Augmented reality environments for the interactive exploration of 3D data

Xiyao Wang

Xiyao Wang. Augmented reality environments for the interactive exploration of 3D data. Human-Computer Interaction [cs.HC]. Université Paris-Saclay, 2020. English. NNT : 2020UPASG052 . tel-03170289

HAL Id: tel-03170289
https://tel.archives-ouvertes.fr/tel-03170289
Submitted on 16 Mar 2021

Augmented Reality Environments for the Interactive Exploration of 3D Data

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
Unité de recherche: Université Paris-Saclay, CNRS, Inria, LRI, 91400 Orsay, France
Référent: : Faculté des sciences d'Orsay

Thèse présentée et soutenue à Gif sur Yvette, le 16/12/2020, par

Xiyao WANG

Composition du jury:

Damien Rohmer Professeur, Ecole Polytechnique	Président
Raimund Dachselt Professeur, Dresden University of Technology	Rapporteur & examinateur
Christophe Hurter Professeur, Ecole Nationale de l'Aviation Civile	Rapporteur & examinateur
Jian Chen Professeure associée, Ohio State University	Examinatrice
Jeanne Vézien Ingénieure de recherche, Université Paris-Saclay, CNRS, Limsi	Examinatrice
Tobias Isenberg Directeur de recherche, Université Paris-Saclay, CNRS, Inria, LRI	Directeur de thèse
Mehdi Ammi Professeur, Université Paris 8	Co-directeur de thèse & examinateur
David Rousseau Directeur de recherche, Université Paris-Saclay, CNRS, IJ-CLab	Invité

Thèse de doctorat

NNT: 2020UPASG052

ABSTRACT

Keywords: 3D data, tangible interaction, scientific visualisation, interaction/HCI, touch input, augmented reality.

Exploratory visualization of 3D data is fundamental in many scientific domains. Traditionally, experts use a PC workstation and rely on mouse and keyboard to interactively adjust the view to observe the data. This setup provides immersion through interaction—users can precisely control the view and the parameters, but it does not provide any depth clues which can limit the comprehension of large and complex 3D data. Virtual or augmented reality (V/AR) setups, in contrast, provide visual immersion with stereoscopic views. Although their benefits have been proven, several limitations restrict their application to existing workflows, including high setup/maintenance needs, difficulties of precise control, and, more importantly, the separation from traditional analysis tools.

To benefit from both sides, we thus investigated a hybrid setting combining an AR environment with a traditional PC to provide both interactive and visual immersions for 3D data exploration. We closely collaborated with particle physicists to understand their general working process and visualization requirements to motivate our design.

First, building on our observations and discussions with physicists, we built up a prototype that supports fundamental tasks for exploring their datasets. This prototype treated the AR space as an extension to the PC screen and allowed users to freely interact with each using the mouse. Thus, experts could benefit from the visual immersion while using analysis tools on the PC. An observational study with 7 physicists in CERN validated the feasibility of such a hybrid setting, and confirmed the benefits. We also found that the large canvas of the AR and walking around to observe the data in AR had a great potential for data exploration. However, the design of mouse interaction in AR and the use of PC widgets in AR needed improvements.

Second, based on the results of the first study, we decided against intensively using flat widgets in AR. But we wondered if using the mouse for navigating in AR is problematic compared to high degrees of freedom (DOFs) input, and then attempted to investigate if the match or mismatch of dimensionality between input and output devices play an important role in users' performance. Results of user studies (that compared the performance of using mouse, space mouse, and tangible tablet paired with the screen or the AR space) did not show that the (mis-)match was important. We thus concluded that the

dimensionality was not a critical point to consider, which suggested that users are free to choose any input that is suitable for a specific task. Moreover, our results suggested that the mouse was still an efficient tool compared to high DOFs input. We can therefore validate our design of keeping the mouse as the primary input for the hybrid setting, while other modalities should only serve as an addition for specific use cases.

Next, to support the interaction and to keep the background information while users are walking around to observe the data in AR, we proposed to add a mobile device. We introduced a novel approach that augments tactile interaction with pressure sensing for 3D object manipulation/view navigation. Results showed that this method could efficiently improve the accuracy, with limited influence on completion time. We thus believe that it is useful for visualization purposes where a high accuracy is usually demanded.

Finally, we summed up in this thesis all the findings we have and came up with an envisioned setup for a realistic data exploration scenario that makes use of a PC workstation, an AR headset, and a mobile device. The work presented in this thesis shows the potential of combining a PC workstation with AR environments to improve the process of 3D data exploration and confirms its feasibility, all of which will hopefully inspire future designs that seamlessly bring immersive visualization to existing scientific workflows.

SYNTHÈSE

Mots-clés: Données 3D,interaction tangible,visualisation scientifique, interaction/IHM, entrée tactile, réalité augmentée.

La visualisation exploratoire des données 3D est fondamentale dans des domaines scientifiques. Traditionnellement, les experts utilisent un PC et s'appuient sur la souris pour ajuster la vue. Cette configuration permet l'immersion par interaction—l'utilisateur peut contrôler précisément la vue, mais elle ne fournit pas de profondeur, qui limite la compréhension de données complexes. La réalité virtuelle ou augmentée (RV/A), en revanche, offre une immersion visuelle avec des vues stéréoscopiques. Bien que leurs avantages aient été prouvés, plusieurs points limitent leur application, notamment les besoins élevés de configuration/maintenance, les difficultés de contrôle précis et, plus important, la séparation des outils d'analyse traditionnels.

Pour bénéficier des deux côtés, nous avons donc étudié un système hybride combinant l'environnement RA avec un PC pour fournir des immersions interactives et visuelles. Nous avons collaboré étroitement avec des physiciens des particules afin de comprendre leur processus de travail et leurs besoins de visualisation pour motiver notre conception.

D'abord, basé sur nos discussions avec les physiciens, nous avons construit un prototype qui permet d'accomplir des tâches pour l'exploration de leurs données. Ce prototype traitait l'espace RA comme une extension de l'écran du PC et permettait aux utilisateurs d'interagir librement avec chacun d'eux avec la souris. Ainsi, les experts pouvaient bénéficier de l'immersion visuelle et utilisent les outils d'analyse sur PC. Une étude observationnelle menée avec 7 physiciens au CERN a validé la faisabilité et confirmé les avantages. Nous avons également constaté que la grande toile du RA et le fait de se déplacer pour observer les données dans le RA présentaient un grand potentiel. Cependant, la conception de l'interaction de la souris et l'utilisation de widgets dans la RA devaient être améliorés.

Ensuite, nous avons décidé de ne pas utiliser intensivement les widgets plats dans la RA. Mais nous nous sommes demandé si l'utilisation de la souris pour naviguer dans la RA est problématique, et nous avons ensuite tenté d'étudier si la correspondance de la dimensionnalité entre les dispositifs d'entrée et de sortie joue un rôle important. Les résultats des études (qui ont comparé la performance de l'utilisation de la souris, de la souris spatiale et de la tablette tangible couplée à l'écran ou à l'espace de RA) n'ont pas montré que la correspondance était importante. Nous avons donc conclu que la dimensionnalité n'était

pas un point critique à considérer, ce qui suggère que les utilisateurs sont libres de choisir toute entrée qui convient à une tâche spécifique. De plus, nos résultats ont montré que la souris restait un outil efficace. Nous pouvons donc valider notre conception et conserver la souris comme entrée principale, tandis que les autres modalités ne devraient servir que comme complément pour des cas spécifiques.

Ensuite, pour favoriser l'interaction et conserver les informations pendant que les utilisateurs se déplacent en RA, nous avons proposé d'ajouter un appareil mobile. Nous avons introduit une nouvelle approche qui augmente l'interaction tactile avec la détection de pression pour la navigation 3D. Les résultats ont montré que cette méthode pouvait améliorer efficacement la précision, avec une influence limitée sur le temps. Nous pensons donc qu'elle est utile à des tâches de vis où une précision est exigée.

Enfin, nous avons résumé tous les résultats obtenus et imaginé un scénario réaliste qui utilise un poste de travail PC, un casque RA et un appareil mobile. Les travaux présentés dans cette thèse montrent le potentiel de la combinaison d'un PC avec des environnements de RA pour améliorer le processus d'exploration de données 3D et confirment sa faisabilité, ce qui, nous l'espérons, inspirera la future conception qui apportera une visualisation immersive aux flux de travail scientifiques existants.

PUBLICATIONS

Full papers - Accepted

1. **Xiyao Wang**, Lonni Besançon, Mehdi Ammi, Tobias Isenberg. Augmenting Tactile 3D Data Navigation With Pressure Sensing. *Computer Graphics Forum (EuroVis 2019)*.

2. **Xiyao Wang**, Lonni Besançon, David Rousseau, Mickael Sereno, Mehdi Ammi, Tobias Isenberg. Towards an Understanding of Augmented Reality Extensions for Existing 3D Data Analysis Tools. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*.

3. Mickael Sereno, **Xiyao Wang**, Lonni Besançon, Michael McGuffin, Tobias Isenberg. Collaborative Work in Augmented Reality: A Survey. *IEEE Transactions on Visualization and Computer Graphics*

Full papers - Submitted and under review

4. **Xiyao Wang**, Lonni Besançon, Mehdi Ammi, Tobias Isenberg. Understanding Differences between Combinations of 2D and 3D Input and Output Devices for 3D Data Visualization. *International Journal of Human-Computer Studies*

Workshop paper

5. **Xiyao Wang**, Lonni Besançon, Florimond Guéniat, Mickael Sereno, Mehdi Ammi, Tobias Isenberg. A Vision of Bringing Immersive Visualization to Scientific Workflows. *CHI-IA 2019-Workshop on Immersive Analytics*

Extended Abstracts

6. **Xiyao Wang**, Lonni Besançon, Mehdi Ammi, Tobias Isenberg. Augmenting Tactile 3D Data Manipulation With Pressure Sensing. *IEEE VIS2017*

7. **Xiyao Wang**, Lonni Besançon, Mehdi Ammi, Tobias Isenberg. Navigation Tactile 3D Augmentée pour Mobiles. *Journées Visu 2018*

ACKNOWLEDGMENTS

Time flies, the past three years' research experience as a Ph.D. student has passed like a flash. My memories of all unforgettable moments feel as if they occurred just now. Throughout the journey of doing my Ph.D. and the writing of this dissertation, I have received a great deal of support and assistant from people around me. I am grateful to a lot of colleagues, friends, and even unknown volunteer participants for encouraging and helping me to start it, persevere with it and finally finish it. Here, I want to express my sincere gratitude.

First, I would like to express my deepest appreciation to my primary supervisor Dr. Tobias Isenberg. People usually describe the Ph.D. study as a lonely journey because we need to reflect on everything, and quite often resolved difficulties by ourselves. Sometimes we may get lost, but his research expertise has lighted up my way to work on this thesis. He fits my envision of an ideal supervisor in the traditional Chinese culture, as we usually describe a teacher as a person to "propagate the doctrine, impart professional knowledge, and resolve doubts". I have learned from his daily behaviors the quality of being persistent and rigorous. Besides, his supports for me do not only limit to scientific work, but also in general life: he gives me advice in everyday life and guides me to define my future career, and to find potential opportunities. Such supports are not less important than scientific supervision for me, as a foreign student living abroad alone.

Similar gratitude belongs to my second supervisor Dr. Mehdi Ammi. Even though we had less chance to discuss or work together in detail due to his move to another university, our occasional meetings are still very inspirational for me, not to mention all the generous help for finding the resources that I needed.

Besides my two official supervisors, I also want to express a special thanks to Dr. Lonni Besancon. Quite like my primary supervisor Dr. Tobias Isenberg, I also see him as a leading person to bring me really into the fields. His help, even before the start of my Ph.D., has encouraged me and gave me the confidence to pursue a research career. Many of my basic research spirits and methods are learned from him.

Without the support of all my Ph.D. defense jury members, I could not successfully finish my defense and come to this moment. So, I wish to pay my special regards to Dr. Damien Rohmer, Dr. Raimund Dachselt, Dr. Christophe Hurter, Dr. Jian Chen, and Dr. Jeanne Vezien, all of whom have participated in my final defense and have given me valuable feedback for my work, and even suggestions for the future. I would particularly thank the two reviewers Dr. Raimund Dachselt, Dr. Christophe Hurter who have carefully read my thesis and gave detailed feedback.

Specifically, Dr. David Rousseau was my main collaborator in particle physics and was invited to be part of my Ph.D. defense. But he is far more than a general collaborator for me. He has not only patiently explained every basic physical law to help me understand the main application field of my thesis but also has made connections between us and the researchers from CERN. Only with his help, I could have finished one of the most important portions of my thesis and have had the luck to visit CERN to get a better understanding of the domain of high-energy physics.

I spent a wonderful three years' period at the Aviz team at Inria Saclay France. With Inria's funding, I was able to start my research career and participate in many international or national conferences to meet and discuss with different people. Besides my gratitude to my funding institution, I also want to express my regards to all the colleagues I have met or collaborated with, including Jean-Daniel Fekete, Pierre Dragicevic, Petra Isenberg, Frédéric Vernier, Steve Haroz, Catherine Plaisant, Paolo Buono, Luiz Morais, Gaëlle Richer, Paola Valdivia, Tanja Blascheck, Qing Chen, Christophe Kinkeldey, Jiayi Hong, Alaul Islam, Alexis Pister, Mickaël Sereno, Natkamon Tovanich, Lijie Yao, Evanthia Dimara, Yuheng Feng, Sarkis Halladjian, Yvonne Jansen, Nicolas Pezzotti, Kejin Jin, Ranjini Aravind, Yumin Hong, and Yuanyang Zhong.

Particularly, the team leader Dr. Jean-Daniel Fekete has helped me throughout the journey for answering my various questions, related to academia, administration, and general French life. My initial application for a Ph.D. position was not quite straightforward, and I will always remember his help in arguing with Inria to accept me as a student in the team. I also learned a lot from Dr. Pierre Dragicevic for both the new practices in HCI statistic methods and French wine culture, and want to express my gratitude to Dr. Petra Isenberg and Dr. Steve Haroz for helping me revise part of my submissions.

I may or may not know the participants of my experiments in advance, but I sincerely thank them for participating in any of my studies as a volunteer. It is only with their time and help that I can achieve my research work. I wish I could list all of them to express my sincere appreciation, but I would not do it as revealing participants' name is violating the rules of GDPR—a European law that we had struggled to figure it out during the past years.

In addition to the scientific aspects, I would like to acknowledge with gratitude, the support and love of my family and every close friend. Even though unfortunately my father has passed away years ago and could not witness any steps of my life anymore, I have still shared with you thousands of times in my heart about the completion of my Ph.D., as no one will be left behind as a family. Because of the busy work of my father from my childhood and these years' time without him, my mother has taken most of the duty of taking care of my education and my general life. Now with this opportunity of almost finishing my life stage as a student, I want to express my gratitude for the support from her throughout my whole life. Even though my language is pale and powerless, my gratitude is sincere and eternal.

And of course, I appreciate all my friends around me, who provided stimulating discussions as well as happy distractions to rest my mind outside of my research.

Last but not the least, I would like to thank my wife, Jianwen Zhang. There were both wonderful and difficult moments during the past three years. I feel that I was the luckiest person in the world as I have you accompany me all the time. We have shared our joys and pain, and I have deeply felt your supports and trust all along, without what I cannot imagine how I could go through this journey to become a doctor in computer sciences today. Every time I think of the moment when I was seriously sick and when you have stopped working and come from another country to take care of me, my moving tears cannot stop flowing. For the rest of our life, I promise that I will do my best to make a nice life for us and to support you as you have supported me.

If I carefully revise my memories of the past, there are too many people that I need to express my gratitude to, including friends and supervisors from China, the Université de Bretagne Occidental, the Université Pierre et Marie Curie, CentraleSupélec, the Université Paris-Saclay, Inria, and all over the world. I might have encountered you during colleague lessons, during user studies, at conferences, or even at a bus stop in front of our lab. I would need another thesis to describe all my gratitude to all of you, so your names might not appear here, but you help will never be ignored, and you will never be forgotten.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACRONYMS

ALICE	A Large Ion Collider Experiment	33
APA	American Psychological Association	111
AR	Augmented Reality	3
ATLAS	A Toroidal LHC ApparatuS	33
CAD	Computer Aided Design	24
CAVE	cave automatic virtual environment	6
CERN	European Organization for Nuclear Research	2
CD	control-to-display	69
CI	confidence interval	73
CMS	The Compact Muon Solenoid	33
DOF	degree of freedom	26
DOFs	degrees of freedom	14
EVE	evolutive virtual environment	14
HCI	Human-Computer Interaction	9
HEP	high-energy physics	1
HMD	head-mounted display	17
MRI	magnetic resonance imaging	2
NHST	Null Hypothesis Significance Testing	111
PC	personal computer	6
RNT	Rotate and Translate	98
RST	rotating-scaling-translating	98
SDK	software development kit	70
TLX	Task Load Index	72
VR	Virtual Reality	3
VR/AR	virtual or augmented reality	16
VP1	Virtual Point 1	33

1D	one-dimensional	102
2D	two-dimensional	2
3D	three-dimensional	1

1

INTRODUCTION

> *Le véritable voyage de découverte*
> *ne consiste pas à chercher de nouveaux paysages,*
> *mais à avoir de nouveaux yeux.*
>
> *The real voyage of discovery*
> *consists not in seeking new landscapes,*
> *but in having new eyes*
>
> Marcel Proust

1.1 IMMERSIVE VS. TRADITIONAL SETUP FOR 3D DATA EXPLORATION

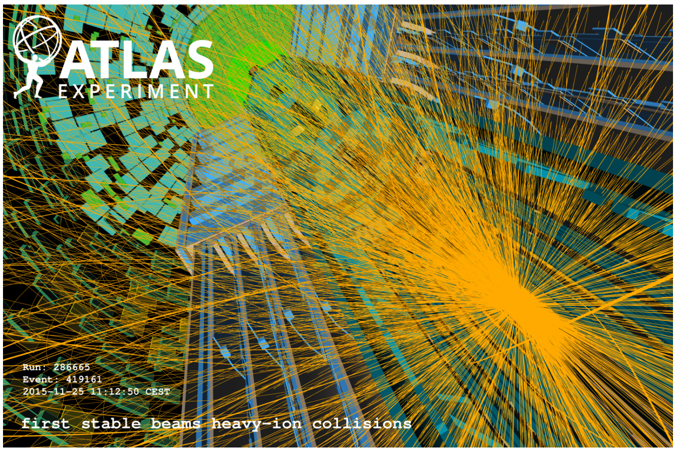

Figure 1.1: Particle collisions visualized by ATLAS VP1, CERN. `https://cds.cern.ch/record/2115422`

Exploratory visualization of three-dimensional (3D) data [Tukey, 1977] is fundamental to many domains in the natural sciences and in medicine. Traditionally, researchers and practitioners use a desktop workstation which is composed of one or several screens, and relies on a mouse and a keyboard as input devices. Such traditional setup provides many controls and exploration tools, and is easily accessible by experts and by general public because it already comes into everyday life and does not require dedicated training, but has major weakness for understanding 3D data. A typical issue is that current datasets used by researchers are growing larger and larger in a short period of time, thus their visualization is becoming more and more complicated. For example, in high-energy physics (HEP), one collision experiment generates more than 10,000 new particle trajectories. Each of them leaves dozens of detection points to analyze, with spatial position and other information, such as the energy. Figure 1.1 shows an image of particle collision visualization used by physicists in European Organization for Nuclear Research (CERN) where each yellow line represents the trajectory of one new-generated particle. This visualization is already a simplified case where only a subset of the dataset is shown, but we can still observe several drawbacks. First, screen-based visualization relies on projection techniques to display 3D content on a two-dimensional (2D) surface (as illustrated in Figure 1.2), thus depth information is lost in any still images. This 3D component can be only understood through interactively manipulating the view, such as rotating and translating. In reality, particles affected by the magnetic fields travel with a curved trajectory in 3D. From this image, we naturally cannot observe such information. Second, after the collisions, generation of particles happens in a limited area, we can notice from this image that it is hard to understand what exactly happens at the center of collision and how the particles start traveling from there, as many of them overlap each other, which introduces more difficulty for understanding.

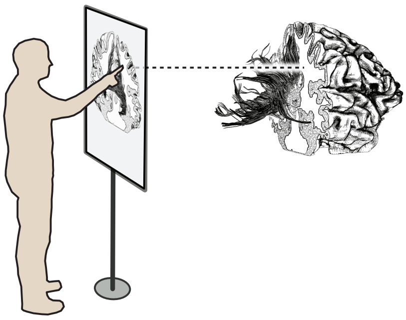

Figure 1.2: 3D content projected on a 2D surface [Isenberg, 2016].

Similar difficulties can also be found in other disciplines dealing with different type of datasets. For example, fluid machinists, biologists and doctors often deal with volumetric datasets. Even with only 128 points per direction of space, one fluid simulation will have $128^3 = 2097152$ nodes. Not to mention the fact that outputs of simulations are usually multiple and time-dependent. To understand its visualization, due to the lack of depth clue on a traditional workstation, experts need to observe and compare several slices at the same time. A typical example is that in a hospital, we can find doctors observe and compare plenty of magnetic resonance imaging (MRI) slices put side by side (Figure 1.3) at the same time.

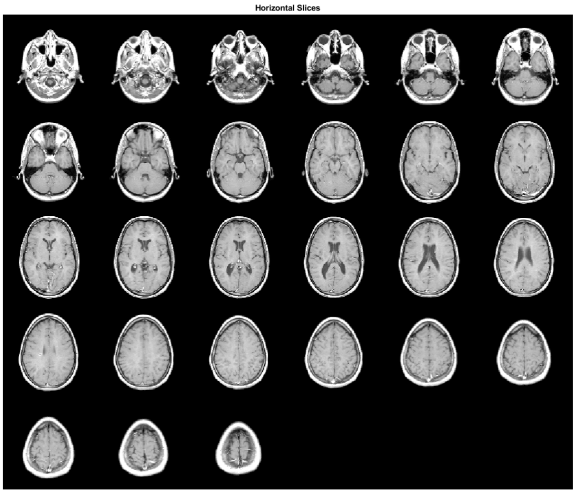

Figure 1.3: Slices of MRI data scan of a human cranium. Image from Matlab.

Nowadays, the rapid development of hardware facilitates the visualization beyond traditional workstation, ranging from small portable devices like smart watches and phones, to large screens or fully virtual environments. For the purpose of improving scientists' experience of exploring and understanding 3D data, the use of immersive environments with Virtual Reality (VR) or Augmented Reality (AR) especially attract researchers' attention. With their stereoscopic displays, 3D data is no longer needed to be projected on a surface, then the spatial understanding is largely enhanced and the occlusion can be reduced. It has already been widely recognized in the literature that visualization enlightens users' understanding of and facilitate the interaction with the large and complex data. For example. many prior work argued that scientific data exploration tasks could take benefits of them (e. g., [Besançon et al., 2017b; Bryson, 1996]). Also, formal studies confirmed the benefits of stereoscopy compared to a normal screen for data understanding (e. g., Figure 1.4). More importantly, VR and AR devices are becoming easily accessible by general public as several commercial products with affordable price are coming into the market, such as Microsoft HoloLens, Oculus Rift, and HTC Vive.

Figure 1.4: Controlled user study proved that Immersion helps data understanding [Prabhat et al., 2008].

Both the traditional workstation and the immersive environments provide a form of immersion in the dataset that is beneficial for the scientists to understand it. The traditional setup provides immersion through interaction: by interactively manipulating the view and exploration tools scientists are able to immerse themselves in the data as they are exploring it. The VR or AR setups (which are referred to immersive environments in this thesis), in contrast, provide immersion through vision as single view can already effectively convey the 3D spatial character of the data, without the need for interactive navigation. Adjusting the view to the 3D tracked position and orientation of the viewer only enhances this effect. However, the immersive environments alone are not without limitation compared to traditional workstations. Major points include, but not limited to the computational power, display resolution, and potential maintenance cost (we present a discussion in Chapter 2). More importantly, scientists or doctors need to run specific analysis tools that are not yet ready with immersive environments (as we discuss in Section 1.2).

To get the best of both worlds, a combination of both types of immersion for the exploratory visualization of 3D data would be highly useful. Some previous literature (e. g., [Besançon, 2018; Isenberg, 2014; Keefe, 2010]) also advocated similar vision, and highlighted that desktops will not be totally replaced by such innovative visualization environments but rather be combined with others to make use of the inherent benefits of each environment.

In this thesis, we attempt to investigate a hybrid setting that combines both interactive and visual immersion, to ultimately bring immersive visualization to existing data analysis workflows.

1.2 FIELDS OBSERVATIONS

The use of a hybrid setting instead of a pure immersive environment is also motivated by our field observations. We closely collaborate with particle physicists to understand their needs and domain-specific requirements (Figure 1.5). We chose the particle physics as our main application domains because it is a typical scientific domain that deals with both spatial and abstract data. Moreover, the visualization of particle collisions has already encountered difficulties using traditional setup with 2D screens as we presented at the beginning of Section 1.1, which makes the physicists themselves also have interests in trying immersive visualization.

Figure 1.5: Our discussions with particle physicists.

In HEP, to explore measured particle collisions event, scientists first need to analyze the data and to eliminate the measuring noise and other irrelevant information related to physical phenomena (for example, the interaction between a particle and detector meshes may introduce useless measures) with statistical tools. And then they relate the measured hit points to reconstruct particle trajectories using specific algorithms. Finally, they explore such processed data to identify and analyze both interesting (for example, several types and a large number of particles could be generated from one collision event, physicists often want to find the few who carry a high energy) and strange events (for example, those particles with a weird constructed trajectory). Once a special part has been identified, they pass to visualization software for data understanding.

We found a major challenge of using immersive environments in their daily workflow is that physicists do not only rely on visualization to explore the data. The initial steps heavily rely on traditional statistical tools (like writing Python scripts) to find and limit their regions of interests. They usually need to regularly switch between the traditional analysis tools and the visualization software to explore the data as illustrated in Figure 1.6.

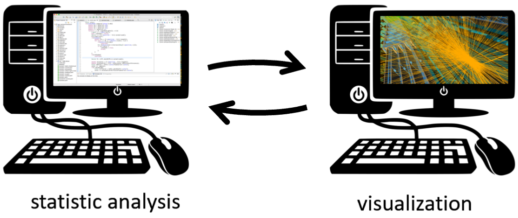

Figure 1.6: Experts need to regularly switch between the analysis tools and visualization tools to explore and understand the data.

According to our observations and discussions, we formulated the following basic requirements (noted as **R**) to bring immersive visualization to their existing workflows.

- **R1** Both traditional analysis and efficient visualization tools are important. The system needs to support easy switch between the two.
- **R2** The immersive systems should be easily integrate to office working environments.

1.3 VISION AND RESEARCH CHALLENGES

R1 largely motivated our vision of using a hybrid system that combines a traditional workstation and an immersive environments. A workstation, will allow experts to continue using the traditional tools that they are familiar with. Based on our experience with different immersive output environments and **R2**, we believe that the use of non-occluded AR headsets (for example, Microsoft's HoloLens) is currently a good solution to provide a common data exploration environment. Such AR headsets immerse users by projecting the data in a stereoscopic view. Compared with large immersive environments like responsive workbench and cave automatic virtual environment (CAVE), AR glasses do not require complex setup and maintenance. Compared with occluding VR headsets, users are not separated from the real word, thus they have more freedom to perform tasks on the desktop, we consequently do not need to replicate existing tools completely in the virtual space, and people can continue to interact with real-world objects (e. g., paper/pen, blackboard). We consider this last point as a major advantage as we observed researchers needing to take traditional notes in their current scientific workflows. The use of AR as an additional output in addition to a personal computer (PC) thus provides an extension of the 2D screen with larger space.

Thus, our general vision is to use a combination of a PC workstation with an AR headset (Figure 1.7). As so, experts can still use their traditional analysis tools on PC while benefiting from the immersive visualization with the AR. With such a setup, several research challenges arise:

Figure 1.7: Our general vision is to combine the traditional workstation with an AR headset to make use of both worlds.

First, for domain experts to be able to perform their data analysis, it is important to clarify which elements can be visualized immersively, and which ones are better used on traditional screens. A challenge is thus how to make the visual transitions between different devices [Isenberg, 2014]. For example, a data exploration and analysis system should provide support to its users to decide what to show on each view as well as how to move a view from desktop to the AR view or in the opposite direction. In addition, existing AR headsets have intrinsic and unchangeable camera parameters and the data display should thus be well adjusted to match these specs. For example, while researchers in both particle physics and fluid dynamics often rely on orthographic projections of their 3D datasets on traditional 2D screens, such views would be equivalent to a flat image in AR. Nonetheless, the projection parameters of AR headsets are equivalent to our normal vision, so that the disadvantages often associated with perspective projection of 3D data may not be as severe as for a general perspective projection. Further investigation is thus needed to understand how to best match the different views between the screen and the AR space.[1]

Second, we do not aim to create a novel environment and to replace current tools. Domains experts are familiar with the data analysis on desktops, the second challenge is thus to design the interaction technique compatible with scientific workflows. We thus need to design ways to interact with popular scientific software (e. g., Python, Paraview, and proprietary software such as MatLab, Virtual Point). While it is easy for experts to interact with them using keyboard and mouse, we need to investigate how to adopt the input to the AR space. More recent forms of input (e. g., tactile/tangible, voice, gestures) are often considered to be intuitive and natural. Yet, the question of whether they are suitable for generic or specific tasks in the context of existing scientific workflows is still open and requires investigation. This research does not only need to address the problems of mapping 2D input to 3D output but also how to create forms of control that is perceived by the intended users as fluid for both desktop and "hologram" representations.

A third challenge lies in the design of dedicated interaction techniques to support such hybrid environments—techniques that make specific use of both views and that seamlessly extend the existing interaction metaphors that the experts are used to on their PC-based tools. We first need to determine which practical tasks require either the traditional or the AR views only, and design appropriate controls (likely captured by the PC) for the AR setting. We also need to support tasks that use both parts of the system and that that allow researchers to easily transition between them. We believe that **R2** does not only apply to output devices, it is also very important while choosing the input devices and designing interaction techniques: we decided to consider only (or mainly) well-established commercial products like

---

[1] For example, the view angle of Microsoft HoloLens is fixed to 18 degrees. In such limited view field, it is hard to visualize the dataset in 3D without proper adjustment of the model.

mice and mobile phones and do not consider those that only exist in lab protocols or that needs complicated assemblage of extra electronic circuits.

Based on the considerations mentioned above, we investigate how to combine both interactive and visual immersions to improve 3D data exploration experience. In this thesis, we narrow down the broad background to specifically investigate the following research questions (noted as **Q**):

Q1 What should a data exploration environment look like that combines an AR display with a traditional workstation?

Q2 How should we treat the AR display compared to a 2D screen and how do we make transitions between the two?

Q3 What should be the appropriate input devices and interaction techniques to work with such a hybrid setting?

Q4 How can data exploration tasks be realized with a hybrid setting?

1.4 THESIS STATEMENT

Based on the discussions above, we argue that using AR technologies to enhance visualization experience involved in scientific process has great potentials. While researchers have extensively investigated different immersive visualization environments in the past, few has focused on their practical application to existing workflows. Studying a practical way to bring immersive visualization to scientific workflows is a key goal of this thesis. Concretely, we investigate the interaction design for a hybrid PC and AR setting with regards to 3D visualization.

We limit the scope with the following restrictions in this thesis:

- **Single user.** While AR provides opportunities for multiple users' collaboration, we limit our investigation to a single-user scenario since it is already a large area to explore, we thus keep the research direction specific and manageable.

- **Human-scale setup.** Even though a lot of immersive visualization environments are large and complex (see discussions in Chapter 2), we limit our considerations to small setups as required by **R2**.

- **Accessible devices.** While many carefully designed input devices have been demonstrated to be efficient for some tasks (e. g., [Fruchard et al., 2019; Klamka et al., 2019]), we first focus on easily-accessible (in another word, commercially-available) devices to fulfill the needs of being easily integrated into scientists' general workflows.

- **Scientific datasets.** We focus on the needs of 3D visualization. We closely collaborated with particle physicists to understand their visualization needs and domain specific requirements, but we also use other scientific datasets like volume data of fluid mechanics to explore more general data exploration questions because our ultimate goal is not to implement a specific tool, but to understand the potentials of such setting and conclude design guidelines that could potentially contribute to the evolution of future working environments. In another word, we expect our findings can be generalized to many scientific domains that deal with a similar type of dataset.

1.4.1 *Methodological approach*

The research questions **Q1–Q4** are, although from different perspectives, still heavily related. Hence, one project could investigate several questions and its results could be a foundation for another one. For an overview, the major research methods of this thesis include:

- **Literature review.** We started by conducting literature review of immersive visualization and related interaction techniques, 3D visualization techniques, and general Human-Computer Interaction (HCI) work. The purpose is to understand state of the art technologies and existing gaps, which would inspire the design of our system.

- **Field observation.** With a special focus on 3D visualization, we closely collaborated with domain experts (particle physicists in this thesis as the particle physics is the main application domain) to understand their traditional working process, interaction needs, and requirements before doing any system design.

- **Interaction design.** An important part of the research was to design the interaction, based on the requirements and previous work.

- **Prototype implementation.** We put considerable effort into implementing prototypes based on our design, including writing shaders to render specific datasets and implementing interaction techniques for different devices/platforms.

- **User evaluation.** We have conducted a series of user studies to access the designed system and to get comments for future improvements. We used an observational study to gather experts' feedback and controlled studies to access the usability of certain techniques. For the observational study, we targeted on domain experts, gave them freedom to use a tool with some tutorials, and gathered their comments and feedback through the think-aloud protocol, interviews, and post-questionnaires. For the controlled studies, we measured quantitative data (e. g., task completion time, accuracy, and workload) and gathered qualitative feedback (e. g., comments and preference).

1.4.2 *Overview*

We first conducted an initial review of the state of the art work in Chapter 2. This chapter provides an overview of existing research of immersive visualization, hybrid visualization system, and interaction techniques for such environments to let us have a global understanding of existing solutions. We reflected on our research goal, summarize limitations of previous works to motivate the work of this thesis.

To start the investigation, we first discussed with several particle physicists to understand their current working procedure and general visualization and interaction needs, and proposed a first prototype of a hybrid AR and PC visualization setting using the mouse to interact with both space. Chapter 3 thus presents the design choices, an observational studies with seven physicists from CERN, and the discussion of results regarding to their comments and feedback. We found that this hybrid prototype is a valid design that can improve their data understanding. We also gathered several different insights that have great potentials to explore such as the large canvas and the support of walking around. However, a major limitation is the interaction part: the use of mouse to control both spaces needs to be revised.

Then, based on the results, especially the problems of using mouse in AR space, we tried to understand if the mouse in 3D is real problematic. Chapter 4 presents an user study to compare users' performance when they use three types of input: mouse, space mouse, and tangible tablet paired with both the screen and the HoloLens. We also tried to understand if the match of dimensionality between the input and output devices play a role to users' performance. Results did not reveal a universal conclusion of which is better, but informed us that it depends a lot on the specific task. We found that the mouse is still a powerful input when accuracy is highly demanded. In fact, this is a quite positive conclusion because we know that the dimensionality mismatch is not critical for designing hybrid visualization systems. Users are flexible to choose different devices depending on the specific tasks. Moreover, it is still good to keep using mouse.

Next, as the use of mobile to visualize data and to control large or virtual spaces increases, Chapter 5 describes a method combing touch and pressure input to augment 3D navigation accuracy.

In Chapter 6, we reflect on the work presented in this thesis and our ultimate goal of paving the way for a continuum of interaction for 3D data visualization [Besançon, 2018; Isenberg, 2014]. Under the scope of this thesis, the continuum means that users would eventually use and interact with different kinds of devices in a seamless manner. Particularly, users would be able to easily switch between traditional workstation and immersive environments. We also reflect on the possible follow-up work that would further extend and develop this concept.

2

RELATED-WORK

Using immersive environments for data exploration has been envisioned for a long time. Bryson [1996], for example, explained that 3D VR environments and scientific visualization naturally match, not only because of the spatial proprieties of scientific data, but also the potential of real-world interactions leveraged by this combination. The research of immersive analytics [Marriott et al., 2018; Dwyer et al., 2018] also promises many advantages for data exploration such as offering spatial visual immersion, and facilitating situated/embodied interaction and collaboration. This chapter provides an overview of the current state of immersive visualization environments and related interaction techniques, while discussing their benefits and limitations with specific regard to scientific workflows.

2.1 VISUAL IMMERSIVE SYSTEMS FOR DATA EXPLORATION

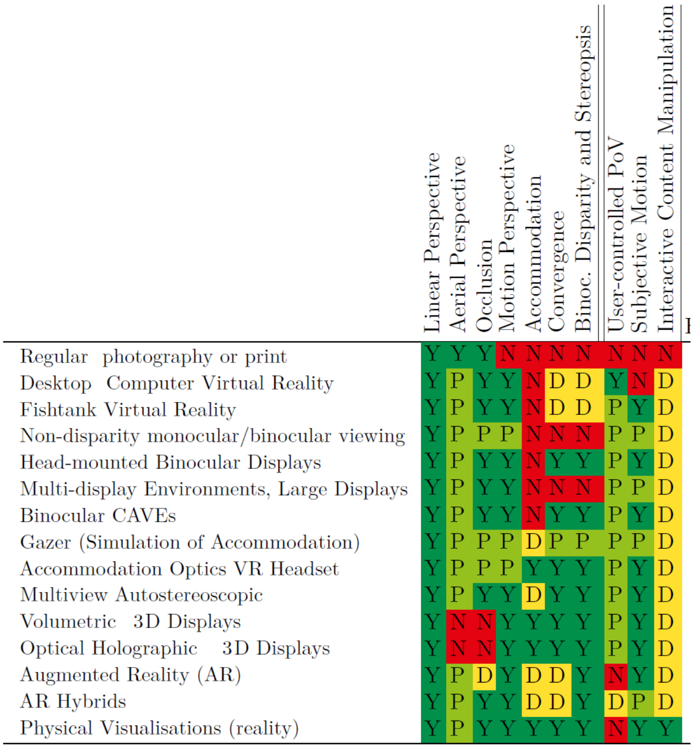

	Linear Perspective	Aerial Perspective	Occlusion	Motion Perspective	Accommodation	Convergence	Binoc. Disparity and Stereopsis	User-controlled PoV	Subjective Motion	Interactive Content Manipulation
Regular photography or print	Y	Y	Y	N	N	N	N	N	N	N
Desktop Computer Virtual Reality	Y	P	Y	Y	N	D	D	Y	N	D
Fishtank Virtual Reality	Y	P	Y	Y	N	D	D	P	Y	D
Non-disparity monocular/binocular viewing	Y	P	P	P	N	N	N	P	P	D
Head-mounted Binocular Displays	Y	P	Y	Y	N	Y	Y	P	Y	D
Multi-display Environments, Large Displays	Y	P	Y	Y	N	N	N	P	P	D
Binocular CAVEs	Y	P	Y	Y	N	Y	Y	P	Y	D
Gazer (Simulation of Accommodation)	Y	P	P	P	D	P	P	P	P	D
Accommodation Optics VR Headset	Y	P	P	P	Y	Y	Y	P	Y	D
Multiview Autostereoscopic	Y	P	Y	Y	D	Y	Y	P	Y	D
Volumetric 3D Displays	Y	N	N	Y	Y	Y	Y	P	Y	D
Optical Holographic 3D Displays	Y	N	N	Y	Y	Y	Y	P	Y	D
Augmented Reality (AR)	Y	P	D	Y	D	D	Y	N	Y	D
AR Hybrids	Y	P	Y	Y	D	D	Y	D	P	D
Physical Visualisations (reality)	Y	P	Y	Y	Y	Y	Y	N	Y	Y

Figure 2.1: A summary of 3D display technologies [Marriott et al., 2018].

Many systems offer stereoscopic view to increase the experience of visual immersion, ranging from large ones that require complex setup to small and portable devices. A recent survey by Fonnet and Prié

[2019] summarized visualization contributions of immersive analytic work for the past few decades and discussed how immersion could be used to better visualize different types of data. Yet, our main focus is how different environments could help the process of exploring and understanding data, rather than studying a specific type of data representation. As for the different immersive displays, Marriott et al. [2018] have made a detailed summary of existing 3D display technologies (Figure 2.1), we do not repeat all the work that has been summarized previously, but quickly go through the several different types of visual immersive systems that are highly investigated for 3D visualization, before discussing their advantages and limitations. Our main focus in this chapter lies on the classic work of different setup on which a lot of novel and specific applications are based.

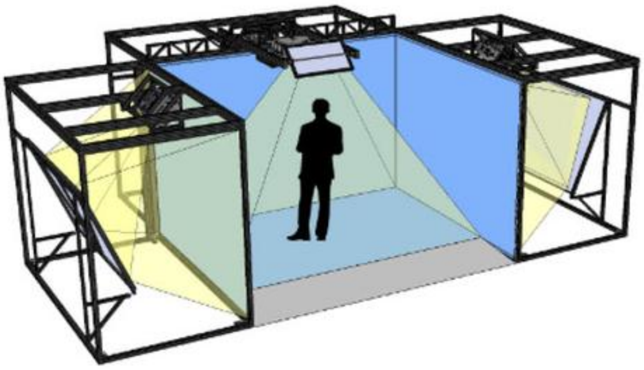

Figure 2.2: A design concept of CAVE, image from `https://www.wavin.ca/vr-cave.html`.

The CAVE [Cruz-Neira et al., 1993b] (Figure 2.2, or other similar setups like the StarCAVE [DeFanti et al., 2009] and the evolutive virtual environment (EVE) [Pierre et al., 2010]) system makes use of a room-size space to project objects. Users, usually wear 3D glasses, perceive objects in stereoscopy and can directly walk into the space, thus being fully immersed inside the visualized data. Since its creation, CAVE has attracted much attention in visualization domains [Cruz-Neira et al., 1993a], as it offers high-quality and large displays, that were not possible with other VR setups at that moment. Also, it allows the integration with different types of device to enhance user experience. For example, even though using a 6 degrees of freedom (DOFs) spatial tracker is the most common input device, there are several attempts to interact with such environments like using a tablet and/or a pen (e. g., the Studierstube project [Szalavári et al., 1998]) and gestures (e. g., [Meulen, 2012]). However, CAVE's stationary installation takes a large space and its complicated setup and maintenance are both time-consuming and costly. Moreover, its reliance on complex hardware setups (usually compute clusters) which make programming for it and running software fairly difficult. A similar, but less complex setup, called 3D power-wall or virtual reality power-wall [Treanor et al., 2009], has attracted a lot of attention in the industrial world, such as automotive manufacturing, and oil and gas industry. As in the CAVE, users wear 3D glasses to perceive content in stereoscopy. However, instead of turning a whole room into display walls, it is usually composed of one large-size display as illustrated in Figure 2.3, or a series of connected LCD displays to form a large screen, what makes itself a less complicated (but still large) setup than the CAVE. This kind of setup, is largely explored and combined with others devices for a hybrid setup (such as in combination with a normal desktop, a tablet, or others), as we will discuss in Section 2.2.

Figure 2.3: An example of 3D power-wall, imagined acquired from `https://commons.wikimedia.org`.

Smaller than the large displays and virtual rooms mentioned above, tabletop environments have also attracted a lot of attention. The responsive workbench [Krueger and Froehlich, 1994] uses a horizontal interactive surface to project stereoscopic images on it. It is also largely applied to visualization domains, like fluid dynamics [Wesche, 1999], battlefield [Durbin et al., 1998], and many others [Wesche et al., 1997]. Such setups also offer a lot of flexibility for users because we can make use of the large space on and around the display surface, especially when we take inspirations from non-immersive tabletop research that make uses of touch input [Lundström et al., 2011] or on-screen controllers [Jordà et al., 2007] to enhance the interaction for visualization. Yet, this setup still requires an important size of space to setup and maintain.

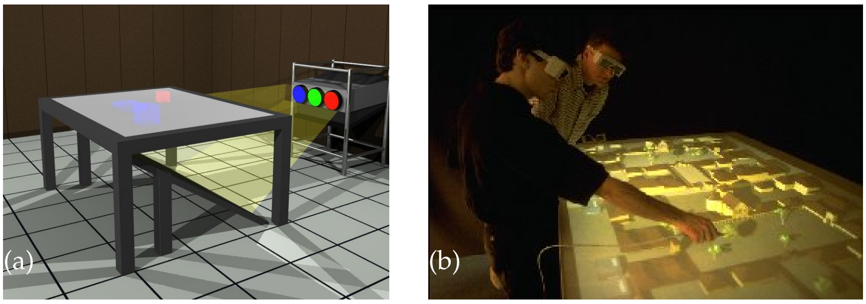

Figure 2.4: The responsive workbench. Images from https://graphics.stanford.edu/projects/RWB/.

Then, to the comparable size of a normal workstation, Fish Tank Virtual Reality [Ware et al., 1993] (Figure 2.5(a)) display stereoscopic images with a normal monitor and 3D glasses while tracking the positions of users' head to offer visual immersion. Compared with larger immersive environments, even though this setup does not necessarily offer superior displays and high computing powers, a study by Demiralp et al. [2006] highlighted that users had higher preferences of using it compared with CAVE. What mores, users were able to achieve the same level of performance for certain scientific visualization applications. This is one of the earliest attempt to bring immersive visualization to a desktop. Nevertheless, such environments are gradually replaced by modern commercial virtual or augmented reality (VR/AR) headsets (Figure 2.5(b) and Figure 2.5(c)) that directly offer both stereoscopic view and head tracking.

2.1.1 *Benefits of visual immersion compared to 2D screen*

All different setups offer a certain level of visual immersion and argued to be beneficial for visual data exploration and understanding. In 1993, Sollenberger and Milgram [1993] conducted three experiments to investigate the effect of using stereoscopic and rotational display by examining accuracy in 3D path-tracing task. Compared with a 2D screen, their results indicated that users' had higher accuracy when visualization was in stereoscopic or when they were using rotational display view. They argued that these elements would help the understanding of complex line cluster and network graphs. Later, follow-up studies confirmed such benefits (e. g., [Ware and Franck, 1996a; Ware and Mitchell, 2005]). Other than that, Prabhat et al. [2008] compared the performance of understanding biological datasets tasks in three environments: desktop, fish tank, and CAVE. The results indicated that CAVE, with the highest visual immersion, yielded the best results for both users' preference and performance of understanding spatial relationship. Similarly, Laha et al. [2012] also performed controlled studies and observed significant benefits of analyzing volume data with immersion presented. Later, the benefits extended to isosurface

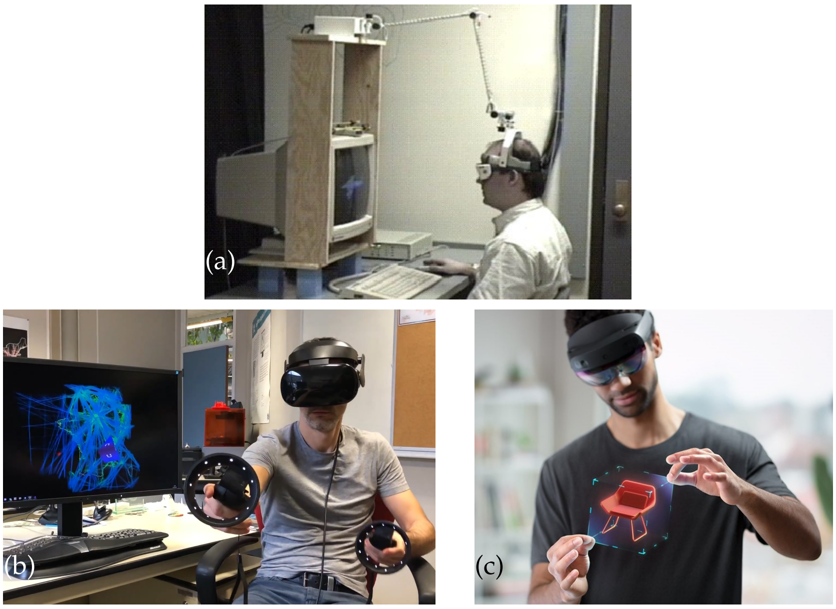

Figure 2.5: (a) the initial Fish Tank design using mechanical tracker, image from [Thabet et al., 2002]. (b) FiberClay visualization systems using VR headset by Hurter et al. [2019]. (c) Microsoft HoloLens 2 AR headset, image from `https://www.microsoft.com/fr-fr/hololens`.

visualization as well [Laha et al., 2014]. Will et al. [2018] explored the use VR environments and 3D interaction techniques for experts to trace neural circuits in brain, finding this system effective and less frustrating compared to traditional tools. They argued that scientists are able to understand large and complex cases better with such setting. Hurter et al. [2019] designed FiberClay (Figure 2.5(b)), a system that visualizes massive 3D airplane trajectories through occluded VR glasses. Its evaluation with experts suggested that it favors the discovery of flying patterns that were not usually noticed, therefore concluding that such immersive systems have benefits for the data sense-making process.

From these studies, we first learned that stereoscopy and motion clues help better understand 3D data, both of which are naturally brought by modern VR/AR headsets. We thus believe that choosing an AR head-mounted display (HMD) is an appropriate choice. Second, our application field of HEP deals with large clusters of particle trajectories (as illustrated in Figure 1.1), which have certain similarities with the studied network clusters or neural circuits. Based on that, we think the advantages of visual immersion will pass to the exploratory visualization of HEP as well.

2.1.2 *Limitations of fully immersive environments*

Despite the benefits offered by different immersive environments, and despite the fact that such environments are gaining popularity in many different areas (for example, VR and AR systems have been already much used in education, art and tourism), it is hard to find them practically applied and integrated into real scientific and engineering workflows. Apart from the issues raised by the technological constraints of the interfaces or physiological characteristics of the human being as summarized by Guillaume [2014], we discuss with regards to the requirements of scientists' general data exploration process.

Many studies pointed out that visual analysis helps researchers and engineers to understand their data, scientific workflows are not limited to spatial aspects only. Abstract data such as statistical results play a pivotal role and the analysis of such data usually requires traditional plots such as histograms, charts, etc. Showing these elements simply as a billboard placed into stereoscopic 3D view is not necessarily always advantageous, and turning the plots into a 3D representation is sometimes argued to be inefficient [Sedlmair et al., 2013]. Also, special 3D spatial input devices, which are usually the default input for VR systems, sometimes fails to meet the interaction requirements of certain scientific tasks, which usually demand a high accuracy. A recent study [Besançon et al., 2017b] using touch and tangible interaction to explore volumetric flow data (Figure 2.6) reported that domain experts mentioned that the traditional scripts-based input is still necessary for accurate and advanced analysis, and would rather combine the studied new interactive approach with traditional mouse and keyboard input.

Figure 2.6: Exploring volumetric flow data with tactile and tangible interaction [Besançon et al., 2017b].

For the particle physics, as we have introduced in Section 1.2, although the physicists agree that stereoscopic output with intuitive and fluid input is inspiring to understand the global event, they expressed the need to find a way to support script writing-loading in such environments as well due to the fact that scientists do not explore the data in a random way. Many predefined views, settings, and interaction (like filtering on different parameters) are specifically designed by and for the scientists to carry out an analysis. Writing scripts with keyboards is still the easiest way to quickly adjust all these parameters with high precision.

In summary, pure immersive environments have their advantages of helping data understanding, but yet several limitations make themselves hard to be applied to scientific workflows. We thus attempt to use a hybrid system to overcome the existing difficulties while benefiting the advantages of each side.

2.2 HYBRID VISUALIZATION SYSTEMS

With the benefits offered by visual immersion and the special requirements of domain experts of using traditional analyze tools, we work towards a hybrid scenario to bring immersive visualization to current scientific workflows as described in Section 1.3. The concept of hybrid environments is not novel. Researchers have proposed many different setups to improve interaction experience and/or to add additional visualization content. For the interaction part, Besançon [2018] discussed the general multi-devices interaction scenarios. In this thesis, we focus on the work related to immersive analytics.

One of the main idea of CAVE2 [Febretti et al., 2013] is to combine immersive visualization with non-immersive content by dividing the large display into two parts. Beside configuring the environments to be fully immersive, users can display immersive and non-immersive side by side. Based on that, they examined a specific collaborative scenario to explore ice covered Lake Bonney using the CAVE2 to display public content while adding a laptop per person to show private information and to control the displayed data (Figure 2.7). Later, researchers studied other visualization applications with the CAVE2 to explore large-scale cosmological simulation [Hanula et al., 2015] and medical images [Marai et al., 2016]. The cosmological application visualized both spatial and non-spatial data, a study shown that such design was efficient. In our application domain, we also need to deal with both abstract and spatial data. We thus hypothesized that a combination of different devices would also help the understanding of scientific data.

Similarly, the Studierstube project [Szalavári et al., 1998] is a co-located collaborative system which combines one AR headset and one hand-held panel for each collaborator. The panel is designed to both facilitate the interaction and to display information. With such a setup, Gröller [2002] have demonstrated several scientific visualization applications.

The interactive Slice Wim [Coffey et al., 2012] combined a table-top display and a wall display. With stereoscopic headsets, users can perceive both the 3D objects in spaces and a projection on the 2D screens. A typical example is that while users are slicing a volume visualization, the volume data is visualized in space with stereoscopy

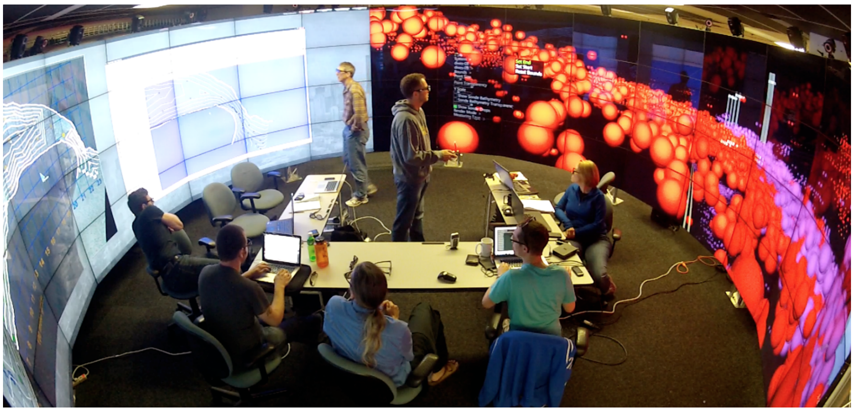

Figure 2.7: A collaborative working scenario using CAVE2 and personal laptops to explore ice covered Lake Bonney, image from Marai et al. [2016].

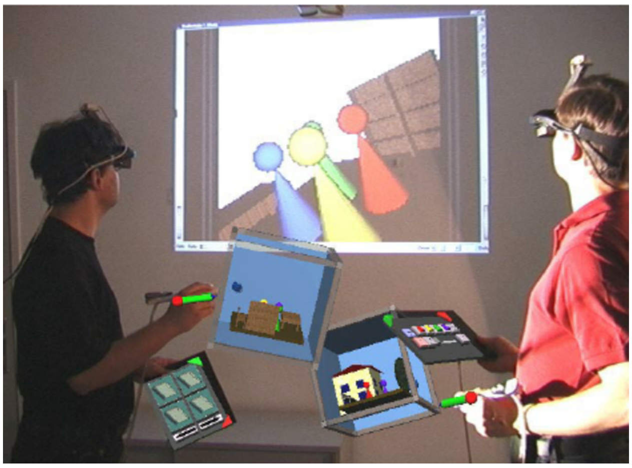

Figure 2.8: Studierstube collaborative application where users see both the 3D visuals in space and other information on the panel [Gröller, 2002].

and the slice plane is projected on the large screen, as illustrated in Figure 2.9.

As we have mentioned here, in most of the traditional setups with a large display, users wear 3D glasses to have stereoscopic views. Such displays project binocular images, thus the 3D objects perceived by the users are directly linked to the content displayed on screens. However, the appeal of small and portable AR headsets has offered another possibility—the immersive displays no longer need to rely on the content on the 2D displays, each of them can have separate views. Thus, research on combining a large 2D display with AR headsets is attracting more and more attention in recent years. For example, Reipschläger et al. [2021] augmented a large visualization screens (Figure 2.11) to address existing challenges that were often encountered with data visualization on large displays, and proposed several ideas to align the visualization of the two displays. Similarly, Büschel et al.

Figure 2.9: The interactive Slice Wim [Coffey et al., 2012].

[2021] explored a in-situ visualization for analyzing spatio-temporal data using a large wall display and anAR headset. These examples are quite aligned with our purpose of extending existing tools with immersive visualization, however, they are generally too large to be setup in a typical office for daily usage.

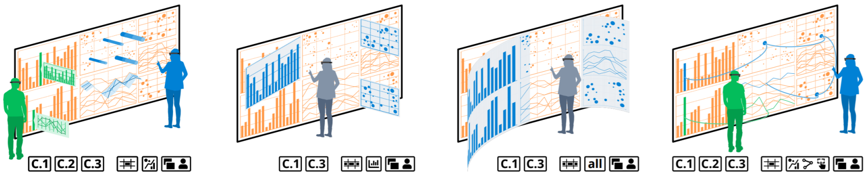

Figure 2.10: Visualizations with AR and large screens by Reipschläger et al. [2021].

Another idea shown by Bornik et al. [2006] is to use a normal tablet PC for showing 2D content and to interact with the stereoscopic view visualized on a large screen (Figure 2.12). They demonstrated this setup with medical datasets. However, an interesting result of their study is that users had some difficulties working with 3D environments, while they can perform things correctly in 2D. They explained this with the important learning cost, we thus further believes that it would be useful to keep a traditional workstation as the experts are used to, instead of replacing with other devices, such as a tablet-like surface used in the Studierstube project. The reason is that we want the new setup can be practically integrated into their normal working environments as requested by **R2**.

Benko et al. [2004] also demonstrated a collaborative case for archaeological data using large surface and VR headsets. Through combining a PC-like device or to large visualization environments, all these environments can fulfill our goal of combing immersive visualization with traditional analysis tools, and can easily display either spatial

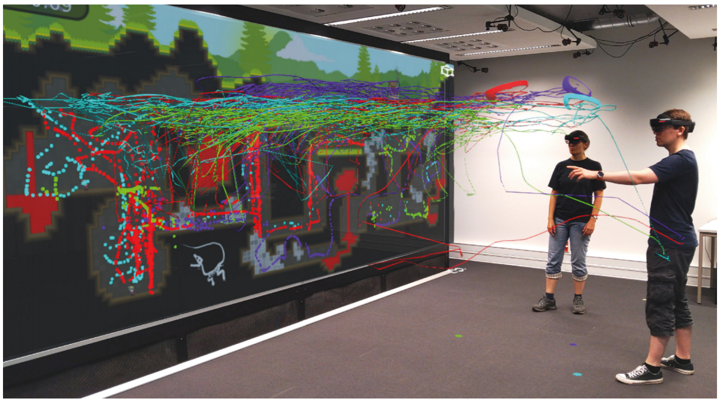

Figure 2.11: Visualizations with AR and large screens by Büschel et al. [2021].

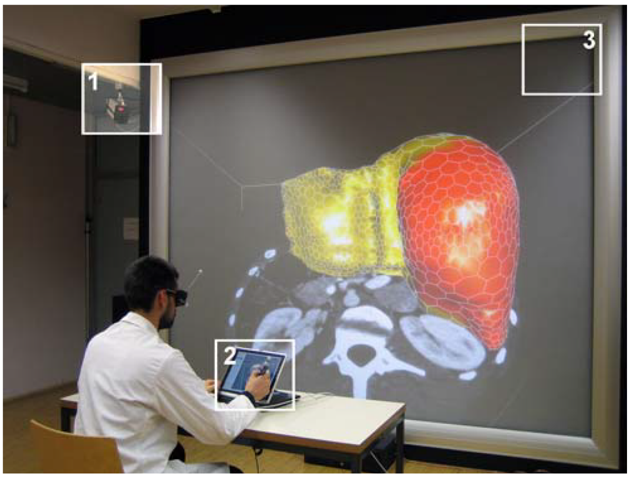

Figure 2.12: A hybrid setup using a tablet and and wall size stereoscopic display for medical data exploration [Bornik et al., 2006].

or abstract information for users to analyze. We believe the major drawback of applying such environments to scientific workflows is still their high cost and complexity for setup and maintenance.

As for the smaller setups, the use of see-through headsets seems to be more common the occluded VR techniques. With fully immersive VR, as users are occluded from the real world, it is required to recreate all needed elements in virtual space, such as a plane surface for visualization and virtual input devices for a reference. A typical example is the VirtualDesk presented by Filho et al. [2019] (Figure 2.13) that created a virtual surface in to facilitate the use of both 2D and 3D content. This method can both achieve **R1** and **R2** since it allows the exploration of all different types of datasets and does not require a considerable size of space to setup. However, as we have presented in Chapter 1, we tend to use non-occluded AR glasses to avoid recreating all the tools that experts need for data analysis, to better make use of real-world objects, and potentially to facilitate the communication among several collaborating users.

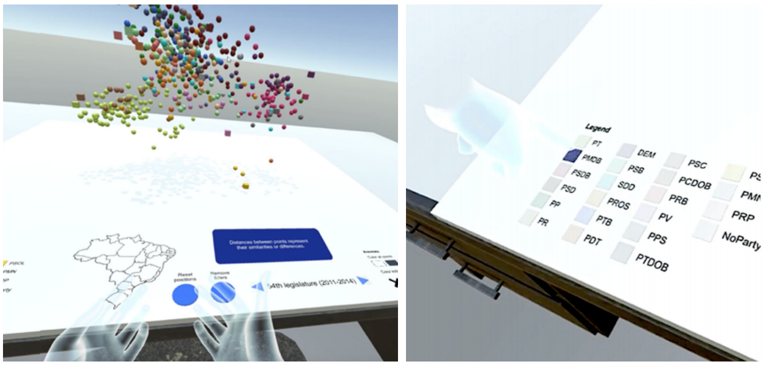

Figure 2.13: VirtualDesk application by Filho et al. [2019].

Previous work combining AR and a normal workstations include, for example, the combinasion of a traditional desktop with zSpace—a 3D stereoscopic screen—for radiologists analysis [Mandalika et al., 2018]. They evaluated this system with both experts and novices. The results shown a higher performance with this hybrid setting compared to both pure 2D and pure 3D setups. The benefits are even greater with non-experts users. Another example is the SpaceTop proposed by Lee et al. [2013]. It remains the size of normal desktop, but allows both traditional 2D and AR displays and interaction. This setup is pretty close to our vision, however, their usage of AR is limited to the space of a traditional space, which we believe could be further explored with modern hardware.

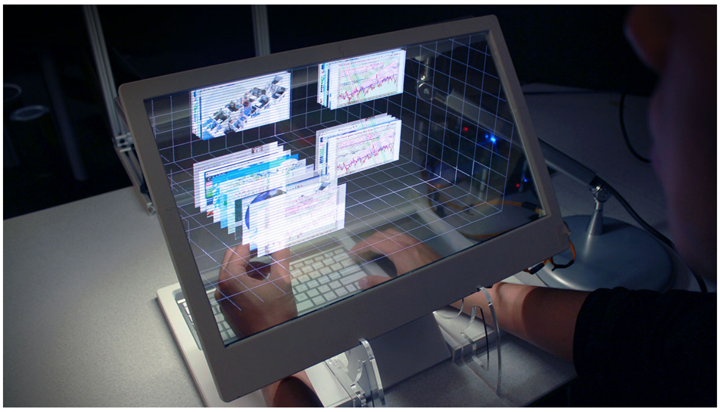

Figure 2.14: SpaceTop by Lee et al. [2013].

Apart from that, another one close to our vision is the DesktopVR (e. g., [Bogdan et al., 2014]). Similar to the fish tank, with the help of 3D glasses, it allows users to switch between displaying the monoscopic or stereoscopic view on the PC screen. Thus, users can switch between the two depending on the analysis they are running. Compared with the use of modern commercial AR devices, this setup still has a few major limitations. For example, the 3D view is limited to the screen space that cannot support large visualization, and the 2D and 3D analysis need to be carried separately—users cannot see the two at the same time. We believe that new hardware brings both new opportunities and challenges for design the interaction that worth to be explored.

Later, Millette and McGuffin [2016] designed a hybrid AR and PC setting to improve the Computer Aided Design (CAD) design process. They also added a smart-phone to control the AR space in 6 DOFs. In fact, it is not rare that people investigated a hybrid setting that one of the device is mainly meant for the interaction. Normand and McGuffin [2018] used AR environments to enlarge the screen of a mobile device, thus López et al. [2016] also proposed to use a tablet to interact with scientific data visualized on a large stereoscopic screen, and proposed a method to address the difference of view between the tablet and the stereo space (Figure 2.15).

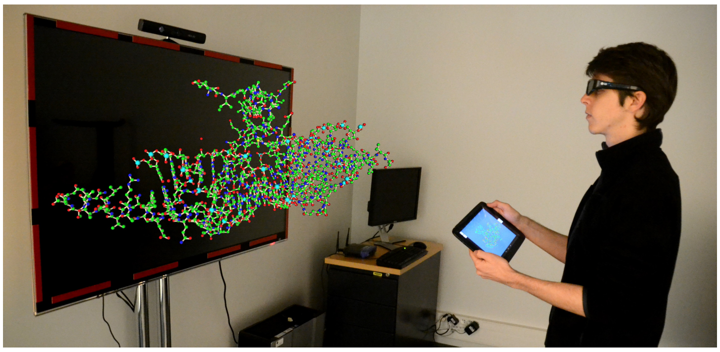

Figure 2.15: A hybrid setup that uses a tablet to control stereoscopic views, by [López et al., 2016].

These past approaches show that the combination of multiple devices can enhance the performance of users and empower them with new types of input by taking the best of each device. Such arguments also motivate our own work of extending traditional workstation with immersive environments. We aim to investigate with special regards to visualization requirements, and study the interaction techniques.

2.3 3D NAVIGATION TECHNIQUES

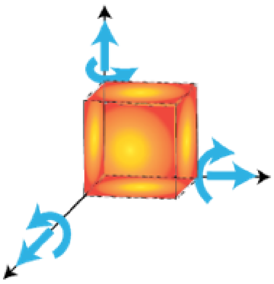

Figure 2.16: Navigating or manipulating in 3D requires 6 DOFs—translating along and rotating around the three axis.

3D navigation (or 3D manipulation), which requires changing position and orientation of the view (or the object) in 6 DOFs (Figure 2.16), is one of the fundamental task in 3D for visualization [Besançon et al., 2021], and is also the main task we need to support in our systems. Some surveys (e. g., [Jankowski and Hachet, 2013; Mendes et al., 2019]) summarized a variety of navigation techniques based on different setups and different input devices, from desktop to virtual immersive environments. And many studies have been performed to compare different input devices for 3D environments. For example, Dang et al. [2009] compared wand, voice, and PC-tablet based interfaces and found out that the wand is the most efficient input, while there are also other work suggested that traditional mouse input is still efficient, such as the one by Sun et al. [2018] that compared positioning tasks. Our main focus here is not to re-summarize all different kinds of input and interaction techniques that can be used for 3D interaction, but to globally discuss a part of them relevant to our hybrid system design. We will add more details in each chapter when a certain type is relevant to our work.

Mouse and keyboard are the default input devices for a traditional workstation. The interaction techniques of using the mouse to navigate in 3D spaces can be roughly classified into two categories. The first is to use widgets, like virtual handles, box, or balls (e. g., Figure 2.17). It allows accurate control of every parameters and sometimes the idea is also transferred to the AR space as done by the designers of the HoloLens (Figure 2.18), even though they do not use the mouse as the input device. However, the design of the mouse cursor in 3D is not easy, drawbacks exist for every kind of existing design [Schemali and Eisemann, 2014]. For example, stereo cursors (which work like a normal 3D object) usually suffer from the occlusion problem as they can be outside of the field of view or hidden by the virtual environments. The one-eye cursor [Ware and Lowther, 1997] improves this aspect by ignoring the depth and keeping the cursor on a 2D screen plane like a traditional mouse cursor on the screen. However, a study by Teather and Stuerzlinger [2013] found that such cursor is not universally beneficial compared with stereo cursor for the selection task, especially when we vary the distance between the objects to select and the display surface. So in the first step, we would avoid choosing interaction techniques that heavily depend on mouse cursor to click on specific widgets or handles, as designing one in 3D environments is an important research question itself.

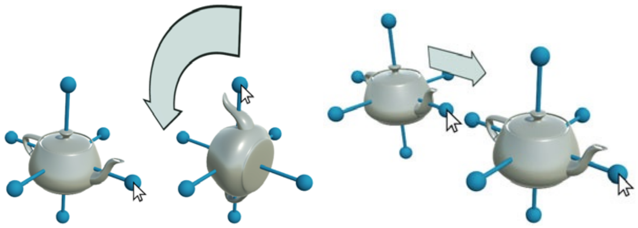

Figure 2.17: Manipulating 3D objects with mouse using handles [Conner et al., 1992; Mendes et al., 2019].

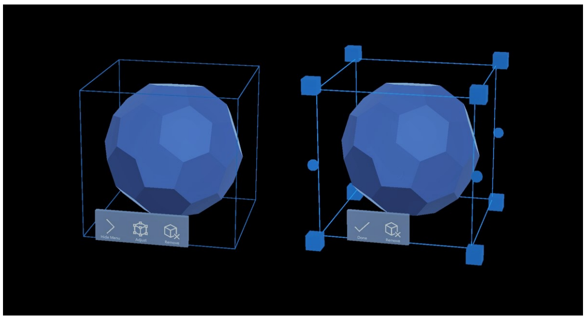

Figure 2.18: The bounding box handles of Micirosoft's mixed reality toolkit.

Another possibility is to directly map the motions of the mouse to manipulating a certain degree of freedom (DOF). This method usually requires the combination with the keyboard, and is supported by most 3D modeling or programming software. An example is that in the inspector windows of Unity3D, while pressing, users can translate their mouse to rotate the camera. Because it does not rely on the absolute position of the cursor, we believe it could be easier to use. In this work, we do not aim to design new interaction mappings (i. e., transfer functions) for the mouse, we thus decided to build upon an existing technique (Figure 2.19), which was found to be used by a large number of 3D modeling software tools in a survey.

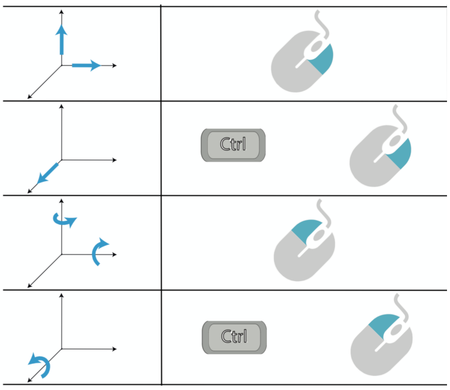

Figure 2.19: An existing mouse-based interaction technique, image from the supplementary materials of Besançon et al. [2017c]

Mid-air gestures are usually used for both VR and AR applications, as it is a natural way to touch and manipulate an object in 3D environments. For example, it is a common approach to capture a user's hand movement and create a visual hand in VR environments to manipulate objects. The first version of the Microsoft HoloLens uses gaze direction and a tap gesture as their default input. For the HoloLens2, the supported gestures are extended to more possibilities (Figure 2.20).

Figure 2.20: Examples of supported in-air gestures supported by HoloLens2, image from Microsoft Mixed Reality Toolkits.

While in-air gestures are often argued to be intuitive and easy to use, it could cause serious tiredness and cannot achieve a high accuracy [Filho et al., 2019]. The tiredness is severe for daily use and a low accuracy usually cannot meet the requirements of visualization tasks. Moreover, by the time we started working on the thesis, the tap gesture is the only supported gesture by the first version of the HoloLens. Without other devices, it is quite hard to meet the requirements of high DOFs' navigation/manipulation in 3D. Even though it is possible to create handles surrounding an object like using the mouse in a desktop, its long-term usage remains difficult and tiring. Apart from that, mid-airs gestures work only for the VR/AR and is hard to be extended to control the PC, which is an important drawback if we want to find one input for both the PC and the AR part of a hybrid system. Based on these considerations, we do not envision that using mid-air gestures is a good choice to offer seamless and unified interaction that could fulfill the needs of visualization tasks—particularly a high accuracy is usually demanded.

Wands or spatial-tracked controllers are also commonly found to interact with VR environments, and are becoming the default input devices for commercial VR headsets. These input devices are very flexibly and can offer many controls through their high input DOFs because they are spatial-aware and they have additional buttons for interaction. Nowadays, a lot of VR-based visualization research is based on commercial controllers. An example is the work by Hurter et al. [2019] as shown in Figure 2.5(b). However, there is still no universal one that have been largely applied in AR. Although some efforts have been made to design special input devices for AR, such as the CHARM (Figure 2.21) by Klamka et al. [2019], a limitation is that self-designed devices are not easily acquired by general public.

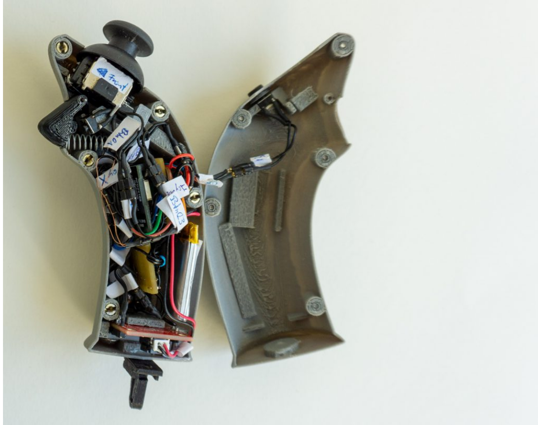

Figure 2.21: A specially designed input device CHARM [Klamka et al., 2019] for AR manipulation.

In addition, a common idea is to add a mobile device or a touch surface to enable the direct touch interaction for immersive environments to form a hybrid system, as we have introduced in Section 2.2. For all the existing solutions, each is beneficial for a certain use case, but also has their limitations. In this thesis, to explore the practical usage of our envisioned hybrid system by general public in their daily workflow, we decided to start from the default input of the PC—the mouse and the keyboard, and gradually revise and improve our design according to obtained results.

2.4 SUMMARY

We have presented in this chapter different immersive visualization environments, and some combinations with a PC-like workstation. Even though there are a lot of possibilities to allow experts to switch between the immersive and non-immersive views to fulfill our basic requirement **R1**, many of them rely on complex setup, or large spaces, which are not compatible with **R2** The one that is most close to our vision is the DesktopVR setup. However, the modern hardware has brought both new opportunities and new challenging for designing the interaction and visualization applications. We thus started to explore the use of a hybrid PC and AR setup for 3D data exploration.

3

UNDERSTANDING AR EXTENSIONS FOR EXISTING TOOLS

Towards an Understanding of Augmented Reality Extensions for Existing 3D Data Analysis Tools

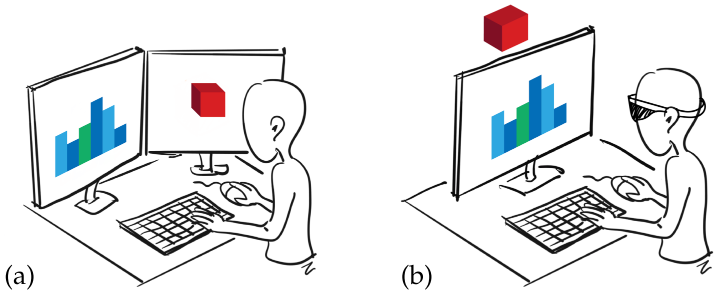

Figure 3.1: Sketch of our vision of transitioning from (a) a traditional workstation to (b) an AR-augmented data analysis environment [Wang et al., 2019b].

To answer **Q1** and **Q2**, we present in this chapter how particle physicists would want to use a hybrid PC and AR setup to explore their data through an observational study. Following our discussions presented in Chapter 1, our goal is to allow researchers to integrate stereoscopic AR-based visual representations and interaction techniques into their tools, and thus ultimately to increase the adoption of modern immersive analytics techniques in existing data analysis workflows. We use Microsoft's HoloLens as a lightweight and easily maintainable AR headset and replicate existing visualization and interaction capabilities on both the PC and the AR view. We treat the AR headset as a second yet stereoscopic screen, allowing researchers to study their data in a connected multi-view manner. We detailed our design choice, implemented prototype, and experiment in this chapter. Through the observational study with 7 physicists from CERN, results indicate that our collaborating physicists appreciate a hybrid data exploration setup with an interactive AR extension to improve their understanding of particle collision events.

Findings of the piece of work presented in this chapter give us fundamental understanding of the feasibility of a PC and AR hybrid visualization system, insights for future improvements, and enlightens us more specific questions to address.

Main portions of this chapter were published at ACM CHI 2020 [Wang et al., 2020]. The term of "we" in this chapter refers to myself, Lonni Besançon, David Rousseau, Mickael Sereno, Mehdi Ammi, and Tobias Isenberg.

3.1 INTRODUCTION

VR allows us to experience remarkably immersive worlds. These environments can be engaging and promise to facilitate tasks that require a high degree of immersion—the psychological state that users experience when they are surrounded by or in an environment that is continuously streaming stimuli [Witmer and Singer, 1998]—into their three-dimensional content. Since the end of 1960s [Sutherland, 1968], a number of technical setups (e. g., [VIVE, 2019; Oculus, 2019; Cruz-Neira et al., 1992; Febretti et al., 2013]) have been introduced and explored by researchers, with recent developments not only coming (visually) close to the vision of a Holodeck [Marks et al., 2014] but also making immersive experiences available to the general public. From the start, VR hardware has also been explored for 3D data visualization (e. g., [Haan et al., 2002; Fröhlich et al., 1999; Hurter et al., 2019; Keefe, 2008; Marks et al., 2014; Sundén et al., 2017; Theart et al., 2017]) and were proven to be more efficient than traditional setups in many different aspects as we have already discussed in Chapter 2.

Compared to fully immersive VR environments, AR offers new opportunities, in addition to also offering immersive 3D stereoscopic data views. First, AR does not transport users to a fully virtual world, allowing them to interact with real-world objects such as traditional input devices (e. g., mouse). Users are thus not forced to use dedicated input devices (e. g., wand, 3D controller) as in most VR settings, resulting in lower learning costs and a large potential to integrate the new environments with existing tools. The latter is essential because domain experts tend to stick to existing analysis tools and are hesitant to transit to new ones, as has been shown in past work [Besançon et al., 2017b] and which we also saw in our field observations as described in Chapter 1. We argue that this is is one of the main reason that current VR-based immersive environments rarely find their way into practical data analysis workflows used by scientists.

In this work we thus investigated a hybrid setup which extends traditional PC-based exploration tools with AR. This setup allows researchers to benefit from the immersion offered by AR technologies, while still being able to use their traditional analysis tools on classical workstations which possess higher computational power. While some past work (e. g., [Millette and McGuffin, 2016; Nagao et al., 2016; Serrano et al., 2015]) and commercial manufacturers [1] already envisioned or explored extensions of 2D screens with AR, we focus on the specific design requirements for scientific visualization. This domain differs from past studied ones such as VR/AR-supported video conferences and game-play, especially because of a high demand for accurate input. We used Microsoft's HoloLens [Microsoft, 2019] as a see-through

---

[1] Examples include the work of https://www.youtube.com/watch?v=0NogltmewmQ and https://www.youtube.com/watch?v=q0K3n0Gf8mA.

AR HMD so that users can seamlessly switch from the PC view to a stereoscopic data representation, and back. We duplicated in both platforms a set of visual data analysis features specific to particle physics to achieve a comparable level of functionality. The features are adjusted with respect to the constraints of their rendering space. We then treat both visual spaces as connected views [Wills, 2008], and let users control them using mouse and keyboard devices to avoid a repetitive switching of input devices. To better understand such AR-supported scientific visualization, we then present an observational study about how scientists want to make use of such a hybrid system, with special focus on particle visualization in HEP. We then discuss the potential usage and future design of such settings with respect to the feedback gathered from our study. Our main contribution is thus not the system design but our study. To our best knowledge, we are the first to examine the practical use of immersive visualization to satisfy real needs of physicists. Our results will guide the design of future hybrid visualization systems needed by physicists and scientists with similar 3D data.

3.2 RELATED WORK

Since we have discussed immersive and hybrid visualization environments in Chapter 2, we only quickly summarize a few main points in this section before discussing approaches that facilitate the interaction between different visual environments. We conclude this section with a small survey of visualization in particle physics—our chosen application domain.

3.2.1 *Immersive and hybrid visualization environments*

In the past, the responsive workbench [Krueger and Froehlich, 1994], occluded virtual reality glasses [Shibata, 2002], and CAVEs [Cruz-Neira et al., 1992] have been extensively studied because, compared to traditional 2D screens, they better support visual data immersion. Prior work argued that such environments can foster and facilitate interaction with large datasets (e. g., [Forsberg et al., 2009; Ware and Franck, 1996b; McIntire et al., 2014a]) as well as improving comprehension (e. g., [Filho et al., 2019; Hurter et al., 2019]). For example, Prabhat et al. [2008] concluded that, for 3D biological data understanding, users preferred and performed best using immersive environments, compared to non-immersive ones. Will et al. [2018] found VR environments with effective 3D input makes the task of tracing neural circuits in brain more effective and less frustrating compared to traditional tools. FiberClay [Hurter et al., 2019] visualized massive airplane trajectories combined with the geological map in VR headsets. Its evaluation with experts suggested that it favors the discovery of flying patterns that were not usually noticed, leading the authors to conclude that such immersive systems have benefits for the data sense-making process.

With the benefits offered by visual immersion, researchers also suggested using hybrid visualization environments that combine both 2D screen and immersive environments to benefit from both 3D stereoscopic view and traditional analysis information (e. g., 2D slicing and abstract data). For example, Mandalika et al. [2018] combined a traditional desktop with zSpace—a 3D stereoscopic screen—for radiologists. Marai et al. [2016] studied collaborative systems combining a CAVE2 displaying public contents—divided into a 2D wall size display part and an immersive stereoscopic display—and one laptop per user to both display private contents and send data to the CAVE. The Studierstube system [Szalavári et al., 1998] is a co-located collaborative system which combines one AR headset and one personal hand-held panel. Benko et al. [2004] also demonstrated a collaborative case for archaeological data. These past approaches show that the combination of multiple devices can enhance the performance of users and empower them with new types of input by taking the best of each device. Such arguments motivate our own work of extending traditional workstation with immersive environments.

3.2.2 *Cross-device interaction*

The communication among multiple devices and the interaction techniques for each device have been studied extensively. Brudy et al. [2019], e. g., surveyed papers from the ACM DL up to May 2018 about cross-device computing taxonomies and gave a detailed list of interaction techniques for different input modalities. As our ultimate vision of a seamless integration of the AR extension into the scientific workflow, we are interested in understanding the possibility of using a common interaction technique to control both sides without switching the devices.

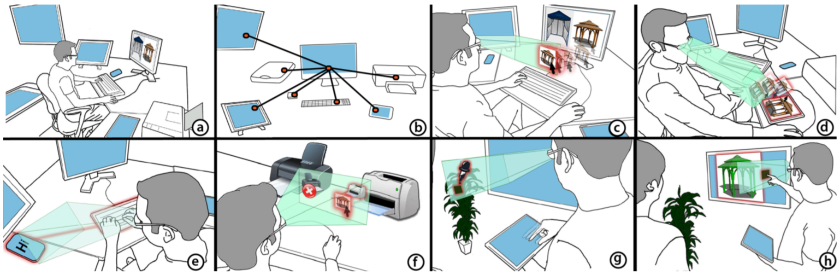

Figure 3.2: Illustration of Gluey [Serrano et al., 2015].

Using a head-mounted AR device as a workstation extension, Serrano et al. [2015] designed Gluey (Figure 3.2) to unify the different devices' interactions for general workflows. In 2015, Microsoft envisioned using their HoloLens to extend the 3D modeling tool Autodesk Maya (Figure 3.3). They allowed users to control the data in both the desktop and the AR space using the mouse. Millette and McGuffin [2016], although they added a smartphone, also kept mouse in a hybrid system for 3D CAD working scenario. In another example, a bank envisioned a scenario where users sit on a desk, pull things from the screen to space, and interact with them using gestures and voice command. While these possibilities have been demonstrated, their benefits and limitations remain unclear as well as how to properly design the interaction, specifically with respect to 3D data exploration needs. We thus based our prototype on mouse control for both the 2D and 3D views to better understand how that could benefit scientists and how it should be implemented.

Figure 3.3: The combination of a traditional workstation and a Microsoft HoloLens to augment the 3D modeling design process, by Microsoft.

3.2.3 *Visualization and data exploration in particle physics*

In our application domain of particle physics, visualization is essential for both collision exploration and public education [Bellis et al., 2018]. For example, experts use statistical tools to identify both strange (whose trajectories are hard to explain by current physics laws) and interesting (those who carry a high energy) particle traces, and visualization is needed to understand both. Various tools already support the interactive visualization of particle collisions for different tasks (Figure 3.4 and Figure 3.5(a)). For example, at CERN, Virtual Point 1 (VP1) [Kittelmann et al., 2010] displays experiments happened inside the A Toroidal LHC ApparatuS (ATLAS) detector for searching elusive dark matter particles, AliEve [Niedziela and Haller, 2017] visualizes events of A Large Ion Collider Experiment (ALICE), and iSpy [Alverson et al., 2012] is used for The Compact Muon Solenoid (CMS) detectors. Most of these tools are traditional PC or web-based with several identified defaults: a typical difficulty is that dense events are difficult to visualize due to the overlapping of trajectories after the projection on the 2D screen. VR recently attracted CERN researchers' attention: it is being gradually recognized by physicists that stereoscopic views can help them to understand their data. Yet the few existing tools, e. g., ATLASrift [Vukotic et al., 2015] and Belle II [Duer et al., 2018], are mostly limited to public education. We, in contrast, investigated the use of AR headsets for data analysis as they do not occlude users from the real world and thus can augment the 3D views of current analysis tools.

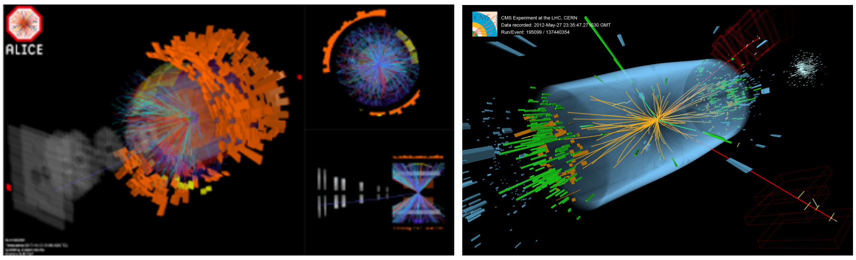

Figure 3.4: Examples of visualization tools used in CERN.

3.3 DESIGN CHOICES

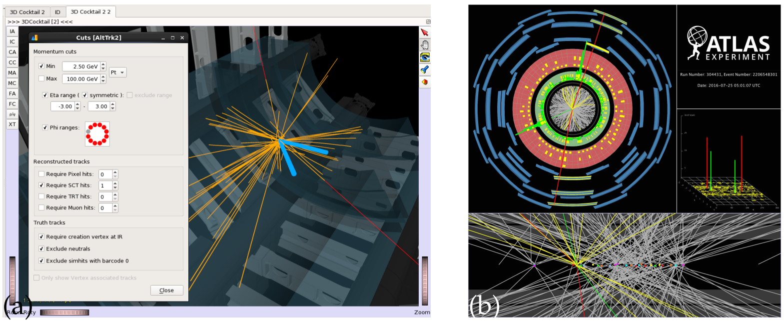

Cuts [AltTrk2]

Momentum cuts

☑ Min 2.50 GeV Pt

☐ Max 100.00 GeV

☑ Eta range (☑ symmetric): ☐ exclude range

-3.00 - 3.00

☑ Phi ranges:

Reconstructed tracks

☐ Require Pixel hits: 0

☑ Require SCT hits: 1

☐ Require TRT hits: 0

☐ Require Muon hits: 0

Truth tracks

☑ Require creation vertex at IR

☑ Exclude neutrals

☑ Exclude simhits with barcode 0

☐ Only show Vertex associated tracks

Figure 3.5: Screen shots of Virtual Point 1 visualization software. Images from the ATLAS Experiment © 2019 CERN, used with permission.

Benefits of visual immersion demonstrated in previous literature and physicists' interests motivated our study. We based our prototype on discussions with our collaborators at CERN (from the Atlas project) to understand their workflow, current tools, and interaction requirements. Through our discussions we learned that their analysis does not only rely on visualization software, they need to switch between data analysis to find and limit the region of interests, and visualization to observe and understand the exact phenomenon. We thus need to consider a system that allows both immersive visualization and traditional data analysis. Moreover, it needs to be easy to setup such a system in an office to be used in the scientists' regular workflows. Based on these considerations, we propose a hybrid setting that combines a PC and a HoloLens Figure 3.1 [Wang et al., 2019b]: We do not aim to replace the existing tools of scientists but rather to propose to use a 3D stereoscopic extension that allows them to better perceive their 3D data, seamlessly combining the benefits of the stereoscopic view with traditional analysis tools. We now explain our choice in detail and the relevance of our study to the domain.

3.3.1 *Input*

We rely on mouse and keyboard as input devices for both PC and the HoloLens. Physicists' analysis (in contrast to visualization) heavily relies on script writing where mouse and keyboard are essential. These devices are thus important to keep in our hybrid system as we want to integrate the 3D extension into their workflow.In addition, previous work pointed out that experts still prefer traditional input even if new forms of intuitive input exist, for example, studies with fluid dynamics researchers [Besançon et al., 2017b] and doctors [Mandalika et al., 2018]—similar to the well-known *Legacy Bias* in interaction design where *"users resort to well-known interaction styles even when more effective and novel techniques are available"* [Brudy et al., 2019].

For AR input, even though mid-air gestures are popular means, it has been argued that they could introduce user confusion, error, and fatigue [Filho et al., 2019]. We thus do not envision its use for scientific visualization where high interaction precision is required. In addition, we are interested in unifying interaction design such that users do not need to switch between different input devices as others argued in the past, including for the HoloLens. But such interaction remains a challenge in purely virtual spaces [Grubert et al., 2018]. Our work is thus a step toward better understanding how hybrid virtual environments can enhance scientific analysis and how to improve scientists' workflow with the commonly used input of mouse and keyboard, as well as how such systems should be realized.

3.3.2 *Output*

We selected an AR HMD because it is light-weight and can easily be used in an office, without occluding the real world as a VR HMD. Users can seamlessly use their traditional analysis tools and benefit from stereoscopic rendering. We thus do not need to recreate all analysis tools as in VR, nor do we need to introduce additional VR-specific input devices. Although AR devices with an additional mouse and keyboard input may fully replace the PC one day, we still explore and study the equivalent PC-AR hybrid setting due to its currently higher fidelity. Moreover, relying on a PC likely will always have merit due to its its high computational power and superior high-resolution screens. We excluded large environments like a CAVE due to their high demand for space and maintenance, which limit a practical integration to regular workflows. We do not use static 3D screens because the virtually unlimited canvas of AR HMDs can provide additional advantages for scientific visualization as such tasks often require multi-view analysis. They are also more flexible to allow users to arrange views to perceive both the 3D space and non-spatial information at the same time. While 3D screens with appropriate interaction techniques can

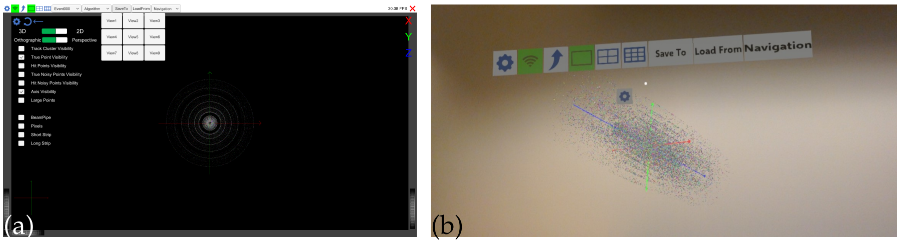

Figure 3.6: Example of user interface on (a) a PC and (b) the HoloLens. Both of them use a menu bar on the top of the view.

also offer large canvases, we believe that gaze-based view-switching has potentials that we should study. Although we did not investigate collaboration in this work, an AR HMD may also facilitate collaborative data analysis [Sereno et al., 2019] which should be studied in the future.

3.3.3 *Study relevance*

Although research in VR/AR with multi-view settings exist [Mahmood et al., 2018]), the interaction requirements of 3D data visualization usually differ from those in other use cases, specifically the demand of high-precision work. For 3D selection, e. g., common methods like ray-casting are unsuitable for scientific datasets because they usually do not natively define objects or regions [Besançon et al., 2019]. Another example is the common use of orthographic projection for precise comparison of parallel structures. It is thus important to study such tasks with domain experts to understand the needs and to conduct design guidelines to bring immersive visualization environments into their workflow. Another reason is that, while some have envisioned hybrid PC plus AR interaction as mentioned above, it still remains unclear how domain experts would want to use such a hybrid system, how an AR extension can support data exploration, and how the interaction/UI should be designed to support their needs. To study it, a working prototype is necessary because paper prototypes with flat images remove the immersion provided by AR, thus several features of AR like walking around would be impossible to investigate.

3.4 PROTOTYPICAL IMPLEMENTATION

We implemented a prototype to serve as an initial tool to understand the potential use case of an AR extension with PC data exploration tools—we did not developed it to replace existing software and settings in usability, interaction details, or computing power, all of which are fast changing according to experts' feedback, yet are not the key points we discuss.

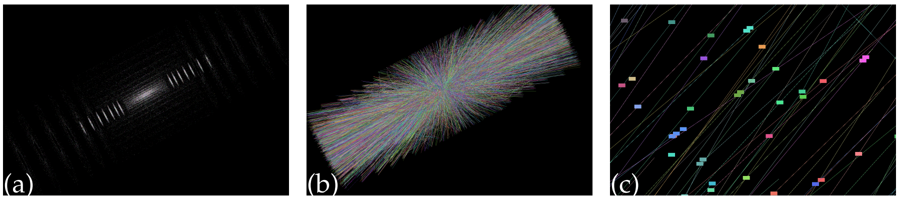

Figure 3.7: Visualization of (a) hit/true points, (b) constructed/true trajectories, and (c) zoomed trajectories with points in our PC prototype.

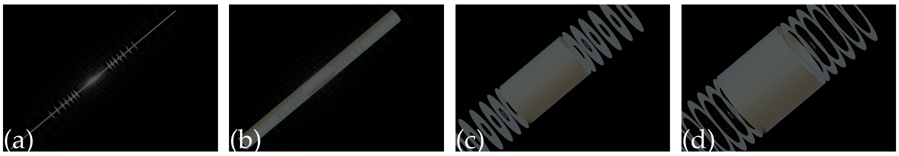

Figure 3.8: Visualization of (a) beam line as well as (b) pixel, (c) short strip, and (d) long trip detectors with particle hit points in our PC prototype.

Our prototype consists of two parts: one on the PC and one on the HoloLens (Figure 3.6), both inspired by our collaborating particle physicists' regular work environments. We envision the metaphor of using a two-screen environment, in which the content of each screen can be defined individually and the mouse can travel from one screen to the other. We then replace one of these screens by the AR environment (Figure 3.1). The users can remain seated and continue to work with their traditional tools as usual on their PC or laptop, but can also transition to the AR environment when needed or go back to the PC at any time. The communication is based on WiFi using the UDP protocol and is bi-directionnal, i. e., motions happening in the AR environment are also transmitted to the PC. We created all implementations in Unity with C# and its framework .NET. We transmit data via UDP due to its simple implementation and small processing overhead. In a controlled network, data usually arrives in order and without loss.

We then created comparable functionality on both the PC and the AR platform, including the user interface and the interaction logic, i. e., all tasks described in this section can be performed both on PC and on HoloLens. Our general idea with this prototype is that both views share the same dataset but can be configured differently (views, settings, manipulations, etc.)—just like multi-view environments on traditional PC settings. Users can pull the current configuration from one side to another using UI actions. They can also switch the real-time synchronization between both device on and off, to understand how users prefer to use such a function.

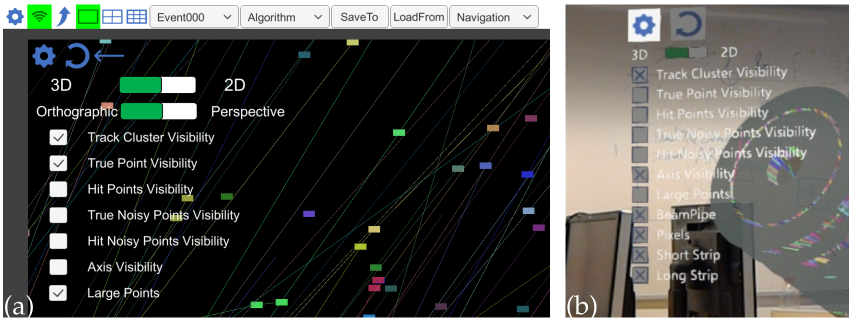

(a)
☑ Track Cluster Visibility
☑ True Point Visibility
☐ Hit Points Visibility
☐ True Noisy Points Visibility
☐ Hit Noisy Points Visibility
☐ Axis Visibility
☑ Large Points

(b)
☒ Track Cluster Visibility
☐ True Point Visibility
☐ Hit Points Visibility
☐ True Noisy Points Visibility
☐ Hit Noisy Points Visibility
☒ Axis Visibility
☐ Large Points
☒ BeamPipe
☒ Pixels
☒ Short Strip
☒ Long Strip

Figure 3.9: UI Interface: (a) PC and (b) HoloLens.

3.4.1 *Data*

We use simulated proton collision events from the MLTrack Challenge [Amrouche et al., 2019]. A single event contains information about (1) the true hit points (collision with detector hardware, including position and momentum data), computed through physics laws, (2) the simulated measured hit points (with simulated measuring error), (3) information on the particle' trajectory, to which we refer as a track (Figure 3.7). One event contains about 10K tracks with 100K points. Basing our visualization only on simulated measurements, we connect the points' positions to get the particles' *measured* trajectories. To reduce the rendering cost, we simply connect the points with straight lines which is also done by the physicists, without introducing ambiguity for understanding the true trajectories.

3.4.2 *Mouse transitions*

Mouse and keyboard input are captured first on the PC, and then transmitted to the HoloLens. On the PC we use the mouse as usual. The mouse switches between the 2D screen and the AR space by pressing the Tab key. We did not use implicit transitions when the mouse crosses the screens' borders as done with two 2D screens to avoid unintended switches between the two platforms. Indeed, the borders' area of their traditional tools usually contains UI elements to perform manipulations or to change system setting (Figure 3.5(a)) causing users to frequently manipulate this area. We did not use gaze focus to control the cursor either like others [Serrano et al., 2015] because we leave users the possibility to see both the 3D and the 2D views at the same time, instead of forcing them use only one. In AR space, the mouse can move—in addition to its 2D motion—along the depth axis with the scroll wheel, while the Shift key is pressed. We we decided to not reinitialize the mouse's depth position after releasing the key (i. e., the mouse will not be back to the default position "in front" of the 3D box). Yet even though the UI widgets in AR are fixed on the box' front, users can still click on them while the mouse is behind them based on ray-casting (without drawing the ray). We added a visual feedback (color highlighting) while the mouse "hovers over" the button.

In the remainder of this section, we list and discuss how we solved the requirements pointed by the experts. We also report our insights from analyzing their current tools.

3.4.3 *Following the track*

Physicists want to follow a single track while exploring, thus each of them need to have a unique color. Their current tool does not provide a standard color map for all tracks, we thus created one to make sure that they are not too bright to hurt eyes in either space.

3.4.4 *Abstraction for data and detector*

It is essential for our collaborators to visualize both the detector structure and the collision data. The Atlas detector comprises three main parts: inner detector, calorimeter, and spectrometer. In this study, we use a simplified model with only the meshes of the beam line (where collision happens) and inner detector (which includes pixel detectors, short strip detector, and long strip detector, see Figure 3.8). Switching on and off different detectors is triggered using UI widgets (Figure 3.9), similar to their existing tools.

3.4.5 *Visualization of different perspectives*

Domain experts need to compare different views of the same data. As shown in Figure 3.5, the current tool uses multiple windows or several views on the same screen. We chose the latter as AR has a larger canvas that supports the simultaneous rendering of multiple views. We implemented 2×2 views (Figure 3.10) including the front, top, left, and perspective views as commonly found in many 3D tools. Those views have different transformation (position, rotation, and scaling), but share the same dataset: manipulations such as filtering are applied to all of them. Those views are not restrained to pre-defined settings, each of the them is configurable. Users first click on a specific view, and then can change its setting (in Figure 3.10, e. g., the right bottom view's is highlighted to indicate that users interact with just this view).

3.4.6 *Save and jump to specific views*

Domain experts also need to be able to save specific states (including transformation, filtering, and abstraction level) of the dataset they

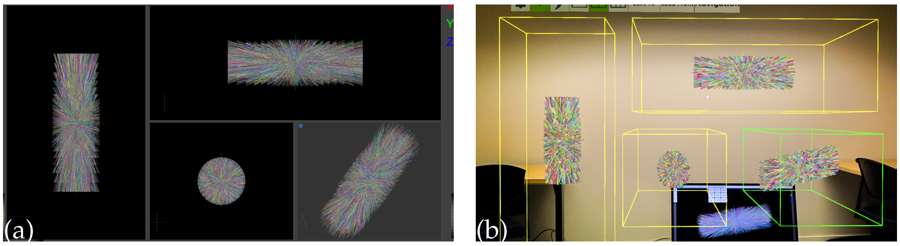

Figure 3.10: Visualization of 2 × 2 views: (a) PC and (b) HoloLens.

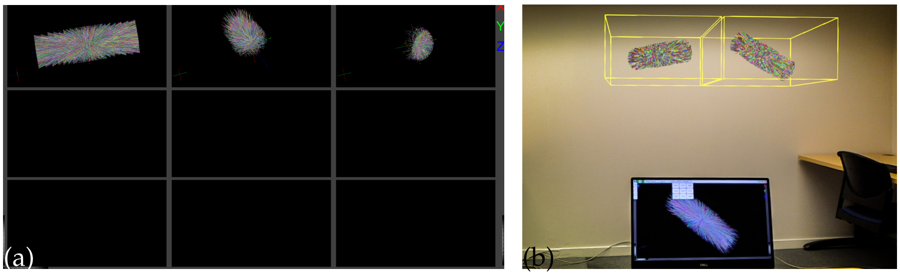

Figure 3.11: Visualization of 3 × 3 views: (a) PC and (b) HoloLens. Only saved views are displayed.

are exploring and may jump back to a former state later. We thus implemented a 3 × 3 dump board (Figure 3.11) that carries the saved states. The dump board is always synchronized between the PC and the HoloLens, thus users can perform interaction on any view and freely switch to the desired setting on the other.

3.4.7 *3D navigation*

3D navigation allows physicists to explore and understand the spatial aspects of the dataset at hand. Their current tool includes interaction with 5 DOFs: x- and y-rotations, x- and y-translations, and uniform scaling. However, they are insufficient in an AR setting where users need to translate the data along the z-axis to specify its position in space. We thus defined a 7 DOFs navigation mapping using mouse and keyboard, derived from one of their used tools and previous work [Besançon et al., 2017c] as follows: the left button triggers a rotation around the x- and y-axes, the right button triggers a translation along the x- and y-axes, the scroll wheel translates along or rotates around the z-axis (a single click on the wheel switches from one to the other), and scrolling while Ctrl is pressed triggers zooming.

3.4.8 *Selection by parameters*

Each particle has many parameters. Experts usually plot histograms and find a region of interest. They then use their tools (VP1 or special Python packages) and focus directly on the special particles by select-

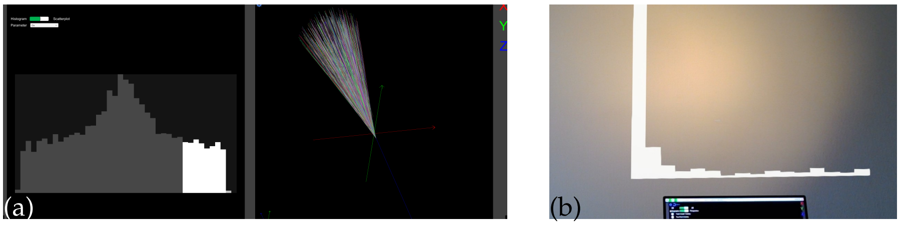

Figure 3.12: Histogram filtering: (a) PC (by η) and (b) HoloLens (by p_t).

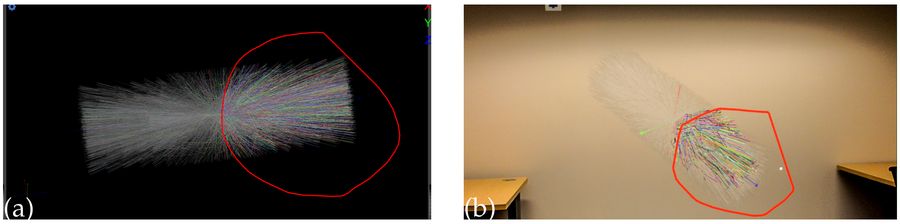

Figure 3.13: Spatial selection: (a) the PC and (b) the HoloLens.

ing them based on the target parameter values. Based on the previous habit of using histogram, we make them interactive. Users can highlight and filter particles through clicking/sliding on the histograms (e. g., left side of Figure 3.12(a)). Explicit filtering is triggered by direct clicking/sliding, while track highlighting is triggered with Ctrl button. In our prototype, we support histograms of following properties: *azimuthal angle* (phi, ϕ) in cylindrical coordinates, *pseudorapidity* (eta, η) related to the the dip angle in cylindrical coordinates, *transverse momentum* (p_t), the momentum of the generated particle projected onto the transverse plane, the radius of the production point of the particle (r_0), and the distance of closest approach to the z-axis of the trajectory of the particle when extrapolated (d_0).

3.4.9 *Spatial selection*

Domain experts sometimes need to select tracks based on their positions in their visualization software. We implemented a lasso tool (Figure 3.13) which is often found in 3D exploration tools (e. g., [Yu et al., 2016]). Users can apply Boolean operations to intersect, unite, or delete the selected tracks with the visualized ones. Using screen widgets, users can also specify if they want to select particles with all trajectories inside the lasso or with at least one part inside it, thus keeping the complete trajectories of the particles. Selected regions are first highlighted, then filtered after confirmation (Figure 3.6).

3.5 OBSERVATIONAL STUDY

To better understand the implications of combining a traditional workstation setting with an AR view and how to develop interaction mechanism for such hybrid environments, we conducted a *preregistered* (`https://osf.io/7qegs/`) observational study with seven experts in the domain of particle physics. While this number of participants may appear low, it is not an unusually low number when conducting observational studies to understand the needs of ***domain experts***. We were interested in their general opinion on such a system and their feedback on how the interaction should be designed to better answer the tasks that they have in their domain. We used an observational strategy that has been used by several other researchers and research work before when designing for and reflecting on specific domain experts needs (e. g., [Fu et al., 2010; Hurter et al., 2019; Klein et al., 2012; López et al., 2016; Lundström et al., 2011; Sultanum et al., 2011]). We specifically decided not to focus on usability studies and time or error measures for several reasons. First, our goal is to understand how can we extend the current data analysis tools with an AR extension instead of presenting new technique/system, we do not aim at proving a hybrid environment that works "better" than a traditional PC-based 3D analysis tools. Also, we do not want to miss meaningful critique and ideas which could be prevented by using quantitative studies as pointed out by Carpendale [Carpendale, 2008].

3.5.1 *Participants*

We recruited 7 CERN researchers as unpaid participants, all working on HEP and denoted as P1–P7 (6 males, 1 female; ages 26–52 years). They had 2–30 years of post-Master's research experience (mean: 12.4, median: 12, and SD: 9.1). All were used to interact with 3D datasets in their work using typical mouse+keyboard interaction (one reported 1–2 times a week, while all others reported several times a day or that it was basically their daily work). Six of them had knowledge about VR glasses, three had limited experience with immersive environments (only VR glasses), and none of them had experience working with Microsoft's HoloLens.

3.5.2 *Apparatus and Setup*

Our prototype used the first version of Microsoft's HoloLens (development edition) and a Dell XPS 9570 laptop (3840 × 2160 15′ screen) running Windows 10 with its integrated keyboard and a Bluetooth mouse. We connected the laptop to a local TP-Link AC750 router via Ethernet, to which the HoloLens connected via WiFi. We ran the study in a meeting room at CERN in Switzerland. We placed the laptop on one side of a big meeting table and let the Hologram initially focus on the center of the table, at around 2.5 meters away from the participant and vertically a little higher than the PC screen (near the other edge of the table).As shown in Figure 3.14, the room had whitish wall and was

Figure 3.14: Experimental environment.

lighted as their normal office (the windows on the other side of the wall were closed by blinds). Nobody walked into or out of this room during the experiment. Participants sat on a fixed-leg chair. While workstation screens in an office have a bigger size than the one used in our experiment, we believe that using a laptop screen is enough for the purpose of the experiment as we aim to understand the potential of the hybrid concept, rather than specific hardware. Apart from that, we have also observed that experts, in many situations, work on their laptop. Even in their office, they still sometimes make use of their laptop screen. Also, due to the needs of setting up this experiment at CERN, using a laptop was the most realistic solution.

3.5.3 *Study design and procedure*

The same dataset was used by all participants in our observational study. Each participant performed the experiment individually, in the mentioned room, alone with the observer. Participants were video-recorded (participants are not always visible but audio is fully recorded) for analysis. We started by explaining to participants the purpose of this study. We then asked participants to read through and sign, if they agreed, a consent form, a media-release form, and a questionnaire collecting their demographic information and past experience with 3D data exploration, 3D interaction, and immersive environments. We then began our three-part experiment:

3.5.3.1 *Explanations and tutorials*

We first re-stated our purpose with this study and emphasized that the goal was not to show a more impressive system with better usability or functionality compared to the traditional one. We then introduced them with the apparatus. We told them that we would first introduce the user interface and explain the interaction. We also explained that the task was to use all possible interaction techniques to explore the data, before a semi-guided interview to understand how they used our system. After they had put on the HoloLens and adjusted it to a comfortable state, we began the introduction. We did not stream their view to avoid the drop of performance, we asked them to confirm their understanding after each presentation and encouraged them to ask questions themselves. In our design scenario, users sit down and work on their data as in their office. However, walking around is a feature of AR HMDs that is under-explored in visualization. We thus told participants to feel free to get up and walk around.

3.5.3.2 *Free exploration and thinking aloud protocol*

Once we finished the explanations and the participant was ready, we left full control to the participant. We asked them to interact with the system using all implemented functions and freely explore the described simulated collision dataset. Participants were allowed to ask for help or pose questions while exploring and were encouraged to think aloud, i. e., to express directly what they observed and thought. The experimenter observed the whole experiment next to the participant, took notes, and guided the participant if they felt that the participant forgot something at the end (e. g., if the participant used only some of the features, we encouraged the participant, without forcing them, to try others as well). There was no set time limit.

3.5.3.3 *Questionnaires and semi-guided interviews*

When a participant reported that they had finished exploring the dataset, we asked them to take off the HoloLens and take a short break. We then conduced a semi-guided interview asking their general feedback about the AR extension and the interaction with mouse and keyboard, as well as any potential improvements they envisioned. We finally asked participants to answer 5- or 7-point Likert questions to quantify some of their ratings. To fairly discuss the potential usage, we highlighted again, before the interview, that our system was a prototype whose purpose was to understand how to use an AR extension, instead of presenting a new system. We also explained that many current hardware limitations would be improved with the release of new AR devices, thus highlighting that they should focus more on the interaction and visualization aspects.

3.6 RESULTS

Participants took 30–60 min to finish the exploration, and 1.5–2 hours for the whole study. We report the collected quantitative data and qualitative insights next. Because they complement each other, we report these two types of results together, organized into several categories.Many of the questions are exploratory (i. e., we ask about their ideas without letting them try other environments) as several have experienced VR before and our goal is to learn about their general

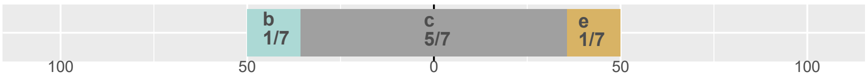

Figure 3.15: "For practical data analysis in my future daily work, I would prefer:" (a) only a PC interface, (b) in-between, (c) half-half hybrid interface, (d) in-between, or (e) only an AR/VR interface.

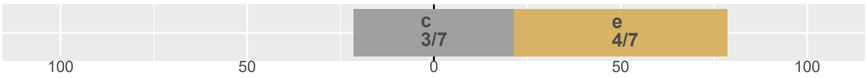

Figure 3.16: "In an interface that uses AR/VR stereoscopic elements, I would prefer:" (a) a VR interface that only shows virtual elements and where the real world is completely invisible, (b) in-between, (c) I have no preference, (d) in-between, or (e) an AR interface in which the virtual elements are shown in addition to the real world.

opinions of the most suitable immersion experience for such analysis tasks instead of precisely comparing the usability.

We based the plot of Likert-scale results on an online tutorial (`http://rcompanion.org/handbook/E_03.html`). The horizontal axis represents the percentages instead of exact numbers, which indicates the trend (towards the left or the right) of all voting results. While we actually use number instead of percentage due to the small number of total participants, we kept the horizontal axis as the percentage number since it is a standard for such plots.

3.6.1 *Interface*

We asked participants about their general impression of such system and if they would prefer an interface with only the PC, an interface using only the HoloLens or another VR headset, or a combined one as we showed. P1 reported a high preference of using only an immersive interface and P2 suggested a hybrid system but with more focus on the PC side, while all other participants preferred a half PC and half AR/VR hybrid system for their future tools as shown in Figure 3.15 and Figure 3.16. We report specific comments below.

p1 did not see the point of keeping the background environment visible. Then, with proper input, a pure VR environments would be fully capable for exploring such datasets. Background information may also be source of disturbance, imaging people walking around. However, the use of AR facilitated the combination with the laptop which is highly advantageous, and there is no sickness feeling as in VR.

p2 agreed that the PC can be used to manage expensive computation, while visualizing the results in AR space.

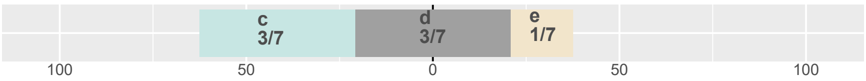

Figure 3.17: "In general for 3D data analysis (not necessarily your own work):" (a) the PC is the best platform, (b) the HoloLens is an interesting yet not particularly useful addition the PC, (c) the HoloLens is a nice addition to the PC that is sometimes useful, (d) a balanced combination of PC and HoloLens is best, (e) the PC is a nice addition to the HoloLens that is sometimes useful, (f) the PC is an interesting yet not particularly useful addition the HoloLens, or (g) the HoloLens is the best platform. P1 did not vote while P4 voted for 2 options.

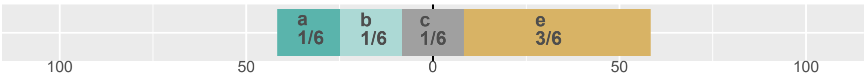

Figure 3.18: "For my typical data analysis at CERN/of particle collision data, the forced perspective view on the HoloLens is" (a) not useful at all, (b) in-between, (c) somewhat usable yet has some perspective projection issues, (d) in-between, or (e)completely equivalent to orthographic views on the PC. P4 did not vote.

p4 thought that the HoloLens is less disturbing than VR headsets by not occluding the real world. The PC worked better for precise interaction and abstract data visualization, while AR offers better depth perception for 3D visualizations.

p5 preferred a hybrid system, but if a future HoloLens becomes powerful enough he can imagine the scenario of using only the HoloLens for both 3D and 2D visualization, potentially with a virtual keyboard.

p6 preferred a hybrid system because laptop could be kept for practical data analysis tasks.

p7 was more interested in stereoscopic visualization than regular 2D screen projections.

Besides, P7 noted that scenarios may exist in which the PC and the HoloLens would better be used separately depending on the tasks. She explained that switching the focus both for visualization and interaction between the 2D and 3D space can be annoying, while all other participants did not share the same feeling. Nonetheless, P3 mentioned that switching was uncomfortable at beginning, since he is not used to look up and down because his office screens are aligned horizontally.

We then asked about the roles of each platform, see the results in Figure 3.17. P1 did not give any preference because he would personally prefer a VR environment instead of an AR, PC, or hybrid system. P3 and P7 thought that the HoloLens is a nice addition to the PC that is sometimes useful, but the major tasks would still be performed on PC. P3's justification is that all the same things can be displayed on a 2D monitor and that experts are quite trained to understand perspective there. P5 and P6 prefer a balanced combination of PC and HoloLens. P4 reported two preferences (c) and (d). He thought that, in principle, a balanced combination is perfect for data interaction, but that it would depend on future performance of the hardware: Today's limitations of the HoloLens mean that it can only be seen as an addition to the PC to use occasionally.

3.6.2 *Perception and data understanding*

We asked about the difference of perceiving data between the HoloLens and the 2D screen. P7 reported that the AR space is similar to an additional 2D screen while other thought that AR provides more than an additional screen, especially emphasizing its added depth perception. We report specific comments below.

p3 appreciated the HoloLens's low resolution in some cases as it makes certain data elements bigger, such as particle hits and curve trajectories which are often tiny. In addition, AR systems would allow a better understanding of detector structures and their spatial arrangement with the particle trajectories, which will largely help at understanding events.

p6 expressed that the direct interaction with the data to see how tracks go from one vertex to another is impossible or, at least, hard to achieve on the PC.

p7 could not perceive the data close to her due to the narrow field of view of the HoloLens. She then needed to move the data further away, which limits the 3D immersion.

Data visualized in AR space is forced to be shown in a perspective view (to maintain stereoscopy), but 3D visualization software often relies on both orthographic and perspective views. We tried to understand how experts feel about this forced yet physically correct perspective projection and the mismatch to a potentially shown orthographic projection on the PC. Figure 3.18 summarizes their general opinions. P1–P6 reported that perspective projection on HoloLens is just natural, they did not see why orthographic projection is useful in AR space for experts. We report other comments below.

p3, p4 said that as domain experts who understood the data well, the link between the orthographic projection on PC and the forced perspective projection in AR is easy to make.

p4 did not vote because he thought the perspective in AR is more useful than orthographic view, but he would keep the orthographic views on the PC side.

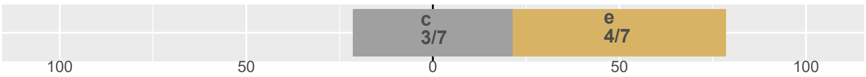

Figure 3.19: "In an improved interface, I would prefer:" (a) the PC and HoloLens views always in sync, showing the same exact views and selections, (b) in-between, (c) the PC and the HoloLens showing different 3D/2D views yet with the same dataset and the same set of filtered/highlighted tracts, (d) in-between, or (e) the PC and HoloLens views completely separated, showing different views, different selections, maybe even different data, like two independent applications.

p5 understands better events with visualization, yet he typically does not need to do measurements in visualization software, thus keeps perspective views which are enough.

3.6.3 *Synchronization*

Understanding the synchronization between the AR and PC views gives insights on how to design and use such a system. P1–P4 would prefer to have the two spaces totally separated and self-configurable, like two different applications. P5–P7, instead, would prefer keep the two different views yet which both reflect the same state and dataset. We present other individual preferences below.

p1 prefers to have two different views/interfaces. While not synchronized, it looks weird if the two views have the same interface but do not show the same content.

p3 has multiple preferences, depending on the application. If he works alone, he would prefer both sides to be configured separately, while he would keep them synchronized for public presentation and collaboration.

p4, p5 prefer to have switchable configuration by users. In one case, one can work on the PC and see the changes directly on the HoloLens; in another situation, they would keep the one state on one side and to then be able to easily compare the different states.

p6, p7 were confused when the two views had the same interface but were not synchronized.

3.6.4 *Input*

All participants agreed that the AR space should not be used only as a static screen, but be interactive using better interaction techniques as Figure 3.20 shows. In our study, however, we saw that certain implementation limitations can play a big role in the perception of

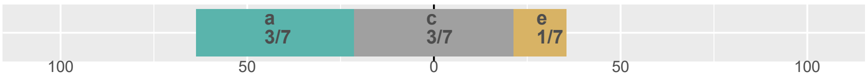

Figure 3.20: "For interacting with the HoloLens/AR view (e. g., selections), I would like to use" (a) completely gesture/hand based input, (b) in-between, (c) input that combines hand and device actions such as a pen to point directly at a track, standing next to it, (d) in-between, or (e) fully hardware-based input where the device is separated from the 3D HoloLens view and that only uses a virtual pointer.

user input devices. Specifically, our participants did not appreciate the current state of the mouse and keyboard input because, in particular, the mouse movement was sampled and transmitted to the HoloLens at the frame rate of the program. This caused the mouse movement to not appear smooth and instantaneous on the HoloLens, compared to the PC. Besides these solvable technical issues,

p1 is familiar with the mouse which is good for precise interaction, but seems not interesting to be imported to AR.

p1, p4, p6 want a real 3D mouse for AR.

p2 may be interested in spatial 3D trackers and joysticks.

p3 is more used to touchpad or touch screen input.

p5 is more interested in 3D mice, while moving 2D ones in 3D is useful as we can easily click on something.

Both P3 and P7 thought that 2D gestures on touch pad or screen could facilitate the interaction from their previous experience, especially for zooming. We also observed that all participants used a lot of zooming while exploring the dataset. Although we saw a preference for investigating the use of a 3D mouse, P6 envisioned other problems for such input because a 3D mouse can be hidden behind or inside the data. Another difference compared to PC is that an AR view has almost unlimited space, so that the cursor can get lost due to a fast interaction (but also because of the narrow field of view).

All participants were willing to use hand gesture input in AR:

p2 thinks that they are less precise but natural.

p4 felt the urge to touch the hologram in space.

p5 would not use them all the time as they can be fatiguing. However, he wants to enable gestures for certain tasks, complementary to a hand-held device.

p6 is interested to navigate and select with hands and fingers, but only if they are accurate enough.

p7 would use them on the HMD and use a mouse for the PC.

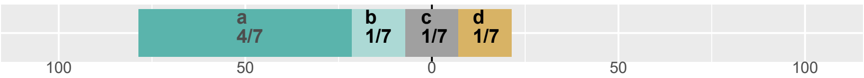

Figure 3.21: "Seeing the AR stereoscopic data from different sides and perspectives, " (a) I prefer getting up and walking around the data view, looking at the data from different locations, (b) in-between (c) I like both, (d) in-between, or (e) I prefer remaining seated, using rotation/translation to see the data from different vantage points.

No participant wanted to use voice commands due to, e. g., the problems in office settings and possible accent issues. Also, no one revealed specific comments on the gaze control.

3.6.5 *Walking around or remaining seated*

Getting up and walking around the data using the HoloLens or remaining seated and using a rotation/translation interactions are often discussed with AR setup (e. g., [Filho et al., 2019]), we present results in Figure 3.21. During the experiment, participants mostly sat on the chair, but had a few attempts to get up. P3–P6 thought that walking around or into the data could be quite helpful for data understanding. We report other individual comments below.

p1 sees no point of walking around, which is especially limited in an office. For him, regular interactions are sufficient.

p2 is wiling to walk, but it may be useless unless other input is supported since mouse&keyboard are not carryable.

p4 prefers moving his head to walking or mouse-based data manipulation. He sees the gaze-based exploration as a main difference to the PC, yet it is less practical in an office.

p6 sees it as the HoloLens' main advantage. However, analysis will be interrupted due to disconnections with the PC.

p7 was disappointed due to the limited field of view.

3.6.6 *Application and collaboration*

We asked participants about their envisioned application and collaboration scenarios, supposing there are no hardware barriers. P1, P2, and P4–P6 stated that a hybrid system could be interesting for collaborative meetings. Further comments included that

p1 finds it interesting to view the data as well as messages from others. Unlimited space could help collaboration.

p2 envisioned the system to be useful to present to others.

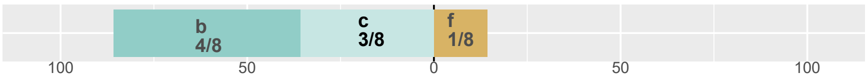

Figure 3.22: "In the future I could envision to use a hybrid AR interface" (a) 8h a day (b) a few times a day (c) a few times a week (d) 1–2× per month (e) 1–2× per year (f) not at all. P4 voted for 2 options.

p3 similarly, specified that the system would be useful for general public presentations but not for experts.

p4 thought that, during collaboration, only one person should interact with the data at a time, while others only observe.

p6 would allow a larger but finite number of users to interact.

3.6.7 *Realistic usage in daily work in the future*

We summarize our participants' envisioned future amount of use of a hybrid setting in Figure 3.22. P7 sees no realistic usage in her daily work since visualization is not her primary task while others disagree. However, to make such hybrid system realistic, P1, P3, and P6 would expect other specific functions. Even though they all work in high-energy physics, their work requirements differ significantly. P4 voted for two options because he thought the realistic usage would depend on whether the task requires more analysis or visualization. He detailed, e. g., that it should have a way to import, export, or communicate the model and settings with other software.

3.6.8 *Comments of each platform*

We report some other comments regarding the HoloLens and the PC.

p1 appreciated the large AR canvas which facilitates working on several things simultaneously. However, the PC has easier accessibility and usability: everyone knows how to use the PC mouse and interact with standardized UI elements.

p2 sees that the additional spatial dimension shown by the HoloLens allows people to see the data in a more intuitive way (e. g., needing less rotation that on the PC). However, people are well trained to use and understand data on PCs.

p3 thinks that the HoloLens improves 3D perception because we can "walk into things" and arrange elements in space.

p4 commented that the AR visual immersion improves data understanding, and that the PC is good for precise control and the display of abstract information.

p6 believes that the PC is good for quick input such as typing, its high resolution allows us to display more details. Its familiarity is also an advantage.

p7 appreciated the HoloLens' better depth perception, but thinks that interaction is easier on PC.

3.7 DISCUSSION AND LIMITATIONS

While our participants' responses to our set of questions as reported in the previous section already provide a lot of inspiration for the future development of AR-supported data exploration environments, we now discuss them in the context of our overall vision of a hybrid system.

3.7.1 *Lessons learned*

We observed that due to our design of the AR-part of the interface to match that of known tools, all our participants quickly understood how to work with both parts of our hybrid setting—no participant expressed a fundamental uneasiness about the new design. While this is not necessarily surprising, it suggests that such a design may lead to higher adoption rates than VR-only setups. Five of our seven participants stated that they would use a hybrid system instead of pure VR/AR or pure desktop systems. However, the placement of the stereoscopic views needs more consideration to avoid disturbing existing (horizontally aligned) screen layouts.

The AR view was clearly seen as complementing the PC—most of our participants, like us, do not expect it to fully replace traditional workstations. Some participants expressed that they would carry out certain types of analysis (script writing, abstract data visualization) on the PC, while they would prefer 3D inspection on the AR view. For the synchronization between PC and AR the opinions were diverse, people suggested scenarios where constant synchronization could be useful as well as other situations were the displays should be akin to separate applications. The possibility of getting up and "walking around or into the data" was evidently quite novel to our participants, but we suspect that people hesitated due to the novelty of this form of data inspection and the environment of an office not inviting such actions. When we developed the system, however, we had observed that one of our collaborating domain experts did get up on his own accord to look at the data "more closely" in an intuitive way. So we believe that this is an exciting possibility for analyzing 3D data.

The main advantage of AR to our participants seemed to be its virtually "unlimited space"—not only for 3D content but also for the visual comparison of other views. Future work should thus investigate how to best make use of this space. Following our initial vision, the AR view could potentially extend existing tools by providing both screen space and stereoscopic views of the 3D data, yet in a fundamentally different way than another 2D or auto-stereoscopic screen on the desk.

Nonetheless, the AR extension is not seen as simply a static 3D stereo view, instead people feel a strong desire to interact with it. Our simple replication of the 2D mouse in AR space did not feel comfortable to people—partially due to the mentioned technical limitations. Yet even if these problems were resolved, it seems that a dedicated input device such as a 3D or space mouse may be more useful. After all, we also change from keyboard to mouse and back during interactions with regular workstations so another dedicated input device may not feel as disturbing as one may expect. This device would need to provide similarly precise input like a regular mouse in 2D to support the precise interactions needed for data analysis. We also would need to understand better how to use a 3D cursor (or a cursor specifically designed for the HoloLens) —it should be inspired by the 2D counter-part yet may need special functions to avoid, for instance, it getting lost in the mentioned large AR interaction space and to always be visible, even in dense data situations.

While our participants mentioned several other possible forms of input such as hand gestures or joysticks, we are hesitant about such designs without empiric evidence that these would provide as much flexibility, control, and precision as a mouse. In particular gestures in empty space—even if envisioned by our participants to be intuitive and "natural"—can quickly become tiring due to the gorilla arm syndrome [Hincapié-Ramos et al., 2014]. One interesting and promising idea, however, is the use of gestures on a potentially existing (laptop) touchpad as they are currently used to augment the interaction in traditional interfaces. Certain well-defined multi-finger gestures for 3D navigation (not only but including two-finger pinch-to-zoom) could be an excellent form of input for the AR space.

Much to our surprise we found that the lack of orthographic projection in AR did not bother any of our participants, despite the prevalence of orthographic views in traditional 3D data analysis tools like the ones used in particle physics. In the future we thus would be interested in studying whether AR views with correct perspective are similarly precise and efficient as orthographic projections (within the domain of particle physics and elsewhere) for solving tasks in 3D space because this is a fundamental prerequisite for an effective 3D data analysis. This apparently "correct perception of 3D shapes" may also shed light on the limitations and benefits of 3D representations of abstract (i. e., non-spatial) datasets [Brath, 2014].

Finally, our participants' suggestions to use an AR-augmented environment for collaboration appears to be straight-forward, yet also raises numerous interesting research questions. In addition to known challenges of collaborative work, we would be interested in how people would actually physically immerse themselves into 3D data representations by walking around in the views, and to what degree this could support data analysis tasks in our application domain of particle physics.

3.7.2 *Generalization*

Our study was based on a PC and a HoloLens using particle physics datasets. However, we believe that our findings can be generalized. Any setup that combines a PC with some sort of stereoscopic 3D display can benefit from our discussion on how to add immersion to existing workflow, but with some limitations regarding the walk-around feature and the management of multiple users in collaborative scenarios. As we stated in our design choices, our results also hold in a possible pure-AR environment which makes use of physical mouse and keyboard. The question remains whether our observations for our specific application domain generalize to other domains, specifically since the daily tasks of our seven participating experts already differed significantly. Yet our participants and anybody dealing with some form of 3D data have to carry out at least the same fundamental 3D manipulations techniques and use similar visualization tools. The tasks for analyzing spatial data are comparable to other applications such as air traffic control [Hurter et al., 2019], and even 3D visualizations of non-spatial data [Bach et al., 2018] require similar forms of interaction. We thus believe that our observations can generalize to or at least inform the interface design in such related domains. Practitioners can thus build on our findings to build hybrid systems more specifically adjusted to specific visualization and data exploration needs.

3.7.3 *Limitations*

In addition to the known limitations of the hardware of the HoloLens 1.0, our prototypical implementation had limited functionality and the communication between the input devices and the AR view exhibited the discussed lag. However, the hardware and software setup only served as a basis to investigate the future design of a more complete solution, futures iteration will focus on a more specific set of tasks within our application domain. We expect that improved AR hardware will remove some of the known technical shortcomings (e. g., limited field of view, low resolution), yet its handling will be similar to the present version. Naturally, the specific experiment design and, in particular, the use of *think-aloud* in the presence of an observer has the potential to bias participants. Yet our study design has established way to conduct observational experiments to extract people's feelings and ways of thinking about technology. We believe that with this approach we obtained much richer input for creating novel hybrid interaction designs than we would have with any controlled speed-and-error experiment. Finally, while we emphasized during the experiment that we did not want the physicist to focus on the technical achievements that the Hololens represent, it is still possible that this novelty effect might have biased their subjective answers. To address this issue we will continue to collaborate with a number of experts in particle physics and work toward a more refined design of a hybrid system that is more usable for actual data analysis and through this continued collaboration we expect to be able to conduct more controlled experiments after our collaborators have gotten used to the new devices.

3.8 CONCLUSION

With the results from our observational study, we provide insights on how to bridge the chasm between the potential benefit of immersive environments and the lack of adoption in the domain sciences. Results suggest that, first, scientists strongly favor hybrid AR setups where the AR complements the PC. Second, content in the two environments should not constantly be linked. Third, walking through the data is fundamentally more intuitive than view manipulations. And the view-based access to lots of virtual screen space is one of AR's main advantages. We thus open up a completely new form of immersive system design for visualization: we no longer need to decide between immersion and existing tools, but we can use the best of both worlds. Such insights are not limited to particle physics, practitioners working with similar 3D dataset can benefit from our results to extend their data analysis tools with immersive views.

This chapter provides some basic answers for the research questions **Q1** and **Q2**. Experts confirmed that using the AR display as an extension to the traditional 2D displays is feasible and has many benefits. We thus confirms that a data exploration system that combines an AR display with traditional workstation that be used practically in an office is simply that users wear an AR HMD while normally working with the PC using traditional mouse and keyboard, even though the input devices needed further consideration for the AR space. As for the transition between the two environments, event though most participants preferred that the two spaces could be configured separately with a function of synchronization on request, it is still be useful to let users define flexibly define it, to benefit different exploration tasks.

4

UNDERSTANDING DIFFERENCES BETWEEN COMBINATIONS OF 2D AND 3D INPUT AND OUTPUT DEVICES FOR 3D DATA VISUALIZATION

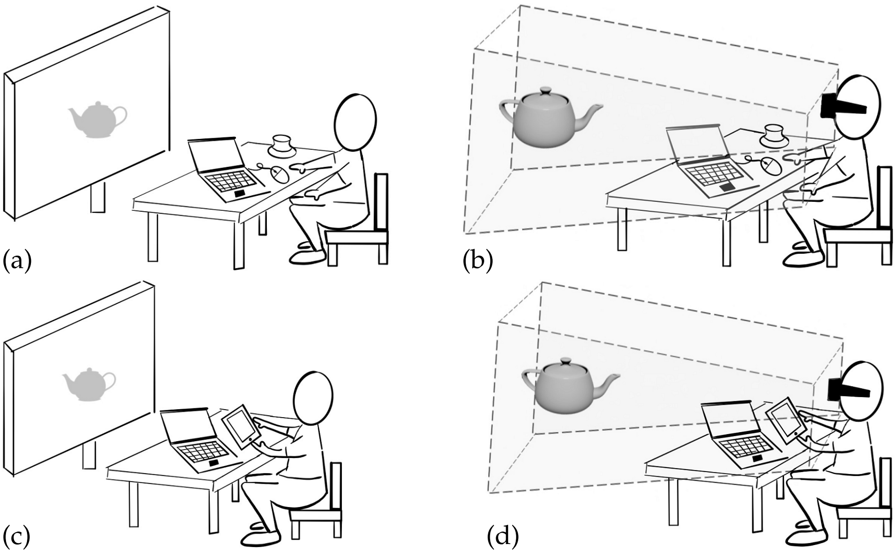
(a) (b) (c) (d)

Figure 4.1: Illustration of experiment setup. (a) users use a mouse or a space mouse with a screen; (b) users use a mouse or a space mouse with a HoloLens; (c) users use a tablet with a screen; (d) users use a tablet with a HoloLens.

Chapter 3 confirmed the feasibility and the advantages could be brought by a hybrid AR and PC visualization system. However, an interaction question arises with our design: the mouse interaction needs to be revised, which is also related to **Q3**. From the results of Chapter 3 and according to previous literature, imitating desktop UI widgets in AR space seems not to be a good choice, thus we do not encourage the heavy use of screen widgets in our application fields. However, it is still important to offer some levels of control with mouse in AR, especially for view and object manipulating. We thus want to understand if using mouse to navigating in 3D causes a severe problem in AR space, to ultimately design the interaction for the hybrid system.

Thus, in this chapter, focusing on interaction needs for scientific data exploration, we evaluated people' performance using 2D (mouse) or 3D (space mouse and tangible tablet) input devices to interact with visualizations shown on 2D screens or stereoscopic AR head-mounted displays. The increasing availability and power of immersive displays drives us to try to understand how to choose input devices, interaction techniques and output displays for the visualization of scientific data, thus to finally help us guide the interaction design for a hybrid AR and PC visualization systems. With a docking task and a clipping plane placement/orientation task, we measure our participants' performance (completion time and accuracy) with each of the different combinations of input and output devices. We also report on their perceived workload, their preference, and on other qualitative feedback. Results show that the mouse remains good with any display, especially for tasks that require a high accuracy. Our results highlight the potential to retain the mouse as a primary input device, and to complement it with other 3D interaction devices for specific uses.

Findings of the piece of work presented in this chapter serve as the guidelines for the interaction design of hybrid AR and PC visualization systems.

Main portions of this chapter were submitted to Elsevier and are under review. The use of we in this chapter refers to myself, Lonni Besançon, Mehdi Ammi, and Tobias Isenberg.

4.1 INTRODUCTION

Exploring and understanding three-dimensional spatial data is key for many scientific disciplines (e. g., computer-aided design, biology, and mechanics). Recent technological developments are allowing users to break free from traditional workstations that consist of 2D screens, a mouse, and a keyboard. For example, analysts can now use devices that range from small-scale portable displays like mobile phones to large surfaces like a wall-size screens and fully immersive environments. With the releases of affordable commercial devices, immersive environments with stereoscopic output are attracting more and more attention in the field of visualization and interaction. Already in 1996, Bryson [1996] pointed out that the visualization of scientific data with a stereoscopic view has a huge potential due to the natural match between the dataset's inherent spatial proprieties and the 3D visual output space as well as the possibility of integrating input devices that differ from the traditional mouse. Later studies also highlighted that stereoscopic views are beneficial for tasks related to understanding scientific data with volume or isosurface visualization, both of which are inherently spatial (e. g., [Prabhat et al., 2008; Laha et al., 2014; Murray et al., 2017]. Recently, the field of immersive analytics [Dwyer et al., 2018; Marriott et al., 2018] has emerged to specifically investigate such settings.

Besides exploring the interaction and visualization techniques with purely virtual environments, visions on integrating them with other setups exist as well (e. g., [Zielasko et al., 2017; Fulmer et al., 2019; Surale et al., 2019; Wang et al., 2020]). We follow the idea of using AR to extend traditional workstations because it allows us to add stereoscopic views to existing analysis tools, thus improving the data exploration workflows. While the benefits of visual immersion have been demonstrated (e. g., [Prabhat et al., 2008; Laha et al., 2012; McIntire et al., 2014b; Hurter et al., 2019]), questions of choosing input devices and designing interaction techniques remain unanswered. Especially with a hybrid AR and PC setup, users need to continuously switch between 2D and 3D views as well as control the two types of views simultaneously or synchronously. The latter leads to interaction needs that differ from the Desktop-VR metaphor [Tait, 1992] where users are forced to work with only one space. Due to their difference in spatial dimensionality, the two output spaces usually rely on fundamentally different input devices and interaction techniques. It is indeed possible to use one input device for one space and a different one for the other. For example, Millette and McGuffin [2016] proposed to use the mouse to control the PC, but switch to a tablet for an AR view. However, we believe that it is essential to provide a seamless interaction experience if our final goal is to design a hybrid system for daily use. Another important consideration is that visualization tasks usually require precise interaction that makes popular input devices in other VR/AR research areas (e. g., 3D gaming) less useful. For example, 6-DOFs tangible devices [Ishii, 2008] are often explored in VR/AR environments since they can map themselves directly to the objects in the virtual world. Previous studies (e. g., [Bérard et al., 2009]) suggest, however, that for certain tasks the mouse remains the most efficient input device when compared with other high-DOF input devices. It was also suggested that adding input DOFs to a traditional desktop is not always useful [Mendes et al., 2019]. Highlighting the importance of the mouse, a recent study by Wang et al. [2020] used it to control both a traditional workstation and an AR headset to visualize particle traces in high-energy physics. Although the mouse was seen as necessary for precise and familiar control, participants expressed to be willing to try other 3D input devices as the primary input for the hybrid space. Motivated by these results, we set out to understand how to choose an appropriate input device to fulfill the interaction needs of exploring data with a hybrid PC and AR setup. Our main question here is whether user performance varies with different devices. More specifically, we want to understand, first, if user performance of using the mouse changes when they perceive the data in different views and, second, if 3D input devices have an advantage over the mouse in such environments. While general interaction needs are complex and can include pointing, selecting, and drawing (and each of these tasks would require a different interaction design), we focus on the first step—manipulating object positions and orientations in 3D space; we do not investigate the design of specific 2D or 3D cursors for the interfaces.

We thus conducted a study to understand these questions with a docking task and a clipping plane positioning/manipulation task, both of which are highly relevant in the scientific visualization domain. We explored how user performance changes while using 2D (mouse) or 3D (space mouse and tangible tablet[1]) input devices to work with visualizations shown on a 2D screen or using a AR HMD. We report the measured performance (we use completion time and accuracy) as well as qualitative feedback reported by our participants. The results of our study do not suggest a universal correlation of how user performance changes with different input paired with the screen or the AR output, but we observe that, in some cases, the two different ways of viewing data influence the performance, even though users may not feel it. Moreover, for our second task of manipulating a clipping plane where high accuracy is needed, the mouse still performs well, despite the dimensionality disparity with the stereoscopic view. We thus conclude that the 2D mouse continues to be a strong contender for 3D spatial tasks, likely due to people's general familiarity with this input device. It thus remains a good choice to use the mouse as the

---

[1] A tangible tablet is a tablet with both touch and tangible (spatial-aware) input.

primary input when designing hybrid visualization systems, while 3D input devices can serve as complementary controls for specific use cases.

4.2 RELATED WORK

3D object manipulation techniques have been extensively studied in the past. As have mentioned in Section 2.3, previous surveys summarized a variety of interaction techniques for environments with either 2D screens or stereoscopic views. Besides the traditional mouse-based interaction, many new forms of input have been created and investigated as well. For example, although mid-air gestures are argued to be natural for immersive environments, severe challenges exist for their long-term adoption in scientific analysis, especially due to their low precision and the fatigue they cause [Filho et al., 2019]. In addition, touch-based and tangible interactions have been explored extensively and a common setting relies on a tablet (with touch and sometimes space-aware input) to allow people to interact with the virtual environment. For example, Büschel et al. [2019] designed a pan and zoom technique for data exploration in AR space based on a phone using both spatial and touch interaction. Millette and McGuffin [2016] also used a spatially-aware mobile phone to interact with a hybrid AR-PC system. Finally, there are devices specifically designed for immersive analytics settings such as the CHARM input device [Klamka and Dachselt, 2015; Klamka et al., 2019] and on-body devices [Fruchard et al., 2019]. While such devices show a lot of potential, they cannot be easily integrated into daily workflows due to not being commercially available yet. In this work, we first focus on existing and well-established input devices to guarantee the accessibility of the general public.

While many input devices offer high DOF counts, it is unclear if they are well suited for 3D visualization tasks. Some researchers believe, in particular, that high-DOFs input devices have advantages compared to 2DOF input for 3D tasks. For example, Schultheis et al. [2012] argued that such devices have inherent advantages compared to the mouse for 3D tasks. Another study [Besançon et al., 2017c] also discovered that users were faster with tangible devices in 6 DOFs manipulation tasks to achieve the same level of accuracy. Spindler et al. [2014] also revealed the value of spatial-aware devices. However, a study by Bérard et al. [2009] showed that the mouse remains the most efficient input device for a 3D placement task, when compared with high-DOFs input devices. This controversy may be explained by several reasons. First, the *legacy bias* [Morris et al., 2014], which emphasizes that users stick to the familiar, well-known interaction techniques such as the mouse, could partly explain the results. Second, it is commonly recognized that different input devices suit different tasks. For example, Sundin and Fjeld [2009] showed through several studies that input devices with isotonic or softly elastic position control and with softly or stiffly elastic rate control are appreciated for different tasks, as the commonly recognized concept of *"everything is best for something and worst for something else"* explained by Buxton [2007]. Thus, researchers often rely on empirical studies to specially compare different input devices for a certain task (e. g., [Zhai and Milgram, 1998]). According to previous studies, users sometimes express high preferences for decomposing manipulations to accomplish complicated tasks (e. g., separating the motions of translation and rotation). Such DOF-separation, which can easily be achieved with low-DOF input devices, favors tasks that require a high amount of accuracy [Bérard et al., 2009; Veit et al., 2009; Wang et al., 2019a]. Even for immersive environments, it has been argued that the mouse and keyboard are still essential for data exploration as we cannot totally separate the visualization from other analysis process that require the use of traditional PC-based tools [Wang et al., 2020]. For a hybrid AR and PC environment, a mouse is thus often kept to interact with the PC. However, it is unclear if its advantages for 3D interaction (such as the separation of DOF for high accuracy) hold if we perceive the data in stereo AR space. To understand if it is a good practice to use the mouse, important questions are, e. g., whether we should keep the mouse as the main input device for immersive analytics where a high accuracy is required and whether it is useful to extend desktop-based mouse interaction techniques to stereo AR displays. Although 3D mouse input is often argued to be problematic such as in the VR-Desktop metaphor [Tait, 1992] due to the difficulties of designing a 3D cursor [Teather and Stuerzlinger, 2013; Schemali and Eisemann, 2014], we focus on the performance changes caused by the mapping from input to virtual space.[2]

A study that especially inspired our own work was conducted by Schultheis et al. [2012] who compared mouse, wand, and two-handed input for docking and construction tasks. Comparing the speed, their results indicated a significant difference in terms of interface, but they did not find significant differences in mono- or stereoscopic views. Different from this work, we want to investigate tasks relevant to 3D visualization, and use input devices that could potentially be integrated into scientists' existing working environments. Another relevant study by Bach et al. [2018] compared different interactive visualization environments, and concluded that there is no universal method that outperforms others but rather depend on specific tasks. Even though they included some discussions on the dimensionality match between the input and output devices, their study used setups

---

[2] We use the term "mapping" to indicate that the input device's action is transmitted to the virtual object's movement, which has the same meaning as a "transfer function" used in other literature (e. g., [Sundin and Fjeld, 2009]).

that all have different input and output devices, thus did not compare the difference of performance while one device paired with others (for example, the mouse paired with a screen and a HoloLens). Also, their focus was to investigate scenarios where users can directly touch and interact with the data, which is different from our hybrid scenario where the AR serves as an extension of the PC.

4.3 OVERVIEW AND RESEARCH HYPOTHESES

As we discussed, we need to study how different input devices compare to each other and which one (or which combination) would best be suited to a hybrid PCAR setting. To support our ultimate goal of seamlessly integrating an AR extension into a current workstation that already is equipped with a mouse and a keyboard, we are particularly interested in investigating input devices that can also be easily integrated with PCs in general working scenarios. As noted, we need to keep mouse and keyboard for controlling existing data analysis software and for script-writing. Since mice only provide 2 DOFs input and rely on a horizontal surface, we decided to also examine a tangible space-aware tablet (from Google's Project Tango, now part of the AR core) as a representative of 6 DOFs devices with zero-order positional input. Third, we also include the 6 DOFs SpaceNavigator[3] (we refer to it as a space mouse in this paper), which uses elastic rate-control [Zhai, 1998; Bérard et al., 2009], because it has been an established input device for 3D manipulation for a long time, and its major characteristic (elastic rate control) has been demonstrated to be suitable for certain navigation tasks [Sundin and Fjeld, 2009]. It also can be a complementary to the mouse interaction to facilitate 3D manipulation. As we discussed, certain previous studies have shown advantages in 3D manipulations (e. g., [Bérard et al., 2009; Besançon et al., 2017c]), while others contradict such findings and highlight that high-DOF input devices are favorable, especially when people are working with 3D environments. We thus want to further investigate the difference between low- and high-DOFs input on users' performance, paired with both a traditional 2D screen and an AR headset. Taking the high precision requirements of visualization into account, we believe that easy separation of manipulation offered by low-DOFs input helps users to achieve precise control. We thus formulate the first hypothesis as:

h1 For visualization tasks that require high accuracy, 2D input devices such as a mouse will yield better performance (regarding traditional HCI measurements like speed and accuracy) than 3D input devices.

---

[3] `https://www.3dconnexion.fr/spacemouse_compact/`

$$\begin{tikzcd}
& \mathcal{V} \arrow[r, "o"] \arrow[dd, dashed, "vi"] & \mathcal{O} \arrow[dr, "u_p"] \arrow[dd, dashed, "om"] & \\
\mathcal{D} \arrow[ur, "v"] & & & \mathcal{U} \arrow[dl, "u_a"] \\
& \mathcal{I} \arrow[ul, "i"] & \mathcal{M} \arrow[l, "m"] &
\end{tikzcd}$$

Figure 4.2: Model of interaction directness by Bruckner et al. [2019].

In addition, many studies (e. g., [Prabhat et al., 2008; Laha et al., 2012; McIntire et al., 2014a; Hurter et al., 2019]) demonstrated that a stereoscopic view helps users to understand spatial data. Yet most of them used fully immersive VR environments. An AR headset, in contrast, does not occlude users from the real world and an thus easily be integrated with existing input devices. It is also likely, however, to introduce distraction (surrounding environments) and/or decreases the image quality due to, e. g., the real-world illumination and background color. Nonetheless, we believe that the stereoscopy provided by AR headsets has advantages for 3D data analysis and manipulations and thus state that

h2 For all examined input devices, users' performance in accomplishing 3D visualization tasks is generally better in AR output space than on a 2D screen.

We then take a step back to examine the potential reasons for these hypothesized differences. Here we base our considerations on the model of spatial directness in interactive visualization by Bruckner et al. [2019] (Figure 4.2). It describes the transformations from data (space $\mathcal{D}$) to visual representations (space $\mathcal{V}$), to an output medium (space $\mathcal{O}$), to what the viewer understands or their mental model of the visual representation (space $\mathcal{U}$), to manipulations of an input device (in space $\mathcal{M}$), to interpretations of this input in form of interaction mappings (space $\mathcal{I}$), and back to the data space $\mathcal{D}$. Transformations between these spaces include projections from higher- to lower-dimensional spaces or *vice versa*, and they facilitate the discussion of the spatial directness of interaction with data representations. Most relevant for our own discussion is the right-most "triangle" of the model, i. e., the mappings between output space $\mathcal{O}$, user space $\mathcal{U}$, and manipulation space $\mathcal{M}$. In our application domain of 3D data analysis, we typically deal with data that has an inherent mapping to 3D space (e. g., simulations in fluid mechanics; Figure 4.3(c)). Scientists analyze this data based on a 3D mental model of the volume in space $\mathcal{U}$. In traditional workstation setups, both space $\mathcal{M}$ using a mouse as input and $\mathcal{O}$ using a 2D screen are two-dimensional (interestingly, the space $\mathcal{M}$ is even rotated by 90 degrees with respect to space $\mathcal{O}$ as explained by López et al. [2016]). This setup creates a mismatch between the dimensions used by the input device, by the output space, and by the mental model of the user (which, in this case, resembles the 3D data space $\mathcal{D}$). If we use 3D input devices to interact with the AR spaces, however, all these spaces have matching dimensionality and this setup should lower people's workload as they explore the data. Nonetheless, in our hybrid working scenario users may sometimes interact with a 2D input device, while focusing on the AR display. Or they they may use 3D input, while interacting with 2D displays. This mismatch may create additional distances in people's mapping between output ($\mathcal{O}$), user ($\mathcal{U}$), and manipulation ($\mathcal{M}$) spaces of the model. We thus formulate our final hypothesis as

h3 We hypothesize that (mis-)match of input and output dimensionality plays a role to users' performance. When input and output dimensionality match each other, user performance will be higher than if there is a mismatch.

At a first glance, ***h3*** seems to be inconsistent with, in particular, ***h1***. The apparent contradiction is that, while working with AR, the 2D mouse creates a mismatch with the output space, while 3D input does not, and the 3D input devices should thus outperform the mouse with 3D output. Our assumption here is, however, that our application domain requires high accuracy and that the DOF separation remains the most important advantage compared to other factors. Nevertheless, the mismatch will still play a role. We thus expect that users' performance using 3D input (like 6 DOFs tangible device) will largely increase paired with AR, while such improvement may be limited with 2D input.

4.4 STUDY

We designed an experiment to examine these research hypotheses. This study was approved by the ethics review board of our institution and pre-registered on `osf.io/7gsk8` to follow latest guidelines to make research more robust [Cockburn et al., 2018; Cockburn et al., 2020]. As we note in this repository, while our preregistration mentioned a specific date to end data collection for 24 participants, time constraints of potential study participants and limited hardware availability led us to extend the deadline for data collection, without changing anything about the study setup. We describe the selected tasks and devices for our experiment as well as the design and the procedure of the study next.

4.4.1 *Tasks*

In general HCI, a variety of tasks have been used to evaluate input devices and 3D interaction techniques, such as positioning/placement tasks, selection tasks, and navigation tasks. As our application domain

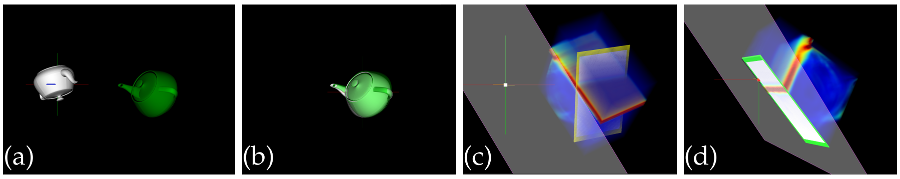

Figure 4.3: Examples of tasks shown on 2D screen: (a) docking task initial status (blue z-axis highlighted for illustration), (b) docking task expected validation status, (c) clipping plane initial status, and (d) clipping plane expected validation status.

is 3D visualization, we specifically needed to select tasks that are highly relevant to it. For this reason, many tasks that are studied in HCI are not well suited for us. For example, the steering task [Accot and Zhai, 1999; Cohen et al., 1993] asks users to follow a narrow path. While it has been been studied in detail with a well-established trade-off model, its application in the visualization field is limited. Similarly, while pointing tasks were also investigated and even specific devices have been designed to improve the performance such as a touch stick [Fallot et al., 2006], we do not include it in this study because we wan to first focus on essential navigation tasks. These include the specification of a volumetric view and object manipulation including clipping, which together were highlighted as one of the fundamental groups of spatial interaction techniques for visualization by Besançon et al. [2021] in their recent survey. Other work that analyzed interaction tasks for 3D data exploration for both general 3D spatial data (e. g., [Keefe and Isenberg, 2013; Besançon et al., 2017b]) and domain specific datasets (e. g., [Laha et al., 2014; Murray et al., 2017]) also stated that specifying 3D positions and orientations is essential because derived tasks such as object manipulation, view navigation, and selection and higher-level classification with a specific context like pattern searching/identification and feature comparison rely on such basic interaction. We thus believe it is important to first compare the different input device and output platform parings in a general 3D interaction scenario, and we thus used a 3D docking task in our experiment. We also asked participants to perform a follow-up task of positioning a clipping plane inside a volume rendering to align our experiment better to a data-visualization context. Both tasks require interaction in 6DOFs (translation along and rotation around x-/y-/z-axis). Although clipping plane manipulations can technically be done with 3DOF input, the use of 6DOF input makes them easier as we argue below.

4.4.1.1 *3D docking task*

A 3D docking task is commonly used to evaluate 3D interaction techniques (e. g., [Chen et al., 1988; Hancock et al., 2007; Glesser et al., 2013; Besançon et al., 2017b; Wang et al., 2019a]). It requires participants to manipulate an object's position and orientation in 6 DOFs to match a target's spatial transformation. We used the Utah teapot in this experiment for its simple shape. Simple objects are commonly used to evaluate interaction performance; e. g., Chen et al. [1988] and Hinckley et al. [1997] used a house, while Zhai and Milgram [1998] used a tetrahedron. We used the Utah teapot in this experiment because its shape is easily understood by non-experts, without orientation ambiguity. We colored the teapot that participants manipulate opaquely gray, while we show the target in semi-transparent green (Figure 4.3(a)). Participants thus get direct feedback about their current relative depth when both objects intersect (the target does not occlude the object). After pilot studies with three members from our lab, we added a representation of the transformation axes of the manipulated teapot to assist the participants during their object manipulation, especially to clearly indicate the z-axis in perspective projection. To illustrate this assistance, we highlighted the z-axis (blue) in Figure 4.3(a) by making it thicker; in reality it is as thin as the x- and y-axes.

For this docking task, we measure accuracy (both Euclidean distance and difference in orientation) and task completion time, without favoring one over the other. Instead, we ask participants to balance their interaction speed and accuracy and to decide when they should finish each trial.

4.4.1.2 *Clipping plane task*

For the second clipping plane task, we asked participants to manipulate a semi-transparent plane to clip a volumetric dataset. We used a fluid mechanic dataset in which color represents the fluid velocity (Figure 4.3(c)). We render a hidden white plane inside the data and ask users to manipulate the semi-transparent clipping plane to align both and to be as close to each other as possible. Such a task would require participants to first find and understand the orientation of the target, before actually aligning the clipping plane with it. Our pilot studies showed, however, that using opaque volume is much too difficult, and frustrates the users to give up quickly. We thus simplified the task by making the data semi-transparent, coloring the borders of the hidden plane to yellow, and showing these borders on the outside of the data volume. In addition, we colored those parts of the target green when they are within a small threshold of the clipping plane at a given time. While this task may differ somewhat from those used in real scientific data exploration where scientists do not have a clear target before exploring the data, our compromise avoids frustrating the participants and we believe that it still allows us to study the dimensionality match between input and output. In addition, we show the virtual pivot point of the clipping plane and its axis of manipulation as in the docking task—pilot studies also showed us that participants would have problems understanding how a flat plane can be manipulated compared to a rigid 3D object if they are not aware of the interaction pivot and axis. We show the virtual pivot point initially as a small white sphere and change it to red when it is on the hidden surface. We believe that such modifications still allow us to examine our research questions as we want to evaluate user performance with different input and output devices, rather than the task, and they equally affect all conditions.

Theoretically, a plane with an infinite surface can be described with only 3 DOFs (rotations around two basis vectors lying on it and a distance along the normal direction from the coordinate origin). Nonetheless, we still offer a 6 DOFs interaction to make the interaction mapping consistent with that used in the docking task and to provide more flexible input. Although adding more DOFs may appear to complicate the task, pilot study participants said such compromise makes interaction actions more predictable. It thus becomes easier compared to 3-DOFs manipulation because the task requires participants to mentally map their manipulations from screen space to model space.

Even though large and complex scenes are important for 3D visualization, experts in many scientific domains work with the visualization of a small/medium size of volume, such as a human brain [Everts et al., 2015] or other biology tissues [e. g., Laha et al., 2012]. Moreover, we believe that, to enable the exploration of large datasets, it is essential to first understand the interaction with small datasets. In addition, we received the volume visualization dataset we used in our second task from our collaborators in fluid mechanics, and which they use in their practical work. We thus believe that investigating such tasks reflects real needs and is not less important than large and complex scenes.

4.4.2 *Apparatus and interaction mapping*

We used a Microsoft HoloLens (1st generation) and a 55″ display as output devices, each of them representing a different visualization output dimensionality (3D vs. 2D). Although intuitively a stereoscopic 3D screen may appear better suited for a controlled study to understand the dimensionality mismatch, our ultimate goal is to bring an AR extension to existing data analysis workflows. We thus directly use the Microsoft's HoloLens to investigate our research hypothesis with our envisioned setup as stated in Section 4.1. Nonetheless, we balanced the view between the two types of devices as much as possible, with the exception of the inherent difference w.r.t. depth clues that we wanted to study as well as some factors that are prescribed by the hardware, such as differences in screen resolution. We thus manually adjusted the parameters of the rendering camera on the 2D screen and the objects' sizes to ensure that participants have a similar perceived feeling with both devices. We specifically ensured

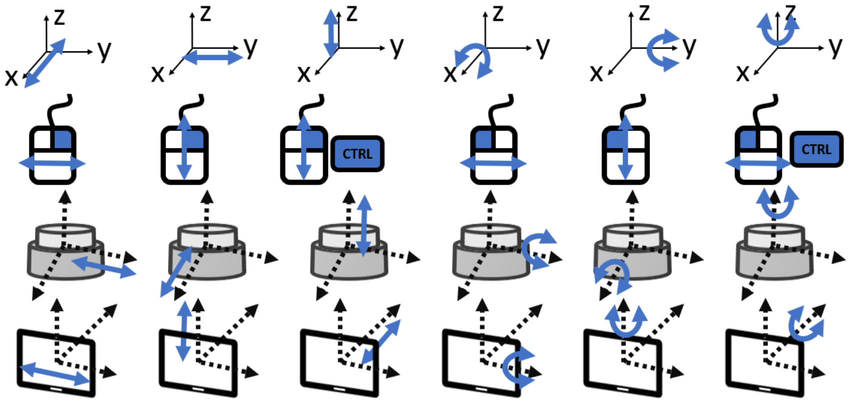

Figure 4.4: Illustration of interaction mapping.

that participants can see the same objects within their field of view with both devices and that the color of the virtual objects on different spaces is as close as possible. In contrast to the fixed camera of the 2D screen, the HoloLens camera can be changed fast by moving the head or walking around. We thus allow participants to manipulate both the object and the camera in 2D screen conditions, as well as the target volume for the clipping plane task. While participants manipulate the target object or camera, we also manipulate the object to dock or the clipping plane accordingly, such that their relative position and orientation does not change. Another inherent difference is that, in the HoloLens condition, users are able to change views faster using head movements and can thus better perceive the depth with the view-motion parallax depth cues, in addition to the stereoscopic view. We do not consider this aspect to be an important issue because we focus on the comparison between the AR display and the screen (as stated in ***H2***), rather than simply comparing mono- with stereoscopic views. We use the input devices descried in Section 4.3 that include a mouse, a space mouse, and a tablet.

We chose an interaction mapping for mouse (and keyboard) input which maps the left button to affect rotation around x- and y-axes, the right button to affect translation along x- and y-axes, with the control button pressed the left button rotates around the z-axis and the right button translates along the z-axis. According to Besançon et al. [2017c], this mapping is frequently used in PC-based 3D modeling and visualization software. For the translation with the mouse, we captured the translation of mouse cursor on the PC screen and adjusted the control-to-display (CD) ratio such that the translation distance of the virtual object in screen is the same as the translation of mouse cursor. For rotation, we implemented the well established Arcball interaction [Shoemake, 1992]—translating the mouse cursor from one border to the opposite rotates the virtual object by 180 degrees.

Based on our tests and pilot studies, we choose to use a similar interaction mapping for the space mouse. In particular, our mapping uses the plane of the table as a reference space, similar to the regular mouse—a mapping that used the screen plane as a reference space turned out to be too confusing. The movement captured by the software development kit (SDK) is not in real-world units but gives a small value each time a movement is captured. There is thus no direct transfer from the real world to the virtual space, and we thus adjusted the scale factors based on our pilot studies to make the users feel as natural as possible.

Finally, the interaction mapping for the space-aware tablet is a link between the tablet's movement and that of the virtual objects. While a direct match between the tablet's spatial position/orientation and the virtual object might be natural, such mapping would cause several problems due to both human and environmental constraints summarized by Wozniak et al. [2014]. First, participants cannot comfortably reach all places with their arms to control the object. Second, there could also be objects in the surrounding environments that limit the users' activity. We thus use clutching [Jacob et al., 1993] such that the tangible interaction is only triggered when participants touch a huge virtual button on the tablet's screen. In thus way, engaging the clutch is based on the touch input of the tablet and can be triggered at any time without affecting the tangible input, in compliance with the principles stated by Wozniak et al. [2014]. We thus make sure that participants can clearly see and interact with the tablet's screen, not only because of the needs of this study, but also it is important to support practical visualization tasks. In addition, we transfer the translation/rotation of the tablet to the virtual object's movement with a 1:1 scale.

Our interaction mappings for the three devices are also illustrated in Figure 4.4. We realized our experiment with Unity and C#. We also used the Google TangoSDK to capture the translation and the rotation of the tablet. Furthermore, we used Activiz library in Unity to read and create the textures of the fluid volume, which initially is in a VTK format. For the first task, we rendered the teapot with a single color, illuminated by a light placed at the camera position. We used the diffuse component of Phong's [Phong, 1975] model to make clearly show the teapot as a 3D object, as illustrated in Figure 4.3(a). For the second task, we used slice-based volume rendering to visualize the fluid volume. We used the same rendering technique for both the screen and the AR display.

We illustrate the general experiment setup in Figure 4.1. While performing tasks with the mouse and the space mouse, participants validate a trial or switch between object manipulation and view changes by either clicking on buttons or use pre-defined hotkeys on on the keyboard, as preferred. With the tablet, participants use virtual buttons on its screen. Even though in realistic hybrid visualization scenarios users would observe data on both the PC and in virtual AR space, we hid the visualization on the PC and on the tablet to have a fair comparison between 2D screens and the HoloLens. The interfaces are illustrated in Figure 4.5.

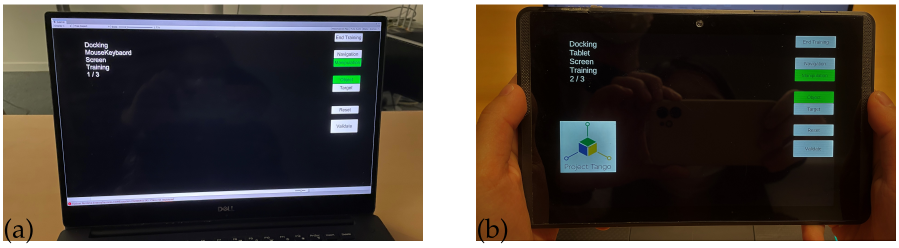

Figure 4.5: Photos to illustrate the interfaces on the PC and on the tablet. They are only used to control the experiment and show additional information. We do not visualize tasks on them.

4.4.3 *Design and measurements*

We asked our participants to first perform the docking task, followed by the clipping plane manipulation task. We did not switch the task order as the docking task served as an additional training for the more realistic and contextual clipping plane manipulation task. For each task, we used a within-subject design with 6 conditions (3 input × 2 output devices). Among the three input devices, the mouse provides 2D input, while the space mouse and the tangible tablet are 3D input devices. As for the output, the screen is a 2D output device and the AR headset (HoloLens) is a 3D output device. We counter-balanced the order of input devices, resulting in six different sequences. We then also counter-balanced the order of output devices, so that we have 12 different sequences overall. For each of the different conditions, we asked participants to perform six trials with different starting conditions. All participants get the same trials in the end (but with a different order, counter-balanced with Latin-Square) for a total of $6 \times 2 \times 6 = 72$ trials per participant per task.

For each task, we measured our participants' performance (task completion time and accuracy), without favoring one over the other. We asked participants to balance their interaction speed and accuracy and to decide when they are done with each trial. In Section 4.5, we detail our methods of computing the accuracy for each task. While precision also can be an index of performance [Albinsson and Zhai, 2003], we did not include it in our study as the accuracy and time are sufficient to access users' 3D manipulation performance [e. g., Chen et al., 1988]. In addition, we collected users' perceived workload with NASA-TLX, their preference, and any comments they had.

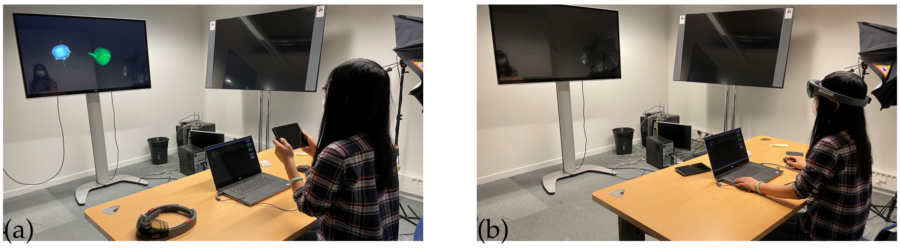

Figure 4.6: Photos to illustrate the experiment setup and conditions.

4.4.4 *Procedure*

We illustrate the experiment conditions in Figure 4.6. One of the authors was present during the experiment as the experimenter. The experimenter first introduced the general goal of the study and all the devices to use in this experiment. If the participant agreed to participate, we asked them to read and sign a consent form. We then asked the first to fill in a questionnaire that collected their basic demographic information and their experience with 3D visualization, 3D interaction, stereoscopic views, etc. We then started the experiment with the docking task. For each condition, we used three phases: demo, training, and trials. During the demo phase, the experimenter first demonstrated one trial to the participant, while explaining the interaction technique and providing ideas on how to solve the task. In case the participant began with the 2D screen, the experimenter and the participant looked at the big screen together. Otherwise, the participant got to wear the HoloLens and saw the trial with the HoloLens, while the experimenter looked at the view from a live streaming.[4] Next, we asked the participant to start their training. We allowed the participant as much time as needed to complete their exploration, and we did not record data during training session. Then, we asked the participants to complete six trials for each condition as described above. On finishing all trials of one condition, the experimenter helped them get on or take off the HoloLens to switch the output. After each task, we asked the participants to fill in a questionnaire that recorded their self-reported workload,[5] preference, and any possible comments before taking a break. Even though thinking-aloud can negatively affect speed measurements, we did not ask them to remain silent during the experiment as this equally affects all conditions. The experimenter took notes during the whole process.

---

[4] We used existing software (`http://www.microsoft.com/en-us/p/microsoft-hololens/9nblggh4qwnx?activetab=pivot:reviewstab`) for the streaming.

[5] We used Hart and Staveland's NASA Task Load Index (TLX) (`http://humansystems.arc.nasa.gov/groups/tlx/downloads/TLXScale.pdf`).

4.4.5 *Participants*

We recruited 24 unpaid participants for this experiment through e-mail announcements (9 female, 15 male; ages 22–32). All of them had at least a bachelor degree, 11 work in visualization or interaction-related domain. 11 had limited knowledge about the Microsoft HoloLens (10 only tried one or twice in their life). 8 had experience with tangible interaction due to mobile games and none of them knew the space mouse before the experiment.

4.5 RESULTS

We analyzed the data with estimation techniques using confidence interval (CI)s and effect sizes instead of p-values. Such methods are now recommended by several research communities, while the use of p-values for dichotomous significance tests have been criticized for their weakness (e. g., [Amrhein et al., 2017; Amrhein et al., 2018; Baker, 2016; Baguley, 2012; Cumming, 2014; Dixon, 2003; Dragicevic et al., 2014; Dragicevic, 2016; Gelman, 2017; Gigerenzer, 2018; McShane and Gal, 2017; Valentine et al., 2015]). We thus discuss the results with a nuanced interpretation by reporting the strength of evidence about the population [Besançon and Dragicevic, 2017; Besançon and Dragicevic, 2019; Cumming, 2014; Dragicevic, 2016; Gigerenzer, 2004; Goodman, 1999; Schmidt and Hunter, 1997] to avoid the dichotomous decision. However, it is still possible to relate the CIs we report to p-values [Dragicevic, 2015; Krzywinski and Altman, 2013].

As all our measurements are strictly positive, we aggregated the data using log-transformed measurements to compute the geometric means and report anti-logarithm forms to decrease the effect of outliers [Dragicevic, 2016; Keene, 1995], as it is common practice (e. g., [Jansen et al., 2013; Le Goc et al., 2016]). We use the t-test with paired variables to analyze the task completion time because participants performed all conditions in a different time range. Our analysis of accuracy is based on bootstrapping CI [Kirby and Gerlanc, 2013], and we report all CIs by default as 95% CIs. To increase transparency, we also visualize our data distributions.

For the plots of absolute mean values, the x-axis represents the absolute value in their unit, which we explain below for each measurement. For the plots of pair-wise ratios, the x-axis represents a ratio without units.

4.5.1 *Docking task*

4.5.1.1 *Completion time*

We report the absolute mean values of task completion time in seconds for each of the different conditions in Figure 4.7, the pair-wise ratios for each input between the two differences output (screen vs. HoloLens), and the pair-wise rations between different input devices for the same output in Figure 4.8.

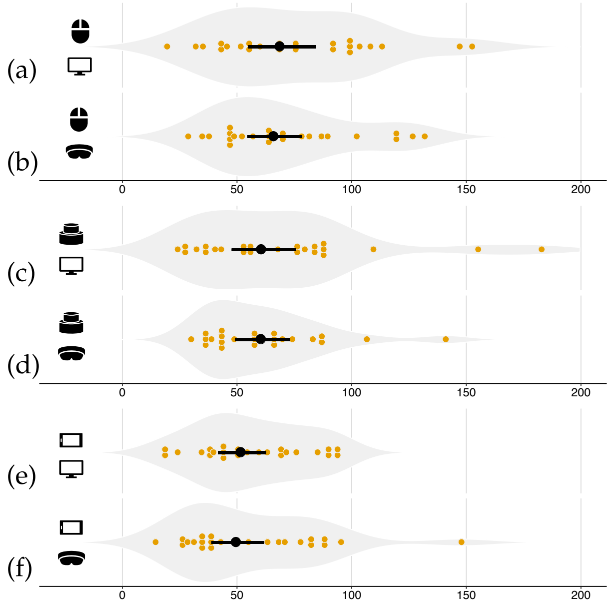

Figure 4.7: Results of docking task completion time (absolute mean value) in seconds. (a) and (b) represent the results using mouse with the screen and the HoloLens; (c) and (d) represent the results using space mouse with the screen and the HoloLens; and (e) and (f) represent the results using tablet with the screen and the HoloLens;

While users use the mouse as input, the average completion time value for a trial is 68.04s (CI [54.91s, 84.32s]) with the screen and 65.33s (CI [54.26s, 78.14s]) with the HoloLens. For the space mouse, the value is 59.95s (CI [47.73s, 75.30s]) with the screen and 59.85s (CI [49.19s, 72.82s]) with the HoloLens. And it is 51.05s (CI [41.67s, 62.54s]) with the screen and 48.97s (CI [38.89s, 61.67s]) with the HoloLens while interacting with the tablet. For any input device, the time difference between the output devices is quite small and the CIs largely overlap each other, so they do not give us enough evidence to conclude effects. We thus further look at the pair-wise ratio (screen/HoloLens). For each of the input devices (mouse, space mouse, and tablet), the ratio is 1.04 (CI [0.93, 1.16]), 1.00 (CI [0.86, 1.16]), and 1.04 (CI [0.91, 1.20]), respectively. As all the values are close to 1, we are unable to find evidence that would prove an effect of time for any input devices paired with 2D or 3D output. We can also estimate that the effect, should it still exist, is relatively small due to the short length of confidence interval.

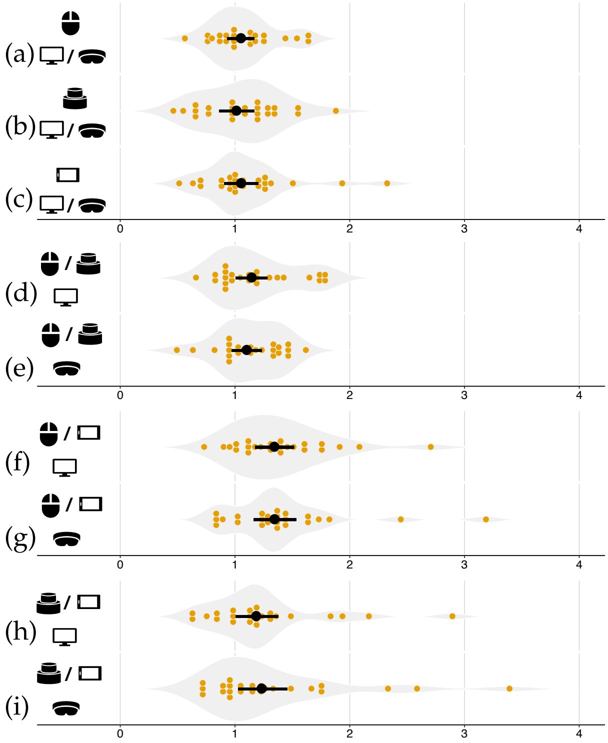

Figure 4.8: Results of docking task completion time effect sizes. (a), (b), and(c) represent pair-wise ratio (screen/HoloLens) the mouse, the space mouse, and the tangible tablet. (d) represents the ratio mouse/space mouse with the screen and (e) with the HoloLens.(f) represents the ratio mouse/tablet mouse with the screen and (g) with the HoloLens.(h) represents the ratio space mouse/tablet with the screen and (i) with the HoloLens.

We also checked the performance among different inputs, with the same output. While users were working with the 2D screen, the measured ratio between mouse and space mouse is 1.13 (CI [1.01, 1.27]). The CI does not overlap with 1 but remains close to it, which signifies that the task completion time is longer with the mouse than the space mouse, but the effect remains small. With the HoloLens, this ratio is 1.09 (CI [0.97, 1.23]), which leads to a similar observation as previously concluded. The ratio between the mouse and the tablet with the 2D screen is 1.33 (CI [1.18, 1.51]), which is also evidence that shows that users are slower using the mouse than the tablet, yet with a relatively small effect. Using the HoloLens, we conclude the same phenomenon as this ratio equals to 1.33 (CI [1.16, 1.53]). As for the ratio between the space mouse and the tablet, the evidence shows that the tablet is faster than the space mouse for both the screen (1.17, CI [1.01,1.37]) and the Hololens (1.22, CI [1.03, 1.45]). We can thus conclude that, for both the screen and the HoloLens, our evidence suggests that the mouse is slower than the space mouse, and both of them are slower than the tablet. However, we did not find evidence that would suggest a difference of task completion time for the different output, while users are working with a specific input device.

4.5.1.2 *Euclidean distance*

We report the absolute mean values of the Euclidean distance for each of the different conditions Figure 4.9, in virtual 3D space units. We also report the pair-wise ratios for each input between the two differences output (2D screen vs. HoloLens), and the pair-wise rations between different input devices for the same output in Figure 4.10. The Euclidean distance is the straight-line difference between the centers of the object (O_o) and the target (O_t) while participants validate one trial

$$d = \|O_o - O_t\| \tag{1}$$

The larger its absolute value is, the less accurate was the trial.

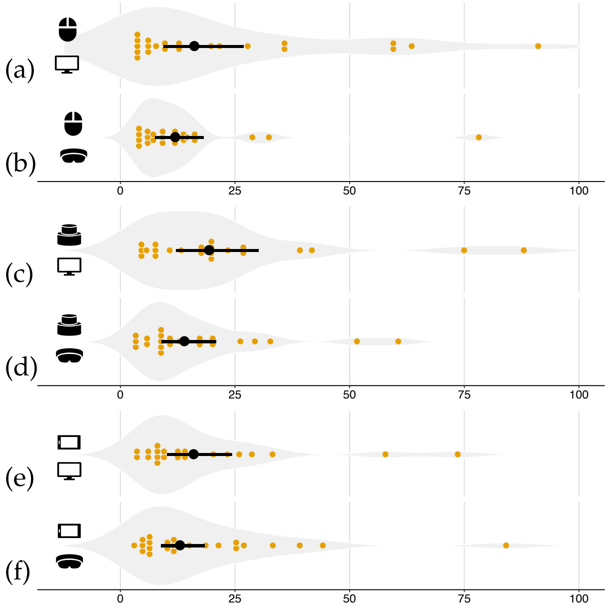

Figure 4.9: Results of docking task euclidean difference between the object and the target, in the unit of virtual space. (a) and (b) represent the results using mouse with the screen and the HoloLens; (c) and (d) represent the results using space mouse with the screen and the HoloLens; and (e) and (f) represent the results using tablet with the screen and the HoloLens;

For the mouse input, the absolute mean values of the distance while users are working with the screen is 15.86 (CI [10.06, 26.50]) and decreases to 11.70 (CI [8.55, 20.08]) for working with the HoloLens. The largely overlapping CIs show only very limited evidence for a difference. However, the pair-wise ratio (screen/HoloLens, 1.355 with CI [0.98, 1.93]) confirms that users are generally more precise while working with the HoloLens using the mouse. Yet, the long CI does not allow us to conclude on the size of this effect. For the space mouse input, the mean value is 19.08 (CI [12.93, 30.19]) with the screen and 13.694 (CI [9.68, 21.15]) with the HoloLens, and the effect size (screen/HoloLens) is 1.39 (CI [1.11, 1.91]) which leads to the similar observation that users perform more accurately with the HoloLens than with the screen. This observation even extends to the tablet as well, which has a mean value of 15.71 (CI [10.92, 24.71]) with the screen and 12.74 (CI [9.24, 19.07]) with the HoloLens. Its pairwise ratio is 1.23 (CI [0.95, 1.61]), which suggests that this difference is supported by only weak evidence. Thus, users are in general more accurate working with the HoloLens than working with the screen for any input devices.

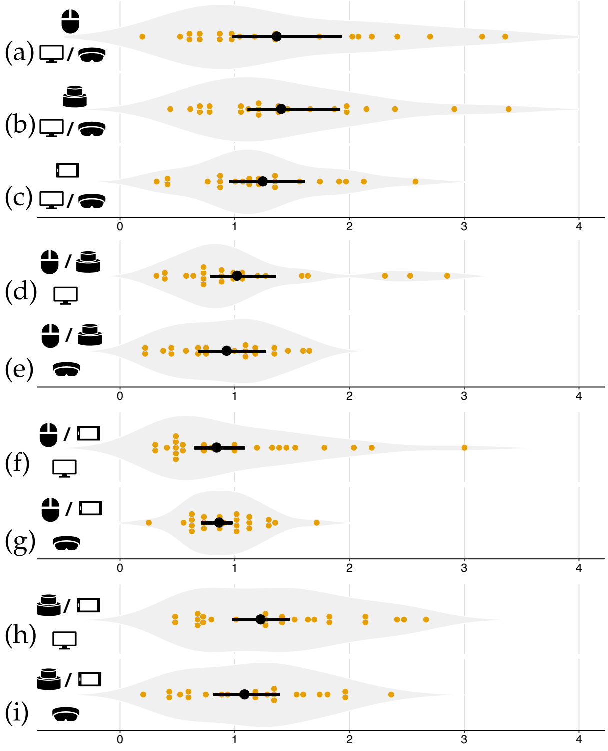

Figure 4.10: Results of docking task Euclidean difference effect sizes. (a), (b), and(c) represent pair-wise ratio (screen/HoloLens) the mouse, the space mouse, and the tangible tablet. (d) represents the ratio mouse/space mouse with the screen and (e) with the HoloLens.(f) represents the ratio mouse/tablet mouse with the screen and (g) with the HoloLens.(h) represents the ratio space mouse/tablet with the screen and (i) with the HoloLens.

We also look into the difference between input devices for the same output. While users are working with the screen, we found no evidence that users perform differently using the mouse or the space mouse, with the ratio being 1.01 (CI [0.79, 1.36]), or with the HoloLens (the ratio is 0.91 with CI [0.69, 1.27]). As for the comparison between the mouse and the tablet, our data shows that mouse is more accurate than the tablet with both types of output (the ratio is 0.83 with CI [0.65, 1.08] using the screen and 0.85 with CI [0.71 0.97] using the HoloLens), though the evidence remains weak. We also find evidence that suggests the space mouse is less accurate than the tablet working with the screen (ratio equals to 1.21 with CI [0.98, 1.48]), but this difference is not observable with the HoloLens which has a ratio of 1.08 (CI [0.81, 1.38]). We found no evidence suggesting a difference of performance between the mouse and the space mouse because the mean value of the pairwise ratio is close to 1 and its CI is large. However, we have weak evidence for the mouse being more accurate than the tablet, and for the tablet being more accurate than the space mouse. This result does not contradict with our observation that no evidence supports an effect between the mouse and the space mouse because a lack of evidence for a difference does not means there is no difference.

4.5.1.3 *Angular distance*

We report the absolute mean values of angular distance in degrees for each of the different conditions in Figure 4.11, the pair-wise ratios for each input between the two differences output (screen vs. HoloLens), and the pair-wise ratios between different input devices for the same output in Figure 4.12. This measure represents the final difference of rotation angle between the object and the target which rotation is respectively represented by the quaternion q_o and q_t. It is then computed by

$$a = 2 \cdot \arccos(q_{d\omega}) \tag{2}$$

where $q_{d\omega}$ is the ω component of

$$q_d = q_o^{-1} \cdot q_t. \tag{3}$$

The bigger its absolute value is, the less accurate was the trial.

First, we found that, for all experiment conditions, the final average angular distances have only little difference and their CIs overlap largely. The average value was 7.29° (CI [5.64°, 9.45°]) using mouse and screen, 7.14° (CI [5.68°, 8.96°]) using mouse and HoloLens, 7.77° (CI [6.80°, 9.51°]) with space mouse and screen, 7.47° (CI [5.99°, 9.30°]) with space mouse and HoloLens, 7.38° (CI [5.84°, 9.35°]) with tablet and screen, and 7.58° (CI [6.01°, 9.71°]) with tablet and HoloLens). We also observe that the pairwise ratios between different conditions remain close to 1. Specifically, the ratio between screen and HoloLens for the mouse was 1.02 (CI [0.84, 1.18]), for the space mouse it was 1.04 with CI [0.90, 1.28], for the tablet it was 0.97 (CI [0.83, 1.13]). While working with the screen, the ratio between the mouse and the

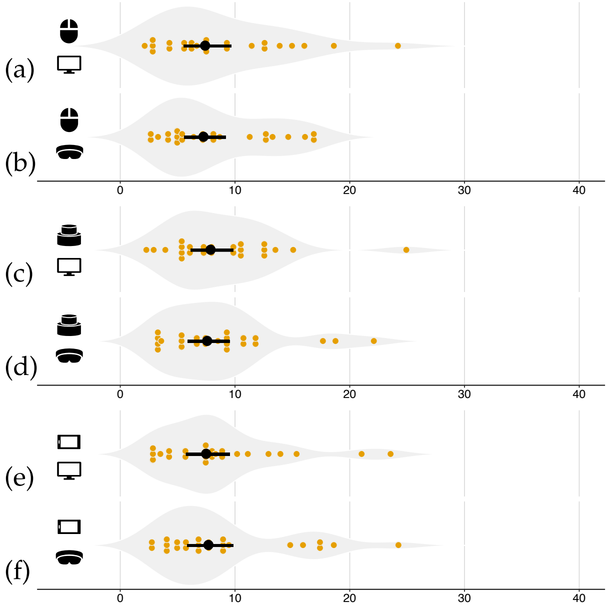

Figure 4.11: Results of absolute mean values of docking task angular difference between the object and the target, measured in degree. (a) and (b) represent the results using mouse with the screen and the HoloLens; (c) and (d) represent the results using space mouse with the screen and the HoloLens; and (e) and (f) represent the results using tablet with the screen and the HoloLens;

space mouse was 0.93 (CI [0.76, 1.15]), between the mouse and tablet it was 0.98 (CI [0.81, 1.16]), and between the space mouse and the tablet it was 1.05 (CI [0.88, 1.37]). As for the HoloLens, these ratios are respectively 0.96 (CI [0.86 1.08]), 0.94 (CI [0.79 1.13]), and 0.96 (CI [0.85 1.17]). In conclusion, we could not find evidence that would suggest that a difference exists for angular accuracy across techniques between the object and target.

4.5.1.4 *Self-reported workload*

We report the participants' ranked workload in Figure 4.13. Our evidence shows that mental and physical demands vary for different input devices, while there is no evidence showing that the output device plays a role. Specifically, users' mental demand is the highest for the mouse as suggested by participants. They said that the mouse interaction mapping needs to be remembered and recalled during the experiment and that it involves using both the mouse and the keyboard identifier. The latter requires users to develop a reflex to use the correct combination right away. The tablet requires the least mental demand because participants feel that the mapping is just natural. For the physical demand, not surprisingly, the tablet requires the most as users need to move it in space, while the space mouse requires the least since it stays on a fixed position on the desk. We

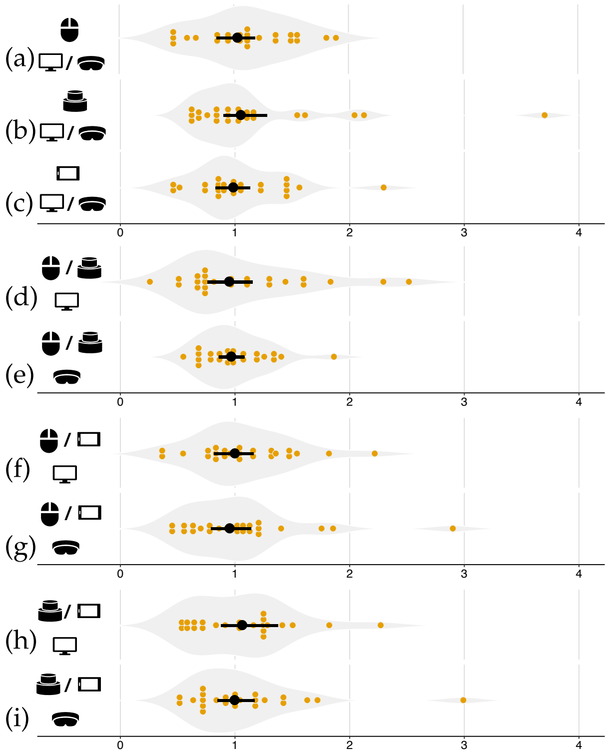

Figure 4.12: Results of absolute mean values of docking task angular difference effect sizes. (a), (b), and(c) represent pair-wise ratio (screen/HoloLens) of the mouse, the space mouse, and the tangible tablet. (d) represents the ratio mouse/space mouse with the screen and (e) with the HoloLens.(f) represents the ratio mouse/tablet mouse with the screen and (g) with the HoloLens. (h) represents the ratio space mouse/tablet with the screen and (i) with the HoloLens.

did not find evidence showing any difference for other workloads, which was confirmed by the verbal comments we received from many participants.

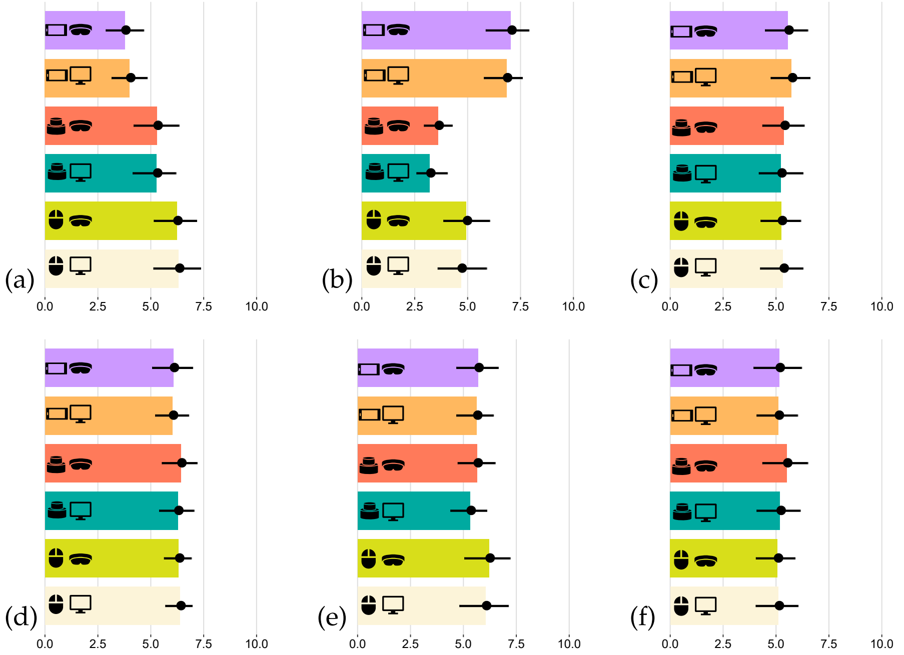

Figure 4.13: Self-reported workloads for the docking task. (a) represents the mental demand; (b) represents the physical demand; (c) represents the temporal demand; (d) represents the users satisfaction of their performance; (e) represents the effort needed to accomplish the task; (f) represents the frustration level.

4.5.1.5 *Preference*

Participants also ranked their preference for input device-output setting combinations from 1 to 6, with 1 meaning least liked and 6 meaning most liked. For our analysis we added all values for a given condition; higher values thus mean higher overall participant preference. We summarize the resulting preference ratings for the docking task in Table 4.1, which shows an almost evenly distributed preference for this task, regardless of input or output device, with a slight preference towards the space mouse and the screen combination.

Table 4.1: Users' serf-rated preference for different input and output combinations, for both the docking task and the clipping plane task.

Input Output	Mouse Screen	Mouse HoloLens	Space Mouse Screen	Space Mouse HoloLens	Tablet Screen	Tablet HoloLens
Docking	83	77	91	86	82	85
Clipping	103	96	90	85	72	60

4.5.2 *Clipping plane task*

4.5.2.1 *Task completion time*

We present the results of mean task completion time in seconds in Figure 4.14, and the pairwise ratios across different conditions in Figure 4.15.

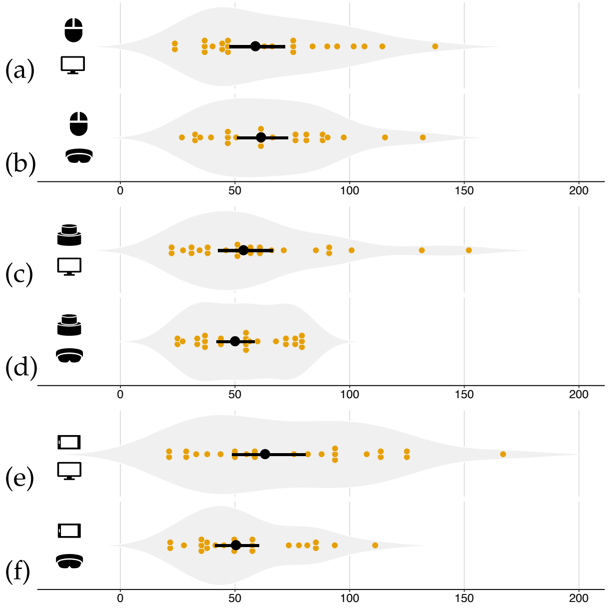

Figure 4.14: Results of clipping task completion time in seconds. (a) and (b) represent the results using mouse with the screen and the HoloLens; (c) and (d) represent the results using space mouse with the screen and the HoloLens; and (e) and (f) represent the results using tablet with the screen and the HoloLens;

We first observe that, for the same input devices, the completion time of different output devices does not vary much. For the mouse, the average time in seconds using the screen is 58.38s (CI [47.58s, 71.63s]) and is 60.82s (CI [50.80s, 72.81s]) using the HoloLens. The pairwise ratio screen/HoloLens is 0.96 (CI [0.86, 1.07]), which being close to 1 does not give us evidence of an effect. For the space mouse, the average time using the screen is 53.18s (CI [42.59s, 66.42s]) and is 49.50s (CI [42.02s, 58.32s]) using the HoloLens. The pairwise ratio screen/HoloLens is 1.07 (CI [0.94, 1.24]), similarly close to 1. For the tablet, the average time in seconds using the screen is 62.60s (CI [48.64s, 80.58s]) and is 49.88s (CI [41.29s, 60.25s]) using the HoloLens. The pairwise ratio screen/HoloLens of the tablet is 1.26 (CI [1.07, 1.47]), thus suggesting an evidence that users are generally faster working with the HoloLens than with the screen while they are using the tablet.

We also looked at the difference of input devices for the same output. While users are working with the screen, the ratio between the mouse and the space mouse is 1.09 (CI [0.93, 1.28]), and the ratio between

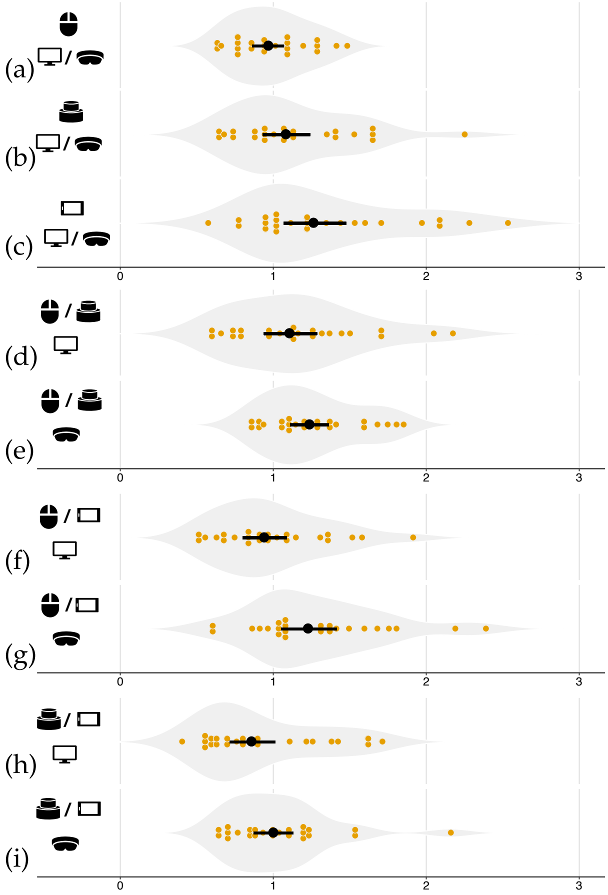

Figure 4.15: Results of clipping task completion time effect sizes. (a), (b), and(c) represent pair-wise ratio (screen/HoloLens) the mouse, the space mouse, and the tangible tablet. (d) represents the ratio mouse/space mouse with the screen and (e) with the HoloLens.(f) represents the ratio mouse/tablet mouse with the screen and (g) with the HoloLens.(h) represents the ratio space mouse/tablet with the screen and (i) with the HoloLens.

the mouse and the tablet is 0.93 (CI [0.80, 1.09]). In both cases, we cannot find enough evidence to claim an effect, but the ratio between the space mouse and the tablet is 0.85 (CI [0.71, 1.01]), which suggests that the space mouse is faster than the tablet, while users are working with the screen. As for the HoloLens, the results are different from those of the screen. We can find evidence suggesting that the use of the mouse resulted in slower interaction times than the space mouse because the ratio between these two conditions is 1.22 (CI [1.11, 1.36]). We can also find evidence showing that interactions with the mouse are slower than with the tablet, given the fact that the ratio is 1.21 (CI [1.05, 1.41]). But we found no evidence for a difference between the space mouse and the tablet with the value of the ratio being 0.99 (CI [0.87, 1.13]), when they are paired with the HoloLens.

4.5.2.2 *Accuracy*

We computed the accuracy as follows. Prior to the study, we first generated a pool of target planes inside the volume with different positions and orientations. A plane-volume intersection can give an intersection plane of different forms (a triangle, a rectangle, a pentagon, etc.). We only kept target planes in a rectangular form (4 corners) for our experiments to lower the participants' mental workflow—they will not be confused about the shape of the target plane. The clipping plane is defined by an artificial (because it is theoretically infinite) center O_{plane} and a normal vector N_{plane}. Each target plane has exactly four corner positions $O_{point_i}, i \in [1, 4]$ (i. e., intersection points of the similarly infinite target plane with the data volume). We then compute the signed distance of each target point O_{point_i} to the manipulated clipping plane as

$$d_i = N_{plane} \cdot (O_{point_i} - O_{plane}). \tag{4}$$

We then determine the absolute values of these distances because the signed information of which side of the plane situates the point is irrelevant. We then average the absolute values to arrive at a final accuracy value of

$$A = \frac{1}{4} \sum_{n=1}^{4} \|d_i\|. \tag{5}$$

We show the results of accuracy in virtual units in Figure 4.16 and 4.17.

For the mouse input, this mean accuracy is 4.26 (CI [3.8, 4.87]) with the screen and 5.39 (CI [4.77, 5.91]) with the HoloLens. The non-overlapping of CIs already suggests that users are more accurate working with the screen. Furthermore, their pairwise ratio is 0.79 (CI [0.70, 0,90]). This CI does not overlap with the value 1, which confirms the effect. This observation partly also applies to the tablet, but the evidence is weaker. With the tablet, the mean value is 5.03 (CI [4.41, 5.86]) with the screen and 5.69 (CI [5.06, 6.27]) with the HoloLens, and the pairwise ratio is 0.88 (CI [0.76, 1.02]). For the space mouse, we did not find evidence that would suggest a difference between the two different output devices. The mean value is 5.06 (CI [4.44, 5.78]) with the screen and is 5.35 (CI [4.82, 5.72]) with the HoloLens, and the pairwise ratio is 0.95 (CI [0.83 1.07]). For users working with the screen, our evidence suggests that the mouse is more accurate than both the space mouse (pairwise ratio being 0.84 with CI [0.72, 0.99]) and the tablet (pairwise ratio being 0.85, CI [0.75, 0.94]), while we found no evidence that would suggest a difference between the two latter devices (pairwise ratio is 1.01 with CI [0.88, 1.19]). For the HoloLens, we only have weak evidence for the mouse being more accurate than the space mouse (ratio 0.95, CI [0.85, 1.04]). We found no evidence to suggest a difference between mouse and tablet (ratio

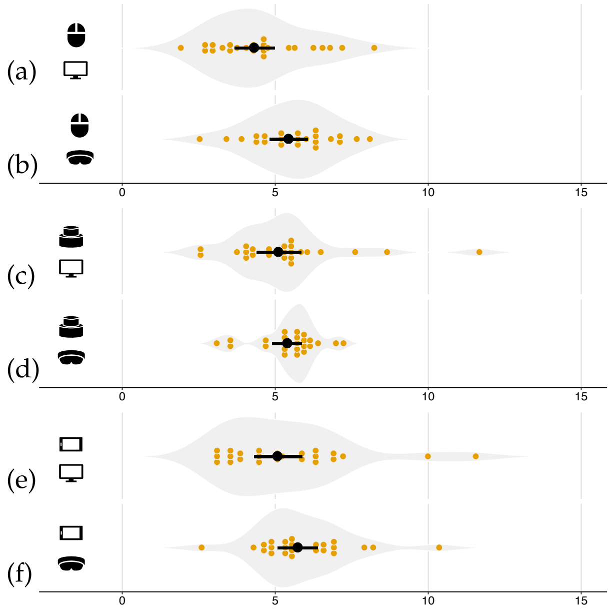

Figure 4.16: Results of clipping task accuracy, measured in virtual units. (a) and (b) represent the results using mouse with the screen and the HoloLens; (c) and (d) represent the results using space mouse with the screen and the HoloLens; and (e) and (f) represent the results using tablet with the screen and the HoloLens;

1.01, CI [0.87, 1.12]). There is also almost no evidence for a possible effect between space mouse and tablet (ratio 0.94, CI [0.84, 1.09]).

4.5.2.3 *Self-reported workload*

For the self-reported workload of clipping plane task (Figure 4.18), the situation differs from the docking task. Although the difference is low, we observed an increasing mental demand, temporal demand, and effort woking with the tablet. For the physical demand, compared with the docking task the difference is that the mouse has a lower average demand than the space mouse. Participants reported that manipulating a clipping plane requires small and precise changes of the plane, where the DOFs separation is advantageous.Other workload factors (temporal demand, effort, and frustration) also revealed that most participants felt that the mouse was easier to work with compared to the tablet, regardless of which output was used.

4.5.2.4 *Preference*

We summarize the preference ratings of our participants for the clipping plane task in Table 4.1, computed the same way as for the docking task. For this task, the participants' preference seems pretty obvious: for the same output device, the mouse was slightly preferred over the space mouse, which in turn was generally preferred over the tablet. Interestingly, in all cases of input, the difference between using the

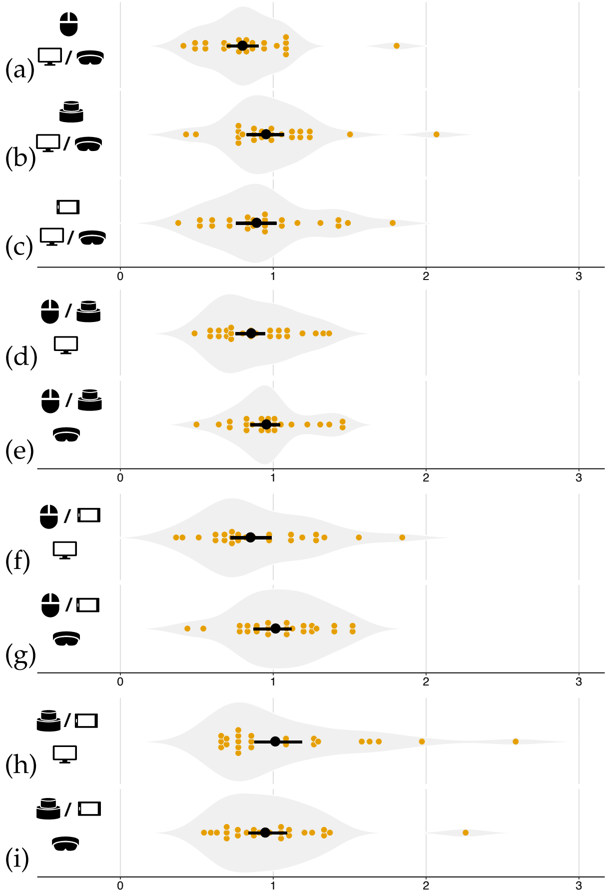

Figure 4.17: Results of clipping task accuracy effect sizes. (a), (b), and(c) represent pair-wise ratio (screen/HoloLens) the mouse, the space mouse, and the tangible tablet. (d) represents the ratio mouse/space mouse with the screen and (e) with the HoloLens.(f) represents the ratio mouse/tablet mouse with the screen and (g) with the HoloLens.(h) represents the ratio space mouse/tablet with the screen and (i) with the HoloLens.

two output devices is low. Most participants also reported that they did not find much difference regarding the output, but that for them the input matters more for the task. Nonetheless, judging from participants' actual recorded ratings, in most cases they preferred the screen over the HoloLens.

4.6 DISCUSSION

First, we look at the difference for input devices paired with the same output. For the 3D docking task that we used to represent general 3D data exploration using object or view manipulation, our participants preferred the 6-DOFs space mouse. Nonetheless, our measurements show that the tablet was the fastest input device. Also, our participants managed to accomplish the task as accurately using

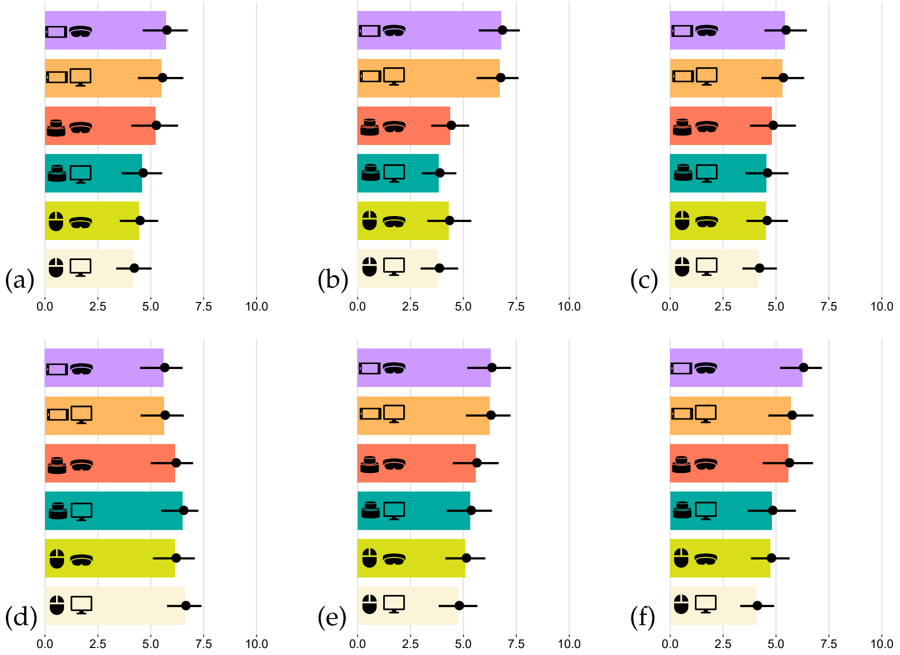

Figure 4.18: Self-reported workloads for the clipping plane task. (a) represents the mental demand; (b) represents the physical demand; (c) represents the temporal demand; (d) represents the users satisfaction of their performance; (e) represents the effort needed to accomplish the task; (f) represents the frustration level.

the mouse as with the space mouse, both being more accurate than the tablet. In the second task, our participants reported a preference toward the mouse (for the same output device), which correlates to their performance—the mouse is generally more accurate than the other two input devices, and faster with the 2D screen. For 3D visualization tasks in which accuracy is a key factor we can thus see our hypothesis ***h1*** to be supported. Even though the mouse was slower than the other two input devices in the clipping task with the HoloLens, the trade-off between time and accuracy seems to be acceptable for most visualization work.

Surprisingly, our participants' preferences differed between the two tasks, even though we used the exact same interaction mapping for both. We see this as an indication that, depending on the specific task and the given application domain, different preferences and also a different performances are to be expected, as previously also concluded by Bach et al. [2018] who used a setup and tasks different from ours to compare visualization and interaction effectiveness. For example, our data shows a clear performance increase over 2D output when performing the clipping task with 3D output. Here, the better spatial understanding of the plane due to the stereoscopic projection makes a clear difference. One reason reported by participants is that the intuitive mapping of the tangible tablet only helps when one can directly understand how the manipulated object or plane is positioned and oriented in space. In other words, since the flat plane is difficult to understand when seen on a 2D screen in the clipping plane task, the mouse is preferred because of its separation of DOF—our participants reported that it allowed them to adjust the data more precisely. The inherent DOF separation of 2D (i. e., 2 DOFs) input devices could thus better match the 2D output space in such situations.

Another interesting and essential point here is that we found, for both tasks with the AR output space, the mouse input to have either equal or better performance when compared to its performance in 2D space. The mismatch between its 2D input to the used 3D output space thus seems not to be a problem for participants when using the mouse. Based on this observation we may answer our initial question: for the types of tasks we are investigating (3D manipulation in 6 DOF), the mouse remains a very effective interaction tool, in particular also if used in combination with AR output. In a hybrid visualization system that relies on mouse input for some non-spatial tasks anyway, it is thus not a bad idea at all to unify the interaction with a mouse to be consistent across devices.

It is also interesting that many participants reported a certain level of indifference about using the 2D screen or the HoloLens, which can also be noticed in the final preference rating that showed that input devices mattered more to them than the output spaces, as we only found a preference toward the 2D screen when compared to the HoloLens. This can be partly explained due to the known hardware limitations of the first generation of the HoloLens, as reported by several participants. One participant explicitly mentioned that the limited field of view resulted in the volume not only being clipped by the clipping plane but also by the HoloLens' field of view, which introduced confusion. Also, other participants mentioned that the headset was still too heavy and it could make them feel worse than working with a screen. Nonetheless, these limitations are not the important issues to discuss here, and we expect them to be resolved in future AR HMDs.[6] However, such self-reported performance is not fully supported by the measurement. For the docking task, our measured values show that they were more precise in terms of Euclidean distance for the docking task using the HoloLens, when compared to the 2D screen, and with similar levels of completion time. For the clipping plane task, we also saw faster task completion times using the tablet with the HoloLens than with the 2D screen, while the mouse and the space mouse showed no evidence for a difference, in both display conditions. Although users seem to be more accurate with the screen for the docking task, the absolute mean difference remains small. Taking the mouse as an example which revealed the biggest effect, the mean value using the screen is 4.26,

---

[6] Already the recent second version of Microsoft's HoloLens improves greatly with respect to issues such as resolution, field of view, and balance.

while it is 5.39 using the HoloLens. Even so, we can consider them as an equal accuracy level compared to the initial stage—a distance of 200 on average. So we conclude that the advantages of stereoscopy for understanding 3D data exist in some cases, even if this may not be evident to the users' mental feeling. Thus, ***h2*** is partly supported. Our results thus do not support the findings drawn by Schultheis et al. [2012], who only used time as a measure. Finding by our experiment indicates that the different output spaces do affect users' performance for precise spatial control.

The basis of our final hypothesis ***h3***—as explained in Section 4.3—was a traditional 3D data analysis setup using mouse and 2D screen creates a mismatch between the dimensions used by the input device, the output space, and the user's mental model. The latter resembles the inherently three-dimensional data space $\mathcal{D}$, while the output space is a vertically displayed 2D projection and the input space is a horizontally oriented 2D space. The inherent dimensionality mismatch of this setup could be resolved with stereoscopic output and 6 DOFs 3D input devices, taking advantage of the many benefits of immersive analytics [Dwyer et al., 2018]. The use of VR technology with its associated dedicated input devices, however, does not always seem to be practical because scientists often rely on analysis tools such as scripts for which a traditional workstation appears much better suited. Instead, our envisioned AR extension to existing workstations uses the best of both worlds: traditional (2D) displays with traditional (2D mouse + keyboard) input, together with an AR-based (3D) display and appropriate input devices.

Focusing only on 3D data and thus a three-dimensional visual representation $\mathcal{U}$, the question we asked in this work is which input devices (operating in a respective $\mathcal{M}$) would be ideal for which output space dimensionality in $\mathcal{D}$. As there is likely no ideal input device, we ask more precisely what the compromises would be, given the choice of a set of input devices and their dimensionality and mapping, with respect to both 2D and 3D output spaces $\mathcal{D}$. We emphasize here that input device manipulations as such are not sufficient to define their effects in a visualization system. Instead, the associated input mapping (i. e., the transformation from space $\mathcal{M}$ to space $\mathcal{I}$ and the resulting manipulations in spaces $\mathcal{D}$ and thus changed $\mathcal{V}$), of course, also play a role. Nonetheless, we simplify the consideration here to the right-most triangle in Figure 4.2, with the input mapping aspects of $\mathcal{I}$ merged into $\mathcal{M}$ and the aspects of $\mathcal{D}$ and $\mathcal{V}$ merged into $\mathcal{O}$. One could hypothesize that a 6 DOFs input device (i. e., one that provides both 3D location and 3D orientation) should be able to better accommodate tools for the exploration of 3D data—it should be able to facilitate, for example, view manipulations, clipping plane adjustments, and other 3D data operations. Nonetheless, as we want an input device that works in a hybrid setup, it should work both for the AR environment as well as for the traditional workstation and even potentially projected 3D views. Moreover, we need to compare these devices with the common mouse as it is established as the primary spatial input device for traditional PCs. While the latter only provides 2 DOFs input, interaction mappings are designed to control 3D spaces with the needed 6 DOFs or more.

Based on our results, unfortunately, we cannot provide a single conclusion for the effects of match or mismatch from our results to support or to reject ***h3***. For the 3D docking task, all input devices—regardless of whether they have 2 DOFs or 6 DOFs—had better performance in 3D output space than in 2D. The increase of performance is only due to the better accuracy of Euclidean distance, while the angular distances did not change fundamentally. But the effect seems to be equal for all input devices, we did not find any evidence suggesting that such increase is greater with 3D input. Also, we did not find evidence that would show if the effect of performance increase is bigger when the dimensionality is matched. The reason for finding higher precision only for Euclidean and not for angular distances is likely that the depth clues offered by stereoscopy, which helped participants to understand the spatial position, had only a limited effect on the understanding of orientation—especially for the simple model of the Utah teapot. With the clipping plane task, we do not find a clear result with respect to ***h3*** either. Here, the performance increase of mouse is generally greater than for the other input devices when we pass from the 2D screen output to the HoloLens. From our results, we thus conclude that the match or mismatch between input and output devices are not crucial for 3D manipulations when we follow the hybrid PC-AR metaphor. However, the properties of input devices matter a lot. Intuitive and fluent high-DOF input like from a tablet is enough for interaction that need fast adjustment and does not require much precise control as our study shows that users finished the docking task with the least amount of time. While it comes to tasks that require more accurate input, low-DOF input devices (like the regular mouse) that could easily separate interaction DOF are generally preferred.

Our failure of finding a simple conclusion of the effect of dimensionality (mis-)match does not decrease the importance of our study, however. Quite in contrast, we now know that, for 3D manipulation, the effect of dimensionality match (if any) is expected to be small that can be neglected. Thus, this consideration can be excluded while choosing input and output devices for 3D visualization. Users are more flexible to choose appropriate devices according to specific tasks and other needs rather than considering if its DOF matches the output. Even more importantly, we see from the results that the mouse performs still well, in particular for accuracy control, and regardless of the used output. We can thus confirm that using the mouse as a primary input for 3D output is a valid choice. This conclusion is important for designing a hybrid visualization systems where the mouse is naturally used for the PC part.

To assist future system designers in choosing appropriate combinations of input and output devices we thus summarize our participants' feedback regarding the main advantages and disadvantages of different input devices as follows. According to them, the mouse is precise but its inherent input mapping for 3D manipulations increases the need for learning and the overall mental demand to recall this mapping, especially as keyboard modifiers are required to control all DOFs. For the space mouse, they liked its fluidity and stated that it requires the least amount of effort, thus causes minimal fatigue levels. Nonetheless, the interaction mapping can be sometimes confusing due to the potentially different frames of reference between the manipulation space $\mathcal{M}$ where the mouse is physically located, and the controlled object or space. Moreover, certain motions are difficult to perform with the space mouse. For example, one participant reported that, while interacting with the space mouse with the right hand, rotating around y-axis is much more difficult than rotating around the x-axis (refer to Figure 4.4) because the latter needs to bend the wrist in an uncomfortable way. Moreover, users see the space mouse as a flexible input device. This propriety, though sometimes it can be considered to be an advantage, increases the error rate and makes its difficult to perform tiny adjustments. However, space mouse is reported to require the minimal interaction efforts among the three devices we chose, users thus believe that they may use it more effectively with more practice. This is partly the reason why users slightly prefer the space mouse for the docking task, even though its measured results are not as good. For the tablet, finally, our participants found its interaction design to be natural in 3D space. Nonetheless, it causes a lot more fatigue than the other devices because users need to hold it and move their arms a lot. Also, users move the tablet in space without a physical reference surface like the table for the mouse and the space mouse, resulting in less accurate manipulation (e. g., making its use more difficult for precise control).

The final aspect to discuss is that of mobility. In our experimental setup we specifically asked participants to sit down to replicate a situation similar to working in an office. Nonetheless, we still observed that participants moved their head to get a better view in the clipping plane task, but did not see similar actions in the docking task. Ultimately, however, it would be good if users can take full advantage of the AR HMD, and so the fact that the tangible tablet exhibited a competitive performance overall suggests that it can be used as an additional input device when mobility is needed or desired. Of course, its ability to also serve as a platform for richer (touch) input only strengthens this point.

4.7 CONCLUSION, LIMITATIONS AND FUTURE WORK

Following our vision of extending traditional workstation with AR for visualization purposes, we discussed the effect of pairing different input and output devices. We reported upon a controlled experiment in which we compared three different, well-established input devices and discussed their suitability to control 3D projected (to 2D) or stereoscopic 3D output spaces. Our results seem to show that performance may vary depending on tasks, but the mouse remains an efficient 3D interaction tool for high accuracy tasks, which are often needed in 3D visualization applications. Our results thus serve as a starting point to guide the interaction design on an AR-extended hybrid visualization system in which mouse and keyboard are still necessary for using existing analysis software. Yet, our studies cannot give a universal conclusion about the effect of the dimensionality match between input and output devices to user performance while working on visualization tasks. This is an essential result of our work: we have established that the match or mismatch of the dimensionality between input and output devices seem to have little importance.

As the mouse remains important for interacting with traditional workstation and it yields good results when a high accuracy is required, we argued that, on the one hand, the mouse can be a primary means of input—even for a hybrid system. High-DOF input devices, on the other hand, could be served as complementary tools for specific cases such as while users walk around and just want to rotate the view roughly, then a tangible tablet would be a good choice.

One of the limitations of our work is that our tasks focused only on the part of 3D manipulation for single objects, without considering scenarios where other motions are required, for example, clicking, typing, selecting, and switching between different tasks. It would thus be interesting to investigate a more realistic scenario to understand how to best combine different input devices, or unifying the interaction with only one device with regards to the needs of 3D visualization. Similarly, we started with simple scenes with few objects. In other scenarios, experts may need to analyze very large, complex visualizations. We also want to understand, in the future, if results change with such data.

Another limitation is that, even though we manually adjusted a lot parameters on both the side to have an equivalent visual and interactive experience, some of our choices could still have affected our measurements and potentially our conclusions. For example, the HoloLens cannot achieve the same level of display quality. Meanwhile, its advantages of moving the head around allows faster understanding that could increase participants' performance. Moreover, although we gave our participants unlimited trials to practice before starting the experiment, many expressed a feeling that their performance would change if they get more used to a device, especially for the space mouse. We thus think another important limitation is that we cannot achieve long-term training for our experiment, due to both the lack of hardware and constraints of experiment length.

Regarding input devices, we chose commercially available devices for our first-stage study. But we do not argue against trying specifically designed ones for dedicated scenarios. For example, the CHARM device by Klamka et al. [2019] may be a good alternative, especially if users want to walk around and lose the access to the mouse. Also, Sundin and Fjeld [2009] investigated a softly elastic spaceCat, and they showed that it favors positioning and docking tasks when compared with elastic rate control input device (like the space mouse we used). In a future vision and besides the mouse as primary input, we would thus want to know how such different input devices can jointly be used in a hybrid visualization setup, or would need to be redesigned and merged, to benefit from both types of input. Apart from that, with a given input device, we implemented the most adapted mapping from previous literature (or the most natural one according to our pilot tests), which does not lead to much liberty for the users. In a future investigation, we would like to allow users more freedom, such as manually adjusting the CD-ratios. As such, it might also be possible to achieve small and accurate adjustment with high-DOF input devices, but we also wonder if the trade-off between time and accuracy is critical and what would users prefer with regards to the requirements of visualization task.

To study the mismatch of dimensionality, we equalized the visuals on screen and on the HoloLens to both be 3D models. It would thus be useful to continue to investigate other types of data in the future, including but not limited to abstract information and graphs because they are also essential forms information—even for 3D data analysis. In addition, to further investigate the design of a hybrid setting, it is also an interesting future work to make use of the benefits of both spaces to visualize different types of data.

Finally, even though our study code and implemented interaction techniques be used directly with a VR headset, we did not actually test it in this study because we wanted to study the question of selecting input devices focusing on a specific setup (AR+PC). Nonetheless, in addition to our hybrid AR and PC setup, pure PC and pure AR visualization environments can benefit from our findings, but we leave the question of how other setups (like VR) can benefit from our findings as an interesting question to be investigated in the future. We envision that due to the inherent differences between VR and AR headsets, we still need to put more efforts. For example, users are occluded form the real world, it may be needed to create visual representations of each input device, and some features like walking-around would be impossible. We also chose basic and representative tasks for general

3D visualization that can be used to various applications, not limited to the datasets used in this our experiment.

This chapter relates to the research question **R3**. While users remain seated, the mouse is still an efficient input, even working AR displays, for 3D visualization tasks, especially for those that require a high accuracy. We thus envision that the mouse should still be kept as the primary input for the hybrid systems, while others might be combined for specific tasks according to different requirements. We need to highlight that, in general HCI, the chosen input devices usually depend on specific tasks and requirements. Thus, even though the tangible tablet was not most efficient nor preferred in our tasks, we do not deny its value in other cases, as its advantages have also been largely explored previously in other scenarios (e. g., [Spindler et al., 2012; Spindler et al., 2014; Sollich et al., 2016]). For example, if a task needs rapid adjustment of the views without being super accurate, we would also encourage the use of a tangible tablet. And the tablet also has other advantages compared to the mouse such as that it does not require on table surface to work, which could make it a good alternative input while users are walking around.

5

AUGMENTING TACTILE 3D DATA NAVIGATION WITH PRESSURE SENSING

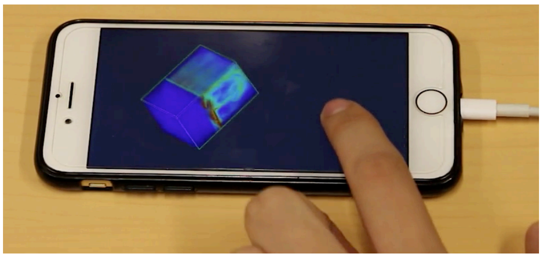

Figure 5.1: We proposed a pressure-augmented tactile interaction technique on mobile phones to improve 3D data navigation.

Chapter 4 gives us basis of how to support interaction with the hybrid setting while users remain seated. We then want to support the interaction while users are walking around as well. As mobile devices are popular means for controlling large and virtual devices with a lot of recent research work has been done to explore the data visualization on them, we thus believe that it is a good and practical tool to use while users are walking around. We then study a touch-based interaction mapping on mobile devices to support 3D navigation task. Specially, this chapter presents a pressure-augmented tactile 3D data exploration technique, specifically designed for small devices, motivated by the need to support the interactive visualization beyond traditional workstations. While tactile interaction has been studied extensively on large screens, current techniques do not scale to small and portable devices. We use readily-available pressure sensing with a binary mapping to allow users to easily select different manipulation. We compare our technique with traditional 3D-RST (rotation, scaling, translation) using a docking task in a controlled study. The results show that our technique increases the precision of interaction, with only little impact on speed. We then discuss the implications for 3D interaction design and verify that our results extend to older devices with pseudo pressure and in realistic usage scenarios of smartphones.

Work presented in this chapter gives a method to increase interaction accuracy based on touch interaction on mobile devices, which we believe is important to accomplish visualization tasks while users walk around with the hybrid setting.

Main portions of this chapter are published on Computer Graphics Forum and presented in EuroVis 2019. The term of "we" in this chapter refers to myself, Lonni Besançon, Mehdi Ammi, and Tobias Isenberg.

5.1 INTRODUCTION

Visualization of 3D content are needed for exploring many types of scientific data, and good interaction design is a fundamental and essential prerequisite for effective visualization tools (e. g., [Keefe, 2010; Yi et al., 2007]). As such, interaction research targeted specifically at data exploration plays an ever increasing role in the field of visualization (e. g., [Sutherland, 1966; Hibbard, 1999; Rheingans, 2002; Johnson, 2004; Tory and Möller, 2004; Keefe, 2010; Keefe and Isenberg, 2013; Munzner, 2014; Isenberg, 2016]). This need for effective and efficient interactive tools is evident, for instance, in the use of navigation for exploratory [Tukey, 1977] 3D spatial data exploration [LaViola Jr. et al., 2017]: by interactively changing the view, scientists are able to immerse themselves in the data [Büschel et al., 2018] to understand its characteristics. Such tasks have long been performed on desktops using mice and keyboards—but here the intuitive and fluent control of the exploration in 3D space is often challenging [Forlines et al., 2007; Besançon et al., 2017c]. Novel environments promise to better support 3D data exploration: they can improve the visual perception and spatial understanding using, e. g., wall-size screens [Bezerianos and Isenberg, 2012], occluded VR glasses [Qi et al., 2006], and CAVEs [Cruz-Neira et al., 1993b]—often combined with dedicated input devices. Researchers also explore novel input metaphors for easy and precise control, e. g., touch-based [Benko et al., 2006] and 3D spatial input [Qi and Martens, 2005].

According to studies of different input metaphors for easy and precise control, tactile input [1] yields important advantages for 3D data exploration (e. g., [Besançon et al., 2017c; Coffey et al., 2012; Fu et al., 2010; Isenberg, 2016; Lundström et al., 2011; Yu et al., 2010]) such as faster completion time or increased 'directness' [Bruckner et al., 2019]. Moreover, the increasing use of smart phones makes this interaction easily accessible due to their good mobility and portability. Yet, mobile devices are restricted to a relatively small screen, possibly making the visualization of and the interaction with complex data impractical. Meanwhile, the emergence of visualization tools for mobile devices (e. g., Arctic Viewer[2], ImageVis3D[3], and KiwiViewer[4]) suggests an increasing demand for mobile visualizations. We thus need to address the challenges of providing interactive data exploration control on these platforms.

Tactile interaction is the primary form of input on smart phones. Performed on a 2D surface, each touch point offers up to 2 DOFs

---

[1] We refer to touch input as 'tactile' **input** as elsewhere (e. g., [Besançon et al., 2017c; Besançon et al., 2017b; Herot and Weinzapfel, 1978; Poupyrev and Maruyama, 2003; Raynal et al., 2010]), we do not mean haptic (tactile) **feedback**.

[2] https://kitware.github.io/arctic-viewer/

[3] https://www.sci.utah.edu/software/imagevis3D.html/

[4] http://www.kiwiviewer.org/

through its translation, but navigation in 3D visualization requires six or more DOFs to specify position and orientation [Isenberg, 2016].[5] To provide the necessary input DOFs, existing techniques thus rely on 'multi-touch' input (e. g., [Hancock et al., 2007; Reisman et al., 2009]) or extra widgets (e. g., [Yu et al., 2010; Zeleznik and Forsberg, 1999]). More touching fingers or widgets on the screen, however, increase occlusion—a critical bottleneck, in particular, on the small surface of mobile devices. Another frequently used solution is 3D *rotation-scaling-translation* (RST) mapping: users control several integrated DOFs [Martinet et al., 2010] with two fingers [Liu et al., 2012]. Although this mapping is popular, a recent study [Besançon et al., 2017c] highlighted its difficulties: participants complained about the lack of separability of different DOFs.

We thus study the augmentation or combination of tactile input with other interaction paradigms to address its limits. In the past, tactile input has been combined with spatially aware tangible devices [Spindler et al., 2012; Spindler et al., 2014; Sollich et al., 2016; Besançon et al., 2017b], mid-air gestures [Kim et al., 2015], and pressure input [Brewster and Hughes, 2009; Heo and Lee, 2012; Corsten et al., 2017] to offer more interaction possibilities. While these techniques have been positively evaluated, most of them rely on custom-made sensing units which limit their adoption. In fact, mobiles' built-in functions are far from fully explored. As one of them, touch pressure sensing[6] is included in a number of recent phones, in this paper we leverage it for the control of different DOFs for 3D data navigation. Specifically, we designed a pressure-augmented scheme to separate different DOFs required for 3D navigation: we use force only to distinguish modes, not as a primary input. Our binary force mapping is easy to execute and to remember and takes inspiration from established tactile interaction mappings. We thus limit the number of touch-points to a maximum of two to leave as much as possible space for data display and allows users to precisely and independently control many of the DOFs involved in 3D navigation. It also does not require the use of additional sensors on the mobile device/display [e. g., Heo and Lee, 2011a; Pelurson and Nigay, 2016; Besançon et al., 2017a].

Our contributions are thus threefold. First, we present the design of our pressure-augmented tactile navigation mapping. Second, we compare our approach to the established RST technique in a controlled study. We found that our technique increases the 3D manipulation accuracy in docking tasks, with only a small increase in interaction time. Finally, we discuss the usage of our technique on devices without pressure sensors based on pseudo-pressure and verify that our technique is also valid in realistic application scenarios.

---

[5] Uniform scaling: 7 DOFs, non-uniform scaling: 9 DOFs [Cohé et al., 2011].

[6] Apple Inc.: '3D-touch'; Huawei Technologies Co. Ltd.: 'force-touch.'

5.2 RELATED WORK

It has been recognized that efficient interaction design plays a key role in data visualization tools, the trend of integrating interaction and visualization research has been argued since 1966 [Sutherland, 1966] until recent days [Keefe, 2010; Keefe and Isenberg, 2013; Besançon, 2018; Büschel et al., 2018; Wang et al., 2019b]. We focus on the effects of combined touch and pressure input for improving 3D dataset navigation on mobile devices.

5.2.1 *Visualization on Mobile Devices*

While the visualization of complex datasets is generally well studied, most work has focused on workstations or environments with large displays. The popularization of personal smart phones and smart watches, however, allows us to put visualizations on small yet readily available displays, and researchers have begun to investigate small-display visualization settings (e. g., [Langner et al., 2017; Pahud et al., 2018]). Mobile devices are also used as a props to interact with large visualization environments, from large screens (e. g., [Jansen et al., 2012; Besançon et al., 2017b]) to AR environments [Schmalstieg et al., 2002; Bornik et al., 2006]. Several examples are also illustrated in Figure 5.2. In our work we augment a mobile's touch input to better support data navigation tasks.

5.2.2 *3D Interaction with 2D Tactile Input*

The navigation in 3D datasets to gain information relies on frequent manipulations of view position and orientation. Due to the many benefits of direct-touch input (e. g., [Fu et al., 2010; Lundström et al., 2011; Coffey et al., 2012; Isenberg, 2016; Yu et al., 2010; Besançon et al., 2017c]), it is now a widely supported means to control 3D environments. While the mapping of finger locations to manipulations is straightforward in 2D (e. g., [Kruger et al., 2005; Hancock et al., 2006]), suitable mappings are less evident in 3D.

Early touch-based techniques like Rotate and Translate (RNT) [Kruger et al., 2005] and rotating-scaling-translating (RST) [Reisman et al., 2009] are often used to control virtual objects in 2D space. Hancock et al. [2007] first extended RST to 3D and provided different mapping possibilities using up to three fingers to support 3D interaction, and later to offer a full 6 DOFs *Sticky Tools* interaction technique [Hancock et al., 2009] which maps two-finger motions to translations (along x, y, and z) and to a rotation around z. A third finger is used for rotations around x and y.Reisman et al. [2009] also proposed 3D RST manipulations by using three or more fingers.

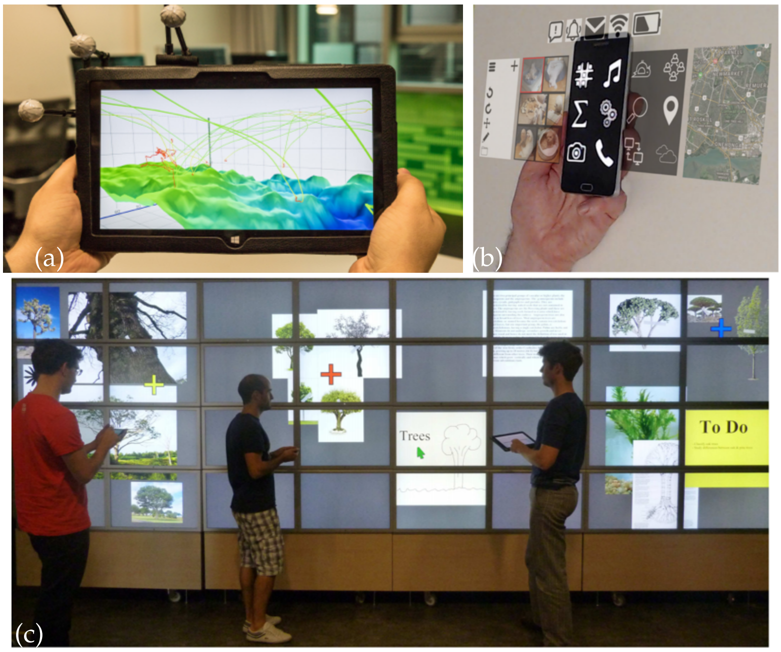

Figure 5.2: Examples of the use of mobile in visualization work: (a) visualizing data on mobile by Büschel et al. [2017]; (b) visualization using a mobile combined with AR environments by Normand and McGuffin [2018]; (c) using mobiles to control the visualization shown on wall-size displays by Chapuis et al. [2014].

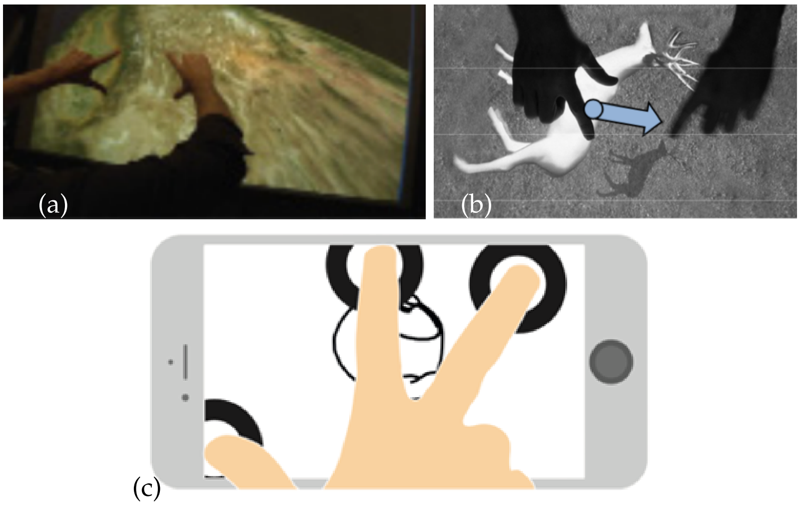

Figure 5.3: (a) 2D RST technique by Reisman et al. [2009]; (b) 3D RST technique by Hancock et al. [2007]; (c) an illustration of occlusion problems caused by increasing the number of touch points.

Another idea frequently found for interactive visualization is to use screen widgets to select a certain manipulation: it reduces the number of fingers needed and preserves a high DOFs count. Cohé et al. [2011], for instance, designed *tBox* which offered direct and independent control of up to 9 DOFs by means of a virtual box around the 3D object. Yu et al. [2010] proposed to reserve the display's borders for mode selection to control different manipulations. Increasing the number of fingers and using screen widgets is efficient in many cases. Yet, on small devices such as phones the former causes occlusion problems and a mismatch with common habits of interacting using at most two fingers, while the latter reduces the limited screen space that is needed for the visualization.

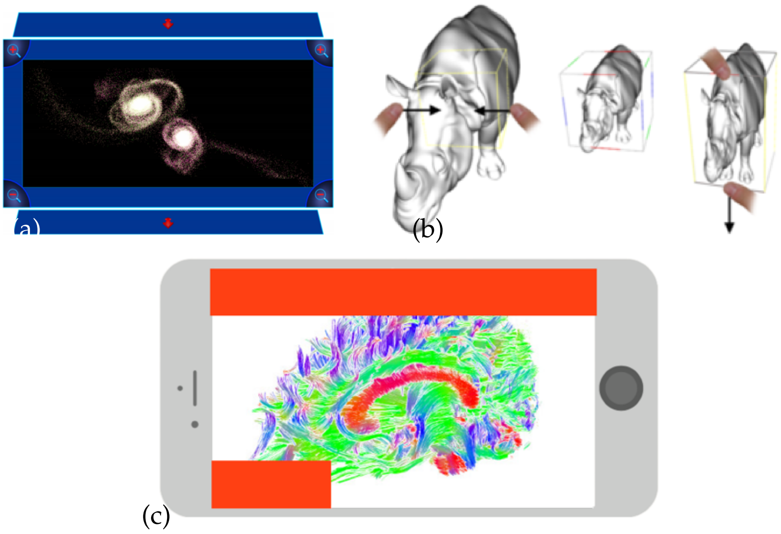

Figure 5.4: (a) FI3D by Yu et al. [2010]; (b) tBox by Cohé et al. [2011]; (c) an illustration of occlusion problems caused by screen widgets.

Reducing the number of fingers while keeping a high DOFs count usually requires the integration of several DOFs. For instance, Liu et al. [2012] controlled 6 DOFs with at most 2 fingers by integrating 3 DOFs (x-/y-translation, z-translation, and z-rotation). Whether it is better to integrate or to separate the control of different DOFs has been discussed in detail (e. g., [Zhai and Senders, 1997a; Zhai and Senders, 1997b; Veit et al., 2009; Martinet et al., 2010]). Researchers found that it depends on the input device: only when many input DOFs are available does it make sense to also provide integrated control [LaViola Jr. et al., 2017]. With limited DOFs of touch input it could thus be beneficial to provide separate DOFs control. The *Depth-Separated Screen Space* [Martinet et al., 2010], for instance, showed that separating rotations and translations led to faster manipulations, while participants in another study [Besançon et al., 2017c] frequently complained about DOFs integration on touch screens. We thus further discuss the effects of DOFs separation with the use of pressure.

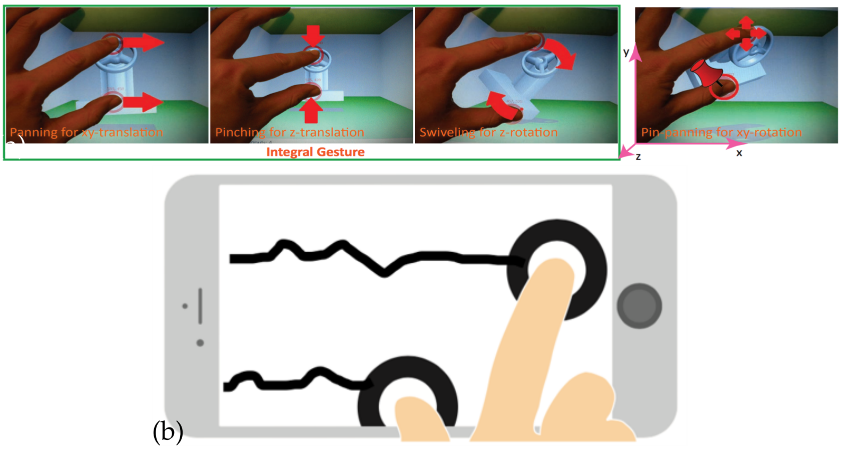

Figure 5.5: (a) an example of 3D manipulation gestures by Liu et al. [2012]; (b) an illustration of inaccurate control caused by integrated mappings.

5.2.3 *Augmenting Tactile Input*

Researchers also tried to combine tactile input with other input paradigms to get best of both worlds (e. g., Figure 5.6). Past research has investigated how to use internal sensor data to combine touch interaction with spatial/tangible interaction (e. g., [Oakley and O'Modhrain, 2005; Sollich et al., 2016; Schwank et al., 2017; Besançon et al., 2017b]). However, we generally discuss the use of another interaction paradigm to augment tactile input instead of being an additional primary input. Chen et al. [2014] proposed *Air+Touch* that combined in-air gestures and touch events using a depth camera, thus providing more interaction possibilities with a single finger. Withana et al. [2015] used infrared sensors to recognize shallow depth gestures to augment spatially limited input devices such as the touch screen. Hinckley et al. [2016] augmented tactile input with pre-touch sensing (fingers above the screen and grip around the mobile). While these ideas have been positively evaluated, they still require special hardware and extra input besides touching the screen.

Figure 5.6: Examples of using complex setups to augment tactile input: (a) the combination of tactile input with in-air gestures by Chen et al. [2014]; (b) combines tactile input with pressure using a self-designed unit [Heo and Lee, 2012]; (c) uses gestures to augment tactile input [Withana et al., 2015].

In our work, we are interested in combing pressure sensing with tactile input. The lack of physical pressure sensing in commercial devices forced researchers in the past to use separate pressure sensors. Pelurson and Nigay [2016] used the non-dominant hand to control pressure and augment navigation in large one-dimensional (1D) data. Heo and Lee [2011a] captured pressure from both sides and the back of a mobile and reported the difficult of maintaining pressure. They later introduced *ForceDrag* and *force lock* [Heo and Lee, 2012] to use pressure as an input modifier: pressure is only used to select the interaction mode before further manipulations are interpreted. They also suggested to use an indicator as virtual feedback. Heo and Lee [2012]'s work inspires our own, but it is still limited to a small number of DOFs which are not directly usable for 3D navigation. Besançon et al. [2017a] captured finger pressure on the back of a tablet to control the interaction's gain factor. They thus used pressure to control an additional variable, and not to improve tactile interaction mapping.

Recent work also investigated the use of direct pressure sensing for 3D manipulations. Wang et al. [2018] designed a 3D positioning technique controlling 3 DOFs with only one finger, providing depth information with pressure input. We are interested, in contrast, in using pressure sensing only to augment tactile input, such as to facilitate the control of 6 DOFs for different 3D navigation modes.

5.2.4 *The Use of Pseudo-Pressure*

Many mobile phones use capacitive sensing, yet without pressure data. Researchers thus explored pseudo-pressure that estimates the applied force based on contact area, temporal changes, or inertial sensors. Initial work assumed the contact area between finger and screen to increase with pressure, or postulated that a harder press usually lasts longer than a light tap. Benko et al. [2006] computed different contact sizes to distinguish the cursor's tracking and dragging state. Boring et al. [2012] used contact area for mode selection. Arif and Stuerzlinger [2013] exposed major challenges for pseudo-pressure: the contact area varies significantly between people and depends on the touch angle, while temporal approaches lengthen completion time. So they combined both touch time and its average surface. They later introduced an authentication system based on key sequences and pseudo-pressure [Arif et al., 2014]. Heo and Lee [2011b] used mobile phones' built-in accelerator and their detection algorithm distinguishes a gentle-tap and a force-tap. Goel et al. [2012] used vibration motor and gyroscope to measure the vibration absorbed by fingers. By using inertial sensors, these last two methods improved pressure prediction rate. In our work we first concentrate on physical pressure sensing but also investigate options to use pseudo-pressure sensing to use our techniques on more devices.And we take inspirations from previous pseudo-pressure based interaction mapping to bring forward our interaction design.

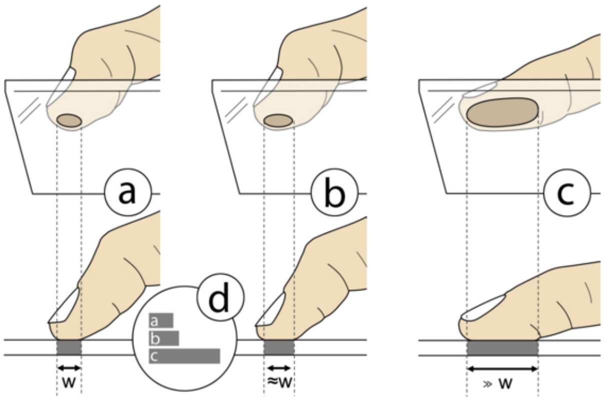

Figure 5.7: The use of contact size to simulate pseudo-pressure [Boring et al., 2012].

5.3 INTERACTION TECHNIQUE

We start by discussing the design goals and resulting interaction designs for our pressure-assisted mode selection based on the mobile's pressure-sensitive screen, without extra sensors. Our general guidelines for supporting data exploration on small screens are as follows (partially from the literature [López et al., 2016; Noirhomme-Fraiture et al., 2005]):

g1 To support effective data navigation, the interaction technique should support 6 DOFs: 3 for x-, y-, and z-translation; and 3 for x-, y-, and z-rotation in perspective projection (or x-/y-rotation and uniform scaling in orthographic projection).

g2 To lower the learning cost and to keep the maximum of intuitiveness, the mapping should not differ too much from currently used techniques on touch screens.

g3 The screen space should be reserved for the visualization. The mapping should thus use as few widget/fingers as possible.

g4 New interaction techniques should provide as much interaction flexibility, preferably more, than existing ones.

Even though many existing techniques support a high-DOFs manipulation, some rely on widgets that obstruct the view onto the manipulated visualizations. Alternatively, an RST-based input mapping can be used. The version of RST used most often for 3D manipulations on mobile devices [Besançon et al., 2017c] works as follows: users perform rotations around the x-/y-axis by moving a single finger, translate along x-/y-axis by moving two fingers in parallel, rotate around the z-axis by rotating two fingers around a given point, and translate along the z-axis by pinching two fingers. With such a mapping, users can control all 6 DOFs with a maximum of two fingers on the screen by integrating several DOFs, without any widgets. We thus chose to combine pressure sensing with RST to fulfill our design goals **g1**, **g2**, and **g3**. Yet, RST is limited because it always integrates x-/y-rotation with uniform scaling or z-translation, which can lead to a misinterpretation of user intents [Besançon et al., 2017c]. We thus use pressure input to separate out the control of different DOFs (supporting **g4**). Due to the known problems of using pressure input together with touch sensing (e. g., imprecise control mentioned by Ramos et al. [2004]), however, we want to treat it as different from other input types:

g5 Pressure should *augment* the touch input, not be a primary input. The pressure control should be as easy as possible to avoid additional workload or input errors.

g5 is important because it has been suggested that keeping a stable level of force during lengthy manipulations is too difficult [Heo and Lee, 2011a; Heo and Lee, 2012]. We thus use pressure sensing only to define spring-loaded modes [7] [Buxton, 1986; Sellen et al., 1992] in a quasi-postural [Isenberg and Hancock, 2012] fashion. Similar to other work (e. g., [Boring et al., 2012; Heo and Lee, 2012]), this means that we *evaluate pressure only at the start* of any input motion to select an interaction mode. Once the mode is selected, the *actual manipulations*

---

[7] Using spring-loaded modes means that a certain mode will be maintained, while users still keep the control, i. e., fingers still on the screen [Hinckley et al., 2006].

can be carried out with a normal force in all cases. This approach is more flexible than using time-based pseudo-pressure because users can reconsider/change their interaction intention after a finger is put down. Even though users can distinguish several discrete force levels [Wilson et al., 2010], a binary mapping (distinguishing only light and hard touch) is enough to provide the precise control that we need (having 6 DOFs and separating rotations from translation). It is also easy to perform since users do not need to worry about disturbing their interaction with accidentally high pressure input.

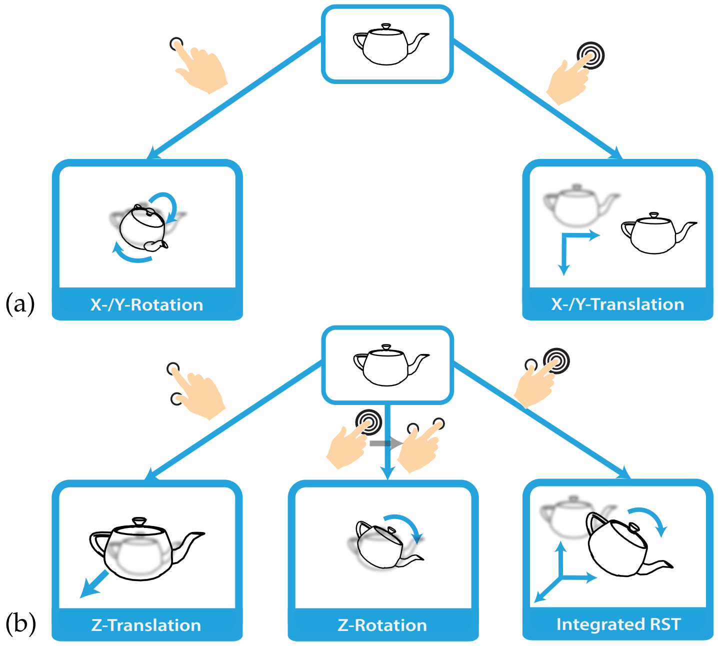

Figure 5.8: Interaction mapping: (a) one-finger motions and (b) two-finger motions.

Based on these considerations, we use the following mappings (Figure 5.8) by default. A single-finger light touch initiates x-/y-rotations, while a single-finger hard touch starts x-/y-translations. A two-finger light touch starts a z-translation, similar to the widely accepted "pinching gesture"—yet by not affecting the scale but the distance to the camera. For rotations around the z-axis, finally, a user first performs a single-finger hard touch and then puts down the second finger to start the manipulation. To avoid unwanted translations (the first finger may move a bit before the second finger is put down), we initialize a timer: if the second finger is put down shortly after the hard touch is performed (we use a threshold of 1 second, based on our pilot studies), we consider that the translation of the first finger is undesired due to mis-operating, then the data will be reset to the status before the first touch is effected. In contrast, if the second finger is put down after the time-out we treat the finger's translation is deliberate and keep the data. The exact mapping could be different according to the specific dataset: for exploring some datasets, translations are more

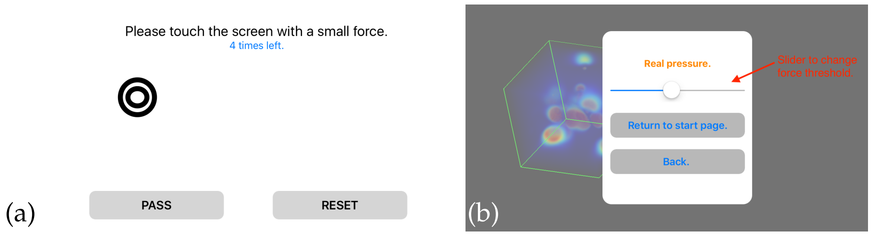
Please touch the screen with a small force.
4 times left.

Figure 5.9: Force threshold setting: (a) initial calibration, users touch at specific positions; (b) slider-based adjustment.

important than rotations, while rotations are key for others. We would map more important motions with light pressure (lower effort), but keeping the overall interaction design. As study the general use of pressure, we do not assume any specific type of data, so our specific mapping should be considered as an example. Finally, we add an integrated two-fingers gesture: putting down two fingers at the same time, with either one being a hard touch, is mapped by default to integrated RST. Aside of this last mapping, our design thus allows users to separate the translations from rotations. Interestingly, we thus go further than the work of Martinet et al. [2010] that studied the effect of separating translations from rotations—we also separate between the x-/y-axes on the one hand and the z-axis on the other hand to understand the impact of a further increase of the separability of DOFs. However, our study is different from theirs as we separate the DOFs with an other input paradigm, and use a different mappings.

We derived this overall mapping based on past experience with touch-based 3D interaction, and the two-level pressure input is likely to be generally applicable. Its specific parameterization, however, will likely depend on people's personal preferences as well as on differences between devices. We thus formulate our next goal as

g6 The choice of different pressure levels should be adjusted for different users and different devices.

To realize this calibration, we distinguish between a *light* and a *hard* touch with a threshold α. We perform an initial calibration to ensure that this threshold is suitable for each user (Figure 5.9(a)). To account for different finger angles [Boring et al., 2012; Roudaut et al., 2009], we record the pressure at several different positions. We average the maximum pressure value that was applied for a single touch gesture in each mode and thus derive α as the midpoint between both pressure averages:

$$\alpha = \frac{\frac{1}{n}\sum_{i=1}^{n} \text{HardMax} + \frac{1}{n}\sum_{i=1}^{n} \text{LightMax}}{2}. \tag{6}$$

We later validate this model of using an averaged single α over the full screen with a small study (Section 5.8). A manual adjustment of α

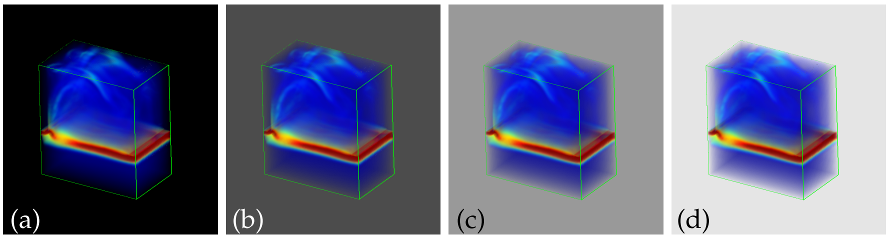

Figure 5.10: Example of visual feedback for volumetric flow data: (a) no touch; (b) light mode; (c) hard mode; and (d) integrated RST mode. We map brighter colors to more force and only change the brightness, not hue, to avoid misperception.

is also possible later (Figure 5.9(b)),[8] the initial position of the slider is associated with a calibrated threshold.

In addition to calibration, it is also essential to make users aware of their input. We thus state

g7 The pressure level has to be evident for effective control.

The importance of visual feedback for pressure input has been argued by previous work (e. g., [Ramos et al., 2004; Heo and Lee, 2011b; Heo and Lee, 2012]). For example, Ramos et al. [2004] recommend "real-time [and] continuous feedback." Since our users do not need to maintain the pressure after a mode has been selected, visual feedback is all the more important. It is the only way to show users their selected interaction mode. We use background color to indicate the current mode, instead of a scroll bar or text to avoid occlusion. It maintains maximum data visibility (to support both **g3** and **g7**). The color can be chosen according to each dataset's properties to ensure that it is not in conflict with the visualization. For example, Figure 5.10 illustrates our choice for visualizing volumetric flow data. Yet, to navigate in large or compact datasets, a common approach is to zoom into the data [Coffey et al., 2012]. In such cases all view space is filled by the data itself, our visual feedback does not have any effects. We thus also added haptic feedback with a short vibration when passing from light to hard touch mode.

We implemented our interaction technique on iOS: we load 3D data with C++ and render it with OpenGL ES 3.0 using our own shaders. We use the VTK 7.0 framework to support some specific scientific datasets (*.vtk, *.vti). We capture and process input events and related data (e. g., touch position, pressure, gesture) with Swift 3 and iOS SDK, and translate the input to change the data view.

---

[8] A dedicated menu can be called with a three-finger touch on the screen, a gesture which we do not used for 3D manipulation.

5.4 EXPERIMENT

To evaluate our technique, we wanted to compare it against other tactile 3D techniques. We excluded techniques using more than two fingers [Hancock et al., 2007; Reisman et al., 2009] because of the occlusion issue that we aimed to avoid. We also did not consider techniques that use screen widgets or interaction zones [Cohé et al., 2011; Yu et al., 2010] because they further reduce the available display size. We thus decided to compare our approach with the frequently used [Besançon et al., 2017c] 3D-RST technique described in Section 5.3.

The goal of our experiment is thus to understand the effects of using pressure for separating DOFs for 3D navigation tasks, based on quantitative data as well as qualitative feedback and observations from the use of both techniques. One of our reasons for separating interaction DOFs was to increase interaction accuracy, so we also wanted to test whether the use of different gain factors (i. e., control-display ratio, CD) had an impact on the performance: With a high gain factor, a small user input results in a large movement, and unwanted operations due to RST's integrated DOFs thus potentially result in a higher frustration. In contrast, a low gain factor may result in these unwanted operations not being noticed.

5.4.1 *Design*

We wanted to compare how participants perform 3D rotations and translations with both techniques. However, generic navigation in large 3D datasets is difficult to control. We thus chose a 3D docking task which comprises translations in 3 DOFs and orientations in 3 DOFs. A docking task consists of bringing a virtual object to a target position and orientation. Such docking tasks are common in the 3D interaction literature [Besançon et al., 2017a; Besançon et al., 2017c; Chen et al., 1988; Glesser et al., 2013; Hancock et al., 2007; Hinckley et al., 1994; Issartel et al., 2016; Vuibert et al., 2015; Zhai, 1998]. We used the Stanford bunny[9]—as done in previous work [Besançon et al., 2017a; Issartel et al., 2016]—due to its easily understood shape without orientation ambiguity. Our docking target was transparent green, while the object to dock was opaque white.

While our technique allows participants to also make use of the RST's integrated DOFs (Figure 5.8(b), bottom-right), we wanted to better understand the advantages and limitations of both approaches (i. e., integrated and separated). We thus removed, only for the experiment, the possibility to use the integrated DOFs with our technique.

Our experiment has two independent variables (3 gain factors and 2 techniques) and we measure two dependent variables (completion

---

[9] https://www.cc.gatech.edu/ turk/bunny/bunny.html

time and accuracy; see Section 5.5.2 for our method of measuring). In addition to the factor of 1 (a 1:1 mapping), we checked for effects of higher (2) and lower (0.5) gain factors. Our experiment thus used a within-subjects design with a total of 6 conditions. To account for variability we used 12 trials per technique per gain factor, resulting in 72 trials per participant. We counter-balanced the trial order using a Latin square to balance learning biases and tiredness. We also validated the target's positions and orientations from randomly generated targets by removing those that were hard to reach, resulting in the same pool of target positions for all participants. We varied the initial angle difference between the object and target from 27° to 180° (mean of 122.2°) and the initial distance difference between 5150 and 7076 units (in virtual space scale; mean of 6433.9).

5.4.2 *Participants*

We recruited 24 unpaid participants (8 female; ages 22–53, mean 30.9, med 26, $SD = 10.4$). 19 had at least a bachelor degree, while 5 had at most an A-level equivalent. 16 were experienced with 3D manipulation through extensive use of video games or 3D modeling software and 7 of them reported frequent use (daily/weekly). All of them reported to use tactile interaction daily on smart-phones or tablets, while 6 of them have used other tactile input devices. All had normal or corrected-to-normal vision. 1 participant was left-handed.

5.4.3 *Apparatus*

We ran the study on an iPhone 7 (4.7″ screen diagonal, 750 × 1334 pixels, 326 ppi, iOS 10.0.0). During the test, the smart phone rested on the table, in landscape mode. While this apparatus does not exactly copy the use of smartphone which are generally handheld, such controlled experiment setups are not uncommon in the field of HCI and visualization with, for instance, non-wrist-worn smartwatches [Blascheck et al., 2018; Blascheck et al., 2019] or touch-based systems [Deber et al., 2015; Knoedel and Hachet, 2011]. We decided to also conduct our study in a similar controlled/constrained way to remove possible confounds (e. g., both how the phone was held by participants and which finger they used could introduce additional noise to the pressure control) as illustrated in Figure 5.11. Moreover, we discuss later in Section 5.8 how our results still generalize to normal phone use.

5.4.4 *Procedure*

We told our participants that we would ask them to perform 3D manipulations on the phone that rested on the table in landscape orientation.

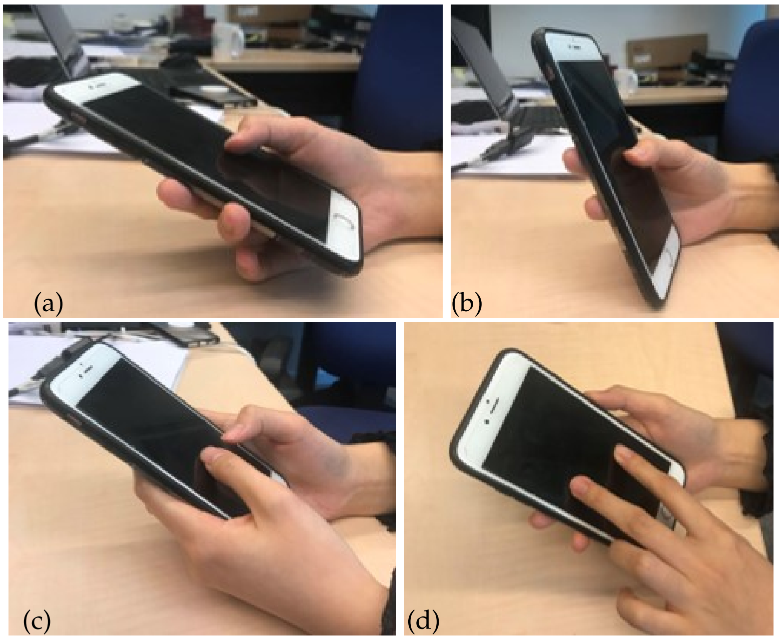

Figure 5.11: Different ways of holding a mobile phone will introduce noise for pressure control.

To avoid participant response bias [Dell et al., 2012], we told them that both methods were state-of-the-art, that none of them was invented by us, and that we wanted to compare their performances.

We began with an initial calibration for α (as described in Section 5.3). We then told participants that they would use one technique first, explained how it worked, and asked them to perform as many training trials as they wanted with different gain factors. We asked them to take good advantage of the training to get familiar with the technique. When participants reported that they were ready, we started the trials. Participants began and validated each trial by touching a button on the corner of the screen. We asked them to balance speed and accuracy but did not reveal their achieved accuracy immediately to avoid a bias toward accuracy [Hinckley et al., 1994]. When they finished all trials for one technique, we introduced the other technique and repeated the process (unlimited training and tests).

We also asked them to fill in questionnaires. Participants reported their age, education background, and how familiar they are with tactile screen and 3D manipulations, before we introduced the task. After each technique, we asked them to fill in a form to assess their workload (NASA's TLX[10]) and their fatigue (based on Shaw's approach [Shaw, 1998]). To avoid seemingly random choices made in the second part of the TLX (often seen as confusing by participants in our pilot studies) that would lead to inconclusive or incorrect results, we removed the second part of the TLX questionnaire to perform a *RAW* TLX. It is, according to Hart's survey [Hart, 2006], equally well suited as a regular

---

[10] http://humansystems.arc.nasa.gov/groups/tlx/downloads/TLXScale.pdf

TLX. At the end of the experiment, we asked all participants to give us their overall preference (the techniques could be named to be equal) as well as advantages and drawbacks that they experienced with both techniques. We recorded think-aloud comments throughout the study.

5.4.5 *Hypotheses*

Based on our pilot studies and results reported in previous work, we formulated a number of hypotheses:

- **h1** Both techniques exhibit a similar performance (accuracy and time) overall as the interaction mappings are close to each other. We do not believe that the separated DOFs has a strong impact on accuracy in general but,
- **h2** we hypothesis that, with a higher gain factor, DOFs integration result in worse performance and we thus expect RST to yield lower accuracy scores than our technique.
- **h3** The overall workload is identical for both technique. While the frustration is higher with the RST due to the integration, the mental workload is probably lower because this technique is already frequently used.
- **h4** The force-touch technique increases the overall fatigue, probably with increased finger and hand fatigue.

5.5 RESULTS

We collected a total of 1728 docking trials from our 24 participants. In addition, we recorded self-stated feedback and answers to subjective preference questions. While data from HCI experiment has usually been analyzed by the Null Hypothesis Significance Testing (NHST) in the past, we choose to report our results using estimation techniques with effect sizes and CIs instead, as recommended by the American Psychological Association (APA) [VandenBos, 2009] and also because the former approach is being increasingly criticized by statisticians [Lai, 2010; Cumming, 2012; Cumming, 2013; Cumming, 2014; Valentine et al., 2015; Baker, 2016], stats practitioners [Amrhein et al., 2017], and HCI researchers [Dragicevic et al., 2014; Dragicevic, 2016; Besançon and Dragicevic, 2017; Besançon and Dragicevic, 2019]. While dichotomous interpretations based on p-values are still extensively used [Besançon and Dragicevic, 2019], approaches relying on effect sizes and more nuanced interpretations are now widely recommended [Cumming, 2012; Cumming, 2013; Cumming, 2014; Valentine et al., 2015; Dragicevic, 2016; Besançon and Dragicevic, 2019]. However, it is still possible to read our results based on p-value by comparing CIs spacing with common p-value spacing [Krzywinski and Altman, 2013]. We detail most of the results in this section by technique and by

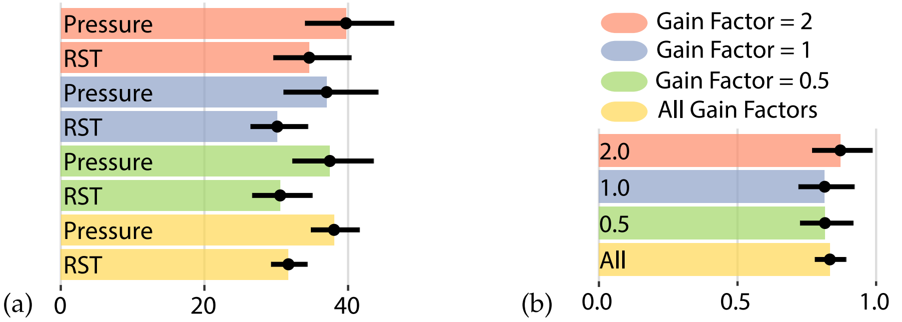

Figure 5.12: Task completion time: (a) absolute values and (b) pair-wise ratios. Error bars: 95% bootstrapped CIs.

gain factor to better understand the latter's potential impact on the former.

5.5.1 *Completion Time*

We analyze log-transformed time measurements to correct for positive skewness and present, in Figure 5.12(a), anti-logged results with geometric means that dampen the effect of extreme completion times that would bias an arithmetic mean, as recommended in such cases [Sauro and Lewis, 2010].[11] Our results show that, overall, there is strong evidence that RST is faster than our technique. For normal (i. e., 1) and low (i. e., 0.5) gain factor values, there is weak evidence that this faster completion time still holds, while for high (i. e., 2) gain factor values our data would suggest that there is no difference. We also report on the effect size in Figure 5.12(b) which corresponds to pairwise ratios,[12] available at `https://aviz.fr/ci/`[Besançon and Dragicevic, 2017].The non-overlap of each confidence interval with the value of 1 clearly shows strong evidence that the completion time is shorter with the RST technique than with our technique in all cases except the 2.0 gain factor value condition. For this one, the evidence is still strong, however. For RST, overall, participants needed a bit longer than 4/5 of the time needed for our technique, so the overall difference in completion time is not large.

---

[11] Arithmetic means use the sum of a set of values, while geometric means use the product of the set's values.

[12] Pairwise ratios are computed for each individual subject, i. e., to divide measurement per each individual subject. Our computation scripts are

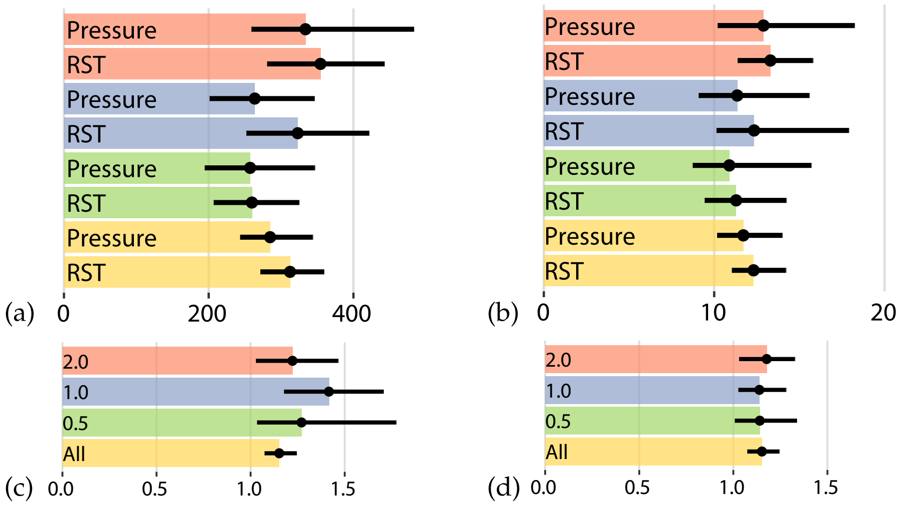

Figure 5.13: Accuracy: (a),(c) Euclidean and (b),(d) angular distances, (a),(b) absolute values and (c),(d) pair-wise ratios. Error bars: 95% bootstrapped CIs. Colors as in Figure 5.12.

5.5.2 *Accuracy*

We determined the Euclidean and angular distances between the manipulated object and the docking target when the participant validated a trial (Figure 5.13), using the objects' centers for the former

$$d_e = \sqrt{(x_t - x_o)^2 + (y_t - y_o)^2 + (z_t - z_o)^2} \tag{7}$$

and deriving the angular difference as

$$d_a = 2 \cdot \arccos(q_{d\omega}) \; ; \quad q_d = q_o^{-1} \cdot q_t \tag{8}$$

with q_o as the manipulated object's quaternion, q_t as the target's quaternion, thus q_d being the difference quaternion, and q_ω being the ω component of an $\omega + xi + yj + zk$ quaternion with $i^2 = j^2 = k^2 = ijk = -1$.

The overlapping Confidence Intervals (CIs) in Figure 5.13(a) suggest that with our data we cannot find evidence of a difference of Euclidean distance between the two techniques. We computed the ratio RST/Pressure for each gain factor to look at the intra-participant difference. The non-overlap of each confidence interval with the baseline 1 provides strong evidence for our technique being more precise than the classical RST mapping w.r.t. Euclidean distance. This observation is strongest for a gain factor of 1: our technique is almost 50% more precise than RST. Similarly, the CIs in Figure 5.13(b) do not provide evidence of a difference between the two technique for angular distances. However, none of the CIs overlaps with 1 in Figure 5.13(d) so our technique is overall more precise than RST.

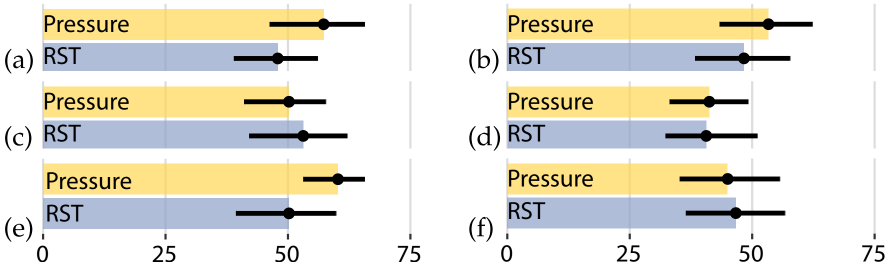

Figure 5.14: Workload in TLX units (lower is better) for (a) physical, (b) mental, and (c) temporal demand, (d) performance, (e) effort, (f) frustration. Error bars: 95% bootstrapped CIs.

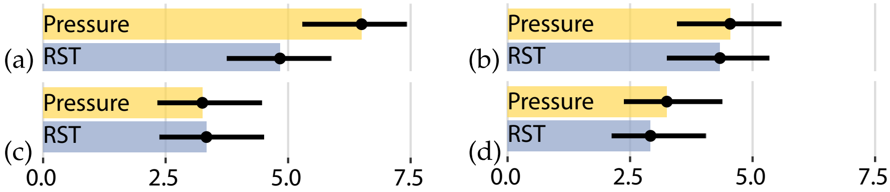

Figure 5.15: Fatigue measurement for (a) fingers, (b) hands, (c) arms, and (d) shoulders. Error bars: 95% bootstrapped CIs.

5.5.3 *Measuring Workload and Fatigue*

We report the TLX results in Figure 5.14. The mental and temporal demand (Figure 5.14(b), (c)) as well as frustration and performance ratings (Figure 5.14(f), (d)) show no evidence of a difference between both techniques. Even though the CIs are overlapping, a difference may exist for physical demand: our technique's physical demand (Figure 5.14(a)) could be higher than that of RST. The even smaller overlaps of the CIs for effort (Figure 5.14(e)) suggest that participants thought it was higher for our technique than for RST. Overall, however, the two differences are small and it appears that the general workload is similar for both techniques.

We report fatigue data in Figure 5.15. The large overlap of CIs for shoulder (Figure 5.15(d)), arm (Figure 5.15(c)), and hand (Figure 5.15(b)) shows that we did not see a difference between the two techniques. A higher finger fatigue caused by our technique, however, is highlighted by the smaller CIs overlap in Figure 5.15(a).

5.6 DISCUSSION

5.6.1 *Completion Time*

We have gathered with our experiment evidence against **h1** because we did notice a difference in completion time between both techniques. It appears through our results that RST can be a bit faster than our technique, thus partially contradicting previous studies [Martinet et al., 2010; Nacenta et al., 2009]. This difference of results can probably be explained by two factors. First, many users are already aware of the gestures used in RST mapping for 2D manipulations on mobile devices. When it comes to our technique, we observed that users chose to take only 2–3 minutes and, at most, 5 minutes for training. This reduced learning (compared to the long experience with 2D RST) resulted in the need to try to recall the interaction mapping during the experiment. This recall leads to slower interaction, which is common with any new technique as explained by MacKenzie and Zhang [1999] (*elusive crossover point*). Second, the tight DOFs integration of RST may also contribute to the faster completion times. Indeed, integrated DOFs allow participants to perform several manipulations in short and fluid succession, which would have to be executed sequentially with our technique. While we believe that this sequential interaction is an asset for accuracy purposes, and the slowing down is likely not a problem for our complete interaction mapping which still provided integrated DOFs manipulations as explained in Section 5.3.

5.6.2 *Accuracy*

Our accuracy results provide evidence to partially validate **h2**. While we wrongly believed that the accuracy would be much more different for a higher gain factor, it remains that, overall and for all gain factors in particular, the accuracy was better with our technique than with the RST mapping. This result can also be explained by the sequential manipulations that are impossible to achieve with the RST mapping, thus making small modifications of single DOFs almost impossible. This is where our technique shines. While it was not critical for a docking task, such an accuracy is very likely to be fundamental in many scientific domains which rely on 3D data visualization and precise interaction.

5.6.3 *Workload and Fatigue*

Our hypothesis **h3** seems to be validated by our experimental data, though our reasoning was wrong. The overall workload seems to be similar for both technique, thus highlighting that having this additional input to augment tactile interaction was not hindering users' workload. The mental workload results clearly highlight the fact that the learning needed to master our technique was not an issue. However, this additional effort in learning probably influenced the effort evaluation done by our participants. This idea is further reinforced by the fact that seven participants actually reported it. One could thus wonder whether, with a longer exposure to pressure-based mode selection, this effort would not show a tendency that is more similar that of RST.

Based on our pilot studies we expected our approach to increase finger and hand fatigue, thus increasing the overall fatigue of our technique compared to the RST mapping. This assumption was partially right: the finger fatigue is indeed increased by the use of our technique, but the hand fatigue measurements are the same for both techniques. This increase in the finger fatigue was easy to predict, as users have to push harder on the screen more often. However, one can notice that the finger fatigue is not much higher with our technique, and that no participant later complained that their finger was hurting, even though the total duration of the study was of 44.6 minutes on average (24.6 min for our technique, a bit more than RST). One should take into account that the overall manipulation time was longer with our technique than with the RST mapping. While this is only a few seconds per trial, it easily represents several minutes when multiplied by the 36 trials per technique. We thus believe that the gained accuracy obtained with our technique is probably worth a bit of extra finger fatigue. Furthermore, the fact that some of our participants had not been exposed to such pressure interaction before did contribute to the higher fatigue: they usually pressed harder than needed to switch modes—such behavior would likely quickly disappear with a longer exposure to our technique.

5.6.4 *Preferences and Qualitative Feedback*

Overall, 14 participants preferred RST, judging it faster and/or more natural, while 9 participants reported our technique as their favorite one. One reported that he liked both equally. None was evaluated as inappropriate for the task. This result contradicts those of Nacenta et al. [2009] for 2D interaction and Martinet et al. [2010] for 3D interaction that integrating DOFs could reduce user satisfaction. We believe that this disparity can be explained by the fact that RST is frequently used for 2D or 3D applications [Besançon et al., 2017c]. Moreover, the studies by Martinet et al. [2010] and Nacenta et al. [2009] were conducted in 2009 or 2010. Since then, RST established itself as predominant. Such a long-term exposure is likely to bias subjective ratings [Burgess and Sales, 1971], and is thus likely to have also biased our participants—16 of them reported to be familiar with 3D manipulations and have surely been exposed to RST before, at least in its 2D form. Moreover, five participants stated that our technique specifically needed to be learned since pressure interaction is but still rarely used today, while RST was deemed as natural and easy to understand by 11 participants. Still, nine participants stated that our technique, though difficult at first, could be understood easily, and once mastered would provide just as good results. The preference for the RST could also possibly be explained by the fact that participants were faster with it. Four participants actually stated so. Taking into account the high number of docking tasks they had to perform, it is likely that there could be a bias toward the fastest technique.

Our fatigue results are also reinforced by participant comments: ten reported that they could feel, at some point, the extra effort on their fingers. This fatigue was probably emphasized by the fact that, for some participants, first pressing hard and then releasing was not natural (two mentioned it). Some continued to press hard throughout all manipulations—spring-loaded modes are also rare in public applications and possibly have to be learned. Nonetheless, participants also highlighted our findings on accuracy. Sixteen of them praised our technique for its better accuracy over RST and reflected on the fact that it helps avoid unintentional DOFs manipulations.

Interestingly, three participants commented on the occlusion issue and mentioned, while performing, that using two fingers for translation (i. e., RST) was problematic because a good part of the screen was hidden. They mentioned that this problem was solved with our technique. This occlusion reduction is an improvement over the RST mapping: most of the interaction (rotation around and translation along the x/y-axes) can be done with a single finger, while the RST technique requires 2 fingers for all but x/y-rotation.

Even though the learning issue was mentioned by nine participants for our technique in contrast with the 'natural' mapping of RST, participants mastered our technique. This finding is also supported by the similar mental workload scores for both techniques and that no participant reported that our technique was too hard to understand, three of them even reported that it was easy to use.

Three participants reported that our visual and haptic feedback helped them. We are still unsure whether haptic feedback would be enough to differentiate between two values only. Since most mobile applications today provide notifications through haptic feedback, however, it does not seem unreasonable to keep a visual feedback.

Though they could change the force feedback and we explained why this is useful, none of our participants felt the need to adjust it. This observation suggest that our initial calibration phase, while simple, is sufficient to calibrate front-of-device input for two values.

Finally, we also noticed an interesting use of the two-finger interaction. Seven of our study participants switched between one-*handed* interaction and two-*handed* interaction, even on such a small device. This possibility was probably exacerbated by our setup since users did not have to hold the device at all, thus having both hands free for interaction.

5.6.5 *Limitations and Future Work*

With our study we evaluated the potential of our pressure-augmented tactile interaction technique on a smart phone for 3D data navigation, in particular to separate out the different DOFs (translations from rotations). While we paid attention to as many factors as we could (e. g., several gain factors), there are other unevaluated aspects. For instance, our study was only done on a single screen size and we can conjecture that the screen size factor could be responsible for an increase/decrease in the occlusion issue as mentioned by our participants. Although there is no reason that our technique cannot be scaled to larger screens, the potential fatigue requires reflection on how to reduce it. A higher gain factor could, for instance, avoid long dragging on large screens.

Similarly, it would be interesting to check whether our technique could be easily remembered and used. We wonder whether our participants would have been faster and/or more accurate after a longer exposure to our technique. Indeed, comparing our mapping to RST was somewhat unfair since most people actually know and understand the RST mapping. As a consequence, we would like to investigate how much better our pressure-based technique could be used with a proper and longer learning phase.

With our design, pressure is used only at the beginning. It would thus also be interesting to investigate if pressure could be used in other phases (e. g., to change interaction mode during the manipulation to avoid frequent shifting of fingers, or to validate/cancel the manipulation at the end).

Docking tasks are abstract, yet they are generalizable to many tasks in 3D (e. g., [Chen et al., 1988; Hinckley et al., 1994]). Even though some interaction intents in 3D visualization go beyond this paradigm, specific tasks (such as 3D subset selections, 3D points manipulations and specifications, temporal navigation, data read-outs, etc. [Keefe and Isenberg, 2013]) could still take inspiration from our initial work. For example, the slicing of volumetric data also requires the precise control of the slicing plane's position and orientation. We thus would like to further investigate how pressure could be used in such demanding contexts. While our study exclusively investigated mobile device use only, such devices can also play a role in larger visualization environments as we noted before. It would thus be interesting to study the use of our technique in combination with such environments.

Also, the prototype we used in our study can naturally be improved. For example, there could be better ways to trigger the menu, we can improve the visual and haptic feedback, and we could adjust the way to manage the exact interaction mapping of the pressure input (e. g., which pressure level triggers what motion), etc.

We used the pressure to select and switch interaction mode because it is an easy way to be realized with commercial mobile devices. Nevertheless, there could be other possibilities to achieve a high accuracy with mobile interaction. For example, some work focuses on back-of-device interaction (e. g., [Maiero et al., 2019; Yadav et al., 2019]),

Table 5.2: Results of pseudo-pressure pilot test.

feature	true positive	false positive	true negative	false negative	accuracy
T(time)	121	15	225	119	72.1%
R(radius)	190	57	183	50	77.7%
T and R	108	12	228	132	70.0%
T or R	203	60	180	37	79.8%

but most of these approaches rely on additional units to capture the interaction. While it would be ideal, for practical usage, to avoid such complex setups, it remains an interesting future research question to study how to combine existing input devices with others, or how to design a new input device that would allow users to interact with both a traditional workstation and stereo immersive environments, as we have discussed in Section 4.7 of Chapter 4.

5.7 EXTENSION TO OLDER DEVICES WITH PSEUDO-PRESSURE

To generalize our approach to devices without physical pressure sensing, we can make use of pseudo-pressure sensing discussed in Section 5.2.4. High-accuracy prediction techniques (e. g., [Hwang et al., 2013; Goel et al., 2012]) usually trigger vibration when a touch is detected. This can disturb normal manipulation and is not compatible with our haptic feedback. We thus investigated previously envisioned and simpler features such as contact area or touch completion time [Arif et al., 2014; Arif and Stuerzlinger, 2013; Benko et al., 2006; Boring et al., 2012]. Indeed, we also observed that the contact area increases with the pressure as described by Arif and Stuerzlinger [2013] and users begin to move their finger(s) very quickly when performing a light touch motion, while a hard mode manipulation exhibits a longer time before fingers start to move as illustrated in Figure 5.16. We propose to establish our model based on these two features.

We then started by conducting a pilot test to parametrize the hybrid model. We asked the same 24 participants to first calibrate the touch sensing by performing five one-finger translations in both light and hard mode. Initial touching target and ending translation targets are shown each time differently on the screen. We computed the average time t that the finger takes to move more than a given distance d (five pixels in our experiment) for both modes, and used the midpoint value for the time threshold. If the finger stays longer than the threshold without shifting outside d, we assume that the time criteria is attained to be classified as a hard touch. Similarly, we computed the radius threshold by using the middle value of the average contact radius.

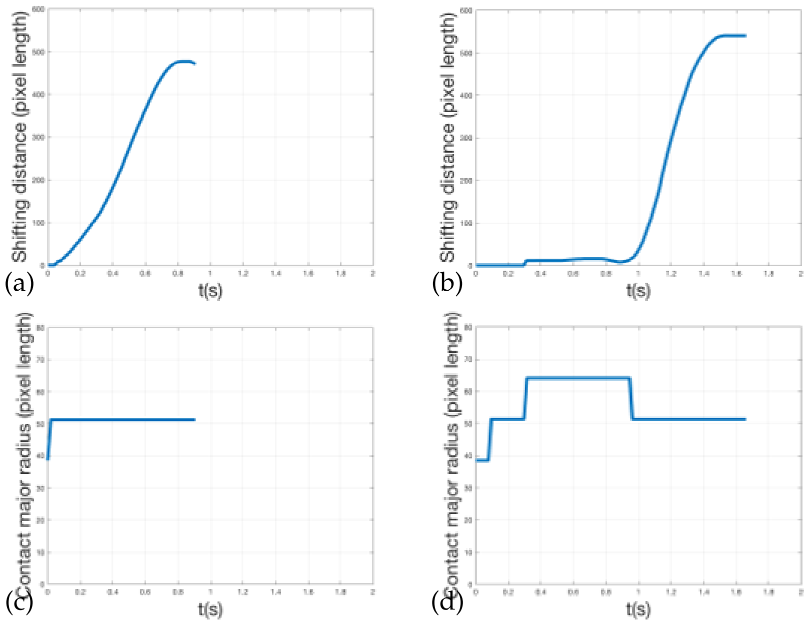

Figure 5.16: Example curves of (a) translation distance vs. time for low pressure; (b) translation distance vs. time for high pressure; (c) touch radius vs. time for low pressure; (d) touch radius vs. time for high pressure.

After calibration, we asked each participant to perform one-finger translation 10 times in each mode, with the order and touch position as well as the trajectory being randomly selected from a pre-defined data pool. We gathered 480 one-finger motions (Table 5.2). Our prediction accuracy is low when compared to the good accuracy of proper pressure sensors (up to 97% for binary level [Goel et al., 2012]). Also, the SDK of iOS 10 does not allow us to directly obtain the continuous contact area on the screen but rather gives us one of the few possible discrete values. However, our captured data showed us that it may be possible to create more precise models. We observed that a hard touch has a peak in its contact area data several milliseconds after the initial touch is made and then suddenly drops. Relying on this, we could improve our accuracy to make our technique efficient on non-pressure sensing devices in the future.

5.8 REALISTIC USAGE OF PHONES

To verify that our technique could also be used in a non-controlled setup when people are holding the phone, we conducted a second test to see if different touching areas and different holding postures would influence pressure. We divided the screen uniformly into nine areas (1 is top-left, 9 is bottom-right), and asked participants to touch these areas both normally and hardly, four times each, in a randomized order, guided by an indication on the screen. We used a within-subjects

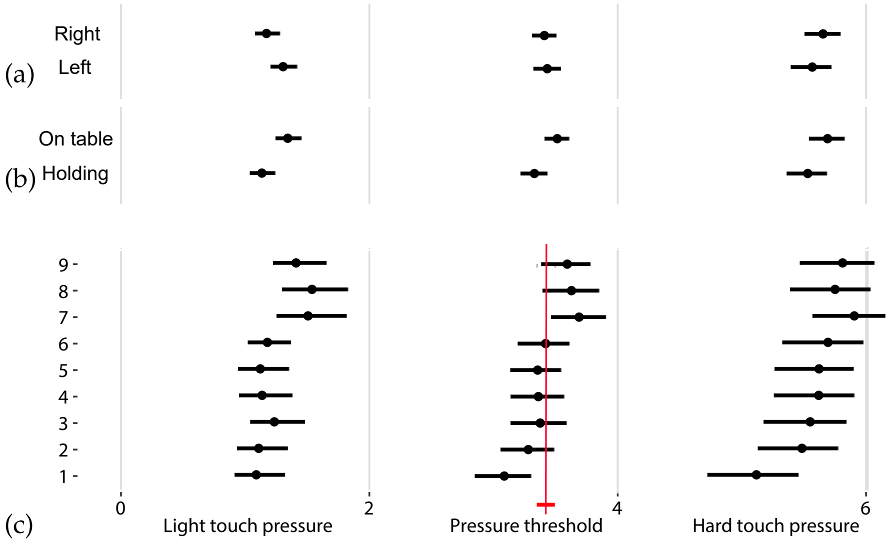

Figure 5.17: Pressure for different study conditions: (a) left and right hand; (b) holding the phone and put it on table; (c) different areas, the red line is the threshold over all screen. Error bars: 95% bootstrapped CIs.

design with a total of 4 counter-balanced conditions (2 different hands (left and right) × 2 postures (phone on the table and phone held by participants)). We told participants to perform the touches only with their index finger as seems to be usual in most 3D applications on mobile devices [Besançon et al., 2017c]. We recruited 8 unpaid volunteers, all right-handed. We ran the experiment on the same device as used in Section 5.4. We gathered 2304 touch events (8 participants × 2 hands × 2 postures × 9 areas × 2 pressure levels × 4 repetitions), see Figure 5.17. The overlapping CIs in Figure 5.17(a) indicate that pressure for light touch, hard touch, and average threshold (α) is similar with both hands. While the overlapping of CIs in Figure 5.17(b) provides evidence that pressure for light touch, hard touch, and average threshold (α) differs according to the mobile's postures, we can see that the absolute value difference is low. The calibration technique we used thus scales to the phone-holding posture. The small value difference can also be compensated by manually adjusting the threshold. The CIs in Figure 5.17(c) show strong evidence of differences between zones for similar pressure input (e. g., between 7, 8, 9, and 1). However, Figure 5.17(c) also shows that it is still possible, for all areas on screens, to distinguish between a hard and a light touch. Based on this data, we are thus confident that our technique scales to realistic hand-held phone usage.

5.9 CONCLUSION AND PERSPECTIVE

We reflected the use of mobile pressure sensing and presented an interaction mapping to separate different DOFs for 3D navigation. This new type of interaction design and input mapping allows us to provide effective visualization exploration tools on new platforms such as mobile phones that have become ubiquitous in today's world yet which have input capabilities and output constraints fundamentally different from established data exploration platforms such as workstations. With our technique, we can reduce, in particular, the occlusion issue often experienced with more-than-two-finger techniques on small devices, do not have to resort to the integration of several DOFs, and can use the entire screen space for data visualization. In our experiment we determined that our technique is more precise than the integrated RST technique which is common in the literature and in mobile applications, with limited impact on speed.

An important insight from our study is that pressure-based spring-loaded moding in quasi-postural fashion is easily understood, used, and does not increase users' mental workload. Also important is the fact that our results and our participants' feedback confirmed that such a separation of DOFs leads to a better accuracy. These results can be of critical importance for some specific scientific domains which rely on interactive visual data exploration to gather insights.

Finally, we reflected on the possibility to use our technique with devices without force input. While we conducted our experiment in a controlled environment to reduce possible confusion, we also verified that our pressure calibration model can be directly applied to realistic application scenarios.

This chapter complement Chapter 4 to answer **Q3**. As walking around to observe the data in AR is seen to have an important potential, we proposed to use a mobile device to support the interaction while walking around because we lose the access to the mouse and the keyboard. With this technique, it is possible to improve the interaction accuracy which is usually highly demanded for visualization tasks.

6

DISCUSSIONS, CONCLUSION, AND PERSPECTIVES

In this thesis, we present the work that combines both visual and interactive immersions to practically bring immersive visualization to existing scientific workflows. In this chapter, we reflect on the basic requirements and examine if we have answered the initial research questions. Meanwhile, based on the insights we gained from our work, we propose and discuss a few possible future research directions.

6.1 THE BASIC REQUIREMENTS AND THE HYBRID DESIGN

We first look back at the basic requirements (**R1** and **R2**) that we had driven based on our observations of and discussions with domain experts which we had stated in Chapter 1 as follows:

- **R1**: Both traditional analysis and efficient visualization tools are important. The system needs to support easy switch between the two.
- **R2**: The immersive systems should be easily integrate to office working environments.

R1 largely motivated our design of using a hybrid system, instead of a purely immersive one, to support the use of traditional analysis tools and meanwhile benefiting from the immersive visualization (Figure 6.1), as we have discussed in Chapter 1 and Chapter 2.

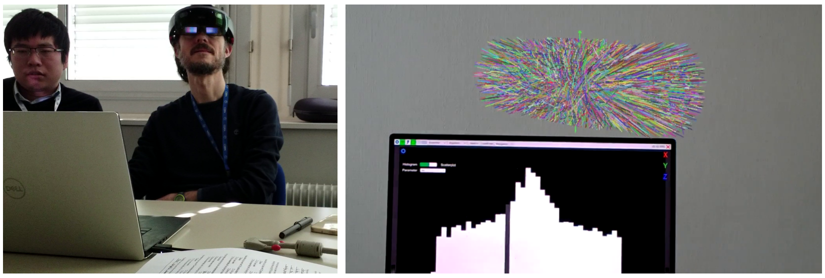

Figure 6.1: The hybrid setup can support both traditional analysis tools on PC and the immersive view in AR.

While existing immersive environments that we reviewed in Chapter 2 can be combined with other devices to fulfill **R1**, most of them are not compatible with **R2**. For example, even though it is possible to add a laptop or a tablet in large environments (e. g., the CAVE2 as illustrated in Figure 2.7), we excluded them because their complex setups and maintenance make themselves too difficult to be applied to an office. For smaller ones, the Desktop-VR metaphor, which allows users to switch between the monoscopic view and stereoscopic view, is already pretty close to our goal, but the new generation of lightweight VR/AR headsets brings new opportunities. For example, the visualization of 3D content can be independent of the computer screen such that it allows users to work on different information at the same time. Moreover, modern VR and AR headsets can now offer a canvas with a large-surface that does not limit the visualization to the existing screen size, even with a small workstation without the need of building a large power screen or a virtual room. We finally chose an AR headset over the VR one for its non-occlusion of the real world, which could largely facilitate the use of a PC. Thus, with a combination of a PC and an AR headset, both **R1** and **R2** are be respected.

The results of our observational study presented in Chapter 3 confirmed that such a hybrid setting is advantageous and is easy to use, thus validated the feasibility of our design. We can answer **Q1** (What should a data exploration environment look like that combines an AR display with a traditional workstation?) that a data exploration system should be a combination of a PC workstation with a lightweight AR HMD to allow experts use their traditional tools on PC while enjoying the immersive visualization with the AR. However, our studied feasibility focused on the mental aspect and weakened the influence of the hardware limitation. Indeed, the existing hardware does have a few important drawbacks to limit the application of our setup to daily workflows. The first version of Microsoft HoloLens had only a very limited field of view, resulting in the difficulties of making full use of its canvas. So users cannot directly perceive the visualization of several different views at the same time, but they need to move their head around to have a larger area of the view. Even though the second version has largely increased its field of view, from our empirical experience, its display quality is still much lower than a normal screen, sometimes even worse than its first version. Under this circumstance, to better use this hybrid setup, we need to carefully consider what to be visualized in which space, especially what we need to put in the AR space. For example, it seems to be not a good idea to put long text or visual representations using large space on it. This is one weakness of our setup compared with large environments that could make use of high-quality displays. Apart from that, the weight of the hardware could not be totally ignored at this stage. A heavy device will decrease users' willingness to use it, we thus think that, our setup is more suitable for occasional uses to deal with a certain visualization task that does not require a long time to perform, instead of wearing the HMD all day long. We believe that an important step towards a long time practical use of such a hybrid setting would require technological breakthroughs to make the HMD more comfortable and less disturbing for the wearers.

For the advantages of using AR compared to VR, besides those reported previously in our analysis for deciding the setup (Chapter 1), experts largely envisioned that this setup could favor collaboration in the future. Being able to see the world around them makes it to possible to first directly interact with their collaborators (for example, non-verbal cues could be used to facilitate the communication) and other tools as they are working normally in their office. Moreover, the use of a hybrid setup does not only favor visualization. We could envision that a possible way to support collaboration among users is to use the PC to visualize private content, while the AR for displaying public content to everyone, which is the same idea as the work of Febretti et al. [2013] using the CAVE2. Although this function could be easily supported by other combinations, such as the work of Reipschläger et al. [2021] and Büschel et al. [2021] that combines a large screen and AR headsets, our setup (with limited requirements of space and setup) supports easily the collaboration within a normal office.

In addition to collaboration, the visibility of surrounding environments benefits the use of real-world objects and traditional analysis tools, but could potentially introduce a certain level of disturbance. For example, the background could be very annoying if one shares an office with several colleagues, or the office might be filled in with large furniture. In this case, the current setup could only reduce this influence by finding other clean places, which could probably be far from the PC screen that makes using our setup less practical. Apart from that, when one's desk is closer to the wall, the visualization will collide with the wall. Even though it is still possible to use our setup, the experience would not be great. To deal with this, we wonder, instead of using the optical see-through environments, whether using a video see-through device would be better suited for a normal office. Generally, users can still perceive everything around them with the video captured through cameras (which we referred to as the AR). When needed, they able to hide the background, and fully immerse themselves inside the data to reduce the potential disturbance of the surrounding environments, and without the needs of regularly taking-off and -on different devices.

From another point of view, using a video see-through device is not without limitations. One of the important advantages highlighted in our experiment is that our chosen hardware did not cause uncomfortable feelings as what is commonly experienced with the VR setups. This is essential towards a long-time utilization of the hybrid setup in their normal workflows, but such advantage might not be passed to video see-through devices due to the possible tiredness of watching videos all the time. We thus think that an interesting future research direction for improving our hybrid setup would be to study the combination of AR and VR setups with a single device, under our general vision of keeping a traditional workstation.

6.2 THE COMBINED USE OF A SCREEN AND AR SPACE

To examine **Q2** (How should we treat the AR display compared to a 2D screen and how do we make transitions between the two?), we presented in Chapter 3 our idea to use the AR as an extension to the PC, in a similar manner that domain experts work with multiple screens in their daily workflows. The results from Chapter 3 informed us that a hybrid system should well consider and make use of the inherent properties of each space. It is not necessary nor advantageous to support everything on both sides.

With such a design, we only took the first step to understand how AR space could be used in combination with a 2D screen and treated them as separated views. However, to continue exploring **Q2**, we could also investigate the use of AR space to augment the 2D screen, similar to the idea presented by Reipschläger and Dachselt [2019]. As such, the two spaces would have totally different roles. In our application domain, we could imagine, for example, use the screen to visualize the results of their traditional analysis results that are often represented as basic bar or line charts. At the same time, a floating 3D view could be placed in front of (or above, depending on how users place the screen) the screen to associate the 2D and 3D representations. We can even go further by combining the two ideas to make better use of the large canvas in AR—the space next to the screen could be used as an augmentation while other space works as an extension, displaying such as different views of the same data, a storyboard stocking all the different states of manipulation, or maybe simply an amplified visualization of a certain region on the screen. This idea of combination (previously examined with other setups, such as the work of Reipschläger et al. [2021] and Büschel et al. [2021]) would also favor the usage of the large canvas of AR.

As for the transition between the two spaces, one type of transition is to investigate how can we transition the design of the traditional setup to the AR space. A 2D screen can be traditionally represented by a flat plane, but the concept is being challenged nowadays. For example, when scientists work with two linked screens, they usually put them side by side, but with a kind of curvature to form a surrounded feeling. In this situation, the screens are represented with two surfaces. Also, some large screens are being curved themselves. Thus, it remains a question of how should the large canvas of AR be designed, and how the objects should be placed. In this thesis, we made use of the canvas simply as a large wall-size display to flatly visualize different views side by side. But would it be better to make it a cylinder, or a set of planes stitched together in a cylindrical way, needs further investigation in the future.

Another point is that the desktop-GUIs, which are initially designed for 2D screens, are not liked in 3D space. Our decision is to reduce its usage in 3D. However, we could still think of if the widgets can be transitioned to 3D in a different form, not just turning a flat image into a textured cube. This is something that we do not know yet how to implement in practice. As for how input cursor and interaction might be transitioned from the PC to the AR space, we discuss it in the next section (Section 6.3).

We also studied the synchronization between the two spaces (screen and AR) to understand how to transition the action happened in one space to another. The way we proposed was to let users' decide whether the synchronization should be immediate or on request, depending on the specific application they are working with. The different envisioned scenario might also suggest that we should not limit the usage of the AR space to only one possibility (augmentation or extension). We should further consider how different spaces should be synchronized. For example, a 3D view linked to a graph on the screen would almost always require instant synchronization, while putting it aside on a storyboard might only need to be triggered on request. We believe that this also needs consideration depending on the specific application scenarios.

We did not study the visual effects of transitioning between the two spaces as investigated by Bogdan et al. [2014] because modern AR headsets handle the visualization independently from the PC screen so that users can observe the two spaces at the same time, without the necessity of transitioning one to another. However, if we want to use the AR to augment the 2D screen as we have discussed above, it might still be useful to understand how the visual animation of transiting between the two spaces help users' understanding.

6.3 INPUT AND INTERACTION

For **Q3** (What should be the appropriate input devices and interaction techniques to work with such a hybrid setting?), our basic decision is to use the mouse to control both parts as it comes naturally with the PC so as that we do not need to reply on special hardware.

In Chapter 3, experts expressed the unpleasant experience of using the mouse in AR while the mouse interaction remains efficient from the results of Chapter 4. The results of these two studies are not contradictory. The unpleasant experience in Chapter 3 was caused by the technical issues that led the mouse movement not very fluent in AR. Also, we represented the mouse cursor as a 3D object that caused some occlusion issues without surprise as previously studied [Schemali and Eisemann, 2014; Teather and Stuerzlinger, 2013] (in fact, all types of 3D cursors have limitations). But in Chapter 4, our focus was the interaction mapping between the input and output dimensionality and we intended to avoid using cursor-based interaction techniques (such as clicking on and dragging a thin handler on the object). Our results suggested the mouse remains efficient when we choose or design a suitable technique to use it. Our general recommendation is to try to map the relative movement of the mouse, in combination with the keyboard, to achieve a certain type of manipulation and to avoid motions that relying on its accurate absolute position in space. With such an idea, the most affected function of our prototype is to filter with a lasso tool. We thus reflect on other 3D filtering/selection techniques. For example, we can imagine a possible way is to turn the mouse cursor into a ray, and let users brushing in space to accomplish the selection, as presented by Hurter et al. [2019]. An advantage of this possibility is that brushing can also be supported on the PC side, without the need of defining two different interaction techniques.

However, we believe that a mouse cursor should still be kept in AR space because we need to, at least, select a certain view to manipulate. Then the question of how to better support a 3D cursor comes. We can envision, for example, the AR canvas is a plane extension of the screen, and makes the cursor stick on a fixed plane, and use ray-casting techniques to select and manipulate different objects. With such a design, the cursor in AR moves exactly like it is on a screen. However, one important problem is that views and objects in AR have a depth position and could be potentially be placed at different levels, a cursor stayed at a fixed-distance plane will be largely affected by the occlusion problems. Another possibility is to make the cursor movable in three directions, including the depth axis, with the help of a keyboard modifier as we have done in our study. We believe this is more appropriate to support tasks in 3D as it can specifically specify a correct position. Or even, it might also be possible to turn the cursor completely into a ray as usually done in VR. But how exactly each way could be useful in our specific application, and for general data exploration with our hybrid design requires further study. As for the transition of the mouse cursor from the screen to the AR, we still believe triggering the transition with a key is the most appropriate way. This is quite important, when we want to investigate how AR could be used to augment the screen because the two spaces overlay each other, it is thus important for users to manually specify which space they are interacting with.

Other than the mouse for accurate input, experts still wanted the system to support some more intuitive input for other tasks that do not need high accuracy. For example, even though they are aware that using mid-air gestures is inaccurate and could cause fatigue, they still suggested that they could serve as a complementary to the mouse, for occasional usage. Inspired by this, we think that a combination of input could be introduced in the future. Users would still perform most tasks with the mouse. Meanwhile, a few manipulations in AR that do not require high accuracy could be supported by simple gestures. For example, in 3D space, we can arrange the position of a specific view with hand gestures. By default, all views are kept at a fixed distance to the users. Users could, for example, use gestures to pull a specific view closer for detailed observation, and push it back later. This rough manipulation of views could also help make use of the large canvas of the AR by allowing users to flexibly adjusting their positions and distance to the users. Also, users could pull the data even in front of the 2D screen to examine it in detail. Another possible way of using hand gestures is to point at something to replace the mouse cursor, this is useful when we do not want to cursor to be transferred from the PC side. For example, we want to do something in AR at the moment waiting for an operation on PC to finish, or when we lose access to the mouse, such as when users are walking around to observe the data.

As reported in Chapter 3, walking around (or directly into) the data to observe it is a great advantage of using AR. Indeed, walking around offers faster interaction than manually adjusting the view, and the outcome of changes is just naturally expected. We think that the latter is a also huge advantage compared to other input modalities that we usually need to memorize the interaction mapping (for example, we need to think about if a translation of mouse would cause a translation or a rotation). We would naturally argue that walking around should be supported. However, walking around can sometimes be difficult in a working environment, facing similar difficulties as we have discussed the limitations in Section 6.1. Even so, experts would still be willing to move their heads around to observe the data, which suggests that it is really beneficial for them to understand the data. Thus we need to investigate, in the future, how this advantage could be better used in an office.

Another problem of walking around means losing control of the mouse and the 2D screen. We thus lose both the interaction and some background information on the screen. At this moment, using in-air gestures could partly replace the mouse, but still have difficulties if we need the control to be accurate. Inspired by previous studies, we proposed to use a mobile device and presented in Chapter 5, a pressure-augmented tactile mapping for 3D navigation that increased users' accuracy. The use of pressure combined with touch input also inspired us to combine different input modalities together to achieve a better user experience. We then believe that the combination of mouse and gestures while seated, and mobile and gestures while walking around would have great interests in the future.

Moreover, the mobile device may not only be served when users are walking around. Results from Chapter 4 suggested that for 3D manipulations that did not require high accuracy, high DOFs input is usually faster. Thus, making use of mobile's inertial sensors, it can also act as a tangible device for fast and intuitive control as the work proposed by Millette and McGuffin [2016]. We think that this could be served at the initial stage, when users just want to change the view to globally observe the data, before going into detailed analysis or manipulations.

To be mentioned, most of the interaction discussions in this thesis are yet limited to specifying a position and an orientation in space. In the future, it would be useful to investigate how to design the interaction techniques to support different data analysis tasks.

6.4 DATA EXPLORATION TASKS

Q4 (How can data exploration tasks be realized with a hybrid setting?) is the major remaining research question that needs to be further explored for this thesis. Even though we have known what experts would expect from the hybrid setting and what advantages do they see from it, we are still unclear how exactly a specific task could be supported. To better understand how can data exploration tasks be realized with such a setting, we are currently working to apply this hybrid setup to a real data exploration scenario. Specifically, we are working on a system to help particle physicists understand and compare the results of different reconstruction algorithms.

First, putting all the conclusions and guidelines mentioned in this chapter, a design of the hybrid visualization systems should have three major parts: a PC, an AR headset, and a mobile device. The PC part deals with all traditional analysis and manipulations that required accurate adjustments to benefit from the mouse as a precise input device. The AR part would be used to visualize spatial data in an immersive manner. To benefit from the large canvas, it should support visualizing and manipulating several different views. Interacting with the AR mainly relies on two ways: the mouse for precise adjustment like manipulating an object as presented in Chapter 4, and gestures for simple and rough, for example, re-arranging the position of a view when we want to closely observe one from several. However, we suggest that interaction with the mouse cursor in 3D should make more use of its relative movement rather than relying on its accurate absolute position in space. In the end, a tablet could be served when users want to walk around, and even more, as a high DOFs input for fast adjustment. As the tablet's screen is relatively small, we do not encourage replicating the PC screen, only keeping important background information.

Then, we want to make use of this setup to deal with a specific application scenario to better understand **Q4**. As mentioned in Chapter 1, a particle trajectory is composed of dozens of hit points detected after the collision. The first step is to run specific algorithms to reconstruct the measured points to tracks. For this process, particle physicists apply different machine learning algorithms which are of course not perfect. Thus, physicists need to analyze the results and compare them between different ones. Traditionally, they plot simple graphs like histograms. A major challenge is that plots of statistic results usually display only global information. It is hard for experts to understand how exactly a trajectory is constructed, either the interaction between different trajectories they are located close to each other. For example, at a certain area in space, the ratio of false construction arises, physicists would want to use our hybrid to understand why this is happening. We hope that with this work, we can better understand how our hybrid setup could support a realistic data exploration task.

Even though we closely collaborated with particle physicists and our main application domain is oriented to HEP, we believe the usage of our setup does not limit to HEP only. All scientific domains dealing with both abstract and spatial data could potentially benefit from the hybrid setup.

6.5 SUMMARY

We have demonstrated how a hybrid PC and AR setup could be used to improve the process of interactive exploration of 3D data, and discussed its interaction design from different perspectives. In conclusion, we have found that the hybrid setup is feasible and truly advantageous for 3D data exploration, despite the existing limitations of the hardware. We also understood that the dimensionality match between input and output devices is not critical while the mouse is still efficient for accurate control, which seems to confirm that the mouse could serve as the primary input for the hybrid setting.

Based on the findings of this thesis, several questions remain to be further discussed. The first is the specific application, with which we expect to better understand how experts would use such a system and how realistic data exploration tasks could be supported with the hybrid setup. The second is how to make better use of the advantages (such as walking around and the large canvas of AR) in an office scenario where the surrounding environments are usually not perfect. Also, it is important to investigate how to combine different input modalities to offer better interaction experience.

Even the universe of immersive visualization is far from being fully understood, we believe that the work presented in this thesis will inspire the creation of novel visualization systems and guide the interaction design to bring immersive visualization to existing scientific workflows, which will ultimately improve the user experience of exploring and understanding data to bring new scientific discoveries.

BIBLIOGRAPHY

Accot, Johnny and Shumin Zhai (1999). "Performance Evaluation of Input Devices in Trajectory-Based Tasks: An Application of the Steering Law." In: *Proc. CHI*. Pittsburgh, Pennsylvania, USA: ACM, pp. 466–472. ISBN: 0201485591. DOI: 10.1145/302979.303133.

Albinsson, Pär-Anders and Shumin Zhai (2003). "High Precision Touch Screen Interaction." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Ft. Lauderdale, Florida, USA: ACM, pp. 105–112. ISBN: 1581136307. DOI: 10.1145/642611.642631.

Alverson, George, Giulio Eulisse, Thomas McCauley, and Lucas Taylor (2012). "iSpy: A powerful and lightweight event display." In: *Journal of Physics: Conference Series* 396.2, p. 22002. DOI: 10.1088/1742-6596/396/2/022002.

Amrhein, Valentin, Fränzi Korner-Nievergelt, and Tobias Roth (2017). "The earth is flat: Significance thresholds and the crisis of unreplicable research." In: *PeerJ* 5. ISSN: 2167-9843. DOI: 10.7717/peerj.3544.

Amrhein, Valentin, David Trafimow, and Sander Greenland (2018). "Inferential Statistics as Descriptive Statistics: There is No Replication Crisis if We Don't Expect Replication." In: *The American Statistician*. DOI: 10.7287/peerj.preprints.26857v4.

Amrouche, Sabrina, Laurent Basara, Paolo Calafiura, Victor Estrade, Steven Farrell, Diogo R Ferreira, Liam Finnie, Nicole Finnie, Cécile Germain, Vladimir Vava Gligorov, and Others (2019). *The Tracking Machine Learning challenge: Accuracy phase*. arXiv preprint 1904.06778. URL: https://arxiv.org/abs/1904.06778.

Arif, Ahmed Sabbir, Ali Mazalek, and Wolfgang Stuerzlinger (2014). "The Use of Pseudo Pressure in Authenticating Smartphone Users." In: *Proceedings of MOBIQUITOUS*. Brussels: ICST, pp. 151–160. ISBN: 978-1-63190-039-6. DOI: 10.4108/icst.mobiquitous.2014.257919.

Arif, Ahmed Sabbir and Wolfgang Stuerzlinger (2013). "Pseudo-pressure Detection and Its Use in Predictive Text Entry on Touchscreens." In: *Proceedings of OzCHI*. New York: ACM, pp. 383–392. ISBN: 978-1-4503-2525-7. DOI: 10.1145/2541016.2541024.

Bach, Benjamin, 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?" In: *IEEE Transactions on Visualization and Computer Graphics* 24.1, pp. 457–467. DOI: 10.1109/TVCG.2017.2745941.

Baguley, Thomas (2012). *Serious stats: A guide to advanced statistics for the behavioral sciences*. Red Globe Press. ISBN: 978-0230577183.

Baker, Monya (2016). "Statisticians issue warning over misuse of P values." In: *Nature* 531.7593, p. 151. DOI: `10.1038/nature.2016.19503`.

Bellis, Matthew, Riccardo Maria Bianchi, Sebastien Binet, Ciril Bohak, Benjamin Couturier, Hadrien Grasland, Oliver Gutsche, Sergey Linev, Alex Martyniuk, Thomas McCauley, Edward Moyse, Alja Mrak Tadel, Mark Neubauer, Jeremi Niedziela, Leo Piilonen, Jim Pivarski, Martin Ritter, Tai Sakuma, Matevz Tadel, and Ben Waugh (2018). *HEP Software Foundation Community White Paper Working Group — Visualization*. arXiv preprint 1811.10309. URL: `https://arxiv.org/abs/1811.10309`.

Benko, Hrvoje, Edward W Ishak, and Steven Feiner (2004). "Collaborative mixed reality visualization of an archaeological excavation." In: *Proceedings of ISMAR*. Los Alamitos: IEEE Computer Society, pp. 132–140. DOI: `10.1109/ISMAR.2004.23`.

Benko, Hrvoje, Andrew D Wilson, and Patrick Baudisch (2006). "Precise Selection Techniques for Multi-touch Screens." In: *Proceedings of CHI*. New York: ACM, pp. 1263–1272. ISBN: 1-59593-372-7. DOI: `10.1145/1124772.1124963`.

Bérard, François, Jessica Ip, Mitchel Benovoy, Dalia El-Shimy, Jeffrey R Blum, and Jeremy R Cooperstock (2009). "Did Minority Report" Get It Wrong? Superiority of the Mouse over 3D Input Devices in a 3D Placement Task." In: *Proceedings of INTERACT*. Berlin/Heidelberg: Springer, pp. 400–414. DOI: `10.1007/978-3-642-03658-3_45`.

Besançon, Lonni (2018). "An Interaction Continuum for 3D Dataset Visualization." Theses. Université Paris-Saclay. URL: `https://tel.archives-ouvertes.fr/tel-01684210`.

Besançon, Lonni, Mehdi Ammi, and Tobias Isenberg (2017a). "Pressure-Based Gain Factor Control for Mobile 3D Interaction using Locally-Coupled Devices." In: *Proceedings of CHI*. New York: ACM, pp. 1831–1842. DOI: `10.1145/3025453.3025890`.

Besançon, Lonni and Pierre Dragicevic (2017). "The Significant Difference between p-values and Confidence Intervals." In: *29ème conférence francophone sur l'Interaction Homme-Machine*. Ed. by AFIHM. AFIHM. Poitiers, France, p. 10. URL: `https://hal.inria.fr/hal-01562281`.

— (2019). "The Continued Prevalence of Dichotomous Inferences at CHI." In: *CHI Extended Abstracts*. New York: ACM. DOI: `10.1145/3290607.3310432`.

Besançon, Lonni, Paul Issartel, Mehdi Ammi, and Tobias Isenberg (2017b). "Hybrid Tactile/Tangible Interaction for 3D Data Exploration." In: *IEEE Transactions on Visualization and Computer Graphics* 23.1, pp. 881–890. DOI: `10.1109/TVCG.2016.2599217`.

— (2017c). “Mouse, Tactile, and Tangible Input for 3D Manipulation.” In: *Proceedings of CHI*. New York: ACM, pp. 4727–4740. DOI: 10.1145/3025453.3025863.

Besançon, Lonni, Mickael Sereno, Mehdi Ammi, Lingyun Yu, and Tobias Isenberg (2019). “Hybrid Touch/Tangible Spatial 3D Data Selection.” In: *Computer Graphics Forum* 38.3, pp. 553–567. DOI: 10.1111/cgf.13710.

Besançon, Lonni, Anders Ynnerman, Daniel F. Keefe, Lingyun Yu, and Tobias Isenberg (2021). “The State of the Art of Spatial Interfaces for 3D Visualization.” In: *Computer Graphics Forum* 40. To appear. DOI: 10.1111/cgf.14189.

Bezerianos, Anastasia and Petra Isenberg (2012). “Perception of Visual Variables on Tiled Wall-Sized Displays for Information Visualization Applications.” In: *IEEE Transactions on Visualization and Computer Graphics* 18.12, pp. 2516–2525. ISSN: 1077-2626. DOI: 10.1109/TVCG.2012.251.

Blascheck, Tanja, Lonni Besançon, Anastasia Bezerianos, Bongshin Lee, and Petra Isenberg (2019). “Glanceable Visualization: Studies of Data Comparison Performance on Smartwatches.” In: *IEEE Transactions on Visualization and Computer Graphics* 25.1, pp. 630–640. DOI: 10.1109/TVCG.2018.2865142.

Blascheck, Tanja, Anastasia Bezerianos, Lonni Besançon, Bongshin Lee, and Petra Isenberg (2018). “Preparing for Perceptual Studies: Position and Orientation of Wrist-worn Smartwatches for Reading Tasks.” In: *Proceedings of CHI Workshop on Data Visualization on Mobile Devices*. URL: https://hal.inria.fr/hal-01744246.

Bogdan, Natalia, Tovi Grossman, and George Fitzmaurice (2014). “HybridSpace: Integrating 3D freehand input and stereo viewing into traditional desktop applications.” In: *2014 IEEE Symposium on 3D User Interfaces (3DUI)*. IEEE. Minneapolis, MN, USA, pp. 51–58. ISBN: 978-1-4799-3624-3. DOI: 10.1109/3DUI.2014.6798842.

Boring, Sebastian, David Ledo, Xiang ‘Anthony’ Chen, Nicolai Marquardt, Anthony Tang, and Saul Greenberg (2012). “The Fat Thumb: Using the Thumb’s Contact Size for Single-handed Mobile Interaction.” In: *Proceedings of MobileHCI*. New York: ACM, pp. 39–48. ISBN: 978-1-4503-1105-2. DOI: 10.1145/2371574.2371582.

Bornik, A, R Beichel, E Kruijff, B Reitinger, and D Schmalstieg (2006). “A Hybrid User Interface for Manipulation of Volumetric Medical Data.” In: *Proceedings of 3DUI*. Los Alamitos: IEEE Computer Society, pp. 29–36. DOI: 10.1109/VR.2006.8.

Brath, Richard (2014). “3D InfoVis is here to stay: Deal with it.” In: *Proceedings of 3DVis*. Los Alamitos: IEEE Computer Society, pp. 25–31. DOI: 10.1109/3DVis.2014.7160096.

Brewster, Stephen A and Michael Hughes (2009). “Pressure-based Text Entry for Mobile Devices.” In: *Proceedings of MobileHCI*. New York:

ACM, 9:1–9:4. ISBN: 978-1-60558-281-8. DOI: 10.1145/1613858.1613870.

Bruckner, Stefan, Tobias Isenberg, Timo Ropinski, and Alexander Wiebel (2019). "A Model of Spatial Directness in Interactive Visualization." In: *IEEE Transactions on Visualization and Computer Graphics* 25.8, pp. 2514–2528. DOI: 10.1109/TVCG.2018.2848906.

Brudy, Frederik, Christian Holz, Roman Rädle, Chi-Jui Wu, Steven Houben, Clemens Nylandsted Klokmose, and Nicolai Marquardt (2019). "Cross-Device Taxonomy: Survey, Opportunities and Challenges of Interactions Spanning Across Multiple Devices." In: *Proc. CHI*. New York: ACM, 562:1–562:28. ISBN: 978-1-4503-5970-2. DOI: 10.1145/3290605.3300792.

Bryson, Steve (1996). "Virtual Reality in Scientific Visualization." In: *Communications of the ACM* 39.5, pp. 62–71. ISSN: 0001-0782. DOI: 10.1145/229459.229467.

Burgess, Thomas D G and Stephen M Sales (1971). "Attitudinal effects of "mere exposure": A reevaluation." In: *Journal of Experimental Social Psychology* 7.4, pp. 461–472. ISSN: 0022-1031. DOI: 10.1016/0022-1031(71)90078-3.

Büschel, Wolfgang, Jian Chen, Raimund Dachselt, Steven Drucker, Tim Dwyer, Carsten Görg, Tobias Isenberg, Andreas Kerren, Chris North, and Wolfgang Stuerzlinger (2018). "Interaction for Immersive Analytics." In: *Immersive Analytics*. Ed. by Kim Marriott, Falk Schreiber, Tim Dwyer, Karsten Klein, Nathalie Henry Riche, Takayuki Itoh, Wolfgang Stuerzlinger, and Bruce H Thomas. Berlin Heidelberg: Springer. Chap. 4, pp. 95–138. DOI: 10.1007/978-3-030-01388-2_4.

Büschel, Wolfgang, Anke Lehmann, and Raimund Dachselt (May 2021). "MIRIA: A Mixed Reality Toolkit for the In-Situ Visualization and Analysis of Spatio-Temporal Interaction Data." In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI '21. Yokohama, Japan: ACM. ISBN: 978-1-4503-8096-6/21/05. DOI: 10.1145/3411764.3445651.

Büschel, Wolfgang, Annett Mitschick, Thomas Meyer, and Raimund Dachselt (2019). "Investigating Smartphone-based Pan and Zoom in 3D Data Spaces in Augmented Reality." In: *Proceedings of MobileHCI*.

Büschel, Wolfgang, Patrick Reipschläger, Ricardo Langner, and Raimund Dachselt (2017). "Investigating the Use of Spatial Interaction for 3D Data Visualization on Mobile Devices." In: *Proceedings of the 2017 ACM International Conference on Interactive Surfaces and Spaces*. New York: ACM, pp. 62–71. ISBN: 978-1-4503-4691-7. DOI: 10.1145/3132272.3134125.

Buxton, Bill (2007). *Multi-Touch Systems that I Have Known and Loved*. Web site: http://www.billbuxton.com/multitouchOverview.html.

Buxton, William (1986). “Chunking and Phrasing and the Design of Human-Computer Dialogues.” In: *Proceedings of IFIP World Computer Congress*, pp. 475–480. DOI: `10.1016/B978-0-08-051574-8.50051-0`.

Carpendale, Sheelagh (2008). “Evaluating Information Visualizations.” In: *Information Visualization: Human-Centered Issues and Perspectives*. Vol. 4950. Berlin/Heidelberg: Springer, pp. 19–45. DOI: `10.1007/978-3-540-70956-5_2`.

Chapuis, Olivier, Anastasia Bezerianos, and Stelios Frantzeskakis (2014). “Smarties: An Input System for Wall Display Development.” In: *Proceedings of the 32nd international conference on Human factors in computing systems*. Toronto, CA: ACM, pp. 2763–2772. DOI: `10.1145/2556288.2556956`.

Chen, Michael, S Joy Mountford, and Abigail Sellen (1988). “A study in interactive 3D rotation using 2D control devices.” In: *ACM SIGGRAPH Computer Graphics* 22.4, pp. 121–129. ISSN: 00978930. DOI: `10.1145/378456.378497`.

Chen, Xiang ‘Anthony’, Julia Schwarz, Chris Harrison, Jennifer Mankoff, and Scott E Hudson (2014). “Air+Touch: Interweaving Touch & In-air Gestures.” In: *Proceedings of UIST*. New York: ACM, pp. 519–525. ISBN: 978-1-4503-3069-5. DOI: `10.1145/2642918.2647392`.

Cockburn, Andy, Pierre Dragicevic, Lonni Besançon, and Carl Gutwin (2020). “Threats of a replication crisis in empirical computer science.” In: *Communications of the ACM* 63.8, pp. 70–79. DOI: `10.1145/3360311`.

Cockburn, Andy, Carl Gutwin, and Alan Dix (2018). “HARK No More: On the Preregistration of CHI Experiments.” In: *Proceedings of CHI*. New York: ACM, 141:1–141:12. DOI: `10.1145/3173574.3173715`.

Coffey, Dane, Nicholas Malbraaten, Trung Le, Iman Borazjani, Fotis Sotiropoulos, Arthur G Erdman, and Daniel F Keefe (2012). “Interactive Slice WIM: Navigating and Interrogating Volume Data Sets Using a Multisurface, Multitouch VR Interface.” In: *IEEE Transactions on Visualization and Computer Graphics* 18.10, pp. 1614–1626. ISSN: 1077-2626. DOI: `10.1109/TVCG.2011.283`.

Cohé, Aurélie, Fabrice Dècle, and Martin Hachet (2011). “tBox: A 3D Transformation Widget Designed for Touch-screens.” In: *Proceedings of CHI*. New York: ACM, pp. 3005–3008. ISBN: 978-1-4503-0228-9. DOI: `10.1145/1978942.1979387`.

Cohen, Oryx, Shawna Meyer, and Erik Nilsen (1993). “Studying the Movement of High-Tech Rodentia: Pointing and Dragging.” In: *Proc. CHI*. Amsterdam, The Netherlands: ACM, pp. 135–136. ISBN: 0897915747. DOI: `10.1145/259964.260149`.

Conner, D., Scott Snibbe, Kenneth Herndon, Daniel Robbins, Robert Zeleznik, and Andries van Dam (Jan. 1992). “Three-Dimensional Widgets.” In: *Computer Graphics (1992 Symposium on Interactive 3D Graphics)* 25. DOI: `10.1145/147156.147199`.

Corsten, Christian, Bjoern Daehlmann, Simon Voelker, and Jan Borchers (2017). "BackXPress: Using Back-of-Device Finger Pressure to Augment Touchscreen Input on Smartphones." In: *Proceedings of CHI*. New York: ACM, pp. 4654–4666. ISBN: 978-1-4503-4655-9. DOI: 10.1145/3025453.3025565.

Cruz-Neira, Carolina, Jason Leigh, Michael Papka, Craig Barnes, S.M. Cohen, S. Das, R. Engelmann, R. Hudson, Trina Roy, L. Siegel, C. Vasilakis, T.A. DeFanti, and Daniel Sandin (Nov. 1993a). "Scientists in wonderland: A report on visualization applications in the CAVE virtual reality environment." In: *Proceedings of 1993 IEEE Research Properties in Virtual Reality Symposium*, pp. 59 –66. ISBN: 0-8186-4910-0. DOI: 10.1109/VRAIS.1993.378262.

Cruz-Neira, Carolina, Daniel J Sandin, and Thomas A DeFanti (1993b). "Surround-screen Projection-based Virtual Reality: The Design and Implementation of the CAVE." In: *Proceedings of SIGGRAPH*. New York: ACM, pp. 135–142. ISBN: 0-89791-601-8. DOI: 10.1145/166117.166134.

Cruz-Neira, Carolina, Daniel J Sandin, Thomas A DeFanti, Robert V Kenyon, and John C Hart (1992). "The CAVE: Audio Visual Experience Automatic Virtual Environment." In: *Communications of the ACM* 35.6, pp. 64–72. DOI: 10.1145/129888.129892.

Cumming, Geoff (2012). *Understanding The New Statistics: Effect Sizes, Confidence Intervals and Meta-Analysis*. New York: Routledge Taylor & Francis Group. ISBN: 978-0415879682. DOI: 10.4324/9780203807002.

— (2013). *Introduction to Statistics 9: Dance of the p Values*. YouTube video: https://www.youtube.com/watch?v=5OL1RqHrZQ8. URL: https://www.youtube.com/watch?v=5OL1RqHrZQ8.

— (2014). "The new statistics: Why and how." In: *Psychological Science* 25.1, pp. 7–29. ISSN: 1467-9280. DOI: 10.1177/0956797613504966.

Dang, Nguyen Thong, Monica Tavanti, Ivan Rankin, and Matthew Cooper (2009). "A comparison of different input devices for a 3D environment." In: *International Journal of Industrial Ergonomics* 39.3. Selected papers from ECCE 2007, the 25th Anniversary Conference of the European Conference on Cognitive Ergonomics, pp. 554 –563. ISSN: 0169-8141. DOI: 10.1016/j.ergon.2008.10.016.

Deber, Jonathan, Ricardo Jota, Clifton Forlines, and Daniel Wigdor (2015). "How Much Faster is Fast Enough?: User Perception of Latency & Latency Improvements in Direct and Indirect Touch." In: *Proceedings of CHI*. New York: ACM, pp. 1827–1836. ISBN: 978-1-4503-3145-6. DOI: 10.1145/2702123.2702300.

DeFanti, Thomas A., Gregory Dawe, Daniel J. Sandin, Jurgen P. Schulze, Peter Otto, Javier Girado, Falko Kuester, Larry Smarr, and Ramesh Rao (2009). "The StarCAVE, a Third-Generation CAVE and Virtual Reality OptIPortal." In: *Future Gener. Comput. Syst.* 25.2, 169–178. ISSN: 0167-739X. DOI: 10.1016/j.future.2008.07.015.

Dell, Nicola, Vidya Vaidyanathan, Indrani Medhi, Edward Cutrell, and William Thies (2012). ""Yours is better!" Participant response bias in HCI." In: *Proceedings of CHI*. New York: ACM, pp. 1321–1330. ISBN: 9781450310154. DOI: 10.1145/2207676.2208589.

Demiralp, C., C. D. Jackson, D. B. Karelitz, S. Zhang, and D. H. Laidlaw (2006). "CAVE and fishtank virtual-reality displays: a qualitative and quantitative comparison." In: *IEEE Transactions on Visualization and Computer Graphics* 12.3, pp. 323–330. DOI: 10.1109/TVCG.2006.42.

Dixon, Peter (2003). "The p-value fallacy and how to avoid it." In: *Canadian Journal of Experimental Psychology* 57.3, pp. 189–202. DOI: 10.1037/h0087425.

Dragicevic, Pierre (2015). *HCI Statistics without p-values*. Research Report RR-8738. Inria, p. 32. URL: https://hal.inria.fr/hal-01162238.

— (2016). "Fair Statistical Communication in HCI." In: *Modern Statistical Methods for HCI*. Ed. by Judy Robertson and Maurits Kaptein. Cham, Switzerland: Springer International Publishing. Chap. 13, pp. 291–330. DOI: 10.1007/978-3-319-26633-6_13.

Dragicevic, Pierre, Fanny Chevalier, and Stéphane Huot (2014). "Running an HCI Experiment in Multiple Parallel Universes." In: *CHI Extended Abstracts*. New York: ACM, pp. 607–618. DOI: http://dx.doi.org/10.1145/2559206.2578881.

Duer, Zach, Leo Piilonen, and George Glasson (2018). "Belle2VR: A Virtual-Reality Visualization of Subatomic Particle Physics in the Belle II Experiment." In: *IEEE Computer Graphics and Applications* 38, pp. 33–43. DOI: 10.1109/MCG.2018.032421652.

Durbin, Jim, J. Edward Swan, Brad Colbert, John Crowe, Rob King, Tony King, Christopher Scannell, Zachary Wartell, and Terry Welsh (1998). "Battlefield Visualization on the Responsive Workbench." In: *Proceedings of the Conference on Visualization '98*. Research Triangle Park, North Carolina, USA: IEEE Computer Society Press, 463–466. ISBN: 1581131062. DOI: 10.1109/VISUAL.1998.745344.

Dwyer, Tim, Kim Marriott, Tobias Isenberg, Karsten Klein, Nathalie Riche, Falk Schreiber, Wolfgang Stuerzlinger, and Bruce Thomas (2018). "Immersive Analytics: An Introduction." In: *Immersive Analytics*. Ed. by Kim Marriott, Falk Schreiber, Tim Dwyer, Karsten Klein, Nathalie Henry Riche, Takayuki Itoh, Wolfgang Stuerzlinger, and Bruce H Thomas. Berlin Heidelberg: Springer. Chap. 1, pp. 1–23. DOI: 10.1007/978-3-030-01388-2_1.

Everts, Maarten H., Eric Begue, Henk Bekker, Jos B. T. M. Roerdink, and Tobias Isenberg (July 2015). "Exploration of the Brain's White Matter Structure through Visual Abstraction and Multi-Scale Local Fiber Tract Contraction." In: *IEEE Transactions on Visualization and*

Computer Graphics 21.7, pp. 808–821. DOI: 10.1109/TVCG.2015.2403323.

Fallot, W., M. Fjeld, C. Speirs, S. Ziegenspeck, H. Krueger, and T. Läubli (2006). "Touch&Type: A Novel Pointing Device for Notebook Computers." In: *Proc. NordiCHI*. Oslo, Norway: ACM, pp. 465–468. ISBN: 1595933255. DOI: 10.1145/1182475.1182538.

Febretti, Alessandro, Arthur Nishimoto, Terrance Thigpen, Jonas Talandis, Lance Long, J D Pirtle, Tom Peterka, Alan Verlo, Maxine Brown, Dana Plepys, and Others (2013). "CAVE2: A hybrid reality environment for immersive simulation and information analysis." In: *The Engineering Reality of Virtual Reality*. Vol. 8649. Bellingham, Washington: SPIE, Article 864903, 12 pages. DOI: 10.1117/12.2005484.

Filho, Jorge A Wagner, Carla M D S Freitas, and Luciana Nedel (2019). "Comfortable Immersive Analytics With the VirtualDesk Metaphor." In: *IEEE Computer Graphics and Applications* 39.3, pp. 41–53. ISSN: 0272-1716. DOI: 10.1109/MCG.2019.2898856.

Fonnet, A. and Y. Prié (2019). "Survey of Immersive Analytics." In: *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–1. DOI: 10.1109/TVCG.2019.2929033.

Forlines, Clifton, Daniel Wigdor, Chia Shen, and Ravin Balakrishnan (2007). "Direct-Touch vs. Mouse Input for Tabletop Displays." In: *Proceedings of CHI*. New York: ACM, pp. 647–656. ISBN: 9781595935939. DOI: 10.1145/1240624.1240726.

Forsberg, Andrew, Jian Chen, and David Laidlaw (2009). "Comparing 3D vector field visualization methods: A user study." In: *IEEE Transactions on Visualization and Computer Graphics* 15.6, pp. 1219–1226. ISSN: 1077-2626. DOI: 10.1109/TVCG.2009.126.

Fröhlich, Bernd, Stephen Barrass, Björn Zehner, John Plate, and Martin Göbel (1999). "Exploring Geo-scientific Data in Virtual Environments." In: *Proceedings of Visualization*. Los Alamitos: IEEE Computer Society, pp. 169–173. ISBN: 0-7803-5897-X. DOI: 10.1109/VISUAL.1999.809884.

Fruchard, Bruno, Arnaud Prouzeau, Olivier Chapuis, and Eric Lecolinet (2019). "Leveraging Body Interactions to Support Immersive Analytics." In: *CHI Workshop on Immersive Analytics*. Glasgow, United Kingdom, 10 pages. URL: https://hal.archives-ouvertes.fr/hal-02095993.

Fu, Chi-Wing, Wooi-Boon Goh, and Junxiang Allen Ng (2010). "Multi-Touch Techniques for Exploring Large-Scale 3D Astrophysical Simulations." In: *Proceedings of CHI*. New York: ACM, pp. 2213–2222. DOI: 10.1145/1753326.1753661.

Fulmer, Willis, Tahir Mahmood, Zhongyu Li, Shaoting Zhang, Jian Huang, and Aidong Lu (2019). "ImWeb: Cross-platform Immersive Web Browsing for Online 3D Neuron Database Exploration." In:

Proceedings of IUI. New York: ACM, pp. 367–378. DOI: 10.1145/3301275.3302319.

Gelman, Andrew (2017). *No to inferential thresholds*. Online. Last visited 04 January 2019. URL: https://andrewgelman.com/2017/11/19/no-inferential-thresholds/.

Gigerenzer, Gerd (2004). "Mindless statistics." In: *The Journal of Socio-Economics* 33.5, pp. 587–606. DOI: 10.1016/j.socec.2004.09.033.

— (2018). "Statistical Rituals: The Replication Delusion and How We Got There." In: *Advances in Methods and Practices in Psychological Science* 1.2, pp. 198–218. DOI: 10.1177/2515245918771329.

Glesser, David, François Bérard, and Jeremy R Cooperstock (2013). "Overcoming limitations of the trackpad for 3D docking operations." In: *CHI Extended Abstracts*. ACM. New York, pp. 1239–1244. DOI: 10.1145/2468356.2468578.

Goel, Mayank, Jacob Wobbrock, and Shwetak Patel (2012). "GripSense: Using Built-in Sensors to Detect Hand Posture and Pressure on Commodity Mobile Phones." In: *Proceedings of UIST*. New York: ACM, pp. 545–554. ISBN: 978-1-4503-1580-7. DOI: 10.1145/2380116.2380184.

Goodman, Steven N (1999). "Toward evidence-based medical statistics. 1: The P value fallacy." In: *Annals of Internal Medicine* 130.12, pp. 995–1004. DOI: 10.7326/0003-4819-130-12-199906150-00008.

Gröller, Eduard (2002). "Insight into Data through Visualization." In: *Graph Drawing*. Ed. by Petra Mutzel, Michael Jünger, and Sebastian Leipert. Berlin/Heidelberg: Springer Berlin Heidelberg, pp. 352–366. ISBN: 978-3-540-45848-7. DOI: 10.1007/3-540-45848-4_28.

Grubert, Jens, Lukas Witzani, Eyal Ofek, Michel Pahud, Matthias Kranz, and Per Ola Kristensson (2018). *Text Entry in Immersive Head-Mounted Display-based Virtual Reality using Standard Keyboards*. arXiv preprint 1802.00626. URL: http://arxiv.org/abs/1802.00626.

Guillaume, Moreau (Nov. 2014). "Visual Immersion Issues in Virtual Reality: A Survey." In: pp. 6–14. DOI: 10.1109/SIBGRAPI-T.2013.9.

Haan, Gerwin de, Michal Koutek, and Frits H Post (2002). "Towards Intuitive Exploration Tools for Data Visualization in VR." In: *Proceedings of Virtual Reality Software and Technology*. New York: ACM, pp. 105–112. ISBN: 1-58113-530-0. DOI: 10.1145/585740.585758.

Hancock, Mark, Sheelagh Carpendale, and Andy Cockburn (2007). "Shallow-depth 3D Interaction: Design and Evaluation of One-, Two- and Three-touch Techniques." In: *Proceedings of CHI*. New York: ACM, pp. 1147–1156. ISBN: 978-1-59593-593-9. DOI: 10.1145/1240624.1240798.

Hancock, Mark, Thomas ten Cate, and Sheelagh Carpendale (2009). "Sticky Tools: Full 6DOF Force-based Interaction for Multi-touch

Tables." In: *Proceedings of ITS*. New York: ACM, pp. 133–140. ISBN: 978-1-60558-733-2. DOI: `10.1145/1731903.1731930`.

Hancock, Mark S, Sheelagh Carpendale, Frederic D Vernier, Daniel Wigdor, and Chia Shen (2006). "Rotation and Translation Mechanisms for Tabletop Interaction." In: *Proceedings of TABLETOP*. Los Alamitos: IEEE Computer Society, pp. 79–88. ISBN: 0-7695-2494-X. DOI: `10.1109/TABLETOP.2006.26`.

Hanula, P., K. Piekutowski, C. Uribe, K. Almryde, A. Nishimoto, J. Aguilera, and G. E. Marai (2015). "Cavern Halos: Exploring spatial and nonspatial cosmological data in an immersive virtual environment." In: *2015 IEEE Scientific Visualization Conference (SciVis)*, pp. 87–99. DOI: `10.1109/SciVis.2015.7429497`.

Hart, Sandra G (2006). "NASA-Task Load Index (NASA-TLX) 20 Years Later." In: *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 50.9, pp. 904–908. ISSN: 1071-1813. DOI: `10.1177/154193120605000909`.

Heo, Seongkook and Geehyuk Lee (2011a). "Force Gestures: Augmented Touch Screen Gestures Using Normal and Tangential Force." In: *CHI Extended Abstracts*. New York: ACM, pp. 1909–1914. ISBN: 978-1-4503-0268-5. DOI: `10.1145/1979742.1979895`.

— (2011b). "Forcetap: Extending the Input Vocabulary of Mobile Touch Screens by Adding Tap Gestures." In: *Proceedings of MobileHCI*. New York: ACM, pp. 113–122. ISBN: 978-1-4503-0541-9. DOI: `10.1145/2037373.2037393`.

— (2012). "ForceDrag: Using Pressure As a Touch Input Modifier." In: *Proceedings of OzCHI*. New York: ACM, pp. 204–207. ISBN: 978-1-4503-1438-1. DOI: `10.1145/2414536.2414572`.

Herot, Christopher F and Guy Weinzapfel (1978). "One-point touch input of vector information for computer displays." In: *ACM SIGGRAPH Computer Graphics* 12.3, pp. 210–216. DOI: `10.1145/965139.807392`.

Hibbard, Bill (1999). "Top Ten Visualization Problems." In: *ACM SIGGRAPH Computer Graphics* 33.2, pp. 21–22. ISSN: 0097-8930. DOI: `10.1145/326460.326485`.

Hincapié-Ramos, Juan David, Xiang Guo, Paymahn Moghadasian, and Pourang Irani (2014). "Consumed Endurance: A Metric to Quantify Arm Fatigue of Mid-air Interactions." In: *Proceedings of CHI*. New York: ACM, pp. 1063–1072. ISBN: 978-1-4503-2473-1. DOI: `10.1145/2556288.2557130`.

Hinckley, Ken, Francois Guimbretiere, Patrick Baudisch, Raman Sarin, Maneesh Agrawala, and Ed Cutrell (2006). "The Springboard: Multiple Modes in One Spring-loaded Control." In: *Proceedings of CHI*. New York: ACM, pp. 181–190. ISBN: 1-59593-372-7. DOI: `10.1145/1124772.1124801`.

Hinckley, Ken, Seongkook Heo, Michel Pahud, Christian Holz, Hrvoje Benko, Abigail Sellen, Richard Banks, Kenton O'Hara, Gavin

Smyth, and William Buxton (2016). “Pre-Touch Sensing for Mobile Interaction.” In: *Proceedings of CHI*. New York: ACM, pp. 2869–2881. ISBN: 978-1-4503-3362-7. DOI: `10.1145/2858036.2858095`.

Hinckley, Ken, Randy Pausch, John C Goble, and Neal F Kassell (1994). “Passive Real-World Interface Props for Neurosurgical Visualization.” In: *Proceedings of CHI*. New York: ACM, pp. 452–458. DOI: `10.1145/191666.191821`.

Hinckley, Ken, Joe Tullio, Randy Pausch, Dennis Proffitt, and Neal Kassell (1997). “Usability Analysis of 3D Rotation Techniques.” In: *Proc. UIST*. Banff, Alberta, Canada: ACM, pp. 1–10. ISBN: 0897918819. DOI: `10.1145/263407.263408`.

Hurter, Christophe, Nathalie Henry Riche, Steven M Drucker, Maxime Cordeil, Richard Alligier, and Romain Vuillemot (2019). “FiberClay: Sculpting Three Dimensional Trajectories to Reveal Structural Insights.” In: *IEEE Transactions on Visualization and Computer Graphics* 25.1, pp. 704–714. DOI: `10.1109/TVCG.2018.2865191`.

Hwang, Sungjae, Andrea Bianchi, and Kwang-yun Wohn (2013). “VibPress: Estimating Pressure Input Using Vibration Absorption on Mobile Devices.” In: *Proceedings of MobileHCI*. New York: ACM, pp. 31–34. ISBN: 978-1-4503-2273-7. DOI: `10.1145/2493190.2493193`.

Isenberg, Tobias (2014). “An Interaction Continuum for Visualization.” In: *Proceedings of the VIS Workshop on “Death of the Desktop: Envisioning Visualization without Desktop Computing” (November 9, Paris, France)*. Ed. by Yvonne Jansen, Petra Isenberg, Jason Dykes, Sheelagh Carpendale, and Dan Keefe. URL: `https://hal.inria.fr/hal-01095454`.

— (2016). “Interactive Exploration of Three-Dimensional Scientific Visualizations on Large Display Surfaces.” In: *Collaboration Meets Interactive Spaces*. Ed. by Craig Anslow, Pedro Campos, and Joaquim Jorge. Berlin/Heidelberg: Springer. Chap. 6, pp. 97–123. DOI: `10.1007/978-3-319-45853-3_6`.

Isenberg, Tobias and Mark Hancock (2012). “Gestures vs. Postures: Gestural’ Touch Interaction in 3D Environments.” In: *Proceedings of 3DCHI*, pp. 53–61. URL: `https://hal.inria.fr/hal-00781237`.

Ishii, Hiroshi (2008). “The Tangible User Interface and Its Evolution.” In: *Commun. ACM* 51.6, pp. 32–36. ISSN: 0001-0782. DOI: `10.1145/1349026.1349034`.

Issartel, Paul, Lonni Besançon, Florimond Guéniat, Tobias Isenberg, and Mehdi Ammi (2016). “Preference Between Allocentric and Egocentric 3D Manipulation in a Locally Coupled Configuration.” In: *Proceedings of Spatial User Interaction*. New York: ACM, pp. 79–88. DOI: `10.1145/2983310.2985750`.

Jacob, Robert, John Leggett, Brad Myers, and Randy Pausch (Mar. 1993). “Interaction styles and input/output devices.” In: *Behaviour & Information Technology* 12.2, pp. 69–79. DOI: `10.1080/0144\-92993\-08924\-369`.

Jankowski, Jacek and Martin Hachet (2013). “A Survey of Interaction Techniques for Interactive 3D Environments.” In: *Eurographics – STARs*. Goslar, Germany: The Eurographics Association, pp. 65–93. DOI: `10.2312/conf/EG2013/stars/065-093`.

Jansen, Yvonne, Pierre Dragicevic, and Jean-Daniel Fekete (2012). “Tangible Remote Controllers for Wall-size Displays.” In: *Proceedings of CHI*. New York: ACM, pp. 2865–2874. ISBN: 978-1-4503-1015-4. DOI: `10.1145/2207676.2208691`.

— (2013). “Evaluating the Efficiency of Physical Visualizations.” In: *Proceedings of CHI*. New York: ACM, pp. 2593–2602. DOI: `10.1145/2470654.2481359`.

Johnson, Christopher R (2004). “Top scientific visualization research problems.” In: *IEEE Computer Graphics and Applications* 24.4, pp. 13–17. DOI: `10.1109/MCG.2004.20`.

Jordà, Sergi, Günter Geiger, Marcos Alonso, and Martin Kaltenbrunner (2007). “The ReacTable: Exploring the Synergy between Live Music Performance and Tabletop Tangible Interfaces.” In: *Proceedings of the 1st International Conference on Tangible and Embedded Interaction*. Baton Rouge, Louisiana: Association for Computing Machinery, 139–146. ISBN: 9781595936196. DOI: `10.1145/1226969.1226998`.

Keefe, Daniel F (2008). “Free-Form VR Interactions in Scientific Visualization.” In: *Proceedings of IEEE Visualization Workshop on Scientific Workflow with Immersive Interfaces for Visualization*. URL: `http://ivlab.cs.umn.edu/generated/pub-Keefe-2008-FreeFormSciVis.php`.

— (2010). “Integrating Visualization and Interaction Research to Improve Scientific Workflows.” In: *IEEE Computer Graphics and Applications* 30.2, pp. 8–13. DOI: `10.1109/MCG.2010.30`.

Keefe, Daniel F and Tobias Isenberg (2013). “Reimagining the Scientific Visualization Interaction Paradigm.” In: *IEEE Computer* 46.5, pp. 51–57. DOI: `10.1109/MC.2013.178`.

Keene, Oliver N (1995). “The log transformation is special.” In: *Statistics in Medicine* 14.8, pp. 811–819. DOI: `10.1002/sim.4780140810`.

Kim, Kwangtaek, Joongrock Kim, Jaesung Choi, Junghyun Kim, and Sangyoun Lee (2015). “Depth Camera-Based 3D Hand Gesture Controls with Immersive Tactile Feedback for Natural Mid-Air Gesture Interactions.” In: *Sensors* 15.1, pp. 1022–1046. ISSN: 1424-8220. DOI: `10.3390/s150101022`.

Kirby, Kris N and Daniel Gerlanc (2013). “BootES: An R package for bootstrap confidence intervals on effect sizes.” In: *Behavior Research Methods* 45.4, pp. 905–927. ISSN: 1554-3528. DOI: `10.3758/s13428-013-0330-5`.

Kittelmann, Thomas, Vakhtang Tsulaia, Joseph Boudreau, and Edward Moyse (2010). “The Virtual Point 1 event display for the ATLAS experiment.” In: *Journal of Physics: Conference Series* 219, part 3. DOI: `10.1088/1742-6596/219/3/032012`.

Klamka, Konstantin and Raimund Dachselt (2015). “Elasticcon: Elastic Controllers for Casual Interaction.” In: *Proceedings of MobileHCI*. New York: ACM, pp. 410–419. DOI: 10.1145/2785830.2785849.

Klamka, Konstantin, Patrick Reipschläger, and Raimund Dachselt (2019). “CHARM: Cord-Based Haptic Augmented Reality Manipulation.” In: *Virtual, Augmented and Mixed Reality. Multimodal Interaction*. Cham, Switzerland: Springer International Publishing, pp. 96–114. DOI: 10.1007/978-3-030-21607-8_8.

Klein, Tijmen, Florimond Guéniat, Luc Pastur, Frédéric Vernier, and Tobias Isenberg (2012). “A Design Study of Direct-Touch Interaction for Exploratory 3D Scientific Visualization.” In: *Computer Graphics Forum* 31.3, pp. 1225–1234. DOI: 10.1111/j.1467-8659.2012.03115.x.

Knoedel, Sebastian and Martin Hachet (2011). “Multi-touch RST in 2D and 3D spaces: Studying the impact of directness on user performance.” In: *Proceedings of 3DUI*. IEEE Computer Society. Los Alamitos, pp. 75–78. DOI: 10.1109/3DUI.2011.5759220.

Krueger, Wolfgang and Bernd Froehlich (1994). “The Responsive Workbench.” In: *IEEE Computer Graphics and Applications* 14.3, pp. 12–15. DOI: 10.1109/38.279036.

Kruger, Russell, Sheelagh Carpendale, Stacey D Scott, and Anthony Tang (2005). “Fluid Integration of Rotation and Translation.” In: *Proceedings of CHI*. New York: ACM, pp. 601–610. ISBN: 1-58113-998-5. DOI: 10.1145/1054972.1055055.

Krzywinski, Martin and Naomi Altman (2013). “Points of Significance: Error bars.” In: *Nature Methods* 10.10, pp. 921–922. DOI: 10.1038/nmeth.2659.

Laha, Bireswar, Doug Bowman, and John Socha (2014). “Effects of VR System Fidelity on Analyzing Isosurface Visualization of Volume Datasets.” In: *IEEE Transactions on Visualization and Computer Graphics* 20.4, pp. 513–522. DOI: 10.1109/TVCG.2014.20.

Laha, Bireswar, Kriti Sensharma, James D Schiffbauer, and Doug A Bowman (2012). “Effects of Immersion on Visual Analysis of Volume Data.” In: *IEEE Transactions on Visualization and Computer Graphics* 18.4, pp. 597–606. ISSN: 1077-2626. DOI: 10.1109/TVCG.2012.42.

Lai, Jerry (2010). “Dichotomous Thinking: A Problem Beyond NHST.” In: *Data and context in statistics education: Towards an evidence based society*. Voorburg, the Netherlands: International Association for Statistical Education, C101:1–C101:4. URL: https://iase-web.org/documents/papers/icots8/ICOTS8C101LAI.pdf.

Langner, Ricardo, Tom Horak, and Raimund Dachselt (2017). “VisTiles: Coordinating and Combining Co-located Mobile Devices for Visual Data Exploration.” In: *IEEE Transactions on Visualization and Computer Graphics* 24.1, pp. 626–636. DOI: 10.1109/TVCG.2017.2744019.

LaViola Jr., Joseph J, Ernst Kruijff, Ryan P McMahan, Doug A Bowman, and Ivan Poupyrev (2017). *3D User Interfaces: Theory and Practice*. 2nd. Boston: Addison-Wesley. URL: `https://www.vitalsource.com/referral?term=9780134034461`.

Le Goc, Mathieu, Pierre Dragicevic, Samuel Huron, Jeremy Boy, and Jean-Daniel Fekete (2016). "A Better Grasp on Pictures Under Glass: Comparing Touch and Tangible Object Manipulation using Physical Proxies." In: *Proceedings of AVI*. New York: ACM. DOI: `10.1145/2909132.2909247`.

Lee, Jinha, Alex Olwal, Hiroshi Ishii, and Cati Boulanger (2013). "SpaceTop: Integrating 2D and Spatial 3D Interactions in a See-through Desktop Environment." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '13. Paris, France: Association for Computing Machinery, 189–192. ISBN: 9781450318990. DOI: `10.1145/2470654.2470680`.

Liu, Jingbo, Oscar Kin-Chung Au, Hongbo Fu, and Chiew-Lan Tai (2012). "Two-Finger Gestures for 6DOF Manipulation of 3D Objects." In: *Computer Graphics Forum* 31.7pt1, pp. 2047–2055. ISSN: 0167-7055. DOI: `10.1111/j.1467-8659.2012.03197.x`.

López, David, Lora Oehlberg, Candemir Doger, and Tobias Isenberg (2016). "Towards an Understanding of Mobile Touch Navigation in a Stereoscopic Viewing Environment for 3D Data Exploration." In: *IEEE Transactions on Visualization and Computer Graphics* 22.5, pp. 1616–1629. DOI: `10.1109/TVCG.2015.2440233`.

Lundström, Claes, Thomas Rydell, Camilla Forsell, Anders Persson, and Anders Ynnerman (2011). "Multi-Touch Table System for Medical Visualization: Application to Orthopedic Surgery Planning." In: *IEEE Transactions on Visualization and Computer Graphics* 17.12, pp. 1775–1784. DOI: `10.1109/TVCG.2011.224`.

MacKenzie, I Scott and Shawn X Zhang (1999). "The Design and Evaluation of a High-performance Soft Keyboard." In: *Proceedings of CHI*. New York: ACM, pp. 25–31. ISBN: 0-201-48559-1. DOI: `10.1145/302979.302983`.

Mahmood, Tahir, Erik Butler, Nicholas Davis, Jian Huang, and Aidong Lu (2018). "Building Multiple Coordinated Spaces for Effective Immersive Analytics through Distributed Cognition." In: *Proceedings of BDVA*. Los Alamitos: IEEE Computer Society, pp. 119–129. DOI: `10.1109/BDVA.2018.8533893`.

Maiero, Jens, David Eibich, Ernst Kruijff, André Hinkenjann, Wolfgang Stuerzlinger, Hrvoje Benko, and Gheorghita Ghinea (Oct. 2019). "Back-of-Device Force Feedback Improves Touchscreen Interaction for Mobile Devices." In: *EEE Trans. Haptics* 12.4, 483–496. ISSN: 1939-1412. DOI: `10.1109/TOH.2019.2911519`.

Mandalika, Veera Bhadra Harish, Alexander I Chernoglazov, Mark Billinghurst, Christoph Bartneck, Michael A Hurrell, Niels de Ruiter, Anthony P H Butler, and Philip H Butler (2018). "A Hybrid

2D/3D User Interface for Radiological Diagnosis." In: *Journal of Digital Imaging* 31.1, pp. 56–73. DOI: 10.1007/s10278-017-0002-6.

Marai, G Elisabeta, Angus G Forbes, and Andrew Johnson (2016). "Interdisciplinary immersive analytics at the electronic visualization laboratory: Lessons learned and upcoming challenges." In: *Proceedings of IA*. IEEE Computer Society. Los Alamitos, pp. 54–59. DOI: 10.1109/IMMERSIVE.2016.7932384.

Marks, Stefan, Javier E Estevez, and Andy M Connor (2014). "Towards the Holodeck: Fully Immersive Virtual Reality Visualisation of Scientific and Engineering Data." In: *Proceedings of IVCNZ*. New York: ACM, pp. 42–47. ISBN: 978-1-4503-3184-5. DOI: 10.1145/2683405.2683424.

Marriott, Kim, Falk Schreiber, Tim Dwyer, Karsten Klein, Nathalie Henry Riche, Takayuki Itoh, Wolfgang Stuerzlinger, and Bruce H Thomas, eds. (2018). *Immersive Analytics*. Berlin Heidelberg: Springer. DOI: 10.1007/978-3-030-01388-2.

Martinet, Anthony, Géry Casiez, and Laurent Grisoni (2010). "The effect of DOF separation in 3D manipulation tasks with multi-touch displays." In: *Proceedings of Virtual Reality Software and Technology*. New York: ACM, pp. 111–118. ISBN: 9781450304412. DOI: 10.1145/1889863.1889888.

McIntire, John P, Paul R Havig, and Eric E Geiselman (2014a). "Stereoscopic 3D displays and human performance: A comprehensive review." In: *Displays* 35.1, pp. 18–26. ISSN: 0141-9382. DOI: 10.1016/j.displa.2013.10.004.

— (Jan. 2014b). "Stereoscopic 3D displays and human performance: A comprehensive review." In: *Displays* 35.1, pp. 18–26. ISSN: 0141-9382. DOI: 10.1016/j.displa.2013.10.004.

McShane, Blakeley B and David Gal (2017). "Statistical Significance and the Dichotomization of Evidence." In: *Journal of the American Statistical Association* 112.519, pp. 885–895. DOI: 10.1080/01621459.2017.1289846.

Mendes, Daniel, Fabio Marco Caputo, Andrea Giachetti, Alfredo Ferreira, and Joaquim A Jorge (2019). "A Survey on 3D Virtual Object Manipulation: From the Desktop to Immersive Virtual Environments." In: *Computer Graphics Forum* 38.1, pp. 21–45. DOI: 10.1111/cgf.13390.

Meulen, Romke (2012). "Prop-free 3D interaction with virtual environments in the CAVE using the Microsoft Kinect camera." PhD thesis. Faculty of Science and Engineering, University of Groningen. URL: https://fse.studenttheses.ub.rug.nl/10575/2/AI.MHMC-2012-VanDerMeulen.pdf.

Microsoft (2019). *Microsoft HoloLens: A new way to see your world*. Web site: https://www.microsoft.com/en-us/hololens/hardware.

Millette, Alexandre and Michael J McGuffin (2016). "DualCAD: Integrating Augmented Reality with a Desktop GUI and Smartphone

Interaction." In: *Proceedings of ISMAR-Adjunct*. Los Alamitos: IEEE Computer Society, pp. 21–26. DOI: 10.1109/ISMAR-Adjunct.2016.0030.

Morris, Meredith Ringel, Andreea Danielescu, Steven Drucker, Danyel Fisher, Bongshin Lee, m. c. Schraefel, and Jacob O Wobbrock (2014). "Reducing Legacy Bias in Gesture Elicitation Studies." In: *Interactions* 21.3, pp. 40–45. ISSN: 1072-5520. DOI: 10.1145/2591689.

Munzner, Tamara (2014). *Visualization Analysis and Design*. Boca Raton, FL, USA: CRC Press. DOI: 10.1201/b17511.

Murray, Paul, Fintan McGee, and Angus G Forbes (2017). "A taxonomy of visualization tasks for the analysis of biological pathway data." In: *BMC Bioinformatics* 18.2. DOI: 10.1186/s12859-016-1443-5.

Nacenta, Miguel A, Patrick Baudisch, Hrvoje Benko, and Andy Wilson (2009). "Separability of Spatial Manipulations in Multi-touch Interfaces." In: *Proceedings of Graphics Interface*. Toronto: Canadian Information Processing Society, pp. 175–182. ISBN: 978-1-56881-470-4. DOI: 10.20380/GI2009.23.

Nagao, Ken, Yucong Ye, Chuan Wang, Issei Fujishiro, and Kwan-Liu Ma (2016). "Enabling interactive scientific data visualization and analysis with see-through HMDs and a large tiled display." In: *Proceedings of IA*. Los Alamitos: IEEE Computer Society, pp. 1–6. DOI: 10.1109/IMMERSIVE.2016.7932374.

Niedziela, Jeremi and Barthélémy von Haller (2017). "Event visualisation in ALICE – Current status and strategy for Run 3." In: *Journal of Physics: Conference Series* 898.7. DOI: 10.1088/1742-6596/898/7/072008.

Noirhomme-Fraiture, Monique, Frédéric Randolet, Luca Chittaro, and Grégory Custinne (2005). "Data visualizations on small and very small screens." In: *Proceedings of Symposium on Applied Stochastic Models and Data Analysis*. ENST. France, pp. 276–285. URL: https://conferences.telecom-bretagne.eu/asmda2005/IMG/pdf/proceedings/276.pdf.

Normand, Erwan and Michael J. McGuffin (2018). "Enlarging a Smartphone with AR to Create a Handheld VESAD (Virtually Extended Screen-Aligned Display)." In: *2018 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. Munich, Germany, pp. 123–133. DOI: 10.1109/ISMAR.2018.00043.

Oakley, I and S O'Modhrain (2005). "Tilt to scroll: Evaluating a motion based vibrotactile mobile interface." In: *Proceedings of World Haptics*. Los Alamitos: IEEE Computer Society, pp. 40–49. DOI: 10.1109/WHC.2005.138.

Oculus (2019). *Oculus Rif: VR Headset for VR Ready PCs*. Web site: https://www.oculus.com/rift/.

Pahud, Michel, Eyal Ofek, Nathalie Henry Riche, Christophe Hurter, and Jens Grubert (2018). *Mobiles as Portals for Interacting with Virtual*

Data Visualizations. arXiv pre-print 1804.03211. arXiv: 1804.03211 [cs.HC].

Pelurson, Sebastien and Laurence Nigay (2016). "Bimanual Input for Multiscale Navigation with Pressure and Touch Gestures." In: *Proceedings of ICMI*. New York: ACM, pp. 145–152. ISBN: 978-1-4503-4556-9. DOI: 10.1145/2993148.2993152.

Phong, Bui Tuong (June 1975). "Illumination for Computer Generated Pictures." In: *Communications of the ACM* 18.6, pp. 311–317. DOI: 10.1145/360825.360839.

Pierre, Cazaux, Vézien Jean-Marc, Martin Pierre, Touraine Damien, Férey Nicolas, and Bourdot. Patrick (2010). "Le système EVE : aperçu des recherches en cours en immersion collaborative et en téléopération." In: *Actes de AFRV 2010*. DOI: 10.1145/234972.234975.

Poupyrev, Ivan and Shigeaki Maruyama (2003). "Tactile Interfaces for Small Touch Screens." In: *Proceedings of UIST*. New York: ACM, pp. 217–220. ISBN: 1-58113-636-6. DOI: 10.1145/964696.964721.

Prabhat, Andrew Forsberg, Michael Katzourin, Kristi Wharton, and Mel Slater (2008). "A Comparative Study of Desktop, Fishtank, and Cave Systems for the Exploration of Volume Rendered Confocal Data Sets." In: *IEEE Transactions on Visualization and Computer Graphics* 14.3, pp. 551–563. ISSN: 1077-2626. DOI: 10.1109/TVCG.2007.70433.

Qi, Wen and Jean-Bernard Martens (2005). "Tangible User Interfaces for 3D Clipping Plane Interaction with Volumetric Data: A Case Study." In: *Proceedings of ICMI*. New York: ACM, pp. 252–258. ISBN: 1-59593-028-0. DOI: 10.1145/1088463.1088507.

Qi, Wen, Russell M Taylor II, Christopher G Healey, and Jean-Bernard Martens (2006). "A Comparison of Immersive HMD, Fish Tank VR and Fish Tank with Haptics Displays for Volume Visualization." In: *Proceedings of APGV*. New York: ACM, pp. 51–58. ISBN: 1-59593-429-4. DOI: 10.1145/1140491.1140502.

Ramos, Gonzalo, Matthew Boulos, and Ravin Balakrishnan (2004). "Pressure Widgets." In: *Proceedings of CHI*. New York: ACM, pp. 487–494. ISBN: 1-58113-702-8. DOI: 10.1145/985692.985754.

Raynal, Mathieu, Guillaume Gauffre, Cédric Bach, Bénédicte Schmitt, and Emmanuel Dubois (2010). "Tactile Camera vs. Tangible Camera: Taking Advantage of Small Physical Artefacts to Navigate into Large Data Collection." In: *Proceedings of NordiCHI*. New York: ACM, pp. 373–382. ISBN: 978-1-60558-934-3. DOI: 10.1145/1868914.1868958.

Reipschläger, Patrick and Raimund Dachselt (Nov. 2019). "DesignAR: Immersive 3D-Modeling Combining Augmented Reality with Interactive Displays." In: *Proceedings of the 2019 ACM International Conference on Interactive Surfaces and Spaces*. Daejeon, Republic of Korea: ACM, pp. 29–42. ISBN: 978-1-4503-6891-9. DOI: 10.1145/3343055.3359718.

Reipschläger, Patrick, Tamara Flemisch, and Raimund Dachselt (Feb. 2021). “Personal Augmented Reality for Information Visualization on Large Interactive Displays.” In: *IEEE Transactions on Visualization and Computer Graphics* 27 (2), pp. 1182–1192. DOI: 10.1109/TVCG.2020.3030460.

Reisman, Jason L, Philip L Davidson, and Jefferson Y Han (2009). “A Screen-space Formulation for 2D and 3D Direct Manipulation.” In: *Proceedings of UIST*. New York: ACM, pp. 69–78. ISBN: 978-1-60558-745-5. DOI: 10.1145/1622176.1622190.

Rheingans, P (2002). “Are we there yet? Exploring with dynamic visualization.” In: *IEEE Computer Graphics and Applications* 22.1, pp. 6–10. ISSN: 0272-1716. DOI: 10.1109/38.974511.

Roudaut, Anne, Eric Lecolinet, and Yves Guiard (2009). “MicroRolls: Expanding Touch-screen Input Vocabulary by Distinguishing Rolls vs. Slides of the Thumb.” In: *Proceedings of CHI*. New York: ACM, pp. 927–936. ISBN: 978-1-60558-246-7. DOI: 10.1145/1518701.1518843.

Sauro, Jeff and James R Lewis (2010). “Average task times in usability tests: What to report?” In: *Proceedings of CHI*. New York: ACM, pp. 2347–2350. DOI: 10.1145/1753326.1753679.

Schemali, Leïla and Elmar Eisemann (2014). “Design and evaluation of mouse cursors in a stereoscopic desktop environment.” In: *Proceedings of 3DUI*. Los Alamitos: IEEE Computer Society, pp. 67–70. DOI: 10.1109/3DUI.2014.6798844.

Schmalstieg, Dieter, Anton Fuhrmann, Gerd Hesina, Zsolt Szalavári, L Miguel Encarnação, Michael Gervautz, and Werner Purgathofer (2002). “The Studierstube Augmented Reality Project.” In: *Presence: Teleoper. Virtual Environ.* 11.1, pp. 33–54. ISSN: 1054-7460. DOI: 10.1162/105474602317343640.

Schmidt, Frank L and John E Hunter (1997). “Eight common but false objections to the discontinuation of significance testing in the analysis of research data.” In: *What If There Were No Significance Tests?* New York: Lawrence Erlbaum Associates Publishers. Chap. 3, pp. 37–64. DOI: 10.4324/9781315629049.

Schultheis, Udo, Jaso Jerald, Fernando Toledo, Arun Yoganandan, and Paul Mlyniec (2012). “Comparison of a two-handed interface to a wand interface and a mouse interface for fundamental 3D tasks.” In: *Proceedings of 3DUI*. Los Alamitos: IEEE Computer Society, pp. 117–124. DOI: 10.1109/3DUI.2012.6184195.

Schwank, Johannes, Franca-Alexandra Rupprecht, and Achim Ebert (2017). “Waggle – Orientation-based Tablet Interaction.” In: *CHI Extended Abstracts*. New York: ACM, pp. 2042–2048. ISBN: 978-1-4503-4656-6. DOI: 10.1145/3027063.3053151.

Sedlmair, Michael, Tamara Munzner, and Melanie Tory (2013). “Empirical Guidance on Scatterplot and Dimension Reduction Technique Choices.” In: *IEEE Transactions on Visualization and Computer Graphics* 19.12, pp. 2634–2643. ISSN: 1077-2626. DOI: `10.1109/TVCG.2013.153`.

Sellen, Abigail J, Gordon P Kurtenbach, and William A S Buxton (1992). “The Prevention of Mode Errors Through Sensory Feedback.” In: *Human Computer Interaction* 7.2, pp. 141–164. DOI: `10.1207/s15327051hci0702_1`.

Sereno, Mickael, Lonni Besançon, and Tobias Isenberg (2019). “Supporting Volumetric Data Visualization and Analysis by Combining Augmented Reality Visuals with Multi-Touch Input.” In: *Computer Graphics Forum*, 3 pages. URL: `https://hal.inria.fr/hal-02123904`.

Serrano, Marcos, 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 MobileHCI*. New York: ACM, pp. 161–171. ISBN: 978-1-4503-3652-9. DOI: `10.1145/2785830.2785838`.

Shaw, Christopher D (1998). “Pain and Fatigue in Desktop VR: Initial Results.” In: *Proceedings of Graphics Interface*. Toronto: Canadian Information Processing Society, pp. 185–192. DOI: `10.20380/GI1998.23`.

Shibata, Takashi (2002). “Head Mounted Display.” In: *Displays* 23.1, pp. 57–64. DOI: `10.1016/S0141-9382(02)00010-0`.

Shoemake, Ken (1992). “ARCBALL: A User Interface for Specifying Three-Dimensional Orientation Using a Mouse.” In: *Proc. Graphics Interface*. Vancouver, British Columbia, Canada: Morgan Kaufmann, pp. 151–156. ISBN: 0969533810. DOI: `10.20380/GI1992.18`.

Sollenberger, Randy L and Paul Milgram (1993). “Effects of stereoscopic and rotational displays in a three-dimensional path-tracing task.” In: *Human Factors* 35.3, pp. 483–499. DOI: `10.1177/00187\-208\-9303\-500306`.

Sollich, Hendrik, Ulrich von Zadow, Tobias Pietzsch, Pavel Tomancak, and Raimund Dachselt (2016). “Exploring Time-dependent Scientific Data Using Spatially Aware Mobiles and Large Displays.” In: *Proceedings of ISS*. New York: ACM, pp. 349–354. ISBN: 978-1-4503-4248-3. DOI: `10.1145/2992154.2996779`.

Spindler, Martin, Wolfgang Büschel, and Raimund Dachselt (2012). “Use Your Head: Tangible Windows for 3D Information Spaces in a Tabletop Environment.” In: *Proceedings of ITS*. New York: ACM, pp. 245–254. ISBN: 978-1-4503-1209-7. DOI: `10.1145/2396636.2396674`.

Spindler, Martin, Martin Schuessler, Marcel Martsch, and Raimund Dachselt (2014). “Pinch-Drag-Flick vs. Spatial Input: Rethinking Zoom & Pan on Mobile Displays.” In: *Proceedings of CHI*.

New York: ACM, pp. 1113–1122. ISBN: 978-1-4503-2473-1. DOI: 10.1145/2556288.2557028.

Sultanum, Nicole, Sowmya Somanath, Ehud Sharlin, and Mario Costa Sousa (2011). "Point it, Split it, Peel it, View it": Techniques for Interactive Reservoir Visualization on Tabletops." In: *Proceedings of ITS*. New York: ACM, pp. 192–201. DOI: 10.1145/2076354.2076390.

Sun, Junwei, Wolfgang Stuerzlinger, and Bernhard E. Riecke (2018). "Comparing Input Methods and Cursors for 3D Positioning with Head-Mounted Displays." In: *Proceedings of the 15th ACM Symposium on Applied Perception*. Vancouver, British Columbia, Canada: Association for Computing Machinery. ISBN: 9781450358941. DOI: 10.1145/3225153.3225167.

Sundén, Erik, Ingemar Lundgren, and Anders Ynnerman (2017). "Hybrid Virtual Reality Touch Table – An Immersive Collaborative Platform for Public Explanatory Use of Cultural Objects and Sites." In: *Proceedings of Eurographics Workshop on Graphics and Cultural Heritage*. Goslar, Germany: Eurographics Association, pp. 109–113. ISBN: 978-3-03868-037-6. DOI: 10.2312/gch.20171300.

Sundin, Martin and Morten Fjeld (2009). "Softly Elastic 6 DOF Input." In: *International Journal of Human–Computer Interaction* 25.7, pp. 647–691. DOI: 10.1080/10447310902964124.

Surale, Hemant Bhaskar, Aakar Gupta, Mark Hancock, and Daniel Vogel (2019). "TabletInVR: Exploring the Design Space for Using a Multi-Touch Tablet in Virtual Reality." In: *Proceedings of CHI*. New York: ACM, 13:1–13:13. ISBN: 978-1-4503-5970-2. DOI: 10.1145/3290605.3300243.

Sutherland, Ivan E (1966). "Ten unsolved problems in computer graphics." In: *Datamation* 12.5, p. 22.

— (1968). "A Head-mounted Three Dimensional Display." In: *Proceedings of the Fall Joint Computer Conference, Part I*. New York: ACM, pp. 757–764. DOI: 10.1145/1476589.1476686.

Szalavári, Zsolt, Dieter Schmalsteig, Anton Fuhrmann, and Michael Gervautz (1998). ""Studierstube": An environment for collaboration in augmented reality." In: *Virtual Reality* 3.1, pp. 37–48. ISSN: 1434-9957. DOI: 10.1007/BF01409796.

Tait, Andy (1992). "Desktop virtual reality." In: *IEE Colloquium on Using Virtual Worlds*. London. IET, 5:1–5:5. URL: https://ieeexplore.ieee.org/document/168239.

Teather, Robert J. and Wolfgang Stuerzlinger (2013). "Pointing at 3D Target Projections with One-Eyed and Stereo Cursors." In: *Proceedings of CHI*. Paris, France: ACM, 159–168. ISBN: 9781450318990. DOI: 10.1145/2470654.2470677.

Thabet, Walid, Mohd Fairuz Shiratuddin, and Doug Bowman (Jan. 2002). "Virtual Reality in Construction: A Review." In: *Engineering Computational Technology*, pp. 25–52. DOI: 10.4203/csets.8.2.

Theart, Rensu P, Ben Loos, and Thomas R Niesler (2017). "Virtual reality assisted microscopy data visualization and colocalization analysis." In: *BMC Bioinformatics* 18.2. DOI: `10.1186/s12859-016-1446-2`.

Tory, Melanie and Torsten Möller (2004). "Human Factors in Visualization Research." In: *IEEE Transactions on Visualization and Computer Graphics* 10.1, pp. 72–84. DOI: `10.1109/TVCG.2004.1260759`.

Treanor, Darren, Naomi Jordan-Owers, John Hodrien, Jason Wood, Phil Quirke, and Roy A Ruddle (2009). "Virtual reality Powerwall versus conventional microscope for viewing pathology slides: an experimental comparison." In: *Histopathology* 55.3, pp. 294–300. DOI: `https://doi.org/10.1111/j.1365-2559.2009.03389.x`.

Tukey, John W (1977). *Exploratory Data Analysis*. Reading, Mass.: Pearson.

Valentine, Jeffrey C, Ariel M Aloe, and Timothy S Lau (2015). "Life After NHST: How to Describe Your Data Without "p-ing" Everywhere." In: *Basic and Applied Social Psychology* 37.5, pp. 260–273. DOI: `10.1080/01973533.2015.1060240`.

VandenBos, Gary R, ed. (2009). *Publication Manual of the American Psychological Association*. 6th ed. Washington, DC: American Psychological Association. URL: `http://www.apastyle.org/manual/`.

Veit, Manuel, Antonio Capobianco, and Dominique Bechmann (2009). "Influence of Degrees of Freedom's Manipulation on Performances During Orientation Tasks in Virtual Reality Environments." In: *Proceedings of Virtual Reality Software and Technology*. New York: ACM, pp. 51–58. ISBN: 978-1-60558-869-8. DOI: `10.1145/1643928.1643942`.

VIVE (2019). *VIVE: Discover Virtual Reality Beyond Imagination*. URL: `https://www.vive.com`.

Vuibert, Vanessa, Wolfgang Stuerzlinger, and Jeremy R Cooperstock (2015). "Evaluation of Docking Task Performance Using Mid-air Interaction Techniques." In: *Proceedings of Spatial User Interaction*. ACM. New York, pp. 44–52. DOI: `10.1145/2788940.2788950`.

Vukotic, Ilija, Edward Moyse, and Riccardo Maria Bianchi (2015). *ATLAS – A Virtual Reality application*. arXiv preprint 1511.00047. URL: `https://arxiv.org/abs/1511.00047`.

Wang, Lu, Laikuan Wong, Yajie Xu, Xiao Zhou, Siyuan Qiu, Xiangxu Meng, and Chenglei Yang (2018). "The Design and Empirical Evaluations of 3D Positioning Techniques for Pressure-based Touch Control on Mobile Devices." In: *Personal Ubiquitous Computing* 22.3, pp. 525–533. ISSN: 1617-4909. DOI: `10.1007/s00779-018-1147-0`.

Wang, Xiyao, Lonni Besançon, Mehdi Ammi, and Tobias Isenberg (2019a). "Augmenting Tactile 3D Data Navigation With Pressure Sensing." In: *Computer Graphics Forum* 38.3, pp. 635–647. DOI: `10.1111/cgf.13716`.

Wang, Xiyao, Lonni Besançon, Florimond Guéniat, Mickael Sereno, Mehdi Ammi, and Tobias Isenberg (2019b). "A Vision of Bringing Immersive Visualization to Scientific Workflows." In: *Proceedings of the CHI Workshop on Immersive Analytics: Interaction Design and Prototyping for Immersive Analytics*. Ed. by Barrett Ens, Maxime Cordeil, Ulrich Engelke, Marcos Serrano, Wesley Willett, and Benjamin Bach. URL: https://hal.archives-ouvertes.fr/hal-02053969/.

Wang, Xiyao, David Rousseau, Lonni Besançon, Mickael Sereno, Mehdi Ammi, and Tobias Isenberg (2020). "Towards an Understanding of Augmented Reality Extensions for Existing 3D Data Analysis Tools." In: *CHI 2020 - ACM Conference on Human Factors in Computing Systems*. Honolulu, United States. DOI: 10.1145/3313831.3376657.

Ware, Colin, Kevin Arthur, and Kellogg S. Booth (1993). "Fish Tank Virtual Reality." In: *Proceedings of the INTERACT '93 and CHI '93 Conference on Human Factors in Computing Systems*. Amsterdam, The Netherlands: Association for Computing Machinery, 37–42. ISBN: 0897915755. DOI: 10.1145/169059.169066.

Ware, Colin and Glenn Franck (1996a). "Evaluating Stereo and Motion Cues for Visualizing Information Nets in Three Dimensions." In: *ACM Trans. Graph.* 15.2, 121–140. ISSN: 0730-0301. DOI: 10.1145/234972.234975.

— (1996b). "Evaluating Stereo and Motion Cues for Visualizing Information Nets in Three Dimensions." In: *ACM Transactions on Graphics* 15.2, pp. 121–140. ISSN: 0730-0301. DOI: 10.1145/234972.234975.

Ware, Colin and Kathy Lowther (1997). "Selection Using a One-Eyed Cursor in a Fish Tank VR Environment." In: *ACM Trans. Comput.-Hum. Interact.* 4.4, 309–322. ISSN: 1073-0516. DOI: 10.1145/267135.267136.

Ware, Colin and Peter Mitchell (2005). "Reevaluating Stereo and Motion Cues for Visualizing Graphs in Three Dimensions." In: *Proceedings of the 2nd Symposium on Applied Perception in Graphics and Visualization*. A Coroña, Spain: Association for Computing Machinery, 51–58. DOI: 10.1145/1080402.1080411.

Wesche, Gerold (1999). "Three-dimensional visualization of fluid dynamics on the Responsive Workbench." In: *Future Generation Computer Systems* 15.4, pp. 469 –475. ISSN: 0167-739X. DOI: 10.1016/S0167-739X(99)00003-5.

Wesche, Gerold, Jürgen Wind, M Gobel, L Rosenblum, J Durbin, R Doyle, D Tate, R King, B Frohlich, M Fischer, et al. (1997). "Applications of the Responsive Workbench." In: *IEEE Computer Graphics and Applications* 17.4, pp. 10–15. DOI: 10.1109/38.595260.

Will, Usher, Klacansky Pavol, Federer Frederick, Bremer Peer-Timo, Knoll Aaron, Yarch Jeff, Angelucci Alessandra, and Pascucci Valerio (2018). “A Virtual Reality Visualization Tool for Neuron Tracing.” In: *IEEE Transactions on Visualization and Computer Graphics* 24.1, pp. 994–1003. ISSN: 1077-2626. DOI: `10.1109/TVCG.2017.2744079`.

Wills, Graham (2008). “Linked Data Views.” In: *Handbook of Data Visualization*. Ed. by Chun-houh Chen, Wolfgang Härdle, and Antony Unwin. Berlin Heidelberg: Springer. Chap. II.9, pp. 217–241. DOI: `10.1007/978-3-540-33037-0_10`.

Wilson, Graham, Craig Stewart, and Stephen A Brewster (2010). “Pressure-based Menu Selection for Mobile Devices.” In: *Proceedings of MobileHCI*. New York: ACM, pp. 181–190. ISBN: 978-1-60558-835-3. DOI: `10.1145/1851600.1851631`.

Withana, Anusha, Roshan Peiris, Nipuna Samarasekara, and Suranga Nanayakkara (2015). “zSense: Enabling Shallow Depth Gesture Recognition for Greater Input Expressivity on Smart Wearables.” In: *Proceedings of CHI*. New York: ACM, pp. 3661–3670. ISBN: 978-1-4503-3145-6. DOI: `10.1145/2702123.2702371`.

Witmer, Bob G and Michael J Singer (1998). “Measuring Presence in Virtual Environments: A Presence Questionnaire.” In: *Presence* 7.3, pp. 225–240. ISSN: 1054-7460. DOI: `10.1162/105474698565686`.

Wozniak, Pawel, Morten Fjeld, and Shengdong Zhao (2014). “Limiting Trial and Error: Introducing a Systematic Approach to Designing Clutching.” In: *Proc. Chinese CHI*. Toronto, Ontario, Canada: ACM, pp. 35–39. ISBN: 9781450328760. DOI: `10.1145/2592235.2592241`.

Yadav, Amit, Alexander Keith Eady, Sara Nabil, and Audrey Girouard (2019). “JoyHolder: Tangible Back-of-Device Mobile Interactions.” In: *Proceedings of the 2019 ACM International Conference on Interactive Surfaces and Spaces*. ISS ’19. Daejeon, Republic of Korea: Association for Computing Machinery, 343–346. ISBN: 9781450368919. DOI: `10.1145/3343055.3360748`.

Yi, Ji Soo, Youn ah Kang, and John Stasko (2007). “Toward a Deeper Understanding of the Role of Interaction in Information Visualization.” In: *IEEE Transactions on Visualization and Computer Graphics* 13.6, pp. 1224–1231. ISSN: 1077-2626. DOI: `10.1109/TVCG.2007.70515`.

Yu, Lingyun, Konstantinos Efstathiou, Petra Isenberg, and Tobias Isenberg (2016). “CAST: Effective and Efficient User Interaction for Context-Aware Selection in 3D Particle Clouds.” In: *IEEE Transactions on Visualization and Computer Graphics* 22.1, pp. 886–895. DOI: `10.1109/TVCG.2015.2467202`.

Yu, Lingyun, Pjotr Svetachov, Petra Isenberg, Maarten H Everts, and Tobias Isenberg (2010). “FI3D: Direct-Touch Interaction for the Exploration of 3D Scientific Visualization Spaces.” In: *IEEE Transactions on Visualization and Computer Graphics* 16.6, pp. 1613–1622. DOI: `10.1109/TVCG.2010.157`.

Zeleznik, Robert and Andrew Forsberg (1999). “UniCam—2D Gestural Camera Controls for 3D Environments.” In: *Proceedings of I3D*. New York: ACM, pp. 169–173. ISBN: 1-58113-082-1. DOI: 10.1145/300523.300546.

Zhai, Shumin (1998). “User Performance in Relation to 3D Input Device Design.” In: *ACM SIGGRAPH Computer Graphics* 32.4, pp. 50–54. ISSN: 0097-8930. DOI: 10.1145/307710.307728.

Zhai, Shumin and Paul Milgram (1998). “Quantifying Coordination in Multiple DOF Movement and Its Application to Evaluating 6 DOF Input Devices.” In: *Proc. CHI*. Los Angeles, California, USA: ACM, pp. 320–327. ISBN: 0201309874. DOI: 10.1145/274644.274689.

Zhai, Shumin and John W Senders (1997a). “Investigating Coordination in Multidegree of Freedom Control I: Time-On-Target Analysis of 6 DOF Tracking.” In: *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 41.2, pp. 1249–1253. DOI: 10.1177/107118139704100211 7.

— (1997b). “Investigating Coordination in Multidegree of Freedom Control II: Correlation Analysis in 6 DOF Tracking.” In: *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 41.2, pp. 1254–1258. DOI: 10.1177/1071181397041002118.

Zielasko, Daniel, Benjamin Weyers, Martin Bellgardt, Sebastian Pick, Alexander Meißner, Tom Vierjahn, and Torsten W. Kuhlen (2017). “Remain seated: Towards fully-immersive desktop VR.” In: *Proceedings of WEVR*. Los Alamitos: IEEE Computer Society, 1:1–1:6. DOI: 10.1109/WEVR.2017.7957707.

Titre : Environnements de réalité augmentée pour l'exploration interactive de données 3D

Mots clés : Données 3D,interaction tangible,visualisation scientifique, interaction/IHM, entrée tactile, réalité augmentée

Résumé : La visualisation exploratoire des données 3D est fondamentale dans des domaines scientifiques. Traditionnellement, les experts utilisent un PC et s'appuient sur la souris pour ajuster la vue. Cette configuration permet l'immersion par interaction---l'utilisateur peut contrôler précisément la vue, mais elle ne fournit pas de profondeur, qui limite la compréhension de données complexes. La réalité virtuelle ou augmentée (RV/A), en revanche, offre une immersion visuelle avec des vues stéréoscopiques. Bien que leurs avantages aient été prouvés, plusieurs points limitent leur application, notamment les besoins élevés de configuration/maintenance, les difficultés de contrôle précis et, plus important, la séparation des outils d'analyse traditionnels.

Pour bénéficier des deux côtés, nous avons donc étudié un système hybride combinant l'environnement RA avec un PC pour fournir des immersions interactives et visuelles. Nous avons collaboré étroitement avec des physiciens des particules afin de comprendre leur processus de travail et leurs besoins de visualisation pour motiver notre conception.

D'abord, basé sur nos discussions avec les physiciens, nous avons construit un prototype qui permet d'accomplir des tâches pour l'exploration de leurs données. Ce prototype traitait l'espace RA comme une extension de l'écran du PC et permettait aux utilisateurs d'interagir librement avec chacun d'eux avec la souris. Ainsi, les experts pouvaient bénéficier de l'immersion visuelle et utilisent les outils d'analyse sur PC. Une étude observationnelle menée avec 7 physiciens au CERN a validé la faisabilité et confirmé les avantages. Nous avons également constaté que la grande toile du RA et le fait de se déplacer pour observer les données dans le RA présentaient un grand potentiel. Cependant, la conception de l'interaction de la souris et l'utilisation de widgets dans la RA devaient être améliorés.

Ensuite, nous avons décidé de ne pas utiliser intensivement les widgets plats dans la RA. Mais nous nous sommes demandé si l'utilisation de la souris pour naviguer dans la RA est problématique, et nous avons ensuite tenté d'étudier si la correspondance de la dimensionnalité entre les dispositifs d'entrée et de sortie joue un rôle important. Les résultats des études (qui ont comparé la performance de l'utilisation de la souris, de la souris spatiale et de la tablette tangible couplée à l'écran ou à l'espace de RA) n'ont pas montré que la correspondance était importante. Nous avons donc conclu que la dimensionnalité n'était pas un point critique à considérer, ce qui suggèreque les utilisateurs sont libres de choisir toute entrée qui convient à une tâche spécifique. De plus, nos résultats ont montré que la souris restait un outil efficace. Nous pouvons donc valider notre conception et conserver la souris comme entrée principale, tandis que les autres modalités ne devraient servir que comme complément pour des cas spécifiques.

Ensuite, pour favoriser l'interaction et conserver les informations pendant que les utilisateurs se déplacent en RA, nous avons proposé d'ajouter un appareil mobile. Nous avons introduit une nouvelle approche qui augmente l'interaction tactile avec la détection de pression pour la navigation 3D. Les résultats ont montré que cette méthode pouvait améliorer efficacement la précision, avec une influence limitée sur le temps. Nous pensons donc qu'elle est utile à des tâches de vis où une précision est exigée.

Enfin, nous avons résumé tous les résultats obtenus et imaginé un scénario réaliste qui utilise un poste de travail PC, un casque RA et un appareil mobile. Les travaux présentés dans cette thèse montrent le potentiel de la combinaison d'un PC avec des environnements de RA pour améliorer le processus d'exploration de données 3D et confirment sa faisabilité, ce qui, nous l'espérons, inspirera la future conception qui apportera une visualisation immersive aux flux de travail scientifiques existants.

Title : Augmented reality environments for the interactive exploration of 3D data

Keywords : 3D data, tangible interaction, scientific visualisation, interaction/HCI, touch input, augmented reality

Abstract : Exploratory visualization of 3D data is fundamental in many scientific domains. Traditionally, experts use a PC workstation and rely on mouse and keyboard to interactively adjust the view to observe the data. This setup provides immersion through interaction---users can precisely control the view and the parameters, but it does not provide any depth clues which can limit the comprehension of large and complex 3D data. Virtual or augmented reality (V/AR) setups, in contrast, provide visual immersion with stereoscopic views. Although their benefits have been proven, several limitations restrict their application to existing workflows, including high setup/maintenance needs, difficulties of precise control, and, more importantly, the separation from traditional analysis tools.

To benefit from both sides, we thus investigated a hybrid setting combining an AR environment with a traditional PC to provide both interactive and visual immersions for 3D data exploration. We closely collaborated with particle physicists to understand their general working process and visualization requirements to motivate our design.

First, building on our observations and discussions with physicists, we built up a prototype that supports fundamental tasks for exploring their datasets. This prototype treated the AR space as an extension to the PC screen and allowed users to freely interact with each using the mouse. Thus, experts could benefit from the visual immersion while using analysis tools on the PC. An observational study with 7 physicists in CERN validated the feasibility of such a hybrid setting, and confirmed the benefits. We also found that the large canvas of the AR and walking around to observe the data in AR had a great potential for data exploration. However, the design of mouse interaction in AR and the use of PC widgets in AR needed improvements.

Second, based on the results of the first study, we decided against intensively using flat widgets in AR. But we wondered if using the mouse for navigating in AR is problematic compared to high degrees of freedom (DOFs) input, and then attempted to investigate if the match or mismatch of dimensionality between input and output devices play an important role in users' performance. Results of user studies (that compared the performance of using mouse, space mouse, and tangible tablet paired with the screen or the AR space) did not show that the (mis-)match was important. We thus concluded that the dimensionality was not a critical point to consider, which suggested that users are free to choose any input that is suitable for a specific task. Moreover, our results suggested that the mouse was still an efficient tool compared to high DOFs input. We can therefore validate our design of keeping the mouse as the primary input for the hybrid setting, while other modalities should only serve as an addition for specific use cases.

Next, to support the interaction and to keep the background information while users are walking around to observe the data in AR, we proposed to add a mobile device. We introduced a novel approach that augments tactile interaction with pressure sensing for 3D object manipulation/view navigation. Results showed that this method could efficiently improve the accuracy, with limited influence on completion time. We thus believe that it is useful for visualization purposes where a high accuracy is usually demanded.

Finally, we summed up in this thesis all the findings we have and came up with an envisioned setup for a realistic data exploration scenario that makes use of a PC workstation, an AR headset, and a mobile device. The work presented in this thesis shows the potential of combining a PC workstation with AR environments to improve the process of 3D data exploration and confirms its feasibility, all of which will hopefully inspire future designs that seamlessly bring immersive visualization to existing scientific workflows.

Université Paris-Saclay
Espace Technologique / Immeuble Discovery
Route de l'Orme aux Merisiers RD 128 / 91190 Saint-Aubin, France