



HAL
open science

Selection Techniques for 3D Extended Desktop Workstation with AR HMD

Carole Plasson, Renaud Blanch, Laurence Nigay

► **To cite this version:**

Carole Plasson, Renaud Blanch, Laurence Nigay. Selection Techniques for 3D Extended Desktop Workstation with AR HMD. 2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR), Oct 2022, Singapore, Singapore. pp.460-469, 10.1109/ISMAR55827.2022.00062. hal-03928037

HAL Id: hal-03928037

<https://hal.science/hal-03928037v1>

Submitted on 6 Jan 2023

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

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

Selection Techniques for 3D Extended Desktop Workstation with AR HMD

Carole Plasson*

Renaud Blanch†

Laurence Nigay†

Univ. Grenoble Alpes, CNRS, Grenoble INP, LIG, F-38000 Grenoble, France

ABSTRACT

Extending a standard desktop workstation (i.e. a screen, a mouse, a keyboard) with virtual scenes displayed on an Augmented Reality Head-Mounted Display (AR HMD) offers many identified advantages including limited physical space requirements, very large and flexible display spaces, and 3D stereoscopic views. While the technologies become more mainstream, the remaining open question is how to interact with such hybrid workstations that combine 2D views displayed on a physical monitor and 3D views displayed on a HoloLens. For a selection task, we compared mouse-based interaction (standard for 2D desktop workstations) and direct touch interaction in mid-air (standard for 3D AR) while considering different positions of the 3D scene according to a physical monitor. To extend mouse-based selection to 3D views, we experimentally explored different interaction metaphors where the mouse cursor moves either on a horizontal or a vertical plane in a 3D virtual scene. To check for ecological validity of our results, we conducted an additional study focusing on interaction with a 2D/3D Gapminder dataset visualization. The results show 1) that the mouse-based interaction, as compared to direct touch interaction in mid-air, is easy and efficient, 2) that using a vertical plane placed in front of the 3D virtual scene to mimic the double screen metaphor outperforms other interaction techniques and 3) that flexibility is required to allow users to choose the selection techniques and to position the 3D virtual scene relative to the physical monitor. Based on these results, we derive interaction design guidelines for hybrid workstations.

Index Terms: Human-centered computing—Visualization—Interaction techniques—Hybrid 2D/3D workstation—AR

1 INTRODUCTION

Extending a standard workstation (i.e. a screen, a keyboard, a mouse) with virtual scenes displayed on an Augmented Reality Head-Mounted Display (AR HMD) offers several advantages. The resulting hybrid workstation (1) enables us to reduce the quantity of information on the physical screen, (2) requires limited physical space on the desk avoiding the need for additional physical monitors, (3) defines a fully flexible display space (e.g., virtual scenes can be placed anywhere), (4) provides 3D stereoscopic views and (5) extends the standard work environment. These advantages have motivated several studies on hybrid workstations in different domains including 3D modeling [19] and information or scientific visualization (financial data [7], particle physics [36]).

In this work, we focus on selection techniques for hybrid workstations that combine 2D views displayed on a physical monitor and 3D views displayed on an AR HMD HoloLens [20]. In AR, mid-air interaction is commonly used: for instance the users directly touch the object of interest with their finger to select it. This interaction is intuitive and direct, but also inaccurate and tiring. The tremor of the hand in the air has a strong impact on interaction precision [27, 35] and interacting in the air causes arm pain, preventing extensive use.

*e-mail: carole.plasson@gmail.com

†e-mail: firstname.lastname@univ-grenoble-alpes.fr

In addition, in hybrid workstations, switching back and forth between mouse interaction (direct manipulation in the 2D Graphical User Interface, GUI, displayed on the physical screen) and mid-air interaction (direct mid-air interaction with stereoscopic AR-based 3D scenes) can be costly as modeled by the Homing Keystroke operator between mouse and keyboard [6].

To avoid these problems, mouse-based interaction with stereoscopic AR-based 3D scenes is a promising solution [33, 36]. Indeed, with several properties contributing to its preeminence [2], mouse-based interaction is well known to users, it enables precise selection and avoids switching between interaction techniques in hybrid workstations.

Our goal is to define a continuous interaction space for hybrid workstations, allowing users to freely and seamlessly switch back and forth between the GUI on the physical screen and the stereoscopic AR-based 3D scene. We experimentally compare this continuous space to the de-facto discontinuous space of the standard mouse+mid-air interaction. There are different ways to position the virtual 3D scene according to the physical monitor and different interaction metaphors for selecting a 3D object using the mouse. Our contributions are twofold: 1) Experimental studies of interaction in hybrid workstations, exploring both the positions of the 3D scene according to the physical monitor and the two mouse/mid-air selection techniques; 2) Guidelines for the design of future hybrid workstation systems. To the best of our knowledge, we are the first to experimentally explore this design space of display arrangement and interaction modality (mouse/mid-air interaction) for seamless hybrid workstations.

In this paper, we review related work on hybrid workstations and interaction techniques with hybrid workstations (Section 2). We then present the design space and the corresponding selection techniques that we experimentally explored (Section 3). The experimental research path that we pursued for exploring this design space is the following. First, we conducted a pilot study to evaluate the usability of the selection techniques identified by the design space and to reduce the number of techniques. After describing an experiment comparing the remaining techniques (Sections 4 & 5), we discuss the results (Section 6). Finally, we developed a hybrid system for exploring Gapminder World multivariate dataset to test interaction in a more ecological context (Section 7). We conclude the paper with the guidelines obtained from these studies.

2 RELATED WORK

For 3D data visualization, prior works showed that stereoscopic 3D views can help users to understand the data [15, 28, 32]. Thus, for 2D and 3D data visualization, hybrid workstations combining a 2D display and a stereoscopic display have been proposed. For instance, Mandalika et al. [18] propose a hybrid system for radiological diagnosis that combines a 2D display with a zSpace stereoscopic display. Users interact with a mouse on the 2D display and with a 3D stylus on the stereoscopic 3D display. Their study shows that a hybrid workstation can improve users' performances (e.g., time, accuracy) as compared to a standard 2D display. However, a stereoscopic display also takes up physical space and is not adaptable (size, position) to users' needs and environments. To address these limitations, recent studies investigate the extension of a standard workstation (i.e. 2D screen, mouse, keyboard) with immersive devices instead of the use of an additional stereoscopic display.

2.1 Hybrid workstations: combining 2D and 3D displays.

Several studies focus on virtual workspaces that combine 2D and 3D virtual displays in Virtual Reality (VR) [13, 24, 38]. To avoid the above issues related to mid-air interaction (midair gestures or handheld controllers), Zhou et al. studied mouse-based interaction in such virtual workspaces [38]: the authors propose Depth-Adaptive Cursor, a mouse-based pointing technique for 3D selection with depth adaptation. As opposed to VR, Augmented Reality (AR) allows users to keep their environment in sight. This makes AR suitable for professional environments. Indeed, a user can continue to use the standard workstation, to take notes on paper and talk to colleagues. For these reasons, we focus on the combination of a standard workstation and AR instead of VR. The recent study by Pavanatto et al. [25] further motivates our focus on hybrid workstations that combine a physical screen with AR-based virtual screens. The authors studied the effect of replacing or extending physical screens with virtual ones for productivity work. They compared 3 conditions: (1) a condition with only physical screens, (2) a condition with only virtual screens, and (3) a condition combining a physical screen and virtual screens. Their study highlights the advantages of virtual screens including no need for physical space, adaptable shapes and sizes, unlimited number of screens, portability, and content privacy. The study also reveals limitations. Using only virtual screens was slower than using physical screens: this is explained by the limited field of view and the low resolution of the virtual screens. Also, reading text on a virtual screen causes eye fatigue. Their study finally shows that the hybrid condition can be a good trade-off between performance and adaptability.

Beyond the advantages of combining a 2D screen with 2D virtual screens [13, 24, 25], another promising avenue is to take advantage of the stereoscopic view provided by AR HMDs. These hybrid workstations that combine a 2D view on a standard 2D screen and a 3D virtual view on an AR HMD are the focus of our work. These hybrid workstations can be useful in many application areas. For instance, the finance company Citi [7] shows a hybrid workstation for financial trading combining 2 physical screens and virtual 3D visualizations displayed on HoloLens in between the two physical screens. Another example is the Maya demo of Microsoft [19]. In this demo, a user models a motorcycle on a standard workstation and can see its 3D rendering at the same time. Finally, Wang et al. [36] propose a system for domain experts in particle physics to overcome difficulties to visualize a collision of particles (e.g., overlapping or occluded 3D trajectories) on a standard 2D screen. Thus, they propose an AR extension of an existing tool already used by physicists that combines a 2D screen for data filtering (only 2D information such as 2D histograms or texts) and an AR HMD-based stereoscopic 3D view of particle collision data.

The many benefits and use cases of hybrid workstations discussed above have motivated our study on interaction.

2.2 Interaction for hybrid workstations.

In AR, mid-air interaction is commonly used: for a selection task users touch the object of interest directly with their finger. This interaction is direct, but inaccurate and it is also tiring [27, 35]. In addition such “natural interaction” may not feel that “natural” to users as D. Norman pointed out [23] and is further developed as a grand challenge in immersive analytics [8]: *“Immersive Analytics brings about new, far more complex transitions which can affect interaction fluidity. To solve these challenges, future systems might aim to lower entry barriers to Immersive Analytics interaction. One promising possibility is to combine established interfaces, like keyboard, mouse or tablet input, with immersive HMDs”*. Thus, instead of mid-air interaction, a promising avenue is to use commonly used interaction such as touch and pen-based interaction [5, 11, 21], keyboard and mouse interaction [12, 25, 36]. In our study we focus on keyboard and mouse interaction. This avoids switching between interaction




Figure 1: Mouse motions in a 3D virtual view.

techniques in the hybrid workstation. There are different ways to interact with a hybrid workstation using mouse interaction. In this paper we provide the first experimental study to compare interaction techniques for these hybrid workstations.

3 DESIGN SPACE & TECHNIQUES

3.1 Design space

Hybrid workstations are systems extending a standard workstation composed of a screen, a keyboard, and a mouse. Then, in the same way as the previous hybrid systems, the interaction techniques that we consider use solely the standard mouse which is already used for the monitor and we do not consider additional input devices such as 3D mice [2, 10, 14, 26, 37].

Using mouse interaction, different types of transitions [30] between the two sides (i.e. physical screens and virtual scenes) of a hybrid workstation are possible. The cursor can jump between the two input spaces (2D and 3D) by pressing a key as in [36]. Alternatively, transitions between a physical screen and a virtual screen can be implemented in the same way as for a dual physical screen [25]. Such continuous transition is simple and intuitive since the cursor is moving on a virtual screen as on a physical screen [19, 33]. We thus focus on such interaction closely related to mouse-based interaction on a physical screen to facilitate the transition between the physical screens and the virtual screens. In ARmouse [33], the virtual scene is displayed above the keyboard of a laptop. When the cursor comes out at the bottom of the physical screen of the laptop, it appears virtually on a horizontal plane above the keyboard. If we represent the 3D virtual space as a cube, the cursor moves in 2D within the base of this cube (see Figure 1b). In addition, a ray is sent vertically from the cursor to select objects. This technique is based on raycasting: the first object intersected by the ray is automatically selected. If several objects are intersected by the ray, a second step is needed to adjust the selection. With ARmouse, this adjustment is made by moving the mouse forward or backward while keeping the left button depressed. Similarly, Wang et al. [36] implement mouse interaction with a cursor moving in 2D within the front face of the cube representing the 3D virtual space. Their technique is thus based on a horizontal raycasting (see Figure 1a). The depth adjustment of the selection is done by using the mouse wheel while pressing the shift button. We thus identify two design solutions: the cursor moves either on a vertical plane (as in [36], see Figure 1a), in the same way as on a physical screen; or on a horizontal plane (as in [33], see Figure 1b). As AR systems rely primarily on mid-air interaction (midair gestures or handheld controllers) and we limit the study to standard workstation input devices, we exclude handheld controllers and consider direct touch as a baseline to interact with the 3D virtual scene. Thus, we oppose this direct interaction with the indirect interaction using a mouse.

AR offers users great flexibility as they can place the holograms wherever they want around the physical screen. Another design parameter to consider is thus the relative positions of the two views [17]. For instance, do we place them next to each other? Or one below the other by placing the virtual scene directly on the desk?

By considering the interaction modality (mid-air interaction or interaction with a mouse) and the position of the 2 views, we obtain the design space presented in Figure 2.




Figure 2: Design space.

3.2 Techniques

The design space defines 6 selection techniques for a hybrid workstation: 4 indirect (with a mouse) and 2 direct (mid-air).

3.2.1 Indirect interaction: mouse-based techniques

Let the 3D virtual scene be represented by a cube as in Figure 1. For the 4 mouse-based techniques of the design space, the cursor can move on a vertical plane or a horizontal plane of the cube.

When using a vertical plane for cursor motion. The 3D scene is placed behind a virtual vertical plane (the front face of the cube, see Figure 1a). This plane is semi-transparent to easily see the scene placed behind (Figure 3a). The cursor moves on this plane and a ray perpendicular to the plane is attached to the cursor to select objects in the 3D scene. In most cases, the object of interest can be directly touched by the ray. However, if several objects are intersected by the ray and if the object of interest is fully occluded, the depth of the selection must be adjusted by adding an extra degree of freedom to the raycasting [3]. To do this, our solution, as in [36], is to use the mouse wheel to adjust the depth of the selection. This plane, and hence the 3D virtual scene, is placed either on the right or below the physical monitor. When it is placed on the right, the cursor comes out of the right side of the physical monitor and enters the virtual plane by its left side. When it is placed below, the cursor comes out of the bottom side of the physical monitor and enters the virtual plane by its top side. These continuous transitions are very similar to the transitions that occur between two physical screens.

The techniques *Mouse-Vertical-Right* (Figure 3b) and *Mouse-Vertical-Below* (Figure 3c) implement a vertical plane.

When using a horizontal plane for cursor motion. The virtual 3D scene is placed above a virtual horizontal plane (the base of the cube, see Figure 1b). The cursor moves on this horizontal plane. As in the previous technique, it has a ray perpendicular to the plane attached to it. The horizontal plane is placed either to the right or below the physical monitor. When it is placed to its right, the cursor comes out of the top right side of the physical monitor and enters the virtual plane by its back left side. When the cursor comes out of the bottom right side of the physical screen, the cursor enters the virtual plane by its front left side. In the same way, when the plane is placed below, the cursor comes out of the bottom side of the

physical monitor and enters the virtual plane by its back side. In this case, the plane lies directly on the table. It is interesting to note that when the horizontal plane is placed below, the transitions between the physical monitor and the virtual horizontal plane are visually continuous as with the vertical plane (i.e. the 2 edges considered in the transitions are parallel). This is not the case when the horizontal plane is placed on the right.

The techniques *Mouse-Horizontal-Right* and *Mouse-Horizontal-Below* (Figure 3d) implement a horizontal plane.

For the four mouse-based techniques, the ray is perpendicular to the plane on which the cursor is moving. Other possibilities could have been explored. A first possibility would be to attach the ray to the user's head [29]. The ray would start from the user's head, pass through the cursor and then reach the 3D virtual scene. This would have the advantage of not decoupling the 3D selection task into a 2D (mouse cursor positioning) task and 1D (mouse wheel depth selection) task. However, in addition to adding instability during interaction (when the user moves the head, the selection is modified), head tracking is needed to implement this 3D selection. This makes this approach inconsistent with regular desktop mouse control. The mouse interaction would therefore not be the same on the physical screen and on the virtual 3D scene, defining a non-uniform interaction space of the hybrid workstation with different interaction techniques. A second possibility would be to use the eye tracker for the 3D selection. This solution implies the same limitations as above (i.e. instability, non-uniform interaction). Moreover, like Wang et al. [36], we want the users to be able to see both the 2D screen and the 3D virtual scene at the same time. The extension of the perspective cursor [22] into 3D space, the In-Depth Mouse [38], is another possibility. We discard this solution because the metaphor of the image plane selection, consistent with a desktop mouse control, is not verified. Moreover, with In-Depth Mouse, the position of the cursor when transitioning between spaces is data-dependent. With our mouse-based techniques, the position of the cursor when transitioning between spaces is easy to understand and above all predictive. We have thus not chosen these solutions and the four mouse-based techniques implement a fixed ray, perpendicular to the plane.

Feedback and validation of the selection. For all mouse-based techniques, the objects intersected by the ray are highlighted, and the current selection (changed with the mouse wheel) turns yellow if it is a distractor and green if it is the target. Users validate the selection by pressing the left mouse button.

3.2.2 Direct interaction: mid-air techniques

To select an element in the virtual 3D scene, the users simply touch it with their fingertips (Figure 4). We used this direct mid-air interaction for several reasons: (1) this interaction is used in a very large majority of AR applications, (2) it does not need additional devices (e.g., handheld controllers), and (3) all objects in the scene are easily reachable by hand. The validation of the selection is done by pressing the space bar. The 3D scene is placed either on the right of the physical screen or on the table. As with the other techniques, the current selection turns yellow if it is a distractor and green if it is the target.

In contrast to the previous mouse-based techniques, the user needs to switch between mouse (2D graphical user interface displayed on the physical monitor) and mid-air interaction (3D elements of the virtual scene on the HoloLens). The techniques *Hand-Right* and *Hand-Below* implement this direct touch interaction.

3.3 Implementation

As we can see in Figures 3 & 4, our system makes it possible to simultaneously visualize: 2D information (e.g., a 2D histogram) on a physical monitor and a 3D scene (e.g., a set of 3D spheres) displayed on an AR HMD.

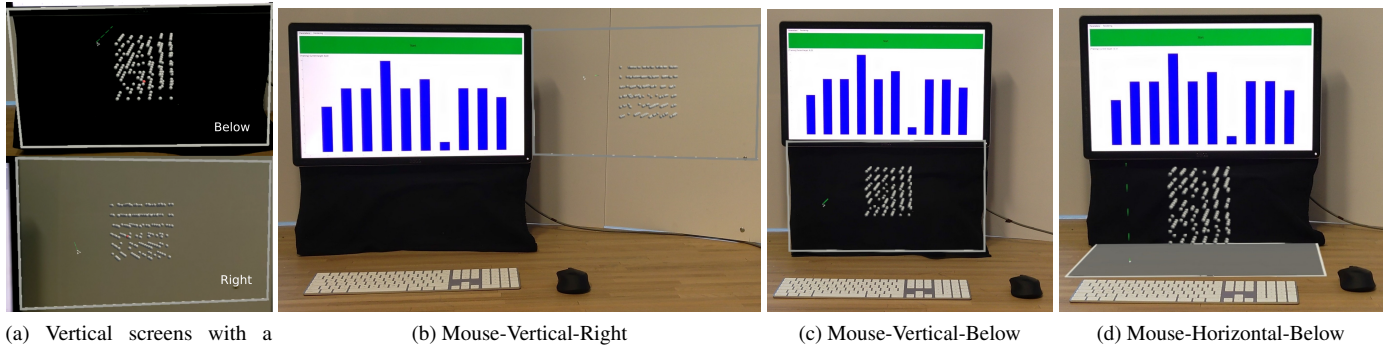


Figure 3: Mouse-based interaction techniques.




Figure 4: Mid-air interaction technique.

We used a laptop and a physical monitor (24”) to display 2D information on the physical monitor. 2D information and the tracking of the mouse cursor are managed by a program written using the Python bridge to the QT toolkit PyQt. When the cursor leaves the screen to enter the virtual space (i.e. the mouse reaching the edge of the physical screen), the cursor is no longer displayed on-screen. Our PyQt software then sends the mouse movements to HoloLens which displays the cursor. The cursor gain is the same on the screen and on the virtual planes. The 3D virtual scene was created with Unity and displayed with the Microsoft HoloLens 2. For direct interaction, the hand mesh and a disk following the users’ index fingertip are displayed for precision. The horizontal and vertical planes have the same physical dimensions as the screen of the physical monitor and their placement is done with a Vuforia marker. Both the laptop and the HoloLens were connected to a router and the mouse positions were sent via this local network.

3.4 Preliminary study: usability and users’ feedback

3.4.1 Protocol

First, we conducted a pilot study to evaluate the usability of the 6 selection techniques, the easiness of the transitions between the physical monitor and the 3D virtual space, and to collect the users’ preferences and feedback. The goal is also to reduce the number of interaction techniques identified by the design space before performing an experimental study. We recruited 4 participants for this qualitative pilot study. All 4 participants were not experts in AR but have occasionally used AR HMD as part of experiments.

The study focuses on the transitions between the two spaces and also the interaction in the virtual scene. In the experimental setting, a simple 2D histogram is displayed on the physical screen (Figures 3 & 4). This setting does not fully utilize the high resolution of the physical screen. Indeed the mouse interaction on the physical screen was not the focus of our work, so we decided to display a simple 2D histogram on it. In a follow-up study (Section 7), we display more complex 2D information on the physical screen (i.e. a map) to illustrate the potential of the combination of a 2D physical display and a 3D virtual view.

Additionally, the experimental setting includes a set of 3D virtual spheres displayed on the HoloLens (Figures 3 & 4). The task was to

successively select a bar of the histogram in 2D and then a sphere in 3D. The only link between the two views is when the participant selects the red bar on the physical screen: Indeed this selection causes the display of a red sphere in the 3D scene. The task then forced the participants to go back and forth between the 2 spaces. Participants performed 21 pairs of selections for each technique. They also filled out two questionnaires: the Raw-TLX questionnaire to evaluate the perceived workload induced by the techniques; and a questionnaire about transitions including questions such as “Are they hard?”, “Are they frustrating?”. Finally we concluded the study with an interview about their preferences after they ranked the techniques.

3.4.2 Results

Transitions between the two spaces. The results showed that the frustration and the difficulty to go back and forth between spaces are quite low with a small advantage for the two metaphors of the double screen with the mouse (Mouse-Vertical-Right, Mouse-Vertical-Below) and the technique with the mouse moving on the table (Mouse-Horizontal-Below). On a Likert scale of 7 points, participants rated the frustration as follows: $mean = 1$ for Mouse-Vertical-Right, Mouse-Vertical-Below and Mouse-Horizontal-Below; $mean = 1.5$ for Hand-Below; $mean = 1.75$ for Mouse-Horizontal-Right; and $mean = 2.5$ for Hand-Right. Participants rated the difficulty of the transitions between spaces as follows: $mean = 1$ for Mouse-Vertical-Right, Mouse-Vertical-Below and Mouse-Horizontal-Below; $mean = 1.75$ for Mouse-Horizontal-Right and Hand-Right; and $mean = 2$ for Hand-Below.

Perceived workload and users’ preferences. The perceived workloads are presented in Figure 5 (Means and 95% confidence intervals are shown in all graphs). The 3 techniques Mouse-Vertical-Right, Mouse-Vertical-Below and Mouse-Horizontal-Below obtained a lower workload than the others. 3/4 participants preferred using the mouse instead of the hand to select the 3D targets.

First, we asked the participants to rank the 4 mouse-based techniques. The participants preferred the 2 techniques with the double screen metaphor: using the mouse and interacting on a vertical plane (Mouse-Vertical-Right: $median_{rank} = 1$; Mouse-Vertical-Below: $median_{rank} = 2$). Mouse-Horizontal-Right is the least preferred (mouse-based) technique ($median_{rank} = 3.5$). For Mouse-Horizontal-Below, it is not clear: some people liked this technique, others not at all ($median_{rank} = 3.5$, ranked from the second to the fourth place). For the mid-air interaction, 3/4 participants preferred it when the virtual 3D scene was on the table. According to them, the table provides support for the arm during the interaction (the arm no longer being in the air but on the table), which facilitates the selection and reduces fatigue.

Users’ feedback. Several participants thought that using a double screen at work could help them with the mouse-based techniques. Participants also compared the mouse-based techniques with their habits: “I’m not used to having two screens next to each other [...]”,




Figure 5: Preliminary study, total workload between 0 and 100 (obtained with Raw-TLX).

I usually have my laptop under my second screen like this technique [Mouse-Vertical-Below]” (P2); “I prefer [Mouse-Vertical-Right] to [Mouse-Vertical-Below] because it is the closest to what I am used to having” (P4).

Overall, participants found the transitions less intuitive with Mouse-Horizontal-Right than with the other mouse-based interaction techniques and sometimes, they lost the cursor with this technique. We thus discarded this technique that is also the mouse-based technique with the highest workload, the highest score of frustration, and the highest difficulty of the transitions. These differences between Mouse-Horizontal-Right and the other mouse-based techniques can be explained by the lack of visual continuity between the 2 spaces (see Section 3.2.1). For mid-air interaction, we decided to remove Hand-Right based on users’ preferences because this mid-air position does not benefit from the table support like Hand-Below. The support of the table was found important by participants to improve accuracy and to reduce fatigue of the interaction. Thus, we discarded 2 of the 6 techniques of our design space by comparing mouse-based techniques on one side and mid-air techniques on the other. The 2 techniques not considered in the following study are: Mouse-Horizontal-Right and Hand-Right.

4 EXPERIMENTAL STUDY

After reducing the number of interaction techniques to 4 (Figure 2), we conducted an objective and subjective experimental study. The goal of this study is to evaluate more precisely the usability and in particular the effectiveness of the 4 remaining selection techniques.

Task. The 3D virtual scene is a set of 160 3D spheres pseudo-randomly placed on a $6 \times 6 \times 7$ regular grid. The task is similar to the pilot study to force users to go back and forth between the 2D graphical user interface and the virtual 3D scene. The task consists of 4 steps:

- 1 Click on the Start&Stop button to start the task. Then, the 2D target appears in the histogram, i.e. one of the bars turning red. The Start&Stop button is deactivated.
- 2 Select the 2D target on the histogram to make the 3D target appear in the 3D virtual scene, i.e. one of the 3D spheres turning red.
- 3 Select the 3D target. The Start&Stop button is reactivated.
- 4 Click the Start&Stop button again to finish the task.

Protocol. We used a within-subject design with the following independent variables:

TECH the technique: Mouse-Vertical-Right, Mouse-Vertical-Below, Mouse-Horizontal-Below, and Hand-Below;

DS the density spacing: To control the density of the environment, we use the method proposed by Vanacken et al. [34]. The closest distractors are placed at a fixed distance from the target. In our setup, the spheres are placed on a grid, this distance represents the spacing between columns and rows. This distance is called the “density spacing” and is equal to 0.5 cm for high density, and 1.6 cm for low density.

SIZE the target size: diameter of the spheres: 0.4 or 0.8 cm.

NSPHERE the number of spheres to be covered with the mouse wheel (mouse-based techniques only) before reaching the target: 2 or 5 spheres. When **NSPHERE** = 2 the target is close

to the plane. Note that for Hand-Below, when **NSPHERE** = 5 the selection is easier because it is on the top part of the set of spheres, and thus easier to reach by hand.

The combinations of **DS** \times **SIZE** \times **NSPHERE** are repeated 5 times in a random order for a total of $2 \times 2 \times 2 \times 5 = 40$ trials for each **TECH** condition. Before each **TECH** condition, the participants performed a training session (6 trials). Participants could take a break after each technique. The order of the techniques was counterbalanced across participants using a Latin square design. The experiment lasted about 50 minutes per participant.

Participants. We recruited 12 unpaid volunteers, ranging from age 23 to 49 (*mean* = 33.6, *std* = 8.7). All participants were right-handed. None of them was an expert in augmented reality.

Measures. For the quantitative analysis, we measured the completion times (times spent between the start (Step 1) and stop (Step 4) buttons) and counted the number of errors when selecting the 3D targets (selection of another 3D object than the target). For the qualitative analysis, participants filled out the Raw-TLX questionnaire after each technique and reported their preferences at the end of the experiment.

5 RESULTS

In the following, we present the qualitative and quantitative results of the experimental study.

Completion times. To compare completion times, we first remove outlier trials. We consider a trial as an outlier if its completion time differs from the mean by more than three standard deviations. We identify and remove 2.6% of the trials (e.g., forgot to press the Start&Stop button). We use ANOVAs and t-tests with Bonferroni adjustment for pairwise comparisons. A Box-Cox transformation with $\lambda = -0.1$ is applied to correct non-normal completion time residuals. When sphericity is violated, we apply Greenhouse-Geisser corrections.

We find a main effect ($F_{1,33,14.61} = 10.35$, $p = 0.004$, $\eta_G^2 = 0.16$) for **TECH** on completion times, see Figure 6a. Pairwise comparisons show that the techniques using a vertical plane (Mouse-Vertical-Below: 5.74s, Mouse-Vertical-Right: 5.75s) are significantly faster ($p < 0.0001$) than the other two techniques (Mouse-Horizontal-Below: 6.43s, Hand-Below: 6.56s). We do not find significant difference between Mouse-Vertical-Below and Mouse-Vertical-Right, as well as between Mouse-Horizontal-Below and Hand-Below.

We find a main effect for **DS** on completion times ($F_{1,11} = 26.89$, $p = 0.0003$, $\eta_G^2 = 0.02$).

We observe significant **TECH** \times **DS** interaction effects ($F_{3,33} = 19.50$, $p < 0.0001$, $\eta_G^2 = 0.06$), see Figure 6b. Pairwise comparisons show that Hand-Below is the only technique significantly impacted by **DS** ($p < 0.0001$). When the environment density is low (**DS**=Large), we observe a significant difference only between Mouse-Horizontal-Below and the two techniques using a vertical plane (Mouse-Vertical-Below: $p = 0.004$, Mouse-Vertical-Right: $p = 0.021$). When the environment density is high (**DS**=Small), Hand-Below is significantly slower than the two techniques using a vertical plane ($p < 0.0001$). Also, Mouse-Horizontal-Below is significantly slower than the two techniques using a vertical plane (Mouse-Vertical-Below: $p = 0.006$, Mouse-Vertical-Right: $p = 0.001$). We observe no significant difference between Hand-Below and Mouse-Horizontal-Below, as well as between Mouse-Vertical-Below and Mouse-Vertical-Right.

We found a main effect for **SIZE** on completion times ($F_{1,11} = 186.85$, $p < 0.0001$, $\eta_G^2 = 0.12$). We do not observe significant **TECH** \times **SIZE** interaction effects ($F_{3,33} = 1.59$, $p > 0.05$).

We find a main effect for **NSPHERE** on completion times ($F_{1,11} = 15.60$, $p = 0.002$, $\eta_G^2 = 0.03$). We observe significant **TECH** \times **NSPHERE** interaction effects ($F_{1,52,16.69} = 21.76$, $p =$




Figure 6: Completion times ...




Figure 7: Experimental study, total workload between 0 and 100 (obtained with Raw-TLX).

0.0001, $\eta_G^2 = 0.07$), see Figure 6c. Pairwise comparisons show that only Mouse-Vertical-Below ($p = 0.032$) and Mouse-Vertical-Right ($p = 0.021$) are (slightly) impacted by *NSPHERE*. When *NSPHERE* = 2 (the target is close to the plane in depth or height), the two techniques using a vertical plane are faster than the others ($p < 0.0001$). We do not observe significant difference between Hand-Below and Mouse-Horizontal-Below, as well as between Mouse-Vertical-Below and Mouse-Vertical-Right. When *NSPHERE* = 5 (the target is far from the plane in depth or height), we observe no significant difference between the four techniques.

Errors. Participants were asked to select targets as quickly as possible, while making as few errors as possible. We observe very few errors during the experiment. Overall, participants did 0.27 wrong selections before selecting the target when using the hand interaction technique, and only 0.01 wrong selections when using the mouse interaction techniques.

Workload. At the end of each technique, the participants filled out the Raw-TLX questionnaire. The results are presented in Figure 7. Hand-Below has the highest workload ($mean = 52.8$), followed by Mouse-Horizontal-Below ($mean = 36.1$). The two techniques using a vertical plane have the lowest workload (Mouse-Vertical-Below: $mean = 24.8$ and Mouse-Vertical-Right: $mean = 24.5$).

Users preferences. To interact with 3D content, participants largely preferred to use the mouse (11/12 participants) rather than the hand (1/12 participant).

At the end of the experiment, each participant was asked to rank the techniques in order of preference. We use a non-parametric Friedman χ^2 test to assess the impact of the technique on users' preferences, and we use a Wilcoxon signed rank test with Bonferroni adjustment for pairwise comparisons. The Friedman test shows a main effect of the technique [$\chi^2(3) = 18, p = 0.0005$]. The median ranks of the techniques are: 1.5 for Mouse-Vertical-Below, 2 for Mouse-Vertical-Right, 3 for Mouse-Horizontal-Below and 4 for Hand-Below. The post-hoc analysis shows a significant difference between Hand-Below and the two mouse techniques using a vertical plane (Mouse-Vertical-Below: $p = 0.002$, Mouse-Vertical-Right: $p = 0.004$). We also observe a smaller significant difference between Mouse-Horizontal-Below and the two techniques using a vertical plane (Mouse-Vertical-Below: $p = 0.02$, Mouse-Vertical-Right: $p = 0.01$). We do not find a significant difference between Mouse-Horizontal-Below and Hand-Below. This is consistent with the results of the workload assessment presented above.

6 DISCUSSION

6.1 Mouse or mid-air interaction

Overall, the objective and subjective results suggest that mouse interaction is slightly more suitable than mid-air interaction (Figures 6b & 7) but participants also indicated that both types of interaction are usable. Switching between mouse interaction with 2D information on the screen and mid-air interaction with a virtual 3D scene could be frustrating. In particular we felt that the most difficult switch would be from mid-air to mouse interaction, i.e. from absolute touch to relative mouse as studied on a laptop [9]. Although the perceived workload and especially the perceived frustration are higher for mid-air interaction than mouse interaction (Figure 7), the participants explained that switching was not the issue and that they are used to finding the mouse without looking at it (proprioception, peripheral vision and habit of switching from keyboard to mouse). They reported that the switch from mid-air to mouse interaction was easy and fast.

When the density is low and/or the target is easily accessible by the hand (e.g., target on top and not in the middle of the 3D scene) all techniques are equivalent (Figure 6b). Participants stated that they liked the directness of the hand interaction, but also indicated that they wanted to keep the mouse interaction.

However, like Pavanatto et al. [25], we observed the limitations of the HMD (e.g., limited field of view) and of the accuracy of the hand detection, especially when the density is high, which is often the case when viewing data with a hybrid workstation. Indeed, when the density increases (Figure 6b) or when the target is surrounded by distractors in the middle of the 3D scene (i.e. when accuracy and especially the accuracy of the perceived depth of the hand is required) the interaction with the hand is negatively affected and has been reported to become tiring. Mouse interaction presents the advantage of being stable under all conditions.

The objective and subjective results as well as the participants' comments lead us to conclude that both mouse and mid-air interaction techniques are desirable in a hybrid workstation.

6.2 Designing mouse-based interaction techniques for hybrid workstations

The objective and subjective results show that for mouse interaction a vertical plane seems the best option: it is fast and appreciated by the participants. Indeed, this condition presents several advantages.

Mimicking the well-known double screen setup. We thought the dual screen metaphor might make cursor transitions easier to understand as it seamlessly extends the physical monitor and implements cursor transitions similar to those between two physical monitors. The experimental results confirm this hypothesis. Participants further indicated that keeping the same orientation as the physical monitor is "easier", "more intuitive" and results in a "lower workload". They also reported that they felt more confident.

An intuitive wheel adjustment. Several participants reported that the depth adjustment with the wheel (vertical plane condition) in contrast to a height adjustment with the wheel (horizontal plane condition) is "more logical" and "more intuitive". One participant

also pointed out that the forward and backward movements of the wheel implement a direct mapping with the movements of the cursor towards the back or front of the 3D scene. The remark is valid if the depth adjustment is done by moving the mouse forward and backward while pressing a mouse button as is the case with AR-Mouse [33].

A better point of view. According to participants, when the virtual 3D scene is positioned on an horizontal plane as with Mouse-Horizontal-Below (Figures 3d), the point of view on the 3D scene and the ray is “not comfortable” and is not “as good as” the techniques based on a vertical plane (Figures 3c & 3b). Since we have not tested Mouse-Horizontal-Right because of the lack of visual continuity between the spaces and the resulting difficulty in cursor transition between them, the remark on the point of view can only be applied to Mouse-Horizontal-Below.

A flexible position of the 3D scene. The objective and subjective results lead us to conclude that both positions (i.e. the 2 views next to each other, or one below the other) should be supported for mouse interaction using a vertical plane. Indeed, we obtained similar time performance and similar users’ perceived workload and preferences. The participants’ comments suggest that the choice depends on users’ habits, and it is probably a trade-off between neck comfort (due to head movements) and harm comfort.

7 STUDY: VISUALIZATION OF A GAPMINDER DATASET

We believe that the benefit of using a mouse interaction with a vertical plane, a horizontal plane, or using mid-air interaction could vary according to (1) the 3D data visualization, (2) the task, and (3) the user. So to increase the ecological validity (one aspect of the external validity) of our work on interaction with a hybrid workstation, we decided to test interaction in a more ecological context.

We conducted a qualitative study where users were put in a data visualization context and should perform elementary tasks (i.e. addressing “individual elements of data”) and synoptic tasks (i.e. requiring “to deal with sets as a whole”) of a Gapminder dataset.

For this study, we recruited 14 unpaid volunteers, ranging from age 21 to 49, to observe the use of the hybrid system and to validate/confirm the observations of our experimental study. None of the participants was an expert in AR (never used an AR HMD or only very occasionally). None of them was an expert in 3D data analysis but they occasionally visualized 3D information in their professional or personal environment.

7.1 Implemented system

Data. We developed a hybrid workstation system, presented in Figure 8. The system simultaneously displays (1) a 2D map of the world on the 2D physical screen, and (2) a 3D scatter plot with the HoloLens. The 2D map is more representative than the simple 2D histogram used in the first experiment of the kind of detailed visualization that is possible on an actual computer screen. The 3D scatter plot shows the number of babies per woman, income, and life expectancy in the world in 1960. Each sphere represents a country, and its diameter is proportional to the population of the country. The two views (i.e. 2D map and 3D scatter plot) are tightly coupled: direct selection with the mouse on the 2D map results in an update of the 3D scatter plot. The user can visualize data of a region of the world and can also compare several regions with each other. The user can select a 3D sphere to obtain information about the corresponding country (i.e. the name of the country, the values of the 3 axes, and the population). This information is displayed at the bottom left of the virtual plane, next to the 3D scatter plot (Field “Current selection”, Figure 8).

Interaction techniques. Our previous results (Sections 5 & 6) suggest that both mid-air and mouse interaction should be kept in a hybrid workstation. Also, we observed that the users’ preferences about the position of the scene varied among participants and were dependent on their habits. As a consequence, we developed a hybrid system offering more flexibility than before.

First, the system offers the flexibility of the position of the scene. With a simple grab movement, the user can switch between two positions of the virtual scene: a position to the right of the physical monitor (similar to Mouse-Vertical-Right, Figure 8a) and a position below the physical monitor (similar to Mouse-Horizontal-Below and Hand-Below, Figure 8b). We decided to remove the Mouse-Vertical-Below condition for two main reasons: (1) this condition had very similar performance and feedback as Mouse-Vertical-Right; and (2) it would require the users to move the physical monitor up during the experiment which could dissuade them to change the scene position during the experiment.

Second, the system offers flexible interaction. In order to maximize this flexibility and also reduce the complexity of the system, both mid-air and mouse interaction can be used to select a sphere of the scatter plot, regardless of the position of the scene (i.e. right or below the screen). In addition to allowing the users to choose mid-air or mouse interaction at any time, we also extend the system to allow the users to use both modalities at the same time. In this case, the mid-air selection is “dominant” which means that only the information of the country selected by the hand is displayed. This last possibility could be useful to quickly compare the values of two countries: one country can be selected with the mouse and the other can be selected in mid-air.

Third, the users can rotate the 3D scatter plot around the Y-axis (up-down axis) by using a virtual slider placed next to it. The users interact with this slider via mid-air interaction (i.e. grab gesture).

7.2 Protocol

We first presented the system and the data to the participants. AR interaction are easy and intuitive but users need practice to handle them correctly [4, 16, 31]. So before starting the experiment, participants had to test the system and especially practice the interaction that can be done (e.g., the interaction with the slider, the change of the position of the scene, the mid-air and the mouse pointing). This training session makes them familiar with all the interaction techniques and the change of the position.


After this training session, they had to perform 15 tasks¹. The current task and the text field for entering the answer were displayed on the physical screen above the 2D map. Participants had to confirm the answer with the button “Validate” and the next task was displayed. We counterbalanced the initial position of the scene, so half of the participants started with the right position, and the others started with the scene below. However, we pointed out to them that they were free to change this position whenever they wished.

Based on the taxonomy of tasks involved in the exploratory data analysis of Andrienko et al. [1], we created elementary tasks (i.e. tasks addressing “individual elements of data”) and synoptic tasks (i.e. tasks requiring “to deal with sets as a whole”) allowing users to identify trends or behaviors of the data. All 15 tasks required participants to use the 2 views (i.e. 2D map and 3D scatter plot). They were related to one axis (e.g., “In Asia, which country has the highest life expectancy?”), two axes (e.g., “Describe if you observe a trend between the life expectancy and the number of babies per woman over the 3 regions of Europe, Africa, and America.”), or the three axes at a time (e.g., “In which region of the world do you observe the least variation on the 3 axes?”). The tasks have been designed for non-experts in data analysis but also to exploit the advantages of the stereoscopic 3D view (e.g., to observe a data trend along 2 or 3 axes).

¹http://iihm.imag.fr/publs/2022/ISMAR2022_appendix.pdf



(a) Right position where both mid-air and mouse interaction techniques are allowed



(b) Below position where both mid-air and mouse interaction techniques are allowed

Figure 8: Hybrid workstation: Visualization of a Gampinder World multivariate dataset.

We did not impose a time limit for completing the tasks and we asked participants to be confident about their answers. After completing the 15 tasks, participants filled out a questionnaire with questions about their usage of the system (e.g., Did they use the rotation of the 3D scatterplot? Did they modify the position of the scene?; Which interaction did they use?; Do they have a favorite position of the scene?). The questionnaire guided the semi-structured interview.

7.3 Results

The sessions with the participants lasted from 35 to 50 minutes per participant. The participants made 17 errors out of 210 answers. In addition to the answers and the interview, we logged several events during the sessions such as changes in scene position, and mouse/mid-air selections.

7.3.1 Interaction for country selection

Based on log events, 6/14 participants used the mouse interaction only, and 8/14 participants used both the mouse and the mid-air interaction to select countries in the 3D scatter plot. Of these 8 participants, 1 participant interacted primarily in mid-air and only occasionally with the mouse (“I used the mouse if two countries were very close to each other”), the other participants mainly used the mouse interaction.

According to the participants, the mouse interaction is easy, stable, not tiring, and accurate. Some participants used the ray attached to the cursor as a “pointer”/a “visual marker” (similar to a cutting plane) to compare a country with others (“I place the cursor on a country to see more easily if the others are above or below”). Finally, some participants declared that in such hybrid systems, using the mouse on the physical monitor makes them automatically keep the mouse on the 3D scene (“Since I had the mouse for the screen, I kept it all along”).

Regarding the mid-air selection, participants reported using it in order to quickly grasp information and usually when the target is big. One participant indicated that the mid-air interaction was more intuitive than the mouse interaction (3D interaction for a 3D task). 3 participants compared 2 countries by selecting one with the mouse and the other with the hand.

7.3.2 Placement of the 3D virtual scene

We also recorded instances where the participants changed the position of the virtual scene during the experiment. 3/14 participants never used it. 7/14 participants changed the position of the scene 2 or 3 times during the experiment. Finally, the remaining 4 participants changed the position more than 6 times. The reasons for these changes reported by the participants are as follows:

- To have the axis of interest in front of them/limit the parallax/use the ray as a visual marker (8 participants);

- To be temporarily closer and more easily select small targets (right position → below) (1 participant);
- To check something by temporarily changing their point of view (2 participants);
- To limit disambiguation when using the mouse (2 participants);
- By comfort (e.g., “at first I was like that, but I changed for comfort”) (1 participant)

The preferred position of the 3D scene in relation to the screen depends largely on the participants (Below: 5/14 participants, Right: 6/14 participants, No preference: 3/14 participants).

8 CONCLUSION

Our work focuses on interaction in hybrid workstations that combine a 2D view displayed on a physical monitor and a 3D view on a HoloLens. The results of an experimental study and also those of a task-driven study provide two main guidelines for the design of hybrid workstation systems:

An efficient continuous mouse interaction. Overall, users reported liking the hybrid workstation, the simple and seamless mouse interaction, and perceived the benefits of this system. With intuitive interaction and smooth transitions between the two spaces, these mouse-based techniques have proven to be effective, easy to learn, easy to use, and well-liked by users. Specifically, we believe that mimicking the dual screen metaphor by using a vertical plane in front of the 3D virtual scene brings users back into a more familiar environment. These results are very encouraging and tend to confirm that such hybrid workstations can help researchers “lower entry barriers” for the adoption of immersive technologies (e.g., Immersive Analytics) [8]. As future work, these techniques could be compared with the In-Depth Mouse technique [38] to study simplicity and seamlessness versus performance in hybrid workstations.

Need for flexibility to suit as many users and tasks as possible. Our results also show the need for flexibility to allow users to choose the interaction technique (mouse/mid-air) and to position the stereoscopic 3D scene relative to the physical monitor.

In this work, we only consider the selection task with discrete data. The impact of other types of data and tasks, such as creating 3D objects from scratch, on the choice of interaction techniques and the position of the virtual scene needs further investigation. Our future work will focus on continuous data and then tasks that require simultaneous manipulation of all three dimensions (e.g., specifying a 3D trajectory). The combined use of mid-air and mouse is a possible extension of our study to be analyzed with regard to the design space of multimodal interaction (synergistic and alternate usage).

Acknowledgments: This work has been partially supported by the MIAI Multidisciplinary AI Institute (ANR-19-P3IA-0003).

REFERENCES

- [1] N. Andrienko and G. Andrienko. *Exploratory analysis of spatial and temporal data: a systematic approach*. Springer Science & Business Media, 2006.
- [2] R. Balakrishnan, T. Baudel, G. Kurtenbach, and G. Fitzmaurice. The Rockin'Mouse: Integral 3D manipulation on a plane. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, CHI '97, p. 311–318. Association for Computing Machinery, New York, NY, USA, 1997. doi: 10.1145/258549.258778
- [3] M. Baloup, T. Pietrzak, and G. Casiez. RayCursor: A 3D pointing facilitation technique based on raycasting. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI '19, p. 1–12. Association for Computing Machinery, New York, NY, USA, 2019. doi: 10.1145/3290605.3300331
- [4] J. Benedict, J. Guliuzo, and B. Chaparro. The intuitiveness of gesture control with a mixed reality device. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 63:1435–1439, 11 2019. doi: 10.1177/1071181319631403
- [5] W. Büschel, A. Mitschick, T. Meyer, and R. Dachsel. Investigating smartphone-based pan and zoom in 3D data spaces in augmented reality. In *Proceedings of the 21st International Conference on Human-Computer Interaction with Mobile Devices and Services*, MobileHCI '19. Association for Computing Machinery, New York, NY, USA, 2019. doi: 10.1145/3338286.3340113
- [6] S. K. Card, T. P. Moran, and A. Newell. The keystroke-level model for user performance time with interactive systems. *Commun. ACM*, 23(7):396–410, July 1980. doi: 10.1145/358886.358895
- [7] Citi. Citi hololens holographic workstation. <https://www.youtube.com/watch?v=0Nog1tmewmQ>, 2016.
- [8] B. Ens, B. Bach, M. Cordeil, U. Engelke, M. Serrano, W. Willett, A. Prouzeau, C. Anthes, W. Büschel, C. Dunne, T. Dwyer, J. Grubert, J. H. Haga, N. Kirshenbaum, D. Kobayashi, T. Lin, M. Olaosebikan, F. Pointecker, D. Saffo, N. Saquib, D. Schmalstieg, D. A. Szafir, M. Whitlock, and Y. Yang. Grand challenges in immersive analytics. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 2021.
- [9] M. J. Foley, Q. Roy, D.-Y. Huang, W. Li, and D. Vogel. Switching between standard pointing methods with current and emerging computer form factors. In *CHI Conference on Human Factors in Computing Systems*, CHI '22. Association for Computing Machinery, New York, NY, USA, 2022. doi: 10.1145/3491102.3517433
- [10] B. Froehlich, J. Hochstrate, V. Skuk, and A. Huckauf. The GlobeFish and the GlobeMouse: Two new six degree of freedom input devices for graphics applications. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '06, p. 191–199. Association for Computing Machinery, New York, NY, USA, 2006. doi: 10.1145/1124772.1124802
- [11] T. Gesslein, V. Biener, P. Gagel, D. Schneider, P. Kristensson, E. Ofek, M. Pahud, and J. Grubert. Pen-based interaction with spreadsheets in mobile virtual reality. pp. 361–373, 11 2020. doi: 10.1109/ISMAR50242.2020.00063
- [12] J. Grubert, E. Ofek, M. Pahud, and P. Kristensson. Back to the future: Revisiting mouse and keyboard interaction for hmd-based immersive analytics. 2020. doi: 10.48550/arXiv.2009.02927
- [13] J. Grubert, E. Ofek, M. Pahud, and P. O. Kristensson. The office of the future: Virtual, portable, and global. *IEEE Computer Graphics and Applications*, 38:125–133, 2018.
- [14] K. Hinckley, M. Sinclair, E. Hanson, R. Szeliski, and M. Conway. The VideoMouse: A camera-based multi-degree-of-freedom input device. In *Proceedings of the 12th Annual ACM Symposium on User Interface Software and Technology*, UIST '99, p. 103–112. Association for Computing Machinery, New York, NY, USA, 1999. doi: 10.1145/320719.322591
- [15] C. Hurter, N. H. Riche, S. M. Drucker, M. Cordeil, R. Alligier, and R. Vuillemot. FiberClay: Sculpting three dimensional trajectories to reveal structural insights. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):704–714, 2019. doi: 10.1109/TVCG.2018.2865191
- [16] L.-H. Lee, T. Braud, S. Hosio, and P. Hui. Towards augmented reality driven human-city interaction: Current research on mobile headsets and future challenges. *ACM Comput. Surv.*, 54(8), oct 2021. doi: 10.1145/3467963
- [17] L. Lischke, S. Mayer, K. Wolf, N. Henze, H. Reiterer, and A. Schmidt. Screen arrangements and interaction areas for large display work places. In *Proceedings of the 5th ACM International Symposium on Pervasive Displays*, PerDis '16, p. 228–234. Association for Computing Machinery, New York, NY, USA, 2016. doi: 10.1145/2914920.2915027
- [18] V. Mandalika, A. Chernoglazov, M. Billinghurst, C. Bartneck, M. Hurrell, N. Ruiter, A. Butler, and P. Butler. A hybrid 2D/3D user interface for radiological diagnosis. *Journal of Digital Imaging*, 31:56–73, 08 2017. doi: 10.1007/s10278-017-0002-6
- [19] Microsoft. Microsoft hololens: Partner demo with maya by autodesk. <https://www.youtube.com/watch?v=q0K3n0Gf8mA>, 2015.
- [20] Microsoft. Microsoft hololens 2. <https://www.microsoft.com/en-us/holoLens/>, 2019.
- [21] A. Millette and M. J. McGuffin. DualCAD: Integrating augmented reality with a desktop gui and smartphone interaction. In *2016 IEEE International Symposium on Mixed and Augmented Reality (ISMAR-Adjunct)*, pp. 21–26, 2016. doi: 10.1109/ISMAR-Adjunct.2016.0030
- [22] M. A. Nacenta, S. Sallam, B. Champoux, S. Subramanian, and C. Gutwin. Perspective cursor: Perspective-based interaction for multi-display environments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '06, p. 289–298. Association for Computing Machinery, New York, NY, USA, 2006. doi: 10.1145/1124772.1124817
- [23] D. Norman. The way I see it: Natural user interfaces are not natural. *interactions*, 17:6–10, 05 2010. doi: 10.1145/1744161.1744163
- [24] E. Ofek, J. Grubert, M. Pahud, M. Phillips, and P. O. Kristensson. Towards a practical virtual office for mobile knowledge workers. 2020. doi: 10.48550/arXiv.2009.02947
- [25] L. Pavanatto, C. North, D. A. Bowman, C. Badea, and R. Stoakley. Do we still need physical monitors? an evaluation of the usability of ar virtual monitors for productivity work. In *2021 IEEE Virtual Reality and 3D User Interfaces (VR)*, pp. 759–767, 2021. doi: 10.1109/VR50410.2021.00103
- [26] G. Perelman, M. Serrano, M. Raynal, C. Picard, M. Derras, and E. Dubois. The Roly-Poly mouse: Designing a rolling input device unifying 2D and 3D interaction. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '15, p. 327–336. Association for Computing Machinery, New York, NY, USA, 2015. doi: 10.1145/2702123.2702244
- [27] C. Plasson, D. Cunin, Y. Laurillau, and L. Nigay. 3D tabletop AR: A comparison of mid-air, touch and touch-mid-air interaction. In *Proceedings of the International Conference on Advanced Visual Interfaces*. Association for Computing Machinery, New York, NY, USA, 2020.
- [28] Prabhat, A. Forsberg, M. Katzourin, K. Wharton, and M. Slater. A comparative study of desktop, fish tank, and cave systems for the exploration of volume rendered confocal data sets. *IEEE Transactions on Visualization and Computer Graphics*, 14(3):551–563, 2008. doi: 10.1109/TVCG.2007.70433
- [29] A. Ramcharitar and R. J. Teather. EZCursorVR: 2d selection with virtual reality head-mounted displays. In *Proceedings of Graphics Interface 2018*, GI 2018, pp. 123 – 130. Canadian Human-Computer Communications Society / Société canadienne du dialogue humain-machine, 2018. doi: 10.20380/GI2018.17
- [30] M. Serrano, B. Ens, X.-D. Yang, and P. Irani. Gluey: Developing a head-worn display interface to unify the interaction experience in distributed display environments. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services*, MobileHCI '15, p. 161–171. Association for Computing Machinery, New York, NY, USA, 2015. doi: 10.1145/2785830.2785838
- [31] W. Sinlapanuntakul, J. Korentsides, A. M. Collard, K. S. Skilton, and B. S. Chaparro. Touching holograms: A preliminary evaluation of mixed reality gestures. 2022.
- [32] W. Usher, P. Klacansky, F. Federer, P.-T. Bremer, A. Knoll, J. Yarch, A. Angelucci, and V. Pascucci. A virtual reality visualization tool for neuron tracing. *IEEE Transactions on Visualization and Computer*

- Graphics*, 24(1):994–1003, 2018. doi: 10.1109/TVCG.2017.2744079
- [33] B. van Klingeren. ARMouse: using the mouse as a foundation for a universal interaction paradigm for hybrid AR environments on desktops. Master’s thesis, 2021.
- [34] L. Vanacken, T. Grossman, and K. Coninx. Multimodal selection techniques for dense and occluded 3D virtual environments. *International Journal of Human-Computer Studies*, p. 237–255, 2009.
- [35] J. Wagner, C. Freitas, and L. Nedel. Comfortable immersive analytics with the virtualdesk metaphor: Case studies and perspectives. *IEEE Computer Graphics and Applications*, 39:41–53, 2019. doi: 10.1109/MCG.2019.2898856
- [36] X. Wang, L. Besançon, D. Rousseau, M. Sereno, M. Ammi, and T. Isenberg. Towards an understanding of augmented reality extensions for existing 3D data analysis tools. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 2020.
- [37] S. Zhai. User performance in relation to 3D input device design. *SIGGRAPH Comput. Graph.*, 32(4):50–54, nov 1998. doi: 10.1145/307710.307728
- [38] Q. Zhou, G. Fitzmaurice, and F. Anderson. In-depth mouse: Integrating desktop mouse into virtual reality. In *CHI Conference on Human Factors in Computing Systems*, CHI ’22. Association for Computing Machinery, New York, NY, USA, 2022. doi: 10.1145/3491102.3501884