



HAL
open science

Understanding Differences between Combinations of 2D and 3D Input and Output Devices for 3D Data Visualization

Xiyao Wang, Lonni Besançon, Mehdi Ammi, Tobias Isenberg

► **To cite this version:**

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*, 2022, 163, pp.102820:1–102820:16. 10.1016/j.ijhcs.2022.102820 . hal-03639831

HAL Id: hal-03639831

<https://inria.hal.science/hal-03639831v1>

Submitted on 13 Apr 2022

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

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

Understanding Differences between Combinations of 2D and 3D Input and Output Devices for 3D Data Visualization

Xiyao Wang^{a,d}, Lonni Besançon^b, Mehdi Ammi^c and Tobias Isenberg^a

^aUniversité Paris-Saclay, CNRS, Inria, LISN, Orsay 91400, France

^bLinköping University, Norrköping 60221, Sweden

^cUniversity of Paris 8, Saint-Denis 93526, France

^dUniversity of Victoria, Victoria V8W 2Y2, Canada

ARTICLE INFO

Keywords:

Augmented reality

3D manipulation

3D visualization


ABSTRACT


Focusing on interaction needs for scientific data exploration, we evaluated people's performance using a 2D mouse, 3D SpaceMouse, or 3D-tangible tablet as input devices to interact with visualizations on 2D screens or stereoscopic augmented reality (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 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. 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.

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 allow 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 an increasing amount of attention in the field of visualization and interaction. 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 properties 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 (Zielasko et al., 2017; Fulmer et al., 2019; Surale et al., 2019; Wang et al., 2020). We follow the idea of using augmented reality (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., 2014; Hurter et al., 2019), questions regarding the choice of 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 both types of views simultaneously or synchronously. The latter leads to interaction needs that differ from the Desktop-VR metaphor (Tait, 1992) where users work with only one space. Due to their difference in spatial

 xiyao.wang23@gmail.com (X. Wang); lonni.besancon@gmail.com (L. Besançon); ammi@ai.univ-paris8.fr (M. Ammi); tobias.isenberg@inria.fr (T. Isenberg)

 xiyaowang.net (X. Wang); lonnibesancon.me (L. Besançon); mehdi-ammi.eu (M. Ammi); tobias.isenberg.cc (T. Isenberg)
 ORCID 0000-0003-2429-6064 (X. Wang); 0000-0002-7207-1276 (L. Besançon); 0000-0003-1763-4045 (M. Ammi); 0000-0001-7953-8644 (T. Isenberg)

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 using a mouse to control the PC, and switching to a tablet for an AR view. We believe, however, that it is essential to provide a seamless interaction experience between interfaces if our final goal is to design a hybrid system for daily use (Ens et al., 2021). 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-degrees-of-freedom (DOF) tangible devices (Ishii, 2008) are often explored in VR/AR environments as they can be mapped directly to 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 DOF 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 were open to trying 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 + AR setup. Our main question here is whether user performance varies with different devices. More specifically, we first want to understand whether user performances change when they perceive the data in different views while they are using a mouse. Second, we want to analyze whether 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 (each of these tasks requiring a different interaction design), we focus on the first step—manipulating object positions and orientations in 3D; we do not investigate the design of specific 2D or 3D cursors for the interfaces.

To answer these questions, we conducted a study to explore how our participants' performance changes when they use 2D (mouse) or 3D (SpaceMouse and tangible tablet¹) input devices to work with visualizations shown on a 2D screen or a stereoscopic AR display. Our results suggest that, for the task in which a high accuracy is needed, 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 primary input when designing hybrid visualization systems, while 3D input devices can serve as complementary controls for specific use cases.

2. Related Work

To understand the practices of choosing input devices for 3D manipulations in different environments, in this section we first summarize previously studied input modalities that have been widely used for 3D manipulations, with respect to their inherent advantages and drawbacks. Taking the special requirements of visualization tasks into consideration, we then further discuss the potential of using them in a hybrid visualization setup. Finally, we introduce existing work that examined different, yet related, questions, which have motivated this work.

3D object manipulation techniques have been extensively studied in the past. Previous surveys (e. g., Jankowski and Hachet, 2013; Mendes et al., 2019; Besançon et al., 2021) summarized a variety of interaction techniques for environments with either 2D screens or stereoscopic views. Early work in HCI recognized the limitations of the 2DOF offered by the classical mouse and, therefore, proposed to either augment its expressivity (e. g., Balakrishnan and Patel, 1998; Cechanowicz et al., 2007) or the number of DOF it can control through the addition of external sensing (e. g., Hinckley et al., 1999), simple shape modifications (e. g., Balakrishnan et al., 1997), or internal sensors (e. g., Zizka et al., 2011; MacKenzie et al., 1997; Cechanowicz et al., 2007; Villar et al., 2009). Besides the traditional or augmented mouse-based interaction, many new forms of input have also been created and investigated. 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. Pioneering work can be found in the approaches by Fitzmaurice (1993), Szalavári and Gervautz (1997), and Poupyrev et al. (1998) who propose to benefit from the tangibility of mobile devices and use them as space-aware devices in several contexts (immersive or not). Some approaches were then developed before touch-screens on mobile devices became ubiquitous. For instance, Henrysson et al. (2005) proposed to manipulate objects in AR via a spatially tracked phone plus keypad.

¹A tangible tablet is a tablet with both touch and tangible (3D-space-aware) input.

Tsang et al. (2002) designed a mechanically tracked display, augmented with a tactile overlay and a microphone. Manipulations of the display make it possible to capture a specific view that can then be annotated through speech. Other approaches proposed to combine touch and tangible input with touch-enabled mobile devices. For example, Spindler et al. (2012) proposed *Tangible Windows*, in which spatially-tracked mobiles act as a peephole into the 3D virtual world and combine tangible interaction, touch interaction, and perspective. Yee (2003) proposed to use a similar concept with pen interaction. Later, 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. This concept was also explored for a collaborative AR environment by Lakatos et al. (2014). Millette and McGuffin (2016) also used a spatially tracked mobile phone to interact with a hybrid AR-PC system. Other paradigm combinations exist (see Besançon et al., 2021), although the approach from Lee et al. (2013) is particularly relevant as it studies how to combine a mouse and keyboard with depth sensing for a see-through desktop environment. Finally, there are devices specifically designed for immersive analytics settings such as the CHARM input device (Klamka and Dachsel, 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 yet being commercially available. 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-DOF input devices have advantages compared to 2 DOF 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., 2017b) also discovered that users were faster with tangible devices in 6 DOF manipulation tasks to achieve the same level of accuracy. 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-DOF 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. 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 in 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 processes 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 setups and concluded that there is no universal method that outperforms others but that it rather depends on the specific task. Even though they included some discussion on the dimensionality match between the input and output devices, their study used setups that all have different input and output devices. They thus did not compare the difference of performance of one device paired with others (for example, the mouse paired with a 2D screen and, later, with a HoloLens). Also, their focus was to investigate scenarios where users can directly touch

²We 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).




Figure 1: Model of spatial interaction directness for visualization by Bruckner et al. (2019).

and interact with the data, which is different from our hybrid scenario where the AR serves as an extension of the PC.

3. Overview and research hypotheses

We study how different input devices compare to each other and which one (or which combination) would best be suited to a hybrid PC-AR setting. To support our ultimate goal of seamlessly integrating an AR extension into a current workstation that is already equipped with a mouse and a keyboard (Wang et al., 2019b), we are particularly interested in studying input devices that can easily be integrated with PCs in general work scenarios. As already noted (Wang et al., 2020), we need to keep the mouse and keyboard for controlling existing data analysis software and for script-writing. Since mice only provide 2 DOF input and rely on a horizontal surface, we decided to also examine a tangible tablet (from Google’s Project Tango, now part of the AR core) as a representative of 6 DOF devices with zero-order positional input. Then, we also include the 6 DOF SpaceNavigator³ (henceforth called SpaceMouse), 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 of elastic rate control has been demonstrated to be suitable for certain navigation tasks (Sundin and Fjeld, 2009). In addition, it can be a complement to the mouse interaction to facilitate 3D manipulations. Some previous studies have shown advantages of traditional mice in 3D manipulations (e. g., Bérard et al., 2009; Besançon et al., 2017b), while others contradict such findings and argue that high-DOF input devices are favorable in 3D environments (e. g., Schultheis et al., 2012; LaViola et al., 2017). We thus further investigate the difference between low- and high-DOF input on users’ performance, paired with both a traditional 2D screen and an AR headset. Taking the high-precision requirement of visualization into account, we believe that easy separation of manipulation offered by low-DOF input may help users to achieve precise control. We thus formulate our first hypothesis as:

H1 For visualization tasks that require high accuracy, 2D input devices such as a mouse will yield better performance (w.r.t. traditional HCI measures: speed, accuracy) than 3D input devices (SpaceMouse, space-aware tablet).

In addition, many studies (e. g., Prabhat et al., 2008; Laha et al., 2012; McIntire et al., 2014; Hurter et al., 2019) demonstrated that stereoscopic views and view-motion parallax depth cues help users to understand spatial data. Yet most of them used fully immersive VR. An AR headset, in contrast, does not occlude users from the real world and can thus easily be integrated with existing input devices. It is likely, however, to introduce distraction (physical environment) and/or to decrease image quality due to, e. g., real-world illumination and the visible background. Still, we believe that AR has advantages for 3D data analysis and manipulations (e. g., Kalkofen et al., 2011; Dwyer et al., 2018) and state that

H2 For all examined input devices, user performance in accomplishing 3D visualization tasks is generally better in the HoloLens’ 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 Bruckner et al.’s (2019) model of spatial directness in interactive visualization (Figure 1). This model is based on research in interactive visualization and exploration of spatial data, where human involvement is an important aspect. It is also relevant to HCI in general when we need to understand the spatial interaction with objects. The model describes the transformations from data (space D) to visual representations (space V), to an output medium (space O), to what the viewer understands or their mental model of the visual representation (space U), to manipulations of an input device (in space M), to interpretations of this input in the form of interaction mappings (space I), and back to the data space 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

³https://www.3dconnexion.fr/spacemouse_compact/

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 2(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 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 here, 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 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 Bruckner et al.'s (2019) model. We thus formulate our final hypothesis as:

H3 We hypothesize that the (mis-)match of input and output dimensionality plays a role on users' performance. When input and output dimensionality match each other, users' 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 in AR, a 2D mouse would create a mismatch with the output space, while a 3D device would 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 the user performance using 3D input (e. g., 6 DOF tangible device) will largely increase paired with the AR, while such improvement may be limited with 2D input.

4. Study

We designed an experiment to examine these hypotheses. This study was approved by the ethics review board of the contact author's institution and pre-registered (available on osf.io/7gsk8) to follow latest guidelines to make research more robust (Cockburn et al., 2018, 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.

4.1. Tasks

We chose a **3D docking task** and a **clipping plane manipulation task** to investigate our research questions. In the former participants manipulate an object's position and its orientation in 6 DOF to match a target's spatial transformation. The latter asks participants to manipulate a semi-transparent plane to clip a volumetric dataset.

A variety of other tasks have been used to evaluate input devices and 3D interaction techniques (such as positioning/placement tasks, selection tasks, and navigation tasks) in HCI, many of those other tasks, are not well suited for our application domain—3D visualization. For example, the steering task (Cohen et al., 1993; Accot and Zhai, 1999) asks users to follow a narrow path. While it has been studied in detail with a well-established trade-off model, its application for visualization is limited. Similarly, pointing tasks were carefully investigated even with specifically designed devices such as the touch stick (Fallot et al., 2006). We do not include it in this study because we want to first focus on essential navigation tasks. Such tasks include the specification of a volumetric view and object manipulations including clipping, which together were highlighted as one of the fundamental groups of spatial interaction techniques for visualization in Besançon et al.'s recent survey (2021). 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., 2017a) 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 both derived tasks (like 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 pairings in a general 3D interaction scenario.

We first asked participants to perform a 3D docking task, which is a commonly used proxy for specifying a generic 3D position in a controlled way (e. g., Chen et al., 1988; Hancock et al., 2007; Glesser et al., 2013; Besançon et al., 2017a; Wang et al., 2019a). We then asked participants to perform a follow-up task of positioning a clipping plane inside a rendered volume to align our experiment better to a data-visualization context. We included this task because clipping




Figure 2: 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.

planes are essential in 3D visualization applications as they allow users to understand the inner structure of volumes or point clouds (Röttger et al., 2011; Song et al., 2011; Keefe and Isenberg, 2013; Lexow et al., 2016; Palomar et al., 2017; Besançon et al., 2021). Even though both tasks are only proxies that can easily be controlled in an experiment, any realistic visualization tool relies on more complex interactions are also based on these basic ones. Although clipping plane manipulations can technically be done with 3 DOF input, we argue below that the use of 6 DOF input makes them easier.

For the docking task, 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 2(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 2(a) by making it thicker; in reality it is as thin as the *x*- and *y*-axes. In this experiment, we varied the initial position and orientation of both the object and the target, but we controlled the initial distance offset between the two to be 600 in virtual space units and their orientation offset to be 60°.

For the second clipping plane manipulation task, we used a fluid mechanic dataset in which color represents the fluid velocity (Figure 2(c)). We rendered a target plane inside the data and asked users to align the semi-transparent clipping plane 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 an opaque volume is much too difficult, and frustrates users to give up quickly. We thus simplified the task by making the data volume semi-transparent, coloring the borders of the hidden plane in yellow, and showing these borders on the outside of the volume. In addition, we colored those parts of the target green when they are within a small distance (we tuned it to be 15 units in our pilot studies) to the clipping plane at a given time. We also decided to fix the initial positions and orientations of the plane and the volume to lower the participants' mental workload while trying to understand the scene. But we did change position and orientation of the target plane for each task. 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. 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. Since they equally affect all conditions, we believe that these modifications still allow us to examine our research questions as we want to evaluate users' performance with different input and output devices, rather than evaluating the task.

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 DOF interaction to make the interaction mapping consistent with that used in the docking task and to provide more flexible input. Although adding more DOF may appear to complicate the task, pilot study participants said such compromise makes interaction actions more predictable. It thus becomes easier compared to 3-DOF 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 visualizations of a small/medium size of volume, such as a human brain (Everts et al., 2015) or other biological




Figure 3: Illustration of interaction mapping. The first row represents the interaction motion, where each column represents a specific interaction DOF. The second row illustrates the interaction mapping for the mouse and keyboard, the third row illustrates the interaction mapping for the space mouse and the fourth row illustrate the interaction mapping for the tangible tablet.

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 for our second task from our collaborators in fluid mechanics, and this is a volume dataset 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.

4.2. Apparatus, interaction mapping, and implementation

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 (Wang et al., 2019b). We thus directly use the Microsoft's HoloLens to investigate our research hypothesis with our envisioned setup as stated in Section 1. Nonetheless, we balanced the view between the two types of devices as much as possible, with the exception of the inherent difference with respect to 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 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 quickly 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 manipulation task. While participants manipulate the target object or camera, we also adjust the object to dock or the clipping plane accordingly, such that their relative position and orientation does not change. Another inherent difference is that, with the HoloLens, users are able to change views faster using head movements and can thus better perceive 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 HoloLens and the 2D screen (as stated in *H2*), rather than simply comparing mono- with stereoscopic views. We also require participants to remain seated in their chair for the study because we want to simulate the envisioned scenario where scientists use the AR extensions with existing tools on desktop in their office. We use the input devices described in Section 3 that include a mouse, a SpaceMouse, 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 affects the rotation around the z -axis and the right button affect the translation along the z -axis. According to Besançon et al. (2017b), this mapping is frequently used in PC-based 3D modeling and visualization software. For




Figure 4: Illustration of the experimental setup. (a) users use a mouse or a SpaceMouse with a screen; (b) users use a mouse or a SpaceMouse with a HoloLens; (c) users use a tablet with a screen; (d) users use a tablet with a HoloLens.

the translation with the mouse, we captured the translation of the mouse cursor on the PC screen and adjusted the control-to-display (CD) ratio such that the translation distance of the virtual object on the screen is the same as the translation of the mouse cursor. For rotation, we implemented the well-established Arcball interaction (Shoemake, 1992)—translating the mouse cursor from one border to the opposite boarder rotates the virtual object by 180 degrees. Please note that, in the remainder of this manuscript, we loosely refer to “mouse” as “mouse and keyboard” input for the sake of simplicity.

Based on our tests and pilot studies, we choose to use a similar interaction mapping for the SpaceMouse. In particular, our mapping uses the plane of the table as a reference space, similar to the regular mouse (illustrated in Figure 3)—a mapping that used the screen plane as a reference space turned out to be too confusing. We translated the 3D mouse input into the translation vector and the rotation quaternion directly based on its SDK. We then tripled the rotation vectors and used one third of the translation vectors, based on our pilot studies, to make the users feel as natural as possible. Please note that, in CAD environments, a SpaceMouse is typically used with the non-dominant hand because it is combined with keyboard and traditional mouse input. This default use of the SpaceMouse, however, is intended for people who work with such setups for extended periods of time and will, therefore, get sufficient training with their non-dominant hand on the SpaceMouse. Our setup is different, we envision use cases where people have had no prior experience with a SpaceMouse or use it infrequently. We thus cannot rely on much training and, consequently, asked participants to use their dominant hand to avoid skewing their performance due to use of the non-dominant hand.

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 this 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 to support also practical visualization tasks. In addition, we transfer the translation/rotation of the tablet to the virtual object’s movement with a 1:1 mapping.

We illustrate our interaction mappings for the three devices in Figure 3. We implemented 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 the 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 (1975) model to clearly show the teapot as a 3D object, as illustrated in Figure 2(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. Our code is available at github.com/xiyaowang/2D3DInOutPutFor3DVis.

We illustrate the general experiment setup in Figure 4. While performing tasks with the mouse and the SpaceMouse, participants validate a trial or switch between object manipulation and view changes by either clicking on buttons or using pre-defined hotkeys 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.

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 inputs \times 2 output devices). Among the three input devices, the mouse provides 2D input, while the SpaceMouse 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 6 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 saw the same trials in the end (but with a different order, counter-balanced with Latin-Square) for a total of 72 trials per participant ($3 \times 2 \times 6 \times 2 = 72$).

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 5, we detail our methods of computing the accuracy for each task. While precision can also 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 assess users' 3D manipulation performance (e. g., Chen et al., 1988). In addition, we collected our participants' perceived workload,⁴ their preference, and any additional comments they had.

4.4. Procedure

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 them 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 stream.⁵ 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 the training session. Then, we asked each participant to complete six trials for each condition as described above. Upon finishing all trials of one condition, the experimenter helped them to get on or take off the HoloLens to switch the output. After each task, we asked each participant to fill in a questionnaire that recorded their perceived workload, preference, and any additional 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.

4.5. 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 worked in visualization or interaction-related domains, 11 had limited knowledge about the Microsoft HoloLens (10 only tried one or twice in their life), and 8 had experience with tangible interaction from mobile games. None of them were familiar with the SpaceMouse before the experiment.

5. Results

We analyzed the data with estimation techniques using confidence intervals (CIs) and effect sizes instead of p -values. Such methods are now recommended by several research communities, while dichotomous significance tests (and thus the use of p -values) have been criticized for their weakness (e. g., Cumming, 2014; Dragicevic, 2016; Gelman, 2017; Gigerenzer, 2018; Amrhein et al., 2019; Helske et al., 2021). We report our results with a nuanced interpretation by reporting the strength of evidence about the population (Cumming, 2014; Dragicevic, 2016; McCook, 2016; Besançon and Dragicevic, 2019) to avoid dichotomous conclusions. However, it is still possible to relate the CIs we report to p -values (Krzywinski and Altman, 2013; Dragicevic, 2015).

⁴We used Hart and Staveland's NASA Task Load Index (humansystems.arc.nasa.gov/groups/tlx/downloads/TLXScale.pdf).

⁵We used existing software (www.microsoft.com/en-us/p/microsoft-hololens/9nblggh4qwnx) for the streaming.

2D and 3D Input/Output Device Combinations for 3D Visualization




Figure 5: Results of docking task completion time (absolute mean value) in seconds. (a) and (d) represent the results using mouse with the screen and the HoloLens; (b) and (e) represent the results using SpaceMouse with the screen and the HoloLens; and (c) and (f) represent the results using tablet with the screen and the HoloLens. The icons that we used to mark each plot in Figures 5–10 and 12–15 represent the input (top) and output (bottom) devices. A division mark (e. g., in Figure 6) means the ratio. For each plot, the big black point represents the mean, error bar indicates the 95% confidence interval, and the violin plot and orange dots represent the actual distribution of measurements.

All our measurements are strictly positive, so we aggregated the data using log-transformed measurements to compute the geometric means and report anti-logarithm forms to decrease the effect of outliers (Keene, 1995; Dragicevic, 2016), as it is common practice (e. g., Jansen et al., 2013; Le Goc et al., 2016). We used 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. Our analysis is mainly based on performance measurements (time and accuracy). We used the perceived workload and preference as a way to gather additional insights, so we only report major findings for them. We still provide all measurements in our OSF repository at osf.io/7gsk8.

5.1. Docking task

Task completion time. We report the absolute mean values of task completion time in seconds for each condition in Figure 5, the pair-wise ratios for each input between the two different outputs (screen vs. HoloLens), and the pair-wise ratios between different input devices for the same output in Figure 6. While users use the mouse as input, the average completion time 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 SpaceMouse, 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 for the tablet, it is 51.05s (CI [41.67s, 62.54s]) with the screen and 48.97s (CI [38.89s, 61.67s]) with the HoloLens. For any input device, the time difference between the output devices is quite small and the CIs largely overlap, so they do not give us enough evidence to conclude effects. We thus further examine the pair-wise ratio (screen/HoloLens). For each of the input devices (mouse, SpaceMouse, 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 the confidence interval. 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 SpaceMouse 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 seems to be longer with the mouse than the SpaceMouse, but the effect remains small. With the HoloLens, this ratio is 1.09 (CI [0.97, 1.23]), which leads to a similar conclusion. 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 1.33 (CI [1.16, 1.53]). As for the ratio between the SpaceMouse and the tablet, the evidence would indicate that the tablet is faster than the SpaceMouse 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 SpaceMouse, and both 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.

Euclidean distance. We report the absolute mean values of the Euclidean distance for each of the different conditions in Figure 7, in virtual 3D space units. We also report the pair-wise ratios for each input between the two different outputs (2D screen vs. HoloLens), and the pair-wise ratios between different input devices for the same output in Figure 8. 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 ($d = \|O_o - O_t\|$). The larger its absolute value is, the less accurate the trial was. 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]) when working with the HoloLens. The largely overlapping

2D and 3D Input/Output Device Combinations for 3D Visualization




Figure 6: Results of docking task completion time effect sizes. (a), (b), and (c) represent pair-wise ratio (screen/HoloLens) the mouse, the SpaceMouse, and the tangible tablet. (d) represents the ratio mouse/SpaceMouse with the screen and (g) with the HoloLens. (e) represents the ratio mouse/tablet mouse with the screen and (h) with the HoloLens. (f) represents the ratio SpaceMouse/tablet with the screen and (i) with the HoloLens.




Figure 7: Results of docking task: Euclidean difference between object and target, in virtual space units. (a) and (d) represent the results using mouse with the screen and the HoloLens; (b) and (e) represent the results using SpaceMouse with the screen and the HoloLens; and (c) and (f) represent the results using tablet with the screen and the HoloLens.




Figure 8: Results of docking task: Euclidean difference effect sizes. (a), (b), and (c) represent pair-wise ratio (screen/HoloLens) the mouse, the SpaceMouse, and the tangible tablet. (d) represents the ratio mouse/SpaceMouse with the screen and (g) the HoloLens. (e) represents the ratio mouse/tablet mouse with the screen and (h) the HoloLens. (f) represents the ratio SpaceMouse/tablet with the screen and (i) the HoloLens.

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 SpaceMouse 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. Our participants were thus generally more accurate when working with the HoloLens than when working with the screen, for any input device. We also examined the difference between input devices for the same output. For working with the screen, we found no evidence that participants performed differently using the mouse or the SpaceMouse with the ratio being 1.01 (CI [0.79, 1.36]), or with the HoloLens, with the ratio being 0.91 (CI [0.69, 1.27]). When comparing the mouse to the tablet, our data shows that mouse is more accurate than the tablet with both types of output (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 SpaceMouse is less accurate than the tablet when working with the screen (ratio is 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 for a performance difference between mouse and SpaceMouse because the mean value of the pairwise ratio is close to 1 and its CI is large. We have weak evidence, however, for the mouse being more accurate than the tablet and for the tablet being more accurate than the SpaceMouse. This result does not contradict our observation that no evidence supports an effect between mouse and SpaceMouse as a lack of evidence does not mean that there is no difference.

Angular distance. We report the absolute mean values of angular distance a in degrees for each of the different conditions in Figure 9, the pair-wise ratios for each input between the two different outputs (screen vs. HoloLens), and the pair-wise ratios between different input devices for the same output in Figure 10. The angle represents the final

2D and 3D Input/Output Device Combinations for 3D Visualization




Figure 9: Absolute mean values of docking task angular difference between the object and the target, in degree. (a) and (d) represent the results using mouse with the screen and the HoloLens; (b) and (e) represent the results using SpaceMouse with the screen and the HoloLens; and (c) and (f) represent the results using tablet with the screen and the HoloLens.




Figure 10: Absolute mean values of docking task angular difference effect sizes. (a), (b), and (c) represent pair-wise ratio (screen/HoloLens) of the mouse, the SpaceMouse, and the tangible tablet. (d) represents the ratio mouse/SpaceMouse with the screen and (g) with the HoloLens. (e) represents the ratio mouse/tablet mouse with the screen and (h) with the HoloLens. (f) represents the ratio SpaceMouse/tablet with the screen and (i) with the HoloLens.




Figure 11: Perceived 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.

difference of orientation between object and target, whose orientations are represented by the quaternions q_o and q_t , respectively. We compute a as $2 \cdot \arccos(q_{d\omega})$ where $q_{d\omega}$ is the ω component of $q_d = q_o^{-1} \cdot q_t$. The bigger its absolute value, 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 SpaceMouse and screen, 7.47° (CI [5.99° , 9.30°]) with SpaceMouse 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 SpaceMouse it was 1.04 (CI [0.90 , 1.28]), and for the tablet it was 0.97 (CI [0.83 , 1.13]). While working with the screen, the ratio between mouse and SpaceMouse 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 SpaceMouse and the tablet it was 1.05 (CI [0.88 , 1.37]). 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 a difference of angular accuracy across techniques between object and target.

Perceived workload. We report the participants' ranked workload in Figure 11. Our evidence shows that mental and physical demands vary for different input devices, while there is no evidence that would show that the output device plays a role. Specifically, our participants' self-reported mental demand was the highest for the mouse. 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 SpaceMouse requires the least since it stays on a fixed position on the desk. We did not find any evidence showing a difference for other workloads, which was confirmed by the verbal comments we received from many participants.

Preference. Participants also ranked their preference for input device-output setting combinations from 1 to 6, with

Table 1

Users' self-rated preference for different input/output combinations, for both docking and clipping plane tasks.

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




Figure 12: Results of clipping task completion time in seconds. (a) and (d) represent the results using mouse with the screen and the HoloLens; (b) and (e) represent the results using SpaceMouse with the screen and the HoloLens; and (c) and (f) represent the results using tablet with the screen and the HoloLens.




Figure 13: Results of clipping task completion time effect sizes. (a), (b), and (c) represent pair-wise ratio (screen/HoloLens) the mouse, the SpaceMouse, and the tangible tablet. (d) represents the ratio mouse/SpaceMouse with the screen and (g) with the HoloLens. (e) represents the ratio mouse/tablet mouse with the screen and (h) with the HoloLens. (f) represents the ratio SpaceMouse/tablet with the screen and (i) with the HoloLens.

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 1, which shows an almost evenly distributed preference for this task, regardless of input or output device.

5.2. Clipping plane manipulation task

Task completion time. We present the results of mean task completion time in seconds in Figure 12, and the pairwise ratios across different conditions in Figure 13. We 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 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 SpaceMouse, the average time using the screen is 53.18s (CI [42.59s, 66.42s]) and 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 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 that users are generally faster working with the HoloLens than with the screen, both with 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 mouse and SpaceMouse is 1.09 (CI [0.93, 1.28]) and the ratio between mouse and tablet is 0.93 (CI [0.80, 1.09]). In both cases, we cannot find evidence to claim an effect, but the ratio between SpaceMouse and tablet is 0.85 (CI [0.71, 1.01]), which suggests that the SpaceMouse is faster than the tablet, while users are working with the screen. The results for the HoloLens are different from those of the screen. We found evidence suggesting that the use of the mouse resulted in slower interaction than the SpaceMouse, the ratio between these two conditions being 1.22 (CI [1.11, 1.36]). We also found evidence for interactions with the mouse being slower than with the tablet, given the ratio being 1.21 (CI [1.05, 1.41]). But we found no evidence for a difference between SpaceMouse and tablet when they are paired with the HoloLens, with the ratio being 0.99 (CI [0.87, 1.13]).

Accuracy. We show the results of accuracy in virtual units in Figures 14 and 15. 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. The intersection of a plane with the data volume box can take different forms (triangle, rectangle,




Figure 14: Results of clipping task accuracy, measured in virtual units. (a) and (d) represent the results using mouse with the screen and the HoloLens; (b) and (e) represent the results using SpaceMouse with the screen and the HoloLens; and (c) and (f) represent the results using tablet with the screen and the HoloLens.




Figure 15: Results of clipping task accuracy effect sizes. (a), (b), and (c) represent pair-wise ratio (screen/HoloLens) the mouse, the SpaceMouse, and the tangible tablet. (d) represents the ratio mouse/SpaceMouse with the screen and (g) with the HoloLens. (e) represents the ratio mouse/tablet mouse with the screen and (h) with the HoloLens. (f) represents the ratio SpaceMouse/tablet with the screen and (i) with the HoloLens.

pentagon, etc.). We only kept target planes in a rectangular form (4 corners) for our experiments to lower the participants' mental workflow—to avoid them being confused about the shape of the target plane. The clipping plane (because it is theoretically infinite) is defined by an artificial 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 $d_i = N_{plane} \cdot (O_{point_i} - O_{plane})$. 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 $\frac{1}{4} \sum_{n=1}^4 \|d_i\|$. For mouse input, this mean accuracy was 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 was 0.79 (CI [0.70, 0.90]). This CI does not overlap with the value 1, which confirms the effect. This observation also partly applies to the tablet, but the evidence is weaker. With the tablet, the mean value was 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 was 0.88 (CI [0.76, 1.02]). For the SpaceMouse, we did not find evidence that would suggest a difference between the two different output devices. The mean value was 5.06 (CI [4.44, 5.78]) with the screen and 5.35 (CI [4.82, 5.72]) with the HoloLens, and the pairwise ratio was 0.95 (CI [0.83, 1.07]). For users working with the screen, our evidence suggests that the mouse was more accurate than both the SpaceMouse (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 was 1.01 with CI [0.88, 1.19]). For the HoloLens, we only have weak evidence for the mouse being more accurate than the SpaceMouse (ratio 0.95, CI [0.85, 1.04]). We found no evidence to suggest a difference between mouse and tablet (ratio 1.01, CI [0.87, 1.12]). There is also almost no evidence for a possible effect between SpaceMouse and tablet (ratio 0.94, CI [0.84, 1.09]).

Perceived workload. For the perceived workload of the clipping plane manipulation task (Figure 16), the situation differs from the docking task. Although the difference is low, we observed an increased mental demand, temporal demand, and effort when working with the tablet. For the physical demand, compared with the docking task, the difference was that the mouse had a lower average demand than the SpaceMouse. Participants reported that manipulating a clipping plane requires small and precise changes of the plane, where the DOF 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.

Preference. We summarize the preference ratings of our participants for the clipping plane manipulation task in Table 1, computed the same way as for the docking task. For this task, the participants' preference seems pretty obvious, the mouse was preferred over the SpaceMouse, which in turn was preferred over the tablet. Interestingly, in all cases of input, the difference between using the 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




Figure 16: Perceived workloads for the clipping plane manipulation 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.

participants' actual recorded ratings, in most cases they preferred the screen over the HoloLens.

6. Discussion

We now summarize our findings and discuss their context, before analyzing the limitations of our empirical experiment.

6.1. Findings

For the 3D docking task that we used as a proxy for general 3D data exploration using object or view manipulation, our participants slightly preferred the 6DOF SpaceMouse. Nonetheless, our measurements show that the tablet was the fastest input device. Also, our participants managed to accomplish the task as accurately using the mouse as with the SpaceMouse, both being more accurate than the tablet. In the second task, our participants reported a preference toward the mouse, 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 performance 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 performance increase over 2D output when performing the clipping task with 3D output. Here, the better spatial understanding of the plane due to the AR display (with both stereoscopic projection and view-motion parallax depth clue) 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 manipulation task, the mouse is preferred because of its separation of DOFs—our participants reported that it allowed them to adjust the data more precisely. The inherent DOF separation of 2D (i. e., 2 DOF) input devices could thus better match the 2D output space in such situations. We note that a clipping plane rendering with stronger support of depth perception (e. g., via gridlines) could have, at least partially, influenced these results. Yet, as we aimed to understand dimensionality mismatch, in particular, between input and output, we believe that these results are still particularly interesting although their effect size could be reduced with more optimal visualization strategies. Another possibility is that many users are generally more familiar with the mouse interaction (see Section 6.2), so many of them may feel that the mouse interaction is more natural and less demanding. For example, expert users may use the mouse only with their fingers and hands while they need to move their arms a lot for the tangible tablet.

Another interesting and essential point is that we found, for both tasks with AR output, mouse input to have either equal or better performance when compared to its performance with 2D projection. The mismatch between its 2D input to the used 3D output space thus does not seem to be a problem for participants. 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 an effective interaction tool, in particular if used in combination with AR output. A hybrid visualization system that relies on mouse input for some non-spatial tasks anyway could thus use the mouse as a single type of input.

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 slight 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 manipulation task, we also saw faster task completion times using the tablet with the HoloLens than with the 2D screen, while the mouse and the SpaceMouse showed no evidence of 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 was 4.26, while it was 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' subjective feeling. **H2** is thus partly supported. As a result, our findings do not support those drawn by Schultheis et al. (2012) who only used time as a measure. Our experimental data seems to indicate that the different output spaces do affect users' performance for precise spatial control.

The basis of our final hypothesis **H3**—as explained in Section 3 as a traditional 3D data analysis setup (using a mouse and a 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 DOF 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 to be 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 1, 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 DOF 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 DOF 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 DOF or 6 DOF—had better performance in the 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 fundamentally change. 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 whether that the 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

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

the clipping plane manipulation task, we do not find a clear result with respect to **H3** either. Here, the performance increase of the 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 such as with a tablet is suitable for interactions that need fast adjustment but do not require precise control, as our study shows that users finished the docking task with the least amount of time. When it comes to tasks that require more accurate input, low-DOF input devices (such as a regular mouse) that could easily separate interaction DOF are generally preferred.

Our failure to find a simple conclusion of the effect of dimensionality (mis-)match does not decrease the importance of our study, however, we now know that, for 3D manipulation, the effect of dimensionality match (if any) is expected to be small, potentially to the point that it can be neglected when choosing input and output devices for 3D visualization. We are free to choose appropriate devices according to specific tasks and other needs, rather than considering if the DOF of the input matches the output. Even more importantly, we see from the results that the mouse still performs 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 system where the mouse is naturally used for the PC part and this finding sheds light on how to assist transitions around immersive environments, an important challenge of immersive analytics (Ens et al., 2021).

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 input devices we have tested as follows. According to our participants, the mouse is precise but its inherently needed 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 SpaceMouse, they liked its fluidity and stated that it required the least amount of effort, thus causing minimal fatigue levels. Nonetheless, the interaction mapping can be sometimes confusing due to the potentially different reference frames between the manipulation space \mathcal{M} , where the SpaceMouse is physically located, and the controlled object or space. Moreover, certain motions are difficult to perform with the SpaceMouse. For example, one participant reported that, while interacting with the SpaceMouse with the right hand, rotating around the y -axis is much more difficult than rotating around the x -axis (refer to Figure 3) because the latter needs to bend the wrist in an uncomfortable way. In general, however, users see the SpaceMouse as a flexible input device. This property, though can sometimes be an advantage, as it also increases the error rate and makes it difficult to perform tiny adjustments. However, SpaceMouse is reported to require the least amount of interaction effort among the three devices we examined, and users thus believe that they may use it more effectively with more practice. This is partly the reason why users slightly prefer the SpaceMouse 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 SpaceMouse, 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 they moved their head to get a better view in the clipping plane manipulation 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.

6.2. Limitations

As is the case with all empirical work, our experiment is subject to a number of limitations. We detail them and analyze their impact on the generalizability of our findings below.

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 only investigated rather simple scenes with few objects, while, in some scenarios, experts need to analyze very large, complex visualizations. While the use of a simple 3D scene as a proxy to more complex, yet fundamentally similar visualization tasks is not uncommon in the visualization literature (e. g., Yu et al., 2010, 2016;

Besaçon et al., 2017b; Büschel et al., 2019), we cannot estimate how our results would fare in such contexts.

Another limitation is that, even though we manually adjusted a lot parameters on both 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 as a regular screen but provides advantages on depth perception. Similarly, as we mentioned in Section 4.2, our use of the SpaceMouse does not reflect how a trained user would use it to perform 3D manipulations. More generally, a key limitation of our experiment also lies in participants' overall expertise with the different devices they used. It is undeniable that participants have more expertise with classical devices such as a mouse or a 2D screen than they would have with a SpaceMouse or a HoloLens. The novelty effect of new devices or the familiarity of older devices are often mentioned as limitations of experimental setups (e. g., Besaçon et al., 2017b). Nonetheless, HCI and visualization study designs can rarely work around this limitation (e. g., Sears and Shneiderman, 1991; Balakrishnan et al., 1997; Forlines et al., 2007; Tuddenham et al., 2010; Yu et al., 2010; Glesser et al., 2013; Besaçon et al., 2017b; Wang et al., 2019a). While more expertise with the SpaceMouse could improve the results that participants obtained with it, we hypothesize based on our acquired data that such expertise would have a moderate influence on our results. Nonetheless, future work should investigate the reliability of our results within a longitudinal study in which participants' expertise would be a direct measure.

On a similar note, as mentioned in Section 4.1, we focused our investigation on 3D manipulations of objects (with a docking task) and clipping planes. We motivated this choice with our need to better understand the consequences of an input/output mismatch for 3D visualization tasks. 3D navigation and clipping planes manipulations are both essential tasks of 3D visualizations (Röttger et al., 2011; Keefe and Isenberg, 2013; Lexow et al., 2016; Palomar et al., 2017; Besaçon et al., 2021). Other tasks, however, are often involved in 3D visualization cases such as drilling (e. g., Klein et al., 2012), peeling (e. g., Sultanum et al., 2011), 3D selection (e. g., Yu et al., 2016; Besaçon et al., 2019), particle placement (e. g., Sobel et al., 2004; Besaçon et al., 2017a), or annotations (e. g., Song et al., 2011). Our work is an initial investigation on the impact of dimensionality mismatch between input and output for the tasks that we considered essential to study first. We hope, however, that our work will inspire more research on the topic for the whole breadth of visualization tasks. In addition, while we think that the visualization community can get inspired by our findings on dimensionality mismatch for the specific devices we selected as representatives, we should also highlight that many other devices could have been studied (e. g., the many prototypes of augmented mouse that we mention in Section 2) and that our results should not be taken as generalizable to all input and output devices, or all of their possible configurations (including mappings, C/D ratios). For example, the CHARM device by Klamka et al. (2019) may be a good alternative, especially if users want to walk around and thus cannot use 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 SpaceMouse we used). Further studies are thus needed to confirm or contrast our findings, not only for the input but also for the output devices. Indeed, even though our study code and implemented interaction techniques can be used directly with a VR headset, we only focused 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) could also leverage them as an interesting question to be investigated in the future. We envision that, due to the inherent differences between VR and AR headsets, our findings would not directly translate to a VR setup where users are occluded from the real world as this may require visual representations of each input device, preventing features such as walking-around.

7. Conclusion and future work

Following our vision of extending traditional workstations 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 that the mouse (and keyboard) remains an efficient 3D interaction tool for high-accuracy tasks, which are often needed in 3D visualization applications. We have also found no evidence that the match or mismatch of dimensionality between input and output devices would have much impact on performance or preferences. Our results thus serve as a starting point to guide the interaction design on an AR-extended hybrid visualization system in which the mouse and keyboard are still necessary for using existing analysis software. Since we started with fundamental, yet simple tasks, it would be beneficial to continue investigating if the results still hold in more complex visualization scenarios. Also, considering the feasibility of a controlled study, we used a fixed interaction mapping, gain factor, and task setups, a future study with different setups

(such as the initial position and orientation offsets for the docking task) would make the conclusion more generalizable. In addition, our results seem to highlight that all devices can provide (sometimes equally) good performance, further underlining that the choice of a device over another could be context-dependent, and that a classical mouse is not particularly out-of-its depth when it comes to immersive 3D manipulations (Wang et al., 2020). Yet, as highlighted in the limitations of our experimental design (Section 6.2), considering the wide variety of possible devices, 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. It would also be useful to continue to study other types of data, including but not limited to abstract information and graphs because they are also essential forms of information—even for 3D data analysis. Nonetheless, as new interaction devices are only slowly adopted by domain experts analyzing their data (e. g., Keefe and Isenberg, 2013; Jackson et al., 2013; Wang et al., 2020; Besançon et al., 2021), and since studying AR and desktop integration and the potential input/output mismatch or transitions has even been recently highlighted as one of the grand challenges of immersive analytics (Ens et al., 2021), we argue that our findings are particularly relevant to real-world applications of visualization.

Acknowledgments

We thank all participants of our experiment for their input. The authors of this article are employed by Inria, France, Linköping University, Sweden, University of Paris 8, France, and Inria, France, respectively. The work was solely supported through these employments. We also thank Yiran Zhang for the sketches and Sabrina Lakhdir for her final proofreading.


References

- Accot, J., Zhai, S., 1999. Performance evaluation of input devices in trajectory-based tasks: An application of the steering law, in: Proc. CHI, ACM, New York. pp. 466–472. doi:10.1145/302979.303133.
- Albinsson, P.A., Zhai, S., 2003. High precision touch screen interaction, in: Proc. CHI, ACM, New York. pp. 105–112. doi:10.1145/642611.642631.
- Amrhein, V., Trafimow, D., Greenland, S., 2019. Inferential statistics as descriptive statistics: There is no replication crisis if we don't expect replication. *The American Statistician* 73(S1), 262–270. doi:10.1080/00031305.2018.1543137.
- Bach, B., Sicat, R., Beyer, J., Cordeil, M., Pfister, H., 2018. The hologram in my hand: How effective is interactive exploration of 3D visualizations in immersive tangible augmented reality? *IEEE Trans. Vis. Comput. Graph.* 24(1), 457–467. doi:10.1109/TVCG.2017.2745941.
- Balakrishnan, R., Baudel, T., Kurtz, G., Fitzmaurice, G., 1997. The Rockin' Mouse: Integral 3D manipulation on a plane, in: Proc. CHI, ACM, New York. pp. 311–318. doi:10.1145/258549.258778.
- Balakrishnan, R., Patel, P., 1998. The PadMouse: Facilitating selection and spatial positioning for the non-dominant hand, in: Proc. CHI, ACM, New York. pp. 9–16. doi:10.1145/274644.274646.
- Bérard, F., Ip, J., Benovoy, M., El-Shimy, D., Blum, J.R., Cooperstock, J.R., 2009. Did “Minority Report” get it wrong? Superiority of the mouse over 3D input devices in a 3D placement task, in: Proc. INTERACT, Springer, Berlin. pp. 400–414. doi:10.1007/978-3-642-03658-3_45.
- Besançon, L., Dragicevic, P., 2019. The continued prevalence of dichotomous inferences at CHI, in: CHI Extended Abstracts, ACM, New York. pp. alt14:1–alt14:11. doi:10.1145/3290607.3310432.
- Besançon, L., Issartel, P., Ammi, M., Isenberg, T., 2017a. Hybrid tactile/tangible interaction for 3D data exploration. *IEEE Trans. Vis. Comput. Graph.* 23(1), 881–890. doi:10.1109/TVCG.2016.2599217.
- Besançon, L., Issartel, P., Ammi, M., Isenberg, T., 2017b. Mouse, tactile, and tangible input for 3D manipulation, in: Proc. CHI, ACM, New York. pp. 4727–4740. doi:10.1145/3025453.3025863.
- Besançon, L., Sereno, M., Yu, L., Ammi, M., Isenberg, T., 2019. Hybrid touch/tangible spatial 3D data selection. *Comput. Graph. Forum* 38(3), 553–567. doi:10.1111/cgf.13710.
- Besançon, L., Ynnerman, A., Keefe, D.F., Yu, L., Isenberg, T., 2021. The state of the art of spatial interfaces for 3D visualization. *Comput. Graph. Forum* 40(1), 293–326. doi:10.1111/cgf.14189.
- Bruckner, S., Isenberg, T., Ropinski, T., Wiebel, A., 2019. A model of spatial directness in interactive visualization. *IEEE Trans. Vis. Comput. Graph.* 25(8), 2514–2528. doi:10.1109/TVCG.2018.2848906.
- Bryson, S., 1996. Virtual reality in scientific visualization. *Commun. ACM* 39(5), 62–71. doi:10.1145/229459.229467.
- Büschel, W., Mitschick, A., Meyer, T., Dachselt, R., 2019. Investigating smartphone-based pan and zoom in 3D data spaces in augmented reality, in: Proc. MobileHCI, ACM, New York. pp. 2:1–2:13. doi:10.1145/3338286.3340113.
- Buxton, B., 2007. Multi-touch systems that I have known and loved. Web site: <http://www.billbuxton.com/multitouchOverview.html>.
- Cechanowicz, J., Irani, P., Subramanian, S., 2007. Augmenting the mouse with pressure sensitive input, in: Proc. CHI, ACM, New York. pp. 1385–1394. doi:10.1145/1240624.1240835.
- Chen, M., Mountford, S.J., Sellen, A., 1988. A study in interactive 3-D rotation using 2-D control devices. *SIGGRAPH Comput. Graph.* 22(4), 121–129. doi:10.1145/54852.378497.
- Cockburn, A., Dragicevic, P., Besançon, L., Gutwin, C., 2020. Threats of a replication crisis in empirical computer science. *Commun. ACM* 63(8), 70–79. doi:10.1145/3360311.
- Cockburn, A., Gutwin, C., Dix, A., 2018. HARK no more: On the preregistration of CHI experiments, in: Proc. CHI, ACM, New York. pp. 141:1–141:12. doi:10.1145/3173574.3173715.

- Cohen, O., Meyer, S., Nilsen, E., 1993. Studying the movement of high-tech Rodentia: Pointing and dragging, in: Proc. CHI, ACM, New York. pp. 135–136. doi:10.1145/259964.260149.
- Cumming, G., 2014. The new statistics: Why and how. *Psychological Science* 25(1), 7–29. doi:10.1177/0956797613504966.
- Dragicevic, P., 2015. HCI Statistics without p-values. Research Report RR-8738. Inria. URL: <https://hal.inria.fr/hal-01162238>.
- Dragicevic, P., 2016. Fair statistical communication in HCI, in: Robertson, J., Kaptein, M. (Eds.), *Modern Statistical Methods for HCI*. Springer International Publishing, Cham, Switzerland. chapter 13, pp. 291–330. doi:10.1007/978-3-319-26633-6_13.
- Dwyer, T., Marriott, K., Isenberg, T., Klein, K., Riche, N., Schreiber, F., Stuerzlinger, W., Thomas, B., 2018. Immersive analytics: An introduction, in: Marriott, K., Schreiber, F., Dwyer, T., Klein, K., Riche, N.H., Itoh, T., Stuerzlinger, W., Thomas, B.H. (Eds.), *Immersive Analytics*. Springer, Berlin/Heidelberg. chapter 1, pp. 1–23. doi:10.1007/978-3-030-01388-2_1.
- Ens, B., Bach, B., Cordeil, M., Engelke, U., Serrano, M., Willett, W., Prouzeau, A., Anthes, C., Büschel, W., Dunne, C., Dwyer, T., Grubert, J., Haga, J.H., Kirshenbaum, N., Kobayashi, D., Lin, T., Olaosebikan, M., Pointecker, F., Saffo, D., Saquib, N., Schmalstieg, D., Szafir, D.A., Whitlock, M., Yang, Y., 2021. Grand challenges in immersive analytics, in: Proc. CHI, ACM, New York. pp. 459:1–459:17. doi:10.1145/3411764.3446866.
- Everts, M.H., Begue, E., Bekker, H., Roerdink, J.B.T.M., Isenberg, T., 2015. Exploration of the brain's white matter structure through visual abstraction and multi-scale local fiber tract contraction. *IEEE Trans. Vis. Comput. Graph.* 21(7), 808–821. doi:10.1109/TVCG.2015.2403323.
- Fallot, W., Fjeld, M., Speirs, C., Ziegenspeck, S., Krueger, H., Läubli, T., 2006. Touch&Type: A novel pointing device for notebook computers, in: Proc. NordiCHI, ACM, New York. pp. 465–468. doi:10.1145/1182475.1182538.
- Filho, J.A.W., Freitas, C.M.D.S., Nedel, L., 2019. Comfortable immersive analytics with the VirtualDesk metaphor. *IEEE Comput. Graph. Appl.* 39(3), 41–53. doi:10.1109/MCG.2019.2898856.
- Fitzmaurice, G.W., 1993. Situated information spaces and spatially aware palmtop computers. *Commun. ACM* 36(7), 39–49. doi:10.1145/159544.159566.
- Forlines, C., Wigdor, D., Shen, C., Balakrishnan, R., 2007. Direct-touch vs. mouse input for tabletop displays, in: Proc. CHI, ACM, New York. pp. 647–656. doi:10.1145/1240624.1240726.
- Fruchard, B., Prouzeau, A., Chapuis, O., Lecolinet, E., 2019. Leveraging body interactions to support immersive analytics, in: CHI Workshop on Immersive Analytics, Glasgow, United Kingdom. URL: <https://hal.archives-ouvertes.fr/hal-02095993>.
- Fulmer, W., Mahmood, T., Li, Z., Zhang, S., Huang, J., Lu, A., 2019. ImWeb: Cross-platform immersive web browsing for online 3D neuron database exploration, in: Proc. IUI, ACM, New York. pp. 367–378. doi:10.1145/3301275.3302319.
- Gelman, A., 2017. No to inferential thresholds. Online blog post. Last visited June 2020. URL: <https://andrewgelman.com/2017/11/19/no-inferential-thresholds/>.
- Gigerenzer, G., 2018. Statistical rituals: The replication delusion and how we got there. *Advances in Methods and Practices in Psychological Science* 1(2), 198–218. doi:10.1177/2515245918771329.
- Glesser, D., Bérard, F., Cooperstock, J.R., 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.
- Hancock, M., Carpendale, S., Cockburn, A., 2007. Shallow-depth 3D interaction: Design and evaluation of one-, two- and three-touch techniques, in: Proc. CHI, ACM, New York. pp. 1147–1156. doi:10.1145/1240624.1240798.
- Helske, J., Helske, S., Cooper, M., Ynnerman, A., Besançon, L., 2021. Can visualization alleviate dichotomous thinking? Effects of visual representations on the cliff effect. *IEEE Trans. Vis. Comput. Graph.* 27(8), 3397–3409. doi:10.1109/TVCG.2021.3073466.
- Henrysson, A., Billinghurst, M., Ollila, M., 2005. Virtual object manipulation using a mobile phone, in: Proc. ICAT, ACM, New York. pp. 164–171. doi:10.1145/1152399.1152430.
- Hinckley, K., Sinclair, M., Hanson, E., Szeliski, R., Conway, M., 1999. The videomouse: A camera-based multi-degree-of-freedom input device, in: Proc. UIST, ACM, New York. pp. 103–112. doi:10.1145/320719.322591.
- Hinckley, K., Tullio, J., Pausch, R., Proffitt, D., Kassell, N., 1997. Usability analysis of 3D rotation techniques, in: Proc. UIST, ACM, New York. pp. 1–10. doi:10.1145/263407.263408.
- Hurter, C., Riche, N.H., Drucker, S.M., Cordeil, M., Alligier, R., Vuillemot, R., 2019. FiberClay: Sculpting three dimensional trajectories to reveal structural insights. *IEEE Trans. Vis. Comput. Graph.* 25(1), 704–714. doi:10.1109/TVCG.2018.2865191.
- Ishii, H., 2008. The tangible user interface and its evolution. *Commun. ACM* 51(6), 32–36. doi:10.1145/1349026.1349034.
- Jackson, B., Lau, T.Y., Schroeder, D., Toussaint, K.C., Keefe, D.F., 2013. A lightweight tangible 3D interface for interactive visualization of thin fiber structures. *IEEE Trans. Vis. Comput. Graph.* 19(12), 2802–2809. doi:10.1109/TVCG.2013.121.
- Jacob, R., Leggett, J., Myers, B., Pausch, R., 1993. Interaction styles and input/output devices. *Behav. Inf. Technol.* 12(2), 69–79. doi:10.1080/01449299308924369.
- Jankowski, J., Hachet, M., 2013. A survey of interaction techniques for interactive 3D environments, in: Eurographics – STARS, The Eurographics Association, Goslar, Germany. pp. 65–93. doi:10.2312/conf/EG2013/stars/065-093.
- Jansen, Y., Dragicevic, P., Fekete, J.D., 2013. Evaluating the efficiency of physical visualizations, in: Proc. CHI, ACM, New York. pp. 2593–2602. doi:10.1145/2470654.2481359.
- Kalkofen, D., Sandor, C., White, S., Schmalstieg, D., 2011. Visualization techniques for augmented reality, in: Furht, B. (Ed.), *Handbook of Augmented Reality*. Springer, New York. chapter 3, pp. 65–98. doi:10.1007/978-1-4614-0064-6_3.
- Keefe, D.F., Isenberg, T., 2013. Reimagining the scientific visualization interaction paradigm. *IEEE Computer* 46(5), 51–57. doi:10.1109/MC.2013.178.
- Keene, O.N., 1995. The log transformation is special. *Stat. Med.* 14(8), 811–819. doi:10.1002/sim.4780140810.
- Kirby, K.N., Gerlanc, D., 2013. BootES: An R package for bootstrap confidence intervals on effect sizes. *Behav. Res. Methods* 45(4), 905–927. doi:10.3758/s13428-013-0330-5.
- Klamka, K., Dachselt, R., 2015. Elasticcon: Elastic controllers for casual interaction, in: Proc. MobileHCI, ACM, New York. pp. 410–419. doi:10.1145/2785830.2785849.
- Klamka, K., Reipschläger, P., Dachselt, R., 2019. CHARM: Cord-based haptic augmented reality manipulation, in: Proc. HCII, Springer International,

- Cham, Switzerland. pp. 96–114. doi:10.1007/978-3-030-21607-8_8.
- Klein, T., Guéniat, F., Pastur, L., Vernier, F., Isenberg, T., 2012. A design study of direct-touch interaction for exploratory 3D scientific visualization. *Comput. Graph. Forum* 31(3), 1225–1234. doi:10.1111/j.1467-8659.2012.03115.x.
- Krzywinski, M., Altman, N., 2013. Points of significance: Error bars. *Nat. Methods* 10(10), 921–922. doi:10.1038/nmeth.2659.
- Laha, B., Bowman, D., Socha, J., 2014. Effects of VR system fidelity on analyzing isosurface visualization of volume datasets. *IEEE Trans. Vis. Comput. Graph.* 20(4), 513–522. doi:10.1109/TVCG.2014.20.
- Laha, B., Sensharma, K., Schiffbauer, J.D., Bowman, D.A., 2012. Effects of immersion on visual analysis of volume data. *IEEE Trans. Vis. Comput. Graph.* 18(4), 597–606. doi:10.1109/TVCG.2012.42.
- Lakatos, D., Blackshaw, M., Olwal, A., Barryte, Z., Perlin, K., Ishii, H., 2014. T(ether): Spatially-aware handhelds, gestures and proprioception for multi-user 3d modeling and animation, in: *Proc. SUI, ACM, New York*. pp. 90–93. doi:10.1145/2659766.2659785.
- LaViola, Jr., J.J., Kruijff, E., McMahan, R.P., Bowman, D.A., Poupyrev, I., 2017. *3D User Interfaces: Theory and Practice*. 2nd ed., Addison-Wesley, Boston. URL: <https://www.vitalsource.com/referral?term=9780134034461>.
- Le Goc, M., Dragicevic, P., Huron, S., Boy, J., Fekete, J.D., 2016. A better grasp on pictures under glass: Comparing touch and tangible object manipulation using physical proxies, in: *Proc. AVI, ACM, New York*. doi:10.1145/2909132.2909247.
- Lee, J., Olwal, A., Ishii, H., Boulanger, C., 2013. Spacetop: Integrating 2D and spatial 3D interactions in a see-through desktop environment, in: *Proc. CHI, ACM, New York*. pp. 189–192. doi:10.1145/2470654.2470680.
- Lexow, G.J., Schurz, D., Gellrich, N.C., Lenarz, T., Majdani, O., Rau, T.S., 2016. Visualization, measurement and modelling of the cochlea using rotating midmodiolar slice planes. *Int. J. Comput. Assisted Radiol. Surg.* 11(10), 1855–1869. doi:10.1007/s11548-016-1374-7.
- López, D., Oehlberg, L., Doger, C., Isenberg, T., 2016. Towards an understanding of mobile touch navigation in a stereoscopic viewing environment for 3D data exploration. *IEEE Trans. Vis. Comput. Graph.* 22(5), 1616–1629. doi:10.1109/TVCG.2015.2440233.
- MacKenzie, I.S., Soukoreff, R.W., Pal, C., 1997. A two-ball mouse affords three degrees of freedom, in: *CHI Extended Abstracts, ACM, New York*. pp. 303–304. doi:10.1145/1120212.1120405.
- Marriott, K., Schreiber, F., Dwyer, T., Klein, K., Riche, N.H., Itoh, T., Stuerzlinger, W., Thomas, B.H. (Eds.), 2018. *Immersive Analytics*. Springer, Berlin/Heidelberg. doi:10.1007/978-3-030-01388-2.
- McCook, A., 2016. We're using a common statistical test all wrong. Statisticians want to fix that. Online blog post. Last visited June 2020. URL: <https://retractionwatch.com/2016/03/07/were-using-a-common-statistical-test-all-wrong-statisticians-want-to-fix-that/>.
- McIntire, J.P., Havig, P.R., Geiselman, E.E., 2014. Stereoscopic 3D displays and human performance: A comprehensive review. *Displays* 35(1), 18–26. doi:10.1016/j.displa.2013.10.004.
- Mendes, D., Caputo, F.M., Giachetti, A., Ferreira, A., Jorge, J.A., 2019. A survey on 3D virtual object manipulation: From the desktop to immersive virtual environments. *Comput. Graph. Forum* 38(1), 21–45. doi:10.1111/cgf.13390.
- Millette, A., McGuffin, M.J., 2016. DualCAD: Integrating augmented reality with a desktop GUI and smartphone interaction, in: *Proc. ISMAR-Adjunct, IEEE Computer Society, Los Alamitos*. pp. 21–26. doi:10.1109/ISMAR-Adjunct.2016.0030.
- Morris, M.R., Danielescu, A., Drucker, S., Fisher, D., Lee, B., schraefel, m.c., Wobbrock, J.O., 2014. Reducing legacy bias in gesture elicitation studies. *Interactions* 21(3), 40–45. doi:10.1145/2591689.
- Murray, P., McGee, F., Forbes, A.G., 2017. A taxonomy of visualization tasks for the analysis of biological pathway data. *BMC Bioinf.* 18(2), 21:1–21:13. doi:10.1186/s12859-016-1443-5.
- Palomar, R., Cheikh, F.A., Edwin, B., Fretland, Å., Beghdadi, A., Elle, O.J., 2017. A novel method for planning liver resections using deformable Bézier surfaces and distance maps. *Comput. Methods Programs Biomed.* 144, 135–145. doi:10.1016/j.cmpb.2017.03.019.
- Phong, B.T., 1975. Illumination for computer generated pictures. *Commun. ACM* 18(6), 311–317. doi:10.1145/360825.360839.
- Poupyrev, I., Tomokazu, N., Weghorst, S., 1998. Virtual notepad: Handwriting in immersive VR, in: *Proc. VRAIS, IEEE Computer Society, Los Alamitos*. pp. 126–132. doi:10.1109/VRAIS.1998.658467.
- Prabhat, Forsberg, A., Katzourin, M., Wharton, K., Slater, M., 2008. A comparative study of desktop, fishtank, and Cave systems for the exploration of volume rendered confocal data sets. *IEEE Trans. Vis. Comput. Graph.* 14(3), 551–563. doi:10.1109/TVCG.2007.70433.
- Röttger, D., Engelhardt, S., Müller, S., 2011. Multimodal visualizations for pre-operative neurosurgical planning, in: *Proc. INFORMATIK, Springer, Berlin, Heidelberg*. p. 438. URL: <https://subs.emis.de/LNI/Proceedings/Proceedings192/438.pdf>.
- Schemali, L., Eisemann, E., 2014. Design and evaluation of mouse cursors in a stereoscopic desktop environment, in: *Proc. 3DUI, IEEE Computer Society, Los Alamitos*. pp. 67–70. doi:10.1109/3DUI.2014.6798844.
- Schultheis, U., Jerald, J., Toledo, F., Yoganandan, A., Mlyniec, P., 2012. Comparison of a two-handed interface to a wand interface and a mouse interface for fundamental 3D tasks, in: *Proc. 3DUI, IEEE Computer Society, Los Alamitos*. pp. 117–124. doi:10.1109/3DUI.2012.6184195.
- Sears, A., Shneiderman, B., 1991. High precision touchscreens: design strategies and comparisons with a mouse. *Int. J. Man Mach. Stud.* 34(4), 593–613. doi:10.1016/0020-7373(91)90037-8.
- Shoemake, K., 1992. Arcball: A user interface for specifying three-dimensional orientation using a mouse, in: *Proc. Graphics Interface, Morgan Kaufmann, San Francisco*. pp. 151–156. doi:10.20380/GI1992.18.
- Sobel, J.S., Forsberg, A.S., Laidlaw, D.H., Zeleznik, R.C., Keefe, D.F., Pivkin, I., Karniadakis, G.E., Richardson, P., Swartz, S., 2004. Particle flurries. *IEEE Comput. Graph. Appl.* 24(2), 76–85. doi:10.1109/MCG.2004.1274065.
- Song, P., Goh, W.B., Fu, C.W., Meng, Q., Heng, P.A., 2011. WYSIWYF: Exploring and annotating volume data with a tangible handheld device, in: *Proc. CHI, ACM, New York*. pp. 1333–1342. doi:10.1145/1978942.1979140.
- Spindler, M., Büschel, W., Dachsel, R., 2012. Use your head: Tangible windows for 3D information spaces in a tabletop environment, in: *Proc. ITS, ACM, New York*. pp. 245–254. doi:10.1145/2396636.2396674.
- Sultanum, N., Somanath, S., Sharlin, E., Sousa, M.C., 2011. “Point it, split it, peel it, view it”: Techniques for interactive reservoir visualization on tabletops, in: *Proc. ITS, ACM, New York*. pp. 192–201. doi:10.1145/2076354.2076390.
- Sundin, M., Fjeld, M., 2009. Softly elastic 6 DOF input. *Int. J. Hum. Comput. Interact.* 25(7), 647–691. doi:10.1080/10447310902964124.

- Surale, H.B., Gupta, A., Hancock, M., Vogel, D., 2019. TabletInVR: Exploring the design space for using a multi-touch tablet in virtual reality, in: Proc. CHI, ACM, New York. pp. 13:1–13:13. doi:10.1145/3290605.3300243.
- Szalavári, Z., Gervautz, M., 1997. The personal interaction panel – A two-handed interface for augmented reality. *Comput. Graph. Forum* 16(3), C335–C346. doi:https://doi.org/10.1111/1467-8659.00137.
- Tait, A., 1992. Desktop virtual reality, in: IEEE Colloquium on Using Virtual Worlds, IET. pp. 5:1–5:5. URL: https://ieeexplore.ieee.org/document/168239.
- Teather, R.J., Stuerzlinger, W., 2013. Pointing at 3D target projections with one-eyed and stereo cursors, in: Proc. CHI, ACM, New York. pp. 159–168. doi:10.1145/2470654.2470677.
- Tsang, M., Fitzmaurice, G.W., Kurtenbach, G., Khan, A., Buxton, B., 2002. Boom chameleon: Simultaneous capture of 3D viewpoint, voice and gesture annotations on a spatially-aware display, in: Proc. UIST, ACM, New York. pp. 111–120. doi:10.1145/571985.572001.
- Tuddenham, P., Kirk, D., Izadi, S., 2010. Graspables revisited: Multi-touch vs. tangible input for tabletop displays in acquisition and manipulation tasks, in: Proc. CHI, ACM, New York. pp. 2223–2232. doi:10.1145/1753326.1753662.
- Veit, M., Capobianco, A., Bechmann, D., 2009. Influence of degrees of freedom's manipulation on performances during orientation tasks in virtual reality environments, in: Proc. VRST, ACM, New York. pp. 51–58. doi:10.1145/1643928.1643942.
- Villar, N., Izadi, S., Rosenfeld, D., Benko, H., Helmes, J., Hodges, S., Ofek, E., Butler, A., Cao, X., Chen, B., 2009. Mouse 2.0: Multi-touch meets the mouse, in: Proc. UIST, ACM, New York. pp. 33–42. doi:10.1145/1622176.1622184.
- Wang, X., Besançon, L., Ammi, M., Isenberg, T., 2019a. Augmenting tactile 3D data navigation with pressure sensing. *Comput. Graph. Forum* 38(3), 635–647. doi:10.1111/cgf.13716.
- Wang, X., Besançon, L., Guéniat, F., Sereno, M., Ammi, M., Isenberg, T., 2019b. A vision of bringing immersive visualization to scientific workflows, in: Proc. CHI Workshop on Immersive Analytics. URL: https://hal.inria.fr/hal-02053969.
- Wang, X., Besançon, L., Rousseau, D., Sereno, M., Ammi, M., Isenberg, T., 2020. Towards an understanding of augmented reality extensions for existing 3D data analysis tools, in: Proc. CHI, ACM, New York. pp. 528:1–528:13. doi:10.1145/3313831.3376657.
- Wozniak, P., Fjeld, M., Zhao, S., 2014. Limiting trial and error: Introducing a systematic approach to designing clutching, in: Proc. Chinese CHI, ACM, New York. pp. 35–39. doi:10.1145/2592235.2592241.
- Yee, K.P., 2003. Peephole displays: Pen interaction on spatially aware handheld computers, in: Proc. CHI, ACM, New York. pp. 1–8. doi:10.1145/642611.642613.
- Yu, L., Efstathiou, K., Isenberg, P., Isenberg, T., 2016. CAST: Effective and efficient user interaction for context-aware selection in 3D particle clouds. *IEEE Trans. Vis. Comput. Graph.* 22(1), 886–895. doi:10.1109/TVCG.2015.2467202.
- Yu, L., Svetachov, P., Isenberg, P., Everts, M.H., Isenberg, T., 2010. FI3D: Direct-touch interaction for the exploration of 3D scientific visualization spaces. *IEEE Trans. Vis. Comput. Graph.* 16(6), 1613–1622. doi:10.1109/TVCG.2010.157.
- Zhai, S., 1998. User performance in relation to 3D input device design. *ACM SIGGRAPH Computer Graphics* 32(4), 50–54. doi:10.1145/307710.307728.
- Zhai, S., Milgram, P., 1998. Quantifying coordination in multiple DOF movement and its application to evaluating 6 DOF input devices, in: Proc. CHI, ACM, New York. pp. 320–327. doi:10.1145/274644.274689.
- Zielasko, D., Weyers, B., Bellgardt, M., Pick, S., Meißner, A., Vierjahn, T., Kuhlen, T.W., 2017. Remain seated: Towards fully-immersive desktop VR, in: Proc. WEVR, IEEE Computer Society, Los Alamitos. pp. 1:1–1:6. doi:10.1109/WEVR.2017.7957707.
- Zizka, J., Olwal, A., Raskar, R., 2011. Specklesense: Fast, precise, low-cost and compact motion sensing using laser speckle, in: Proc. UIST, ACM, New York. pp. 489–498. doi:10.1145/2047196.2047261.



Xiyao Wang is a postdoctoral fellow at the University of Victoria in Canada. He received his PhD degree from the University Paris-Saclay and Inria, France. Currently, his work focuses on exploring the visualization with 3D environments. He also studies novel 3D interaction and visualization techniques beyond traditional workstations (e. g., mobile and VR/AR).



Lonni Besançon is a postdoctoral fellow at Linköping University, Sweden. He received his Ph.D in computer science at University Paris Saclay, France. His thesis “An interaction continuum for 3D dataset visualization” received the second prize of the prix de these GDR-IGRV. He is particularly interested in interactive visualization techniques for 3D spatial data relying on new input paradigms and his recent work focuses on the visualization and understanding of uncertainty in empirical results in computer science.



Mehdi Ammi is Professor in Computer Science at the University of Paris 8. He is engineer in electronics (2000), PhD in robotics (2005) and rehabilitated to supervise research (HDR) in computer science since 2013. His research focuses on the study of pervasive environments and their societal and industrial applications (e-health, smart building, factory of the future, etc.). He addresses this field with a multidisciplinary approach combining artificial intelligence (machine learning, neurosciences, etc.), mechatronic technologies (organic electronics, smart textiles, etc.) and human factors (psychology, behavioral analysis, etc.). Mehdi Ammi was the leader or co-leader of several international working groups (IEEE TCH, EuroVR, Haptic SIG, etc.), and was involved in more than 20 national & international projects and industrial collaborations.

2D and 3D Input/Output Device Combinations for 3D Visualization



Tobias Isenberg is a senior research scientist at Inria, France. He received his doctoral degree from the University of Magdeburg, Germany, in 2004. Previously he held positions as post-doctoral fellow at the University of Calgary, Canada, and as assistant professor at the University of Groningen, the Netherlands. His research interests comprise topics in scientific visualization, illustrative and non-photorealistic rendering, and interactive visualization techniques. He is particularly interested in interactive visualization environments for 3D spatial data that rely on novel interaction paradigms such as augmented reality, tactile screens, and tangible devices.