilles Bailly. 2013. Screenfinity: Extending the Perception Area of Content on Very Large Public Displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, 1719–1728. <https://doi.org/10.1145/2470654.2466227> 11, 16
- [194] Sebastian Schmidt, Miguel A. Nacenta, Raimund Dachselt, and Sheelagh Carpendale. 2010. A Set of Multi-touch Graph Interaction Techniques. In *ACM International Conference on Interactive Tabletops and Surfaces (ITS '10)*. ACM, 113–116. <https://doi.org/10.1145/1936652.1936673> 28, 31
- [195] Tobias Schwarz, Simon Butscher, Jens Mueller, and Harald Reiterer. 2012. Content-aware Navigation for Large Displays in Context of Traffic Control Rooms. In *Proceedings of the International Working Conference on Advanced Visual Interfaces (AVI '12)*. ACM, 249–252. <https://doi.org/10.1145/2254556.2254601> 26, 29
- [196] Stacey D. Scott, M. Sheelagh T. Carpendale, and Kori Inkpen. 2004. Territoriality in Collaborative Tabletop Workspaces. In *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work (CSCW '04)*. ACM, 294–303. <https://doi.org/10.1145/1031607.1031655> arXiv:arXiv:1011.1669v3 12, 14, 56, 58, 59, 64, 65
- [197] Mickael Sereno, Lonni Besançon, and Tobias Isenberg. 2022. Point Specification in Collaborative Visualization for 3D Scalar Fields Using Augmented Reality. *Virtual Reality* (2022). <https://doi.org/10.1007/s10055-021-00614-2> 20, 22
- [198] 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> 19

- [199] Lauren Shupp, Christopher Andrews PhD, Margaret Dickey-Kurdziolek PhD, Beth Yost, and Chris North. 2009. Shaping the Display of the Future: The Effects of Display Size and Curvature on User Performance and Insights. *Human-Computer Interaction* 24, 1-2 (2009), 230–272. <https://doi.org/10.1080/07370020902739429> arXiv:<https://doi.org/10.1080/07370020902739429> 40
- [200] 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). <https://doi.org/10.3389/frobt.2019.00128> 56, 57
- [201] Martin Spindler and Raimund Dachselt. 2009. PaperLens: Advanced Magic Lens Interaction above the Tabletop. In *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces (ITS '09)*. ACM, Article 7, 1 pages. <https://doi.org/10.1145/1731903.1731948> 19, 22
- [202] Priyanka Srivastava, Anurag Rimzhim, Palash Vijay, Shruti Singh, and Sushil Chandra. 2019. Desktop VR is Better Than Nonambulatory HMD VR for Spatial Learning. *Frontiers in Robotics and AI* 6 (2019), 50. <https://doi.org/10.3389/frobt.2019.00050> 84
- [203] Markus Steinberger, Manuela Waldner, Marc Streit, Alexander Lex, and Dieter Schmalstieg. 2011. Context-Preserving Visual Links. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (Dec. 2011), 2249–2258. <https://doi.org/10.1109/TVCG.2011.183> 30
- [204] Andrew Suk. 2012. k-Quasi-Planar Graphs. In *Graph Drawing*, Marc van Kreveld and Bettina Speckmann (Eds.). Springer Berlin Heidelberg, 266–277. https://doi.org/10.1007/978-3-642-25878-7_26 41
- [205] T. Sun, Y. Ye, I. Fujishiro, and K. Ma. 2019. Collaborative Visual Analysis with Multi-level Information Sharing Using a Wall-Size Display and See-Through HMDs. In *2019 IEEE Pacific Visualization Symposium (PacificVis) (PacificVis '19)*. IEEE, 11–20. <https://doi.org/10.1109/PacificVis.2019.00010> 6, 19, 20, 22, 23, 26, 50, 54
- [206] Desney Tan, Mary Czerwinski, and George Robertson. 2003. Women Go with the (Optical) Flow. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '03)*. ACM, 209–215. <https://doi.org/10.1145/642611.642649> 83
- [207] Desney Tan, Darren Gergle, Peter Scupelli, and Randy Pausch. 2003. With similar visual angles, larger displays improve spatial performance. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '03)*. ACM, 217–224. <https://doi.org/10.1145/642611.642650> 83
- [208] Desney S. Tan and Mary Czerwinski. 2003. Information Voyeurism: Social Impact of Physically Large Displays on Information Privacy. In *CHI '03 Extended Abstracts*

- on *Human Factors in Computing Systems (CHI EA '03)*. ACM, 748–749. <https://doi.org/10.1145/765891.765967> 15
- [209] 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> 13, 56
- [210] Vildan Tanriverdi and Robert J. K. Jacob. 2000. Interacting with Eye Movements in Virtual Environments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '00)*. ACM, 265–272. <https://doi.org/10.1145/332040.332443> 29
- [211] Unity Technologies. 2019. *Unity Manual [online]*. <http://docs.unity3d.com/Manual/index.html> Accessed 29 June 2019. 38
- [212] Lucia Terrenghi, Aaron Quigley, and Alan Dix. 2009. A taxonomy for and analysis of multi-person-display ecosystems. *Personal and Ubiquitous Computing* 13, 8 (2009), 583. <https://doi.org/10.1007/s00779-009-0244-5> 1
- [213] Daniel Thalmann, Jun Lee, and Nadia Magnenat Thalmann. 2016. An Evaluation of Spatial Presence, Social Presence, and Interactions with Various 3D Displays. In *Proceedings of the 29th International Conference on Computer Animation and Social Agents (CASA '16)*. ACM, 197–204. <https://doi.org/10.1145/2915926.2915954> 84
- [214] Bruce H. Thomas, Michael Marner, Ross T. Smith, Neven Abdelaziz Mohamed Elsayed, Stewart Von Itzstein, Karsten Klein, Matt Adcock, Peter Eades, Andrew Irlitti, Joanne Zucco, Timothy Simon, James Baumeister, and Timothy Suthers. 2014. Spatial augmented reality — A tool for 3D data visualization. In *2014 IEEE VIS International Workshop on 3DVis (3DVis)*. 45–50. <https://doi.org/10.1109/3DVis.2014.7160099> 19
- [215] James J. Thomas and Kristin A. Cook. 2006. A Visual Analytics Agenda. *IEEE Comput. Graph. Appl.* 26, 1 (Jan. 2006), 10–13. <https://doi.org/10.1109/MCG.2006.5> 1, 25
- [216] C. Tominski, J. Abello, F. van Ham, and H. Schumann. 2006. Fisheye Tree Views and Lenses for Graph Visualization. In *Tenth International Conference on Information Visualisation (IV'06)*. IEEE, 17–24. <https://doi.org/10.1109/IV.2006.54> 28, 31
- [217] 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 (Seoul, Republic of Korea) (CHI '15)*. Association for Computing Machinery, New York, NY, USA, 3255–3264. <https://doi.org/10.1145/2702123.2702129> 15

- [218] 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> 54
- [219] Stéphane Valès, Carole Dupré, Hélène Gaspard-Boulinc, Stéphane Conversy, Claire Ollagnon, Vincent Peyruquéou, and Joel Viala. 2008. MAMMI Phase 3 - Exploring collaborative workspaces for air traffic controllers in the scope of SESAR. In *INO 2008, 7th Eurocontrol Innovative Research Workshop & Exhibition*. Paris, France, pp 15–24. <https://hal-enac.archives-ouvertes.fr/hal-01022265> 12
- [220] Frank van Ham and Adam Perer. 2009. “Search, Show Context, Expand on Demand”: Supporting Large Graph Exploration with Degree-of-Interest. *IEEE Transactions on Visualization and Computer Graphics* 15, 6 (Nov. 2009), 953–960. <https://doi.org/10.1109/TVCG.2009.108> 28
- [221] Luc Vandenamee, Hoorieh Afkari, Johannes Hermen, Louis Deladiennée, Christian Moll, and Valérie Maquil. 2022. DeBORAH: A Web-Based Cross-Device Orchestration Layer. In *Proceedings of the 2022 International Conference on Advanced Visual Interfaces (AVI 2022)*. ACM, Article 58, 3 pages. <https://doi.org/10.1145/3531073.3534483> 16
- [222] Vinoba Vinayagamorthy, Maxine Glancy, Christoph Ziegler, and Richard Schäffer. 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> 19
- [223] Tatiana Von Landesberger, Arjan Kuijper, Tobias Schreck, Jörn Kohlhammer, Jarke Van Wijk, Jean-Daniel Fekete, and Dieter W. Fellner. 2011. Visual Analysis of Large Graphs: State-of-the-Art and Future Research Challenges. *Computer Graphics Forum* 30, 6 (Jan. 2011), 1719–1749. <https://doi.org/10.1111/j.1467-8659.2011.01898.x> 30
- [224] Ulrich von Zadow, Wolfgang Büschel, Ricardo Langner, and Raimund Dachsel. 2014. SleeD: Using a Sleeve Display to Interact with Touch-sensitive Display Walls. In *Proceedings of the Ninth ACM International Conference on Interactive Tabletops and Surfaces (ITS '14)*. ACM, 129–138. <https://doi.org/10.1145/2669485.2669507> 26
- [225] Jorge Wagner Filho, Marina Rey, Carla Freitas, and Luciana Nedel. 2018. Immersive Visualization of Abstract Information: An Evaluation on Dimensionally-Reduced Data Scatterplots. In *Proceedings of the Conference on Virtual Reality and 3D User Interfaces (VR '18)*. IEEE, 483–490. <https://doi.org/10.1109/VR.2018.8447558> 82, 84
- [226] Jorge Wagner Filho, Wolfgang Stuerzlinger, and Luciana Nedel. 2020. Evaluating an Immersive Space-Time Cube Geovisualization for Intuitive Trajectory Data Exploration.

- IEEE Transactions on Visualization and Computer Graphics* 26, 1 (2020), 514–524. <https://doi.org/10.1109/TVCG.2019.2934415> 84
- [227] 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> 103
- [228] 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> 54
- [229] James R. Wallace, Stacey D. Scott, Taryn Stutz, Tricia Enns, and Kori Inkpen. 2009. Investigating Teamwork and Taskwork in Single- and Multi-display Groupware Systems. *Personal Ubiquitous Comput.* 13, 8 (Nov. 2009), 569–581. <https://doi.org/10.1007/s00779-009-0241-8> 26
- [230] 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> 3, 19
- [231] Colin Ware. 2004. *Information Visualization: Perception for Design* (2nd ed.). Morgan Kaufmann, San Francisco, CA, USA. <https://doi.org/10.1016/B978-1-55860-819-1.X5000-6> 44
- [232] Colin Ware and Peter Mitchell. 2008. Visualizing Graphs in Three Dimensions. *ACM Trans. Appl. Percept.* 5, 1, Article 2 (Jan. 2008), 15 pages. <https://doi.org/10.1145/1279640.1279642> 50, 84
- [233] Gordon Wetzstein, Robert Konrad, Nitish Padmanaban, and Hayato Ikoma. 2017. Build Your Own VR System: An Introduction to VR Displays and Cameras for Hobbyists and Educators. In *ACM SIGGRAPH 2017 Courses (SIGGRAPH '17)*. ACM, Article 14, 786 pages. <https://doi.org/10.1145/3084873.3084928> 87
- [234] Daniel Wigdor, Chia Shen, Clifton Forlines, and Ravin Balakrishnan. 2007. Perception of Elementary Graphical Elements in Tabletop and Multi-Surface Environments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*. ACM, 473–482. <https://doi.org/10.1145/1240624.1240701> 13, 17
- [235] 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> 19, 22

- [236] N. Wong, S. Carpendale, and S. Greenberg. 2003. Edgelens: an interactive method for managing edge congestion in graphs. In *IEEE Symposium on Information Visualization (InfoVis '03)*. IEEE, 51–58. <https://doi.org/10.1109/INFVIS.2003.1249008> 28, 31
- [237] Mike Wu and Ravin Balakrishnan. 2003. Multi-Finger and Whole Hand Gestural Interaction Techniques for Multi-User Tabletop Displays. In *Proceedings of the 16th Annual ACM Symposium on User Interface Software and Technology (UIST '03)*. ACM, 193–202. <https://doi.org/10.1145/964696.964718> 15
- [238] Yalong Yang, Tim Dwyer, Bernhard Jenny, Kim Marriott, Maxime Cordeil, and Haohui Chen. 2019. Origin-Destination Flow Maps in Immersive Environments. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2019), 693–703. <https://doi.org/10.1109/TVCG.2018.2865192> 49
- [239] Beth Yost, Yonca Haciahmetoglu, and Chris North. 2007. Beyond Visual Acuity: The Perceptual Scalability of Information Visualizations for Large Displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*. ACM, 101–110. <https://doi.org/10.1145/1240624.1240639> 1, 82, 83
- [240] Beth Yost and Chris North. 2006. The Perceptual Scalability of Visualization. *IEEE Transactions on Visualization and Computer Graphics* 12, 5 (2006), 837–844. <https://doi.org/10.1109/TVCG.2006.184> 14
- [241] Johannes Zagermann, Ulrike Pfeil, Roman Rädle, Hans-Christian Jetter, Clemens Klok-mose, and Harald Reiterer. 2016. When Tablets Meet Tabletops: The Effect of Table-top Size on Around-the-Table Collaboration with Personal Tablets. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, 5470–5481. <https://doi.org/10.1145/2858036.2858224> 16
- [242] 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> 18, 22

APPENDIX



Supplementary Material for Chapter 4

Supplemental Material

Evaluating the Extension of Wall Displays with AR for Collaborative Work

Contents


1	View of the final results	2
1.1	Wall Classification	2
1.2	Wall+AR Classification	3
1.3	Bird View Wall+AR Classification	4
1.4	Wall Story	5
1.5	Wall+AR Story	6
1.6	Bird View Wall+AR Story	7
2	Heatmap of participants and Surfaces	8
2.1	Classification Wall	8
2.2	Classification Wall+AR	9
2.3	Story Wall	10
2.4	Story Wall+AR	11

1 View of the final results


1.1 Wall Classification



(a) Pair 0 Wall Classification



(b) Pair 1 Wall Classification



(c) Pair 2 Wall Classification




(d) Pair 3 Wall Classification



(e) Pair 4 Wall Classification




(f) Pair 5 Wall Classification




(g) Pair 6 Wall Classification




(h) Pair 7 Wall Classification




(i) Pair 8 Wall Classification



(j) Pair 9 Wall Classification




(k) Pair 10 Wall Classification




(l) Pair 11 Wall Classification


1.2 Wall+AR Classification




(a) Pair 0 Wall+AR Classification




(b) Pair 1 Wall+AR Classification



(c) Pair 2 Wall+AR Classification




(d) Pair 3 Wall+AR Classification




(e) Pair 4 Wall+AR Classification




(f) Pair 5 Wall+AR Classification




(g) Pair 6 Wall+AR Classification




(h) Pair 7 Wall+AR Classification




(i) Pair 8 Wall+AR Classification



(j) Pair 9 Wall+AR Classification




(k) Pair 10 Wall+AR Classification




(l) Pair 11 Wall+AR Classification


1.3 Bird View Wall+AR Classification




(a) Pair 0 Wall+AR Classification




(b) Pair 1 Wall+AR Classification




(c) Pair 2 Wall+AR Classification




(d) Pair 3 Wall+AR Classification




(e) Pair 4 Wall+AR Classification




(f) Pair 5 Wall+AR Classification




(g) Pair 6 Wall+AR Classification




(h) Pair 7 Wall+AR Classification




(i) Pair 8 Wall+AR Classification



(j) Pair 9 Wall+AR Classification




(k) Pair 10 Wall+AR Classification




(l) Pair 11 Wall+AR Classification


1.4 Wall Story



(a) Pair 0 Wall Story




(b) Pair 1 Wall Story



(c) Pair 2 Wall Story




(d) Pair 3 Wall Story



(e) Pair 4 Wall Story




(f) Pair 5 Wall Story



(g) Pair 6 Wall Story



(h) Pair 7 Wall Story




(i) Pair 8 Wall Story



(j) Pair 9 Wall Story




(k) Pair 10 Wall Story



(l) Pair 11 Wall Story


1.5 Wall+AR Story




(a) Pair 0 Wall+AR Story




(b) Pair 1 Wall+AR Story




(c) Pair 2 Wall+AR Story




(d) Pair 3 Wall+AR Story



(e) Pair 4 Wall+AR Story



(f) Pair 5 Wall+AR Story




(g) Pair 6 Wall+AR Story



(h) Pair 7 Wall+AR Story

No data


(i) Pair 8 Wall+AR Story



(j) Pair 9 Wall+AR Story




(k) Pair 10 Wall+AR Story




(l) Pair 11 Wall+AR Story


1.6 Bird View Wall+AR Story




(a) Pair 0 Wall+AR Story




(b) Pair 1 Wall+AR Story




(c) Pair 2 Wall+AR Story




(d) Pair 3 Wall+AR Story




(e) Pair 4 Wall+AR Story



(f) Pair 5 Wall+AR Story




(g) Pair 6 Wall+AR Story




(h) Pair 7 Wall+AR Story

No data


(i) Pair 8 Wall+AR Story



(j) Pair 9 Wall+AR Story















(k) Pair 10 Wall+AR Story















(l) Pair 11 Wall+AR Story

2 Heatmap of participants and Surfaces













2.1 Classification Wall




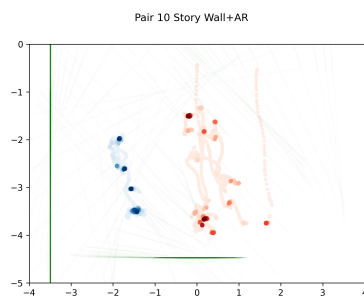









2.2 Classification Wall+AR



2.3 Story Wall



2.4 Story Wall+AR



Titre: Environnements de Réalité Physique et Augmentée utilisés pour l'analyse visuelle collaborative

Mots clés: Interaction homme-ordinateur ; Réalité augmentée ; Data visualisation ; Travail collaboratif ; Coopération (psychologie) ; Systèmes d'affichage

Résumé:

La quantité de données produite chaque jour est en croissance constante. Cette masse d'information demande de trouver des solutions pour les analyser et les visualiser. Les murs d'écrans sont de bon candidat pour résoudre ce problème, car ils permettent d'afficher une grande quantité de données. Cependant, ce sont des écrans partagés, conçus pour plusieurs personnes, ce qui ne permet pas d'avoir des espaces de travail privés. De plus, un mur d'écran reste un périphérique d'affichage physique qui est coûteux à agrandir lorsque la quantité de contenu à afficher devient trop grande.

Dans ce manuscrit, j'explore la combinaison de la Réalité Augmentée (RA) et d'un mur d'écrans, pour premièrement, ajouter un affichage d'informations personnalisées; et deuxièmement, étendre virtuellement l'espace de rendu. Nous concentrons nos efforts sur l'étude de la synergie de ces deux dispositifs à travers deux questions de recherche :

QR1 - Comment la RA peut-elle aider à l'exploration et la navigation d'un réseau sur un mur d'écrans, en ayant un visuel et une navigation personnelle ?

QR2 - Comment l'extension du mur d'écran avec la RA change-t-elle l'exploitation de l'espace par les utilisateurs durant une session de collaboration ?

Pour RQ1, j'étudie le cas de l'utilisation de la RA avec un mur d'écrans pour la navigation de réseaux. Nous présentons 4 techniques, proposant différents mécanismes d'aide visuelle à la navigation. Nous évaluons ensuite ces 4 techniques lors de deux expériences, une de parcours de chemin de faible précision (sélection de chemin) et une autre de parcours de chemin de haute précision (tracé de chemin). Pour la tâche de sélection de chemin, une connexion persistante entre le curseur en RA et le réseau sur le mur permet de meilleur résultat. Dans le cas du tracé de chemin, nous observons

qu'une connexion plus légère offre une meilleure performance. De plus, nous montrons la viabilité d'un système où l'interaction se fait en RA, pour la navigation de réseau sur des écrans partagés.

Pour RQ2, je porte mon attention sur l'espace disponible devant le mur d'écrans, et envisage l'utilisation de cet espace pour étendre le mur d'écran avec de la RA. Je présente un système, étendant le mur d'écrans grâce à la RA en utilisant les dimensions de l'espace disponible devant le Mur. Nous comparons notre système avec un mur d'écran seul avec deux tâches collaboratives. Nous observons qu'avec notre système, les utilisateurs utilisent principalement l'espace virtuel. Bien que cela crée un surcoût pour l'interaction, nous n'observons aucune différence de performance avec le mur d'écran seul, et nous observons un vrai bénéfice à cet espace supplémentaire.

La complexité de mise en place des systèmes précédemment étudiés nous a amené à étudier une manière moins coûteuse d'utiliser un mur d'écran:

QR3 - Peut-on émuler un mur d'écran avec un casque de RV ?

Cette question ayant plusieurs facettes, je me concentre sur la capacité des casques de RV d'atteindre la résolution nécessaire à répliquer l'expérience d'utilisation d'un mur haute résolution. Pour ce faire, nous étudions le modèle optique des casques de RV et le comparons au modèle de la vision humaine. Notre analyse indique que les casques actuels n'ont pas une résolution suffisante pour émuler un mur de haute résolution. Nous confirmons cette analyse avec une expérience pilote qui compare un mur d'écrans, deux casques de RV et une émulation d'un casque de RV parfait.

Je conclus ce manuscrit en discutant plus largement des différentes manières pour les murs d'écran d'être complétés par les casques d'AR et des implications de leur remplacement par des casques de VR/AR, et j'élaborerai sur les directions de recherches qu'ouvre cette thèse.

Title: Augmented and Physical reality environments used in collaborative visual analysis

Keywords: Human-Computer Interaction; Augmented Reality; Data Visualization; Collaboration; Co-operation(Psychology); Displays

Summary:

The amount of data produced every day is constantly growing. This mass of raw information requires finding solutions to analyze and visualize it. Wall displays are a good candidate to solve this issue, allowing to display a large amount of data. However, they are shared devices made to accommodate several people, which does not allow for the existence of a private workspace. Moreover, wall displays are still physical devices that can not be easily extended when the content to display becomes too large.

In this manuscript, I explore the combination of Augmented Reality (AR) and a wall display to: first, add a personalized information space onto a wall display; and second, extend the rendering space. We focus on using an augmented reality headset (AR-HMD) to display content complementary to the context offered by the wall. We study the synergy of these two devices through two research questions:

RQ1 - How can AR assist exploration and navigation of a network on a wall display, through personal view and navigation?

RQ2 - How does extending the screen space of a wall display using AR changes users' space exploitation during a collaborative session?

For RQ1, I study the case of using AR with a wall display for navigating networks. I present four techniques, proposing different visual aid mechanisms for navigation. I then evaluate these four techniques in two experiments, one for low accuracy path following in the network (path selection) and one for high accuracy path following (path tracing). We report that for the path selection task, a persistent connection between the cursor in AR and the network on the wall obtained better results. In the case of path tracing, we ob-

serve that a lighter connection offered better performance. Moreover, we show the feasibility of a system where the interaction is done privately in AR for network navigation on shared screens.

For RQ2, I focus on the space available in front of the wall display and consider using AR to extend the wall display. I present a system extending a wall display with AR, taking advantage of the space in front of the wall. We compare our system that combines a wall display with AR with a wall display alone using two collaborative tasks. We observe that users extensively use the available virtual space with our system. Although this creates an additional cost of interaction, we observe no performance difference and a real benefit of this extra space.

The complexity of setting up the previously studied systems led us to study a cheaper way to use a wall display:

RQ3 - Can we emulate a wall display inside a VR headset?

This question has many aspects, and I focus on the capacity for VR headsets to reach the necessary resolution to replicate the user experience of a high-resolution wall display. For this, we study the optic model of VR headsets and compare it to the human vision model. Our analysis indicates that current headsets need a higher resolution to emulate a wall display. We confirm our analysis by running a pilot study comparing a wall display, two VR headsets, and an emulation of a perfect VR headset.

I conclude this manuscript by discussing the different ways for a wall display to be combined with AR headsets, the requirements for replacing wall displays by VR/AR headsets, and by elaborating on the future work this thesis opens.