



HAL
open science

Comparing Physical and VR-emulated Ultra-walls in a Classification Task

Olivier Chapuis, Raphaël James, Marzieh Rafiei, Anastasia Bezerianos

► **To cite this version:**

Olivier Chapuis, Raphaël James, Marzieh Rafiei, Anastasia Bezerianos. Comparing Physical and VR-emulated Ultra-walls in a Classification Task. Proceedings of the 35th Conference on l'Interaction Humain-Machine, AFIHM; ACM, Mar 2024, Paris, France. pp.1 – 14, 10.1145/3649792.3649797 . hal-04451520

HAL Id: hal-04451520

<https://hal.science/hal-04451520v1>

Submitted on 11 Feb 2024

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.

Comparing Physical and VR-emulated Ultra-walls in a Classification Task

Comparaison entre ultra-murs physiques et ultra-murs émulsés en RV dans une tâche de classification

OLIVIER CHAPUIS, Université Paris-Saclay, CNRS, Inria, LISN, France

RAPHAËL JAMES, Université Paris-Saclay, CNRS, Inria, LISN, France

MARZIEH RAFIEI, Université Paris-Saclay, CNRS, Inria, LISN, France

ANASTASIA BEZERIANOS, Université Paris-Saclay, CNRS, Inria, LISN, France

Due to their large size and high pixel density, ultra-walls are efficient devices for manipulating large datasets. However, they are expensive, require large rooms, and are complex to maintain. This paper explores the possibility of emulating an ultra-wall in a head-mounted display (HMD). We first analyze the ability of an HMD to emulate the high pixel density of an ultra-wall and show that they cannot yet emulate an ultra-wall. Given this limitation, in a first study we emulate two different virtual walls, a scaled wall and a simulation of a HMD that features a pixel density matching human visual capacity. We also consider different virtual navigation techniques to deal with huge virtual walls and situations where users cannot physically travel. In a follow-up study, we compare an actual physical ultra-wall with a virtual emulation of a wall of the same size in two commercial HMDs using physical navigation. Our results suggest that the pixel density of state-of-the-art HMDs does not allow to reach the performance of an ultra-wall but that this might be possible in the future with HMDs that reach human visual capacity (*e.g.*, two displays at 2600 *dpi* with a resolution of 5600×5600 pixels). However, our results also suggest that even with such HMDs, users would most likely prefer to use an ultra-wall.

En raison de leur grande taille et de leur haute densité de pixels, les ultra-murs sont des dispositifs efficaces pour la manipulation de masse de données. Cependant, ils sont complexes, coûteux et nécessitent de grandes salles. Cet article explore la possibilité d'émuler un ultra-mur dans un casque de réalité virtuelle (CRV). Nous analysons d'abord la capacité d'un CRV à émuler la densité de pixels d'un ultra-mur et montrons qu'ils ne peuvent pas encore le faire. Pour explorer cette limitation, dans une première étude, nous émuloons deux types de murs virtuels, un mur agrandi et une simulation d'un CRV doté d'une densité de pixels correspondant à la capacité visuelle humaine. Nous utilisons différentes techniques de navigation virtuelle pour interagir avec de très grands murs virtuels et pour le cas où l'espace est limité. Dans une deuxième étude, nous comparons un ultra-mur physique avec une émulation d'un mur de la même taille dans deux CRVs en utilisant la navigation physique. Nos résultats suggèrent que la densité de pixels des CRVs ne permet pas d'atteindre les performances d'un ultra-mur, mais que cela pourrait être possible à l'avenir avec des CRV atteignant la capacité visuelle humaine. Cependant, nos résultats suggèrent que même avec de tels CRVs, les utilisateurs préféreraient probablement utiliser un ultra-mur.

CCS Concepts: • **Human-centered computing** → **Graphical user interfaces**; **Virtual reality**.

Additional Key Words and Phrases: Head mounted display, Wall-sized display, Pixel density, Pick-and-Drop, Virtual navigation

Mots Clés et Phrases Supplémentaires: Casque de réalité virtuelle, mur d'écrans, densité des pixels, Pick-and-Drop, Navigation virtuelle

Reference:

Olivier Chapuis, Raphaël James, Marzieh Rafiei, and Anastasia Bezerianos. 2024. Comparing Physical and VR-emulated Ultra-walls in a Classification Task.

This is the author's draft version of the work. It is posted here for your personal use. Not for distribution. The definitive version of the record is published in IHM '24: Proceedings of the 35th Conference on l'Interaction Humain-Machine.

© 2024




Fig. 1. Virtual wall: A user moving a disc while performing the classification task of Liu et al. [30] on a $5.5\text{ m} \times 1.8\text{ m}$ virtual wall display that is rendered in an HMD (large letters: low information density condition).

1 INTRODUCTION

Ultra-high resolution wall-sized displays (ultra-walls) are large wall displays with a pixel density similar to that of the screen of a workstation (typically 100dpi). They allow to display a considerable amount of pixels and are appropriate platforms for manipulating a large amount of data [3, 19, 30, 38, 39].

However, ultra-walls are expensive, need large rooms, and are complex to set up and maintain. For instance, the cost of the wall used in Liu et al. [30] to compare a desktop and an ultra-wall, and the one used in our second study, can be estimated to at least 100k€. In addition, ultra-walls need large rooms, and given the number of pixels to be rendered, they should be driven by a cluster of high-end workstations (at least for now). For these reasons, these platforms, although powerful, can be very power-consuming and complex to setup, maintain and program for.

In comparison, Virtual Reality (VR) head-mounted displays (HMDs) are not too expensive, do not necessarily need large rooms, are easy to maintain, and benefit from various available development environments, such as Unity3D. Moreover, work in the recent Immersive Analytics trend [9, 42] suggests using HMDs for data analytics [11, 26, 47]. Although 3D visualization for abstract data has some drawbacks [13, 34], it is nevertheless possible to display them on 2D surfaces rendered in HMDs.

Thus, a natural research question is to compare the performance of an HMD with an ultra-wall when manipulating a large amount of data. This article explores this question, examining in particular whether it is possible to emulate an ultra-wall in an HMD (as is done in Figure 1) with respect to its ability to show large amounts of information.

The size and the pixel density of an ultra-wall determine the quantity of information that can be displayed. Both characteristics are fundamental in the study of ultra-walls [14, 30, 52]. The big size of an ultra-wall allows for displaying a large number of dataset elements concurrently, as its high pixel density allows to display these elements (e.g., images, plots, or text descriptions) in high enough resolution to be easily distinguishable and readable.

As HMDs could emulate a wall display of arbitrary size, in this paper, we start by analyzing, in a theoretical way, the pixel density of a virtual wall rendered in HMD. We take as examples the HTC Vive Pro 1 & 2 and the Varjo XR-1 & XR-3. These HMDs are currently the highest-resolution commercially available HMDs. As we will see, headsets' pixel density is sub-optimal and does not allow, a priori, to emulate the pixel density of an ultra-wall.

Then, to empirically validate the above computational approach, we run two user studies. For this purpose, we used the classification task that has been proposed by Liu et al. [30], in their case to compare an ultra-wall and a desktop. This task has been used in several other works to compare interaction techniques or paradigms [1, 21, 28, 35]. The advantage of this task is that it involves both interaction (pick-and-drop) and navigation (virtual or physical pan & zoom). More importantly, the task operationalizes information density, and thus the quantity of information the device can render. Liu et al. [30] operationalizes this information density by simply considering classes represented by letters of different sizes. Large letters operationalize low information density (see Figure 1), and small letters operationalize high information density.

In our first study, we consider a special emulation of a wall by considering a scaled virtual wall to make the task manageable. We also consider a simulation of an HMD that features a pixel density matching human visual capacity (60 pixels per degree). We use virtual navigation techniques (Teleport and Drag) to deal with our huge scaled virtual wall and situations where users cannot physically travel (no space available). In the second study, we compare an actual physical wall, a Vive Pro, and a Varjo XR-1, and a too-optimistic simulation of a “perfect-resolution” HMD under physical navigation (users can walk in the room). In both studies, we used the exact same task as Liu et al. [30], allowing us to compare the results of our studies and those of Liu et al. [30].

Combined, our results suggest that the pixel density of state-of-the-art HMDs does not allow to reach the performance of an ultra-wall but that this might be possible in the future with HMDs that reach human visual capacity (*e.g.*, two displays at 2600 *dpi* with a resolution of 5600×5600 pixels). However, our results also suggest that even with such HMDs, users would most likely prefer to use an ultra-wall.

2 RELATED WORK

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

2.1 Screen Size and Resolution

There is a long history in HCI about comparing screens of different sizes and resolutions. Pioneer work mainly compares standard monitors with larger screens (*e.g.*, multi-monitor setting) in high level “desktop tasks” [5, 12, 18, 43]. Most of these results suggest that users are more efficient with larger screens. The above works consider users seated in front of the display, and thus physically static (besides eyes and head movement). However, with very large screens of high resolution, like ultra-walls, it is not possible for the users to see (distinguish) all the pixels at a fixed position, and thus users should physically travel the wall room to take advantage of both the size and the high resolution, or should use virtual navigation (*e.g.*, pan & zoom) to bring information closer to them. We consider both such navigation in our work.

Interestingly, it has been shown that, for an ultra-wall (2.7m wide, 96 *dpi*), physical navigation is more efficient than virtual navigation for simple search and pattern finding tasks [3]. Several subsequent works studied the cause of this effect. Ball and North [2] suggests that the field of view is not a key factor, Ragan et al. [37] suggests that the constant positioning of information is important, Jansen et al. [24] studies spatial memory, and Liu et al. [30] advocates that the main cause is the power of head movements to reach information (*i.e.*, immersion when close to the wall). Moreover, Liu et al. [30] insist that such results depend on the “density of information” and the full use of the high resolution.

2.2 HMD versus Other Displays

There are a lot of works that compare HMDs with desktops and CAVE style environments. Santos et al. [44] provides a good overview of such works done before 2009. Most of this work concerns tasks in a 3D virtual world (*e.g.*, navigation in a maze or buildings, volume, visualization, etc). Results are mixed regarding performance, but HMDs are often preferred and some results suggest that users create a better mental model of the environment with HMDs.

Subsequent works are in the same lines: 3D tasks where resolution is secondary, have very few results on performance differences, but find some perceptual advantage for HMDs. For instance, Lhemedu-Steinke et al. [27] suggests a better concentration, involvement and enjoyment with a HMD than with a large display. Also, [46] suggests a better perceived spatial presence of the virtual environment with an HMD than with a desktop. On the other hand, some works report cybersickness issues with HMDs that do not occur with standard displays [8, 45].

2.3 HMD in Immersive Analytics

The interest around immersive visual analytics [9, 42] in the visualization and HCI community has increased in the recent years. In particular, HMDs have become very popular for visualizing abstract data with no obvious 3D representation. Although the use of 3D in this context is disputable [13, 34], some works have been conducted to take advantage of the immersion and the 3rd dimension to visualize multivariate data [4, 10], scatterplots [26, 36, 47], networks [11, 50], small multiples [31], *etc.* Surprisingly, there are very few works that compare HMD with other types of displays in the context of immersive analytics. Cordeil et al. [11] compares an HMD and a CAVE style $9m \times 1.7m$ horseshoe shaped wall display at 34 dpi with stereo 3D rendering for collaborative analysis of connectivity of simple networks. Results suggest that the HMD is more efficient than the CAVE, but collaboration is out of the scope of this report. To evaluate the possible advantage of immersive space-time cube geovisualization, Wagner Filho et al. [48] compares an HMD and a desktop, but there is no decisive results in term of performance.

2.4 Emulating Walls in VR HMDs

Recent work uses virtual reality HMDs to simulate various types of “physical” environments, in particular walls, to evaluate them without the need to build or deploy the full physical environment. A pioneer work in this area is the work of Mäkelä et al. [32] where they suggest conducting field studies on public displays in virtual reality. They observed similar phenomena in real-world and virtual public display setups, such as engagement with the displays. Similar types of work have been conducted to emulate Augmented Reality (AR) [16], real-world authentication prototypes [33], cross-reality systems [15], *etc.*

However, to our knowledge, the only work that focuses on emulating a wall display is the work Mäkelä et al. [32]. Moreover, our goal is different, Mäkelä et al. [32] focus on field studies and high-level user engagement, as we want to compare an ultra-wall and HMDs with a focus on their resolution.

3 BACKGROUND AND GEOMETRICAL MODELS OF VISION

We describe next the well-known geometrical model of vision and through it we examine the possibility of emulating a wall display in an HMD in terms of pixel density.

3.1 Principle of the monocular eye for a physical screen

The angular size of an object is the angle a viewed object creates on the human retina. The relationship between the size of a circle of diameter w seen on a screen positioned at distance d from the eye, and its angular size μ in the human “cyclopean” eye, is given by $w = 2 \cdot d \cdot \tan \frac{\mu}{2}$ (Figure 2). A human with normal vision can distinguish an object that is not smaller than an angular size of $1'$, *i.e.*, one minute of arc $= \frac{1}{60}^\circ$. For example if the circle mentioned above had an angular size of $1'$ then the viewer would be able to tell it apart from other similar circles placed side-by-side. This means that if this circle was a screen’s pixel, the human vision’s optimal resolution would be about 60 pixels per degree (ppd). From this, we can determine the smallest size of such a circle that can be distinguished by the human eye, by setting the angular size to $1'$, getting $w_{min} = 2 \cdot d \cdot \tan \frac{1'}{2} \sim 0.00029 \cdot d$. By abuse of language, we call this smallest entity a *visual pixel*.




Fig. 2. Simplified scheme for the monocular eye looking at a physical screen. $w = 2 \cdot d \cdot \tan \frac{\mu}{2}$.

Device	<i>dpi</i>	pixel size	optimal visual distance
Smartwatch (Apple watch s. 8)	330	0.0769 mm	26.4 cm
Smartphone (OnePlus 10T)	394	0.0640 mm	22.0 cm
Tablet (iPad)	264	0.0960 mm	33.0 cm
Laptop (15", FullHD)	145	0.1750 mm	60.2 cm
Our Wall (4K screens)	140	0.1800 mm	61 cm
Our Wall (8K screens)	280	0.0900 mm	30.9 cm

Fig. 3. Different devices: resolution, corresponding pixel size and optimal viewing distance. A closer distance will result in pixelization; a farther distance will result in details not being perceived.



Fig. 4. The wall that we used in the second study.

Conversely, we can compute the “ideal” visual distance from a computer screen given the size p of its pixels, *i.e.*, the distance where the pixel size is equal to the size of a visual pixel, by just inverting the previous equation $d_{ideal} = \frac{p}{2 \cdot \tan \frac{1}{2}}$. Within this distance, we have a pixelization effect, *i.e.*, users can see the pixels, and beyond it, users might lose information rendered on the screen.

Figure 3 provides examples for standard devices, from a smartwatch to the ultra-wall that we used in the second study (Figure 4). We can observe that, **current screen technology for smartwatches, smartphones, laptops, and ultra-walls has reached human visual capacity and beyond.**

3.2 Principle of the monocular eye for a VR headset

We now turn our attention to HMDs. One particularity of HMDs is that the screens are (and must be) very close to the eyes. They are so close to the eyes that the human eyes cannot accommodate an image rendered on the screen. Thus, common VR HMDs use convex lenses to create a virtual image of the screens at a far distance from the eyes (*e.g.*, 1.5 m), avoiding any need for the eyes to perform an accommodation.

In this context the “encadré” (frame) below explains how to compute (from the characteristics of an HMD) the angular size of a pixel of the virtual image in the users’ eyes, the ppd of the HMD, and the ideal size of a pixel on the display, *i.e.*, the size so that the pixel in the virtual image has an angular size of 1 arc minute.

Table 1 provides examples of these values for different headsets. Note that AR headsets such as HoloLens use technology other than convex lenses, but the optic system of such headsets has the same goal: to build, using the projected images on the glasses, a virtual image at a large distance from the eyes. Unfortunately, this table shows our lack of information, because manufacturers do not disclose technical characteristics of their HMDs such as focal distance and the distance from the lens to the screens (we provide our own estimations for the Vive Pro).

However, we can observe that **current VR headsets are far from being optimal regarding human visual capacity in their full field of view.** Indeed, the bests AR headsets reach 30 ppd (Varjo XR-3, Vive Pro 2), half of what would be needed. The Varjo’s use an interesting technology using an additional 1920×1920 pixels display (one per eye and a semi-transparent mirror system) that is “projected” and concentrated in the center of the field of view (we call this display a “focus display”, FD in Table 1). Varjo claims that this technology allows reaching the optimal resolution of 60 ppd, and even more, on a part of the field of view (*e.g.*, 27 × 27 degree).

However, covering the entire field of view with a resolution of 60 ppd would require screens with a resolution of about 5600×5600 pixels (more or less, depending on the field of view) or focus displays that can move to follow the users’ gaze (and in particular eyes saccades, which seems unlikely, see Kim et al. [25], however).

Let us consider an HMD using convex lenses with focal distance f (and width and height w_l and h_l) located between the eyes and the screen (of width and height w_s and h_s), at a distance d_l from the screen (as seen in Figure 5). As stated by Wetzstein et al. [51] during their SIGGRAPH class, the Gaussian thin lens formulas tell us that the distance from the lens to the virtual image equals to:

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

with d_l the distance from the lens to the screen, and f the focal distance of the lens (Figure 5). From this, we can infer that the distance D from the eye to the virtual image is:

$$(2) \quad D = \left| \frac{1}{\frac{1}{f} - \frac{1}{d_l}} + d_e \right|$$

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

$$(3) \quad M = \frac{f}{f - d_l}$$

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

$$(4) \quad A = 2 \cdot \arctan \frac{M \cdot p}{2 \cdot D}$$

With this same formula, we can compute the (horizontal or vertical) field of view for one eye by taking for p the (horizontal or vertical) display size. Conversely, we can compute the ideal size of the pixel on the display, *i.e.*, the size so that the pixel in the virtual image has an angular size of 1 arc minute:

$$(5) \quad P = \frac{2 \cdot D \cdot \tan \frac{1'}{2}}{M}$$




Fig. 5. Simplified scheme for optics of a VR headset. f : focal distance, d_e : distance between the eye and the screen, d_l : distance between the lens and the virtual image.

model	f	d_l	M/D	display (pixels)	resolution	PPD	ideal resolution	FoV (both eyes)
phone vr	40	39	40/1.55 m	75x75mm	16 ppmm 400 dpi	15	87 ppmm 2213 dpi	80°
Vive Pro	33	32	33/1.06 m	1440x1600	25 ppmm 640 dpi	17	105 ppmm 2668 dpi	107°
Vive Pro2	33	32 ^x	33/1.06 m ^x	2448x2448	35 ppmm 890 dpi	25	105 ppmm 2668 dpi	116°H, 96°V
Varjo-XR1	?	?	?	1440x1600	?	18	?	87°
Varjo-XR1 FD	?	?	?	1920x1080	?	>60	?	28°H, 16°V
Varjo-XR3	?	?	?	2880x2720	?	30	?	115°H, 90°V
Varjo-XR3 FD	?	?	?	1920x1920	?	70	?	27°
HoloLens	?	?	?	1268x720	?	41 [§]	?	30°H, 17°V*
HoloLens 2	?	?	?	1440x936	?	20 [§]	?	43°H, 29°V*

^x Values inferred from the Vive Pro values.

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

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

Table 1. Specification for the focal distance, lens-to-screen distance, FoV (H: horizontal, V: vertical), and resolution of various VR devices (ppmm: pixels per millimeter; dpi (or ppi): pixels per inch; ppd: pixels by degree). All specifications given are for one eye only. Technical characteristics are most of the time not disclosed by the manufacturer, and are either retrieved from 3rd party analysis (see table notes) or marked in the table by a question mark when the values are not known.

3.3 Emulating a Wall

The emulation of a wall display in VR can follow a simple procedure: we use a rectangle of the same size as the wall (the virtual wall) where we render the same scene that is rendered in the physical wall using physical dimensions in the VR space. Next, we are interested in calculating under which circumstances the users can see details rendered on this virtual

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

$$\frac{2 \cdot d_v \cdot \tan \frac{1'}{2}}{ppd} \quad (6)$$

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

For instance, the dpi of a virtual wall at 1 m of the UC , a comfortable distance to work with a wall, is of 24 dpi with the Vive Pro, 26 dpi with the Varjo XR-1 (standard display), of 36 dpi with the Vive Pro 2, of 44 dpi with the Vajo XR-3 (standard display), and should have, a priori, an “optimal” dpi of 87 with the Varjo XR-1 and Varjo XR-3 focus displays.

4 THE CLASSIFICATION TASK

Liu et al. [30] compared an ultra-wall ($5.5 m \times 1.8 m$ at 100 dpi) with a desktop (30" display at 100 dpi) in a classification task that abstracts various concrete classification and scheduling tasks (such as classifying brain scans and schedule the CHI 2013 conference). The scene rendered for the task has the size of the ultra-wall and is divided in 32 containers, each containing at most six discs with a letter, see Figure 1. The task consists of moving the discs so that all the containers contain discs with the same letter. The task starts with a partially classified scene where, in each container, there is a majority of discs with the same letter, which are shown in green, and the participants have to move miss-classified discs, shown in red, so that at the end all discs are green (correctly classified).

The difficulty of the task is controlled by the number of different letters (the classes). Liu et al. [30] consider an *Easy* task with two letters (two classes: “C” and “D”) and a *Hard* task with four letters (four classes: “H”, “K”, “N” and “R”).

The most important factor in Liu et al. [30] is the size of the letters (factor LETTERSIZE) that represents information density. Small letters correspond to a large amount of information (e.g., a lot of text can be displayed), and large letters represent a low amount of information. Liu et al. [30] use three letter sizes:

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

The main result of Liu et al. [30] is that for large letters, where almost no navigation is needed, the desktop is faster than the ultra-wall, and that for small letters, where a lot of virtual or physical navigation is needed, the ultra-wall is faster than the desktop (35% faster for the hard task).

4.1 The classification Task in VR

To “replicate” the classification task in an HMD, we could emulate a wall in virtual reality: we use a rectangle of the same physical size as the wall (the virtual wall) where we render the same scene using physical dimensions (see Figure 1).

For the *Large* letter size condition, letters are large enough to be easily seen when the user camera is at 2 m of the virtual wall, and, a priori, a participant can perform the task without moving from their starting position. We consider

this condition for testing whether the basic input interactions could impact the results: we expected no large differences between the different `DEVICE` conditions we will test.

The *Small* letter size condition is more interesting. Indeed, here a letter starts to be readable when it is rendered (in the VR HMD) in a box with a width of at least 4 pixels and clearly readable when rendered in a 5 pixels wide box (Figure 6). The width of a letter is 1.8 mm, and thus, we need to have pixels of at least $\frac{1.8}{4} = 0.45$ mm, which correspond to a wall at 56 dpi and an optimal distance of 1.55 m.



Fig. 6. Mapping the letter N into a 4x4 and a 3x3 matrix.

With the Vive Pro, Equation 6 shows that the user camera *UC* should be at a distance of about 44 cm of the virtual wall so that a user can see the letters (with a Vive Pro 2 we estimate this distance to 64 cm, and a little bit more with the standard display of the Vajo XR-3).

With an ideal HMD at 60 ppd, this distance is obviously of 1.55 m, and thus, in theory, with the Varjo XR-1 (and XR-3) and its focus display, this distance can be estimated to 1.55 m (as Varjo claim 60 ppd on the focus display). However, this is true only on a small part of the field of view, and for the rest of the field of view, a letter will be seen when the user camera is about 46 cm (XR-1) or 67 cm (XR-3) of the virtual wall.

Thus, we consider a special condition *Simulated* where we try to simulate an optimal VR HMD at 60 ppd. This is possible because of the nature of the classification task that considers isolated letters to represent information density. We, thus, just have to scale up the size of the letters to simulate an HMD with a higher ppd. For instance, to simulate an ideal HMD at 60 ppd we just have to scale the letters by a factor of 3.5, allowing a user to see the letters at a distance of 1.55 m in the Vive Pro. We consider such a condition in our first experiment.

Moreover, to understand the impact of scaling information more globally, we also considered an additional condition *Scaled* where we scale the entire wall (not just the letters), by a factor of 3.5. This results in a virtual wall of $19.25\text{ m} \times 6.3\text{ m}$. The width of a small letter is then of 6.3 mm and is displayed with about 4 pixels when the wall is at a distance of 1.55 m from the *UC*, allowing, a priori, participants to perform the classification at a reasonable distance of the wall. However, given the size of this wall, virtual navigation is needed to perform the task (*e.g.*, to reach the top of the wall).

4.2 Virtual Navigation Techniques

To perform the classification task, the participants must move the discs with miss-classified letters, and navigate in front of the scene that is made up of the disk containers. Following the spirit of Liu et al. [30], we decided to use the standard controller associated with our HMD and ray-casting for pointing. Pick-and-drop [41] is used to move a disc: a click with the pad of the Vive controller picks a disc, then the disc follows the ray on the plan of the virtual wall, and another click drops it in a container. Using pick-and-drop avoids long drags, which are not practical (especially with a wall), and also allows to use the controller to perform other actions during the pick-and-drop.

As explained in the introduction, in VR, we want the users to work without the need for a large physical space (like the big rooms needed to install actual Ultra-walls). we consider users working in more confined spaces sitting comfortably on a chair. Nevertheless, for our seated users to be able to "move" and see the content rendered in the virtual walls, we need virtual navigation techniques.

We consider two navigation techniques. For both techniques, the users can control the distance between them, *i.e.*, the (*UC*), and the virtual wall (*z* translation) by using horizontal scroll on the pad of the Vive controller (Figure 7-(a)), similarly to how users zoom in and out using their mouse wheel. We consider two basic techniques for *xy* navigation (virtual wall plan):




Fig. 7. (a) HTC vive controller. (b) The Drag technique. (c) The Teleport technique.

- *Drag*: by pressing the pad and moving the Vive controller, the users can move the wall on its plan (*i.e.*, dragging in all the directions following the ray-casting of the controller, without modifying their distance to the wall), similarly to how users can drag a scene or a document in a desktop when pressing a mouse button. [Figure 7-\(b\)](#).
- *Teleport*: The users point on the wall and click using the trigger of the Vive controller ([Figure 7-\(a\)-right](#)). This teleports the users (UC) in front of this point without modifying their distance to the wall. This specific teleportation technique is more effective than classical teleportation (*e.g.*, pointing on the floor), especially in the case of a huge wall, as it directly positions the users at the desired position in front of the wall. [Figure 7-\(c\)](#).

Thus, the goal is that, as observed in Liu et al. [30], the users will choose an appropriate distance to the wall to see enough details (the letters) using vertical scroll on the pad and then will navigate at a constant distance of the wall in the same plan of the wall (*Teleport*) or move the wall in its plan (*Drag*). Note that both techniques can be used during a pick-and-drop.

We consider the two above techniques as they are standard for interacting in VR: (i) navigation in the virtual world with *Teleport* is common, efficient [6, 7] and reduces motion sickness [17]; (ii) and moving/dragging objects using drag with *Drag* is also standard in VR and is the basic technique for navigating maps or documents in GUI.

5 FIRST STUDY: VIRTUAL NAVIGATION, SCALE AND SIMULATE

In this first study, we set out to compare how the different virtual walls perform and if the performance conforms to our theoretical calculations. We run the classification task in a Vive Pro (when we ran the experiment, the Vive Pro was the commercial HMD with the higher resolution). with the participants seated on a chair and using the virtual navigation described in the previous section, factor NAVTECH with two levels *Teleport* and *Drag*. However, during early testing we found that performing the task with the Vive Pro for the *Small* letters condition was really tedious (recall that the UC should be at about 44 cm of the virtual wall), and we decided to remove it from the study.

Thus, for the *Small* letters condition, we consider the two conditions *Scaled* and *Simulated* described in the Classification section. They allow to perform the task at a reasonable distance from the wall (1.55 m in theory and 1 m considering the results of Liu et al. [30]). The *Scaled* condition takes advantage of the flexibility of VR to emulate any kind of wall and it allows to overcome the resolution issue by scaling the wall (and thus the letters). The *Simulated* condition allows to jump into the future to have an idea of the performance of an ideal HMD regarding resolution. Thus, our second main factor is DEVICE with two levels *Scaled* and *Simulated*.

We used the exact same tasks as in Liu et al. [30] (same disc-letters layouts), with the goal to compare our results with those of Liu et al. [30]. Thus we do not consider an ultra-wall condition in this first experiment and our virtual wall has the same base size as the wall of Liu et al. [30] (multiplied by 3.5 for the *Small-Scaled* condition).

Note that for *Large* letters, there is no problem performing the task with the Vive Pro with a virtual wall of the same size as the wall of Liu et al. [30] (even with no navigation). So with *Scaled* we will not scale the wall for the *Large* letters. In practice, *Scaled* can be seen as an adaptive virtual wall that scales the wall when needed. Thus, for *Large* letters, the conditions *Scaled* and *Simulated* are equivalent, and thus, we only run it as one single condition (to avoid running both *Scaled* and *Simulated* for *Large* letters and tire our participants). In addition, we are mainly interested in the *Small* letters condition; the *Large* letters are mainly useful to compare the input devices later in this paper (the task is, under this condition, a repetition of pointing tasks on large targets: discs and containers).

5.1 Hypotheses

One important goal of this experiment is to validate the computation presented in the two previous sections. Thus, our first hypothesis relates to these computations and states that:

(H₁) For the *Small* letters conditions, participants will perform the task at a virtual distance of about 1 m from the virtual wall, distance predicted by our calculations and Liu et al. [30].

Regarding navigation techniques, we expect different performance. With a huge wall (like *Scaled*) that needs a large amount of navigation over big distances, we may expect that teleportation will be more efficient. But when the amount of navigation needed is not as much (for example, for *Simulated* the wall size is smaller) we may expect that the better control that *Drag* provides will allow for better performance. However, this is an assumption given the nature of the techniques and navigation needed, and it is not supported by some related work or theory. Thus our next hypothesis is:

(H₂) When considering navigation NAVTECHS, for *Scaled* participants will be faster with *Teleport* than with *Drag*; and for *Simulated*, participants will be faster with *Drag* than with *Teleport*.

As *Scaled* is larger it will require more navigation than *Simulated*, thus we could expect that:

(H₃) When considering DEVICES, participants will be faster with *Simulated* than with *Scaled*. This only applies to *Small* letters (as the *Large* is identical across devices).

Moreover, by design and given the results of Liu et al. [30], we predict that participants will be faster with *Large* than with *Small* and faster with *Easy* than with the *Hard*.

Finally, we expect that when we compare with the results of Liu et al. [30]:

(H₄) We will found only small differences between *Simulated* and the result of Liu et al. [30] for the *WALL* in the *Small* letters condition.

5.2 Participants

We recruited 12 volunteers (4 women, 8 men) aged from 20 to 31 (mean 25.2, median 24.5). They were all right-handed, were not color blind, and had normal vision. Moreover, all the participants had already used a HMD, and they evaluated their familiarity with VR with a median of 3 (on a 1–5 Likert scale).

5.3 Apparatus

We used a HTC Vive Pro with one controller. Participants were seated in a fixed chair to avoid motion sickness, but they could move their hands, their head, and possibly their torso. The experimental software was written using Unity 3D with the SteamVR plugin.

5.4 Design and Procedure

The experiment is a $2 \times 2 \times 2 \times 2$ mixed-factorial within-participants design with factors:

- NAVTECH: *Drag* and *Teleport*;
- DEVICE: *Scaled* and *Simulated*;
- LETTERSIZE: *Large* and *Small*;
- DIFFICULTY: *Easy* and *Hard*.

All factors are crossed but as we discussed, (*Scaled, Large*) and (*Simulated, Large*) are the same condition, and thus DEVICE \times LETTERSIZE leads to 3 conditions: (*Scaled~Simulated, Large*), (*Scaled, Small*) and (*Simulated, Small*)

The experiment was split into two sessions (separated by one or two days), one for *Drag* and the other for *Teleport* and the presentation order was counter-balanced. A session started with four training tasks, and then the participants performed two times the task for each of the six DEVICE \times LETTERSIZE \times DIFFICULTY conditions (*i.e.*, 2 repetitions). The presentation order of these conditions was counter-balanced among the participants using a 6-Latin square.

We use the *exact* same classification tasks (discs layouts) as the one used in Liu et al. [30].

Participants started by signing a consent form that highlighted the risks of cybersickness and the operator verbally warned participants about cybersickness and urged them to immediately stop the experiment if they felt bad (none of them stopped the experiment). Then, the participants provided demographic information and read instructions to perform the task as fast as possible and to not drop a disc in a wrong container. At the beginning of each task, the user camera was positioned in the middle of the virtual wall at 2 meters of distance.

At the end of each session, the participants had to grade, on a 5-scale Likert, their perceived efficiency, enjoyment, physical fatigue, mental load, and motion sickness. Moreover, participants had to rank the two navigation techniques at the end of the second session. A session lasted about 40 minutes.

5.5 Results

We first check (H_1), then we focus on our main measure, the task completion time (*TCT*), and report on the questionnaire.

Virtual distance to the virtual wall. Figure 8 shows the average virtual distance from the wall by condition (in all the bar-charts, error bars show the 95% confidence intervals). We observe that for *Large* letters participants did not have to come closer to the wall. More importantly, for the *Small* letters condition, we can observe that the distance between the *UC* and the virtual wall is close to 1 m. Thus, (H_1) is supported.

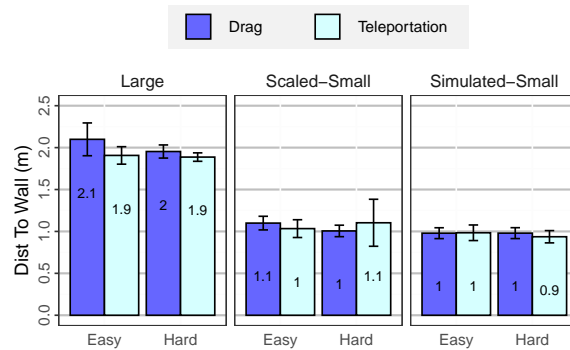
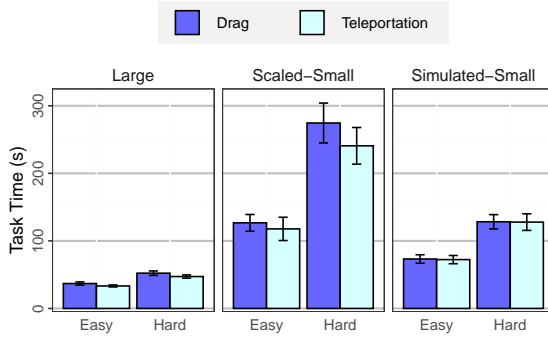


Fig. 8. Virtual distance from the wall for each NAVTECH by DIFFICULTY and grouped by DEVICE \times LETTERSIZE.



Effect	n, d	$F_{n,d}$	p	η_G^2
NAVTECH	1,11	3.68	0.0813	0.02
DEVICE	1,11	103	< 0.0001	0.41
LETTERSIZESIZE	1,11	249	< 0.0001	0.79
DIFFICULTY	1,11	228	< 0.0001	0.51
NAVTECH×DEVICE	1,11	4.60	0.0515	0.01
NAVTECH×LETTERSIZESIZE	1,11	1.13	0.2761	0.01
DEVICE×LETTERSIZESIZE	1,11	104	< 0.0001	0.41
NAVTECH×DIFFICULTY	1,11	2.64	0.1323	0.01
DEVICE×DIFFICULTY	1,11	30.1	0.0002	0.12
LETTERSIZESIZE×DIFFICULTY	1,11	175	< 0.0001	0.36

Fig. 9. (Left): Mean task completion time in seconds for each NAVTECH by DIFFICULTY grouped by DEVICE × LETTERSIZE. (Right): Degree 2 factorial ANOVA with participant as random factor for the TCT measure.

Task Completion Time. We checked for outliers and observed that, as in Liu et al. [30], 95% of the trials were within 15% of the mean completion time per condition. Only one trial was more than 25% slower (namely 39%), and we removed this trial for the analysis (a *Teleport-Scaled-Small-Hard* trial). Normality tests did not show any evidence of non-normality. We used the same statistical method as Liu et al. [30], repeated measure ANOVA, and Holm corrected (paired) t-tests.

Figure 9 shows the task completion time for each NAVTECH, by DIFFICULTY grouped by DEVICE × LETTERSIZE, and the results of the ANOVA for the model (Participant as a random factor; for the analysis we duplicate the *Large* data for *Scaled* and *Simulated* as they are the same condition): $TCT \sim \text{NAVTECH} \times \text{DEVICE} \times \text{LETTERSIZESIZE} \times \text{DIFFICULTY}$

Although *Teleport* seems slightly faster than *Drag*, there is no significant effect of NAVTECH, nor significant interactions involving NAVTECH on TCT. Thus, (H₂) is not supported in terms of performance in terms of navigation technique.

We have a significant effect of DEVICE on TCT, *Simulated* being faster than *Scaled*. As the data are the same for *Large* letters for *Scaled* and *Simulated*, we have a significant DEVICE × LETTERSIZESIZE interaction effect as the difference between *Scaled* and *Simulated* could happen only for *Small* letters. Thus (H₃) is supported. Note also that we have a significant DEVICE × DIFFICULTY interaction on TCT, and indeed, the difference between *Scaled* and *Simulated* is larger for the *Hard* task (Cohen $d = 3.07$) than for the *Easy* task (Cohen $d = 2.18$).

As expected, we observe large significant effects of LETTERSIZESIZE and DIFFICULTY on TCT: (i) participants have been faster with the *Easy* tasks than with the *Hard* tasks for all the NAVTECH × DEVICE × LETTERSIZESIZE conditions (all p 's < 0.001); (ii) participants have been faster with *Large* than with both *Simulated-Small* and *Scaled-Small* (all p 's < 0.001).

We can also observe a significant LETTERSIZESIZE × DIFFICULTY interaction effect. Indeed, the differences between the LETTERSIZESIZE conditions were relatively bigger for the *Hard* tasks than for the *Easy* tasks (Cohen d 's differences of ~ 1).

Subjective Results. Figure 10 shows the results of our questionnaire. Overall, participants found the navigation techniques effective and enjoyable, but reactions were more mixed regarding fatigue and mental load (probably because some found the task challenging). Three participants also said they suffered from motion sickness but did not stop the experiment. We notice that overall *Teleport* has been slightly better rated than *Drag*, in particular regarding fatigue, but, indeed, we could not measure any significant (or marginal) differences (all questions, Wilcoxon paired signed rank tests).

We also asked the participants to rank the navigation techniques for both small letter conditions. In the case of *Simulated-Small*, 8 participants (over 12) preferred *Drag*, as for *Scaled-Small*, 9 participants preferred *Teleport*. Thus, regarding preferences, (H₂) has an agreement score of 17 over 24 ($p = 0.064$).




Fig. 10. Results of the questionnaire for Efficiency, Enjoyment, Physical Fatigue, Mental Load, and Motion Sickness (the four questions of Liu et al. [30] that come from the NASA-TLX, plus a question regarding motion sickness). Values on the right (in green) are “better” for all the questions.

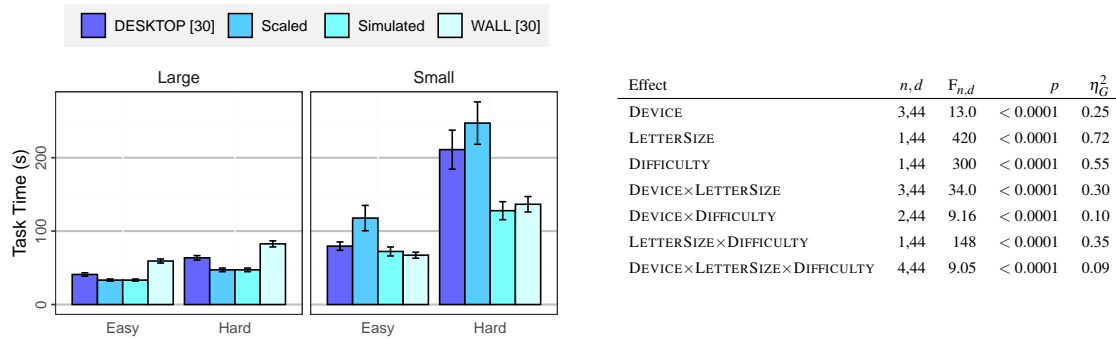


Fig. 11. (Left) Mean task completion time in seconds for each DEVICE by DIFFICULTY grouped by LETTERSIZE. (Right) Mixed factorial ANOVA with participant as random factor for the TCT measure (DEVICE between factor).

5.6 Comparing with Liu et al.

Now, we compare the data of our experiment with the data of Liu et al. [30]. To simplify, we will consider only the data for *Teleport*, as the results were slightly better than for *Drag*. In the above experiment, we recall that we used the same procedure, design, disc layout, order, etc., as in Liu et al. [30]. This allows us to merge and compare our data with that of Liu et al. [30] (note we removed the *Medium* data from Liu et al. [30] as we did not run this condition).

To summarize, we consider a $[4 \times 2 \times 2]$ mixed factorial design with factors:

- DEVICE: *WALL* and *DESKTOP* (from Liu et al. [30]) and *Scaled* and *Simulated* (from our experiment);
- LETTERSIZESIZE: *Large*, *Small* (both experiments);
- DIFFICULTY: *Easy* and *Hard* (both experiments).

with DEVICE considered as a between-participants factor. Note that, obviously, physical navigation was used in Liu et al. [30] experiment for the *WALL* condition, while we used virtual navigation in our VR conditions.

Results for TCT can be observed in Figure 11, as well as the ANOVA for the mixed model (participant as a random factor): $TCT \sim [DEVICE] \times LETTERSIZESIZE \times DIFFICULTY$. Again, obviously, we have a strong effect of LETTERSIZESIZE and DIFFICULTY on TCT, and a significant LETTERSIZESIZE \times DIFFICULTY interaction.

We focus on the significant effect of DEVICE and the *Small* LETTERSIZESIZE condition. Given the effect size of the DEVICE \times LETTERSIZESIZE interaction effect, we also consider DIFFICULTY in each case:

- *Small letters – Easy task*: *Scaled* is slower than all the other DEVICE: *DESKTOP* ($p = 0.0650$, $d = 1.22$), *Simulated* ($p = 0.0096$, $d = 1.44$), and *WALL* ($p = 0.0224$, $d = 1.64$). For *Simulated* we found no significant difference with the *DESKTOP* ($p = 0.612$) and the *WALL* ($p = 0.8166$). Recall that Liu et al. [30] find that the *WALL* was faster than the *DESKTOP* ($p = 0.0396$, $d = 1.11$).
- *Small letters – Hard task*: *Scaled* is also slower than all the other DEVICE, the difference is significant with *Simulated* ($p = 0.0012$, $d = 2.32$) and *WALL* ($p = 0.0012$, $d = 2.17$), but not with *DESKTOP* ($p = 0.8166$). For *Simulated* we find that it is significantly faster than the *DESKTOP* ($p = 0.0050$, $d = 1.77$), but the difference with the *WALL* is not significant ($p = 0.8166$). Again, we recall that Liu et al. [30] find that the *WALL* is significantly faster than the *DESKTOP* ($p = 0.0090$, $d = 1.61$).

Thus, (H₄) is supported, and to summarize, our results suggest that an HMD with a resolution matching human vision can be as efficient as an ultra-wall in tasks that involve high-density information (*Small letters*).

Note that for *Large letters*, participants have been faster in our experiment than in Liu et al. [30] experiment. We explain this result by the difference in the input devices (Vive active tracking in our experiment vs. Vicon passive tracking and mouse) and that our large letters were, in fact, larger than the ones in Liu et al. [30].

6 FOLLOW-UP STUDY: PHYSICAL WALL VS. HMDS, USING PHYSICAL NAVIGATION

In the previous study, we had seated participants and used virtual navigation (on propose), so as to avoid using an exact emulation of a wall, that would require otherwise a lot of walking in the *Small* condition with the Vive Pro (our initial tests caused too much motion sickness). However, when considering physical navigation, some tests suggested that running the *Small* condition with the Vive Pro was reasonable. A second limitation of the first experiment is that we did not compare the results of our HMDs condition directly with an ultra-wall, rather we used the results of a previously published experiment.

Thus, in this second study, we compare a 140 *dpi* ultra-wall of about the same size as the wall of Liu et al. [30] (*Wall*) with emulations of this wall in a Vive Pro (*Vive*) and in a Varjo XR-1 (*Varjo*) using physical navigation. We used physical navigation (which is required for the physical *Wall* as we cannot teleport), to ensure that any effects found are due to the rendering and not the navigation techniques. And to ensure a fair comparison (since walking is required), we emulated walls in the HMDs that are identical in size to the physical *Wall* and not bigger.

We also wanted to consider the *Simulated* condition. **Unfortunately, we made a mistake and scaled the letters by a factor of 4.2 (and not 3.5), making the classification task too easy.** However, as we will see, this led to an interesting result. We denote by *Simu-4.2* this condition (for *Large* we used the same size as Liu et al. [30], *i.e.*, we did not scale).

Thus, in this study, our primary factor is DEVICE with four levels: *Wall*, *Vive*, *Varjo*, and *Simu-4.2*. We introduced the Varjo XR-1, an HMD that includes a focus display with a (claimed) perfect-resolution.

We used again the classification task (see Classification section), with four letter classes: “H”, “K”, “N” and “R” (*i.e.*, we only consider the *Hard* task; we do not consider the *Easy* task as it led to smaller effects than the *Hard* task, in both Liu et al. [30] and our own first study). Thus, our second factor is LETTERSIZE with two levels *Large* and *Small* (as in the first study). We recall that as we are studying the effect of resolution, we are mainly interested in the *Small* size condition. The *Large* size is useful to see whether the different DEVICES perform equally well for the basic interactions (*e.g.*, pick-and-drop): we expected very similar performance for the different DEVICES in the *Large* condition.

We stopped the experiment after 6 participants once we noticed our error in the simulated condition. Nevertheless, we report the results of the 6 participants here, as they present interesting tendencies, and our mistake led to an interesting observation regarding the subjective results.

6.1 Hypothesis

Our analysis in the Background section and the above experiment (H_4) suggest that:

(H_5) For the *Small* letters, the *Wall* will out perform the *Vive* and the *Varjo*.

We can also hope that the focus display of the *Varjo* will be beneficial for small letters:

(H_6) For the *Small* letters, the *Varjo* will be faster than the *Vive*.

Moreover, as explained above, we expect that for the *Large* letters, we will find only small differences among the DEVICE. Finally, our original hypothesis regarding *Simu-4.2* was that we would observe only small differences between the *Wall* and *Simu-4.2*. However, given the mistake we made, it should be expected (at posteriori) that *Simu-4.2* should be faster than the *Wall* (bigger letters thus easier task).

6.2 Participants

We recruited 6 volunteers (5 men, 1 woman) aged from 24 to 45 (mean 30, median 27) with normal or corrected-to-normal vision (over the 12 participants initially planned). Participants were HCI researchers, engineers or graduate students in computer science. Most of them already used an HMD and a wall.

6.3 Apparatus

For the wall display (*Wall*), we used a 5.50×1.80 m display with a resolution of 30720×8640 , composed of $8 \times 4 = 32$ 4K LCD displays and driven by 8 workstations (140 dpi, see Figure 4, note that the displays of this wall also support a resolution 8K, doubling the resolution and the dpi of the wall). For the VR headsets, the Vive Pro (*Vive*) and the Varjo XR-1 (*Varjo*), we ran the experiment in the same room as the wall display room. The Varjo XR-1 was, when we ran the experiment, the best HMD available regarding resolution. We chose the Vive Pro as it is a standard, consumer-grade headset with a good resolution, and we used it in the first study. We ran the erroneous simulation condition (*Simu-4.2*) in the Vive Pro because it is more comfortable to wear than the Varjo XR-1.

For the input, we use one Vive controller (participants used the trigger of the controller for pick-and-drop), but we perform the tracking using a Vicon system (for all the DEVICE conditions). We used physical navigation for all conditions, with the cable for the VR headsets suspended so the participants would not trip on it. The software used was Unity 3D with the same base code for the HMDs and the wall.

6.4 Experiment Design & Procedure

Design. The experiment is a full factorial DEVICE \times LETTERSIZE within-participants design. We blocked by DEVICE, and fixed the LETTERSIZE presentation order to *Large* first, followed by *Small*, as we are mainly interested in effects due to resolution (that we expect with *Small* letters). In the initial 12 participants' design, we counter-balanced the device condition using a Latin square. For each device, participants had 6 trials: 2 training trials with each LETTERSIZE condition in the order defined above, then 2×2 experiment trials in the same order. We used the exact same initial letter dispositions as in Liu et al. [30] and the previous experiment to allow for a comparison of results.

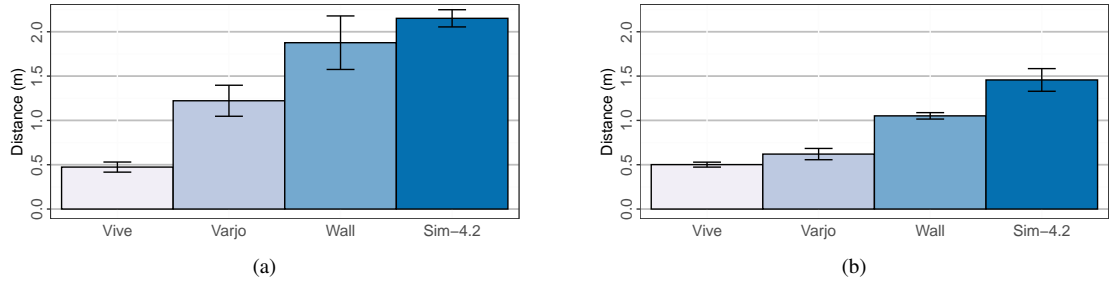


Fig. 12. (a) Result of the vision test: maximal distance (in meters) to the virtual or physical wall needed to see the small letters used in the experiment; (b) average distance to the virtual or physical wall when performing the actual classification task for the *Small* letters.

Procedure. Each participant took part in two sessions, on two separate days, seeing 2 (out of the 4) device conditions in each of the two sessions. A break was enforced in-between conditions inside a session. At the beginning of each VR condition, we warned the participants about possible motion sickness discomfort. At the end of the second session the participants filled-in a questionnaire asking to rank the techniques (with possible ex-aequo) regarding their preferences: overall, enjoyment, effectiveness, mental and physical load, and motion sickness.

Each condition started with a **vision test** to determine what is a comfortable viewing distance for each one of the conditions, and if it is close to our theoretical calculations. The operator asked the participant to look at a letter of *Small* size on the wall, and to then move back, one step at a time, until they hit the limit where they could not tell the letter anymore (at each step, we changed the letter shown, chosen among “H”, “K”, “N” and “R”). The resulting distance was recorded. The operator then explained the interactions and the task’s goal, ensuring the participant felt comfortable.

Next, participants moved onto the **classification** tasks. Before each trial, the participants were placed at a 2 m distance from either the physical wall or the virtual wall in the VR conditions. The 2 m distance was chosen as it was, a priori, a comfortable distance to perform the *Large* letters task.

6.5 Results

Distance to the Wall: Vision Test and Classification Trials. Figure 12-(a) shows the result of the vision test. Participants were able to see the *Small* letters at 50 cm with the *Vive*, which is slightly more than what we predicted (45 cm). For the *Varjo*, the participants needed to be at 1.20 m from the letters to see them, which is 35 cm less than predicted. This suggests that the focus display of the *Varjo XR-1* does not reach the claimed 60 ppd (it is closer to 50 ppd).

The result for the physical *Wall* is about 1.9 m distance, 35 cm more than the expected value of 1.55 m. We have no explanation, but we will see that in practice, the participants performed the classification task for the small letters at about 1 m from the screen. For *Simu-4.2*, the distance is about 2.2 m, more than our expectation of 2 m.

For the classification task, results differed. Figure 12-(b) shows, for the *Small* condition, the average distance of the participant to the virtual or physical wall during a trial. Note that the distance to the wall for the physical *Wall* was of about 1 m (far smaller than the 1.9 m of the vision test and very close to the results of Liu et al. [30]). This trend to have a smaller distance than the vision test is also true for *Simu-4.2*, but the difference is not as strong. In contrast, the distance during the classification task is close to the vision test distance for the *Vive*. Finally, despite its focus display, this distance for the *Varjo* is close to the one of the *Vive*, indicating the focus display did not really help in the classification task.

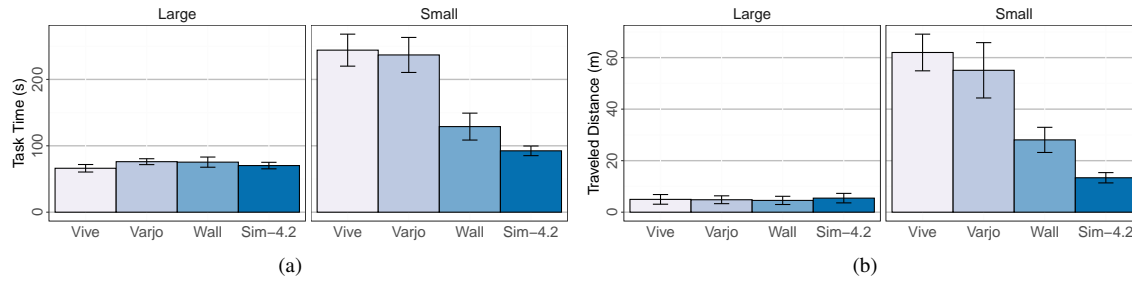


Fig. 13. (a) Task time by DEVICE×LETTERSIZES conditions during the classification task. (b) Average distance traveled by the participants' head by DEVICE×LETTERSIZES conditions (in meters).

Task Time. Figure 13-(a) shows the task time by DEVICE × LETTERSIZE conditions. For the *Large* condition, we can observe, as expected, that the task times are very similar. This suggests that the input system worked equally well for all DEVICE. For the *Small* condition, we can observe that (we did not report statistical tests as we have only 6 participants. However, the results we state are supported by t-tests, as well as by Bayesian analysis).

- As expected, the physical *Wall* was faster than the *Vive*, but also faster than the *Varjo* (this was the case uniformly for the 6 participants, with clear differences between mean times). Thus (H₅) is supported.
- The task time is similar for *Vive* and *Varjo*, making (H₆) unlikely to hold. This suggests that for a task where users have to move in front of a (virtual) wall and move their heads in all directions to see small details, the focus display of the *Varjo* is not really useful.
- *Simu-4.2* was faster than the *Wall* (clearly for 5 participants, but one participant was 2 s faster with *Wall*). This is expected given the size of the letters with *Simu-4.2* (that is by error larger than our model calculation).

Traveled Distance. Figure 13-(b) shows the average traveled distance by the participants during a trial by DEVICE × LETTERSIZE conditions. For the *Large* letter condition, participants did not need to come close to the wall and did not need to travel too much. For the *Small* condition, the results confirm the results on the task time: task time is “positively correlated” with the distance to the wall to perform the task (far away is better), and the traveled distance during the task (shorter is better).

Subjective Results. Regarding their overall preferences, the six participants ranked the physical *Wall* condition 1st and the *Simu-4.2* condition 2nd (4 participants ranked the *Vive* 3rd, and 3 participants ranked the *Varjo* 3rd, ex-aequo were allowed). Given the quantitative results above, it is not surprising that the participants preferred the *Wall* and the *Simu-4.2* conditions over the other virtual walls. However, it is surprising and interesting that the six participants uniformly preferred the *Wall* condition over the *Simu-4.2* condition, given that ex-aequo were allowed and that *Simu-4.2* exhibits better quantitative performance than the *Wall* (uniformly: for the six participants).

We obtain the same results for enjoyment and effectiveness and similar tendencies regarding mental and physical load. Only one participant reported a moderate discomfort regarding motion sickness.

Thus, the subjective results suggest that participants appreciated more the physical aspect of the wall display compared to the virtuality of a headset.

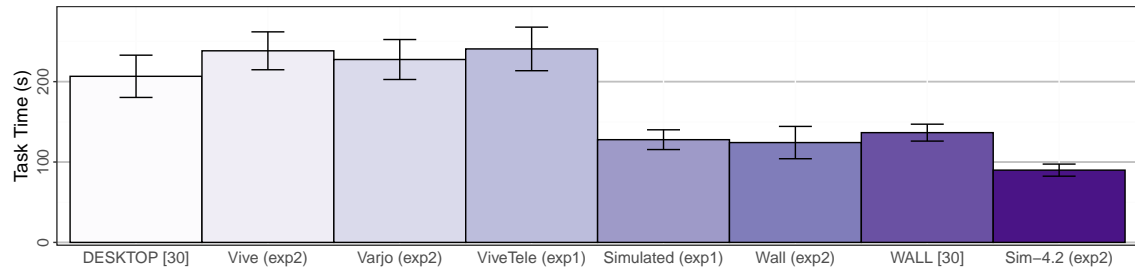


Fig. 14. Task completion time for the *Small* letters conditions (*Hard* task) for the devices conditions from Liu et al. [30] (*DESKTOP*, *WALL*), our 1st study (*ViveTeleport*, *Scaled* and *Simulated* in the *Teleport* condition) and our 2nd study (*Vive*, *Varjo*, *Wall* and *Simu-4.2*).

6.6 Putting everything together

Figure 14 shows the task completion time for the *Large* and *Small* letters conditions (*Hard* task) for all the “device” conditions considered in this paper: Liu et al. [30] and our two studies. Beside the differences already noticed in the two previous sections, we can note that for the *Small* letters, the *Simulated* condition (study 1), the *Wall* condition (study 2), and the *WALL* condition (Liu et al. [30]) exhibit close performances, and clearly outperform the *DESKTOP* (Liu et al. [30]) and our realistic HMDs condition (*Scaled*, *Vive*, *Varjo*). Note that these four last conditions have close performance, too. These results confirm our analysis of the capacity of HMDs to emulate an ultra-wall: HMDs are not yet capable to emulate physical walls, but “perfect” future HMDs will likely arrive at similar performance.

7 CONCLUSION, DISCUSSION AND FUTURE WORK

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

Moreover, our results also suggest that a hypothetical perfect-resolution HMD (60 ppd in the field of view) would be competitive in terms of performance. However, the second study suggests that users may prefer a physical wall anyway. Indeed, participants were faster and traveled less distance with the *Simu-4.2* condition than with the *Wall* condition (*Simu-4.2* that was easier due to our error), but despite that *all* participants preferred the *Wall* condition over the *Simu-4.2* condition.

A limitation of our work is that the second study has only 6 participants. However, we believe that the results are robust. For instance, we find similar results to Liu et al. [30] for the walls, and the results are consistent with our theoretical prediction. Of course, it would be interesting to run the second study with 12 participants and with the Vive Pro 2 and the Varjo XR-3, the current best resolution HMDs (about 30 ppd, only half of what is needed for a perfect-resolution HMD). Nevertheless, given our theoretical analysis and our empirical results that are consistent with this analysis, we believe that even these new HMDs will not reach the performance of a physical ultra-wall, but rather fall somewhere between the performances seen in the HMDs tested in this article and that of the ultra-wall.

In fact, it would be more interesting to test an HMD that *both* provides perfect resolution in its field of view and a view on reality (*e.g.*, a super Hololens or an ultra VR HMDs with video-see-through). Apple announced such an HMD, and the

classification task could be an interesting benchmark to evaluate its resolution and the quality of its view on reality (for example comparing a virtual wall and a physical one, both seen through the HMD).

However, HMDs should not only be tested regarding their resolution. For instance, wall displays are mainly useful for collaborative work, and a view on the reality is a decisive advantage in a co-located collaborative context. This does not necessarily mean that virtual walls are not useful even in these contexts. Results from James et al. [22] show that pairs did not hesitate to interact collaboratively with AR surfaces even in the presence of a wall display. We think that future work should compare ultra-walls with high-resolution AR headsets (e.g., the Apple Vision Pro) in collaborative tasks such as the classification task of [28] and in tasks that require less interaction and more high-level decision making and planning such as those of [20, 29, 49], but also continue to study contexts where the two types of devices can be combined to provide richer and more flexible collaborative environments [23, 40].

SUPPLEMENTAL MATERIAL

Experimental data and the discs layouts of Liu et al. [30] are available as supplemental material and online at <https://www.lri.fr/~chapis/ihm-2024-data/>. We thank the authors of Liu et al. [30] for sharing their data.

ACKNOWLEDGEMENTS

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

REFERENCES

- [1] Ignacio Avellino, Cédric Fleury, Wendy E. Mackay, and Michel Beaudouin-Lafon. 2017. CamRay: Camera Arrays Support Remote Collaboration on Wall-Sized Displays. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, 6718–6729. <https://doi.org/10.1145/3025453.3025604>
- [2] Robert Ball and Chris North. 2008. The Effects of Peripheral Vision and Physical Navigation on Large Scale Visualization. In *Proceedings of Graphics Interface (GI '08)*. CIPS, 9–16. <http://dl.acm.org/citation.cfm?id=1375714.1375717>
- [3] Robert Ball, Chris North, and Doug A. Bowman. 2007. Move to improve: promoting physical navigation to increase user performance with large displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*. ACM, 191–200. <https://doi.org/10.1145/1240624.1240656>
- [4] Andrea Batch, Andrew Cunningham, Maxime Cordeil, Niklas Elmqvist, Tim Dwyer, Bruce H. Thomas, and Kim Marriott. 2020. There Is No Spoon: Evaluating Performance, Space Use, and Presence with Expert Domain Users in Immersive Analytics. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (2020), 536–546. <https://doi.org/10.1109/TVCG.2019.2934803>
- [5] Xiaojun Bi and Ravin Balakrishnan. 2009. Comparing Usage of a Large High-Resolution Display to Single or Dual Desktop Displays for Daily Work. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09)*. ACM, 1005–1014. <https://doi.org/10.1145/1518701.1518855>
- [6] Doug A. Bowman, David Koller, and Larry F. Hodges. 1997. Travel in Immersive Virtual Environments: An Evaluation of Viewpoint Motion Control Techniques. In *Proceedings of the Virtual Reality Annual International Symposium (VRAIS '97)*. IEEE, 45–52. <https://doi.org/10.1109/VRAIS.1997.583043>
- [7] Evren Bozgeyikli, Andrew Raij, Srinivas Katkooi, and Rajiv Dubey. 2016. Point & Teleport Locomotion Technique for Virtual Reality. In *Proceedings of the Annual Symposium on Computer-Human Interaction in Play (CHI PLAY '16)*. ACM, 205–216. <https://doi.org/10.1145/2967934.2968105>
- [8] James Brooks, Riley Lodge, and Daniel White. 2017. Comparison of a head-mounted display and flat screen display during a micro-UAV target detection task. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, Vol. 61. SAGE, 1514–1518. <https://doi.org/10.1177/1541931213601863>
- [9] Tom Chandler, Maxime Cordeil, Tobias Czauderna, Tim Dwyer, Jaroslaw Glowacki, Cagatay Goncu, Matthias, Klapperstueck, Karsten Klein, Kim Marriott, Falk Schreiber, and Elliot Wilson. 2015. Immersive Analytics. In *Big Data Visual Analytics (BDVA '15)*. IEEE, 1–8. <https://doi.org/10.1109/BDVA.2015.7314296>
- [10] Maxime Cordeil, Andrew Cunningham, Tim Dwyer, Bruce H. Thomas, and Kim Marriott. 2017. ImAxes: Immersive Axes As Embodied Affordances for Interactive Multivariate Data Visualisation. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*

- (*UIST '17*). ACM, 71–83. <https://doi.org/10.1145/3126594.3126613>
- [11] Maxime Cordeil, Tim Dwyer, Karsten Klein, Bireswar Laha, Kim Marriott, and Bruce H. Thomas. 2017. Immersive Collaborative Analysis of Network Connectivity: CAVE-style or Head-Mounted Display? *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (Jan 2017), 441–450. <https://doi.org/10.1109/TVCG.2016.2599107>
 - [12] Mary Czerwinski, Greg Smith, Tim Regan, Brian Meyers, George Robertson, and Gary Starkweather. 2003. Toward characterizing the productivity benefits of very large displays. In *Proceedings of IFIP INTERACT: Human-Computer Interaction (INTERACT '03)*. IOS, 9–16.
 - [13] Niklas Elmquist. 2017. 3D Visualization for Nonspatial Data: Guidelines and Challenges. <https://sites.umiacs.umd.edu/elm/2017/10/03/3d-visualization-for-nonspatial-data-guidelines-and-challenges/> Accessed May 6th, 2020.
 - [14] Alex Endert, Christopher Andrews, Yueh Lee, and Chris North. 2011. Visual encodings that support physical navigation on large displays. In *Proceedings of Graphics Interface (GI '11)*. CHCCS, 103–110. <http://dl.acm.org/citation.cfm?id=1992917.1992935>
 - [15] Uwe Gruenefeld, Jonas Auda, Florian Mathis, Stefan Schneegass, Mohamed Khamis, Jan Gugenheimer, and Sven Mayer. 2022. VRception: Rapid Prototyping of Cross-Reality Systems in Virtual Reality. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI '22)*. ACM, Article 611, 15 pages. <https://doi.org/10.1145/3491102.3501821>
 - [16] Kota Gushima and Tatsuo Nakajima. 2021. Virtual Fieldwork: Designing Augmented Reality Applications Using Virtual Reality Worlds. In *Virtual, Augmented and Mixed Reality*. Springer, 417–430. https://doi.org/10.1007/978-3-030-77599-5_29
 - [17] M. P. Jacob Habgood, David Moore, David Wilson, and Sergio Alapont. 2018. Rapid, Continuous Movement Between Nodes as an Accessible Virtual Reality Locomotion Technique. In *Conference on Virtual Reality and 3D User Interfaces (VR '18)*. IEEE, 371–378. <https://doi.org/10.1109/VR.2018.8446130>
 - [18] Dugald Ralph Hutchings, Greg Smith, Brian Meyers, Mary Czerwinski, and George Robertson. 2004. Display Space Usage and Window Management Operation Comparisons between Single Monitor and Multiple Monitor Users. In *Proceedings of the Working Conference on Advanced Visual Interfaces (AVI '04)*. ACM, 32–39. <https://doi.org/10.1145/989863.989867>
 - [19] Mikkel R. Jakobsen and Kasper Hornbæk. 2013. Interactive Visualizations on Large and Small Displays: The Interrelation of Display Size, Information Space, and Scale. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2336–2345. <https://doi.org/10.1109/TVCG.2013.170>
 - [20] Mikkel R. Jakobsen and Kasper Hornbæk. 2014. Up Close and Personal: Collaborative Work on a High-resolution Multitouch Wall Display. *ACM Transactions on Computer-Human Interaction* 21, 2, Article 11 (Feb. 2014), 34 pages. <https://doi.org/10.1145/2576099>
 - [21] Mikkel R. Jakobsen and Kasper Hornbæk. 2015. Is Moving Improving? Some Effects of Locomotion in Wall-Display Interaction. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, 4169–4178. <https://doi.org/10.1145/2702123.2702312>
 - [22] Raphaël James, Anastasia Bezerianos, and Olivier Chapuis. 2023. Evaluating the Extension of Wall Displays with AR for Collaborative Work. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. ACM, Article 99, 17 pages. <https://doi.org/10.1145/3544548.3580752>
 - [23] Raphaël James, Anastasia Bezerianos, Olivier Chapuis, Tim Dwyer, Maxime Cordeil, and Arnaud Prouzeau. 2020. Personal+Context Navigation: Combining AR and Shared Displays in Network Path-following. In *Proceedings of Graphics Interface (GI '20)*. CHCCS/SCDHM, 13. <https://doi.org/10.20380/GI2020.27>
 - [24] Yvonne Jansen, Jonas Schjerlund, and Kasper Hornbæk. 2019. Effects of Locomotion and Visual Overview on Spatial Memory When Interacting with Wall Displays. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '19)*. ACM, Article 291, 12 pages. <https://doi.org/10.1145/3290605.3300521>
 - [25] Jonghyun Kim, Youngmo Jeong, Michael Stengel, Kaan Aksit, Rachel Albert, Ben Boudaoud, Trey Greer, Joohwan Kim, Ward Lopes, Zander Majercik, Peter Shirley, Josef Spjut, Morgan McGuire, and David Luebke. 2019. Foveated AR: Dynamically-Foveated Augmented Reality Display. *ACM Trans. Graph.* 38, 4, Article 99 (jul 2019), 15 pages. <https://doi.org/10.1145/3306346.3322987>
 - [26] M. Kraus, N. Weiler, D. Oelke, J. Kehrer, D. A. Keim, and J. Fuchs. 2020. The Impact of Immersion on Cluster Identification Tasks. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (2020), 525–535. <https://doi.org/10.1109/TVCG.2019.2934395>
 - [27] Quinatei Lhemedu-Steinke, Gerrit Meixner, and Michael Weber. 2018. Comparing VR Display with Conventional Displays for User Evaluation Experiences. In *Conference on Virtual Reality and 3D User Interfaces (VR '18)*. IEEE, 583–584. <https://doi.org/10.1109/VR.2018.8446076>
 - [28] Can Liu, Olivier Chapuis, Michel Beaudouin-Lafon, and Eric Lecolinet. 2016. Shared Interaction on a Wall-Sized Display in a Data Manipulation Task. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, 2075–2086. <https://doi.org/10.1145/2858036.2858039>
 - [29] Can Liu, Olivier Chapuis, Michel Beaudouin-Lafon, and Eric Lecolinet. 2017. CoReach: Cooperative Gestures for Data Manipulation on Wall-sized Displays. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, 6730–6741. <https://doi.org/10.1145/3025453.3025594>
 - [30] Can Liu, Olivier Chapuis, Michel Beaudouin-Lafon, Eric Lecolinet, and Wendy E. Mackay. 2014. Effects of Display Size and Navigation Type on a Classification Task. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. ACM, 4147–4156. <https://doi.org/10.1145/2556288.2557020>
 - [31] Jiazhou Liu, Arnaud Prouzeau, Barrett Ens, and Tim Dwyer. 2020. Design and Evaluation of Interactive Small Multiples Data Visualisation in Immersive Spaces. In *Conference on Virtual Reality and 3D User Interfaces (VR '20)*. IEEE, 588–597. <https://doi.org/10.1109/VR46266.2020.00-23>
 - [32] Ville Mäkelä, Rivu Radiah, Saleh Alsherif, Mohamed Khamis, Chong Xiao, Lisa Borchert, Albrecht Schmidt, and Florian Alt. 2020. Virtual Field Studies: Conducting Studies on Public Displays in Virtual Reality. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing*

- Systems (CHI '20)*. ACM, 1–15. <https://doi.org/10.1145/3313831.3376796>
- [33] Florian Mathis, Joseph O'Hagan, Kami Vaniea, and Mohamed Khamis. 2022. Stay Home! Conducting Remote Usability Evaluations of Novel Real-World Authentication Systems Using Virtual Reality. In *Proceedings of the 2022 International Conference on Advanced Visual Interfaces (AVI 2022)*. ACM, Article 14, 9 pages. <https://doi.org/10.1145/3531073.3531087>
- [34] Tamara Munzner. 2014. *Visualization Analysis and Design*. A. K. Peters Visualization Series, CRC Press.
- [35] Erwan Normand and Michael J. McGuffin. 2018. Enlarging a Smartphone with AR to Create a Handheld VESAD (Virtually Extended Screen-Aligned Display). In *International Symposium on Mixed and Augmented Reality (ISMAR '18)*. IEEE, 123–133. <https://doi.org/10.1109/ISMAR.2018.00043>
- [36] Arnaud Prouzeau, Maxime Cordeil, Clement Robin, Barrett Ens, Bruce H. Thomas, and Tim Dwyer. 2019. Scaptics and Highlight-Planes: Immersive Interaction Techniques for Finding Occluded Features in 3D Scatterplots. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI '19)*. ACM, Article 325, 12 pages. <https://doi.org/10.1145/3290605.3300555>
- [37] Eric D. Ragan, Alex Endert, Doug A. Bowman, and Francis Quek. 2012. How Spatial Layout, Interactivity, and Persistent Visibility Affect Learning with Large Displays. In *Proceedings of the International Working Conference on Advanced Visual Interfaces (AVI '12)*. ACM, 91–98. <https://doi.org/10.1145/2254556.2254576>
- [38] Fateme Rajabiyazdi, Jagoda Walny, Carrie Mah, John Brosz, and Sheelagh Carpendale. 2015. Understanding Researchers' Use of a Large, High-Resolution Display Across Disciplines. In *Proceedings of the International Conference on Interactive Tabletops & Surfaces (ITS '15)*. ACM, 107–116. <https://doi.org/10.1145/2817721.2817735>
- [39] Khairi Reda, Andrew E. Johnson, Michael E. Papka, and Jason Leigh. 2015. Effects of Display Size and Resolution on User Behavior and Insight Acquisition in Visual Exploration. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, 2759–2768. <https://doi.org/10.1145/2702123.2702406>
- [40] Patrick Reipschläger, Tamara Flemisch, and Raimund Dachsel. 2021. Personal Augmented Reality for Information Visualization on Large Interactive Displays. *IEEE Transactions on Visualization and Computer Graphics* 27 (2 2021), 1182–1192. Issue 2. <https://doi.org/10.1109/TVCG.2020.3030460>
- [41] Jun Rekimoto. 1997. Pick-and-drop: a direct manipulation technique for multiple computer environments. In *Proceedings of the 10th annual ACM symposium on User interface software and technology (UIST '97)*. ACM, 31–39. <https://doi.org/10.1145/263407.263505>
- [42] William Ribarsky, Jay Bolter, Augusto Op den Bosch, and Ron van Teylingen. 1994. Visualization and Analysis Using Virtual Reality. *IEEE Comput. Graph. Appl.* 14, 1 (Jan. 1994), 10–12. <https://doi.org/10.1109/38.250911>
- [43] George Robertson, Mary Czerwinski, Patrick Baudisch, Brian Meyers, Daniel Robbins, Greg Smith, and Desney Tan. 2005. The Large-Display User Experience. *IEEE Comput. Graph. Appl.* 25, 4 (July 2005), 44–51. <https://doi.org/10.1109/MCG.2005.88>
- [44] Beatriz Sousa Santos, Paulo Dias, Angela Pimentel, Jan-Willem Baggerman, Carlos Ferreira, Samuel Silva, and Joaquim Madeira. 2009. Head-mounted display versus desktop for 3D navigation in virtual reality: a user study. *Multimedia Tools and Applications* 41, 1 (2009), 161. <https://doi.org/10.1007/s11042-008-0223-2>
- [45] Priyanka Srivastava, Anurag Rimzhim, Palash Vijay, Shruti Singh, and Sushil Chandra. 2019. Desktop VR is Better Than Nonambulatory HMD VR for Spatial Learning. *Frontiers in Robotics and AI* 6 (2019), 50. <https://doi.org/10.3389/frobt.2019.00050>
- [46] Daniel Thalmann, Jun Lee, and Nadia Magnenat Thalmann. 2016. An Evaluation of Spatial Presence, Social Presence, and Interactions with Various 3D Displays. In *Proceedings of the 29th International Conference on Computer Animation and Social Agents (CASA '16)*. ACM, 197–204. <https://doi.org/10.1145/2915926.2915954>
- [47] Jorge Wagner Filho, Marina Rey, Carla Freitas, and Luciana Nedel. 2018. Immersive Visualization of Abstract Information: An Evaluation on Dimensionally-Reduced Data Scatterplots. In *Proceedings of the Conference on Virtual Reality and 3D User Interfaces (VR '18)*. IEEE, 483–490. <https://doi.org/10.1109/VR.2018.8447558>
- [48] Jorge Wagner Filho, Wolfgang Stuerzlinger, and Luciana Nedel. 2020. Evaluating an Immersive Space-Time Cube Geovisualization for Intuitive Trajectory Data Exploration. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (2020), 514–524. <https://doi.org/10.1109/TVCG.2019.2934415>
- [49] James R. Wallace, Nancy Iskander, and Edward Lank. 2016. Creating Your Bubble: Personal Space On and Around Large Public Displays. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, 2087–2092. <https://doi.org/10.1145/2858036.2858118>
- [50] Colin Ware and Peter Mitchell. 2008. Visualizing Graphs in Three Dimensions. *ACM Trans. Appl. Percept.* 5, 1, Article 2 (Jan. 2008), 15 pages. <https://doi.org/10.1145/1279640.1279642>
- [51] Gordon Wetzstein, Robert Konrad, Nitish Padmanaban, and Hayato Ikoma. 2017. Build Your Own VR System: An Introduction to VR Displays and Cameras for Hobbyists and Educators. In *ACM SIGGRAPH 2017 Courses (SIGGRAPH '17)*. ACM, Article 14, 786 pages. <https://doi.org/10.1145/3084873.3084928>
- [52] Beth Yost, Yonca Haciahmetoglu, and Chris North. 2007. Beyond visual acuity: the perceptual scalability of information visualizations for large displays. In *Proceedings of the SIGCHI conference on Human factors in computing systems (CHI '07)*. ACM, 101–110. <https://doi.org/10.1145/1240624.1240639>