

Multiple Monitors or Single Canvas? Evaluating Window Management and Layout Strategies on Virtual Displays

Leonardo Pavanatto, Feiyu Lu, Chris North, and Doug A. Bowman




Fig. 1. Four configurations explored in this study: (a) Hybrid Multiple Monitors, where five virtual monitors are placed in a grid around a central physical monitor, (b) Virtual Multiple Monitors, where six virtual monitors are placed in a grid; (c) Hybrid Canvas, which provides space without boundaries around a physical monitor; and (d) Virtual Canvas, which provides space without any boundaries. Background is shown in blue and curvature is not shown for easier visualization.

Abstract—Virtual displays enabled through head-worn augmented reality have unique characteristics that can yield extensive amounts of screen space. Existing research has shown that increasing the space on a computer screen can enhance usability. Since virtual displays offer the unique ability to present content without rigid physical space constraints, they provide various new design possibilities. Therefore, we must understand the trade-offs of layout choices when structuring that space. We propose a single Canvas approach that eliminates boundaries from traditional multi-monitor approaches and instead places windows in one large, unified space. Our user study compared this approach against a multi-monitor setup, and we considered both purely virtual systems and hybrid systems that included a physical monitor. We looked into usability factors such as performance, accuracy, and overall window management. Results show that Canvas displays can cause users to optimize window layouts more than multiple monitors with snapping behavior, even though such optimizations may not lead to longer window management times. We did not find conclusive evidence of either setup providing a better user experience. Multi-Monitor displays offer quick window management with snapping and a structured layout through subdivisions. However, Canvas displays allow for more control in placement and size, lowering the amount of space used and, thus, head rotation. Multi-Monitor benefits were more prominent in the hybrid configuration, while the Canvas display was more beneficial in the purely virtual configuration.

Index Terms—Window Management, Virtual Displays, Multiple Monitors, Canvas, Layout, Augmented Reality

1 INTRODUCTION

Traditionally, a monitor's physical dimensions limit its width and height and, thus, how much screen space is available. We can reduce these restrictions by changing the scale of the display, switching virtual desktops, and overlapping or minimizing windows. However, these solutions still require the content to be contained within the physical space of the monitor, limiting how much content we can display *simultane-*

ously without cluttering the space and making window management difficult. Previous research on physical monitors has demonstrated that additional screen space can improve task performance [9, 10], making this a critical issue for knowledge workers.

We can achieve larger screen space by placing monitors side-by-side, forming wall-sized or multi-monitor setups. These configurations can show more information but also change how users interact with the system and the constraints they face [10, 43, 45]. For example, while a single monitor is a continuous display, multi-monitor setups have *boundaries* between the individual screens, dividing the space into sections. These are often called bezels, due to the physical frame around a screen. Even if the system integrates the space, these gaps are still visible to users and can change how they organize or partition their content [22, 45]. These divisions can limit users' freedom to place certain content, as they may avoid having windows that cross the boundaries between monitors or are too far apart [22].

Virtual displays offer the unique ability to present content without

• Leonardo Pavanatto, Feiyu Lu, Chris North, and Doug A. Bowman are with the Center for Human-Computer Interaction, Department of Computer Science, Virginia Tech, Blacksburg, VA, United States. E-mail: {lpavanat|feiyulu|north|dbowman}@vt.edu

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxx

rigid physical space constraints. We define them as interfaces that:

1. are displayed through an augmented or virtual reality (AR/VR) head-worn display (HWD);
2. extend or replace a physical monitor of a personal computer;
3. allow access to the computational capabilities of a personal computer; and
4. can display two-dimensional windows or monitors, on either a plane or any 2D manifold (e.g., a cylinder).

Past studies have demonstrated that it is feasible to use virtual displays for productivity work [6, 36]. They can be used by themselves or to extend physical monitors, which might provide benefits regarding resolution, field of view, and readability [36]. While the specifications of the HWD do dictate the system capabilities, the placement and dimensions of the 2D manifolds are virtually defined and, thus, have more flexible properties than physical monitors.

In particular, this flexibility allows us to create virtual displays without boundaries by rendering UI elements directly on a large virtual surface. The implication is that we can achieve much larger screens without subdivisions. However, such a *Canvas* approach is very different than what users are used to, and the extra freedom might cause more problems than it solves. Furthermore, if we use virtual displays to extend a physical display, it is more difficult to eliminate the boundaries of the physical display. We need to investigate the trade-offs that structuring the screen space in different ways causes on productivity and user experience.

Although AR/VR systems are often associated with novel 3D interactions and input devices, in this work, we assume that traditional 2D displays and windows, and traditional desktop input devices (i.e., mouse & keyboard) will be used with virtual displays. This approach allows users to transition seamlessly from today's personal computers to computing with virtual displays, and enables the use of familiar, highly refined interfaces and applications that are already available.

In this paper, we present the design of a single Canvas approach that eliminates boundaries by using virtual displays. We then describe a user study that investigated the effects of replacing multi-monitor setups with Canvas while considering either purely virtual systems or extending an existing physical monitor (which we call a *hybrid* setup). To achieve ecological validity, we developed a prototype that presents a full-featured Windows 10 operating system on virtual displays seen through a Microsoft HoloLens 2 AR HWD. We obtained qualitative and quantitative measures, including performance, accuracy, head movement, focus, comfort, confidence, readability, and user preference. We further looked into window management and how each condition impacted operations such as placement and resizing.

The contributions of this work include: (1) a quantifiable understanding of the trade-offs between multi-monitor and Canvas setups in virtual and hybrid settings, (2) identification of needs and opportunities for enhancing window management in such systems, (3) a formal evaluation method that could be used or modified to consistently evaluate similar systems now and in the future in terms of multi-tasking and window management, and (4) a validation of the use of a Canvas display as a virtual display.

2 RELATED WORK

2.1 Working with Large Physical Monitors

Previous research on large monitors has shown that when conducting cognitively difficult tasks, a display with larger screen space can provide a significant advantage in performance [9, 10]. Enhanced performance was attributed to physical navigation and maintaining an overview context [2]. The location and visual appearance of content in large displays also become valuable clues to keep users aware of the organization of the space as a type of external memory [1, 3]. These findings indicate the importance of having more screen real estate and accessing it using body motion instead of window or desktop switching.

Working with large monitors can also introduce issues such as losing track of the mouse cursor, accessing distant information, dealing with bezels, and the managing windows and tasks in the extra and distant available space [14, 40]. Cursor position awareness can be enhanced by showing temporal cursor trail [10] or an animated circle around the cursor [40]. Deciding where to place a new window, how to quickly move it, and how to organize a space with a large number of windows are some window management issues [40]. Differences in physical configurations can also impact the usage of such systems, where the designer needs to carefully consider placement strategies for the mouse, keyboard, and displays [14]. These factors may also impact large virtual displays, although we believe some of those can be addressed more easily without the physical constraints.

Multi-Monitor displays, sometimes called multi-display systems [19, 44], are a common approach for obtaining large displays as they are built by putting several small displays together in the same physical space. This approach has been shown to enable tasks that are complex and divided between multiple windows, while providing peripheral awareness [22] and increasing immersion [5]. These workspaces allow users to handle a larger number of windows [14], and enable more efficient multi-tasking [10].

2.2 Effects of Display Boundaries

While extending screen space, Multi-Monitor displays introduce physical boundaries between the displays that can influence user experience and interaction [44]. While the physical monitor literature often uses the term *bezels* to describe these physical boundaries, we will use the more general term *boundaries* throughout this paper, since this work focuses on virtual displays.

Multi-Monitor boundaries can present both benefits and issues [44]. Physical discontinuities introduced by boundaries or depth differences have been shown to have less effect on performance than other factors [43, 45]. Boundaries can be used as anchors that allow users to organize windows in a way that benefits their task [40, 45]. On the other hand, crossing boundaries introduces issues. A large window may span multiple monitors and create visual discontinuity at the boundaries [40], making it hard to read or understand patterns. This also makes window management more complicated, as users wish to avoid placing windows crossing them, and may place windows in less optimal locations. Techniques such as *Snapping* and *Bumping* have been proposed to avoid windows that cross boundaries [40]. Moreover, moving a cursor across boundaries may lead to interaction discrepancies, as the space behind the boundary doesn't exist in the computer's representation of the display [40]. There is no universal guideline as to whether we should eliminate boundaries or not, and both options are task- and user-dependent.

Boundary-less large monitors have also been investigated. Projector arrays combine multiple projectors to create a large seamless display panel [44]. They use back projection to align the screens side-by-side without any space between them. While eliminating boundaries, these displays present challenges of their own on presenting a uniform appearance [42], including color, contrast, brightness calibration, and alignment issues [44]. Previous studies have compared such displays against traditional monitors. Czerwinski et al. [10], compared a single 15" monitor against a large 42" boundary-less display called *DSharp*, which used projectors and a curved Plexiglas. Results showed the large display enhanced performance, and was most preferred with higher satisfaction. Ni et al. [34] also found that using large high-resolution projector arrays can improve performance on navigation, search, and comparison tasks in information-rich virtual environments when compared against both a single projected screen and an LCD panel.

Wallace et al. [45] simulated boundaries on a projected large display to measure the effects of various boundary widths, from which they concluded that the impact of bezel presence and width were small to the visual search task, and that the bezels helped users segment content across the space. Those findings were limited to a visual search task, and to a projected screen far from the user. We need to understand better how boundaries influence productivity work, especially on virtual

displays, as those have more flexibility in terms of form factor.

2.3 Window-Based Productivity Work in AR/VR

The idea of displaying windows through an HWD is not new. Feiner et al. [17] explored the issue in the 90s with a system that displayed floating windows through an optical see-through display. They explored different ways of registering windows, including head- and world-fixed, and how to create links between windows and physical objects. Raskar et al. [38] further explored what it would mean to expand an office space using AR and a collection of projectors and cameras. They suggested the creation of spatially immersive display surfaces over physical objects. However, this topic only became more popular in recent years as the concept of pervasive and everyday AR gained traction [4, 20]. Grubert et al. [21] present an interesting vision of future office spaces and the limitations and possibilities introduced by immersive technologies. Existing research has also explored using VR and AR technologies for everyday tasks, including productive work scenarios that include sitting, standing, and walking [4]. McGuffin et al. [31], argued that AR HWDs would be likely to be adopted in office environments over the next decade, and they discussed research opportunities around their use by knowledge workers.

Ens et al. [16] investigated the design and advantages of an immersive multitasking system that supports surrounding the user with windows. Combining virtual displays with laptop/tablet touchscreens was a feasible approach to aid mobile workers [7]. Combining virtual monitors with tablets used for touch input has also been shown to achieve performance and accuracy comparable to touch-controllers [25]. Working with a combination of physical and digital documents through AR proved to be feasible and accepted by users [27]. Existing work has also combined physical displays and virtual representations, such as displaying visualizations over tabletops [8], or large displays [29, 39]. Davari et al. [11] introduced Ganceable AR as an approach to access information at a glance and studied resolving real-world occlusion caused by virtual displays. Lu et al. [28] presented novel approaches to place and summon glanceable virtual content in a less obtrusive manner. Other studies further explore transforming traditional large display environments into immersive environments to facilitate sense-making [24, 26]. While these systems were innovative and showed great promise, we also need to investigate how to use virtual displays to support more traditional forms of working, such as knowledge work from a personal computer.

2.4 Replacing or Extending Physical Monitors with Virtual Displays

Conducting productivity work in HWDs can provide more display flexibility, reducing costs [36] and addressing many challenges such as lack of space, illumination issues, and privacy concerns [21, 35]. In a recent study, VR reduced the distraction of users working in open office environments, induced flow, and was preferred by users [41]. However, traditional anchored input devices such as keyboard, mouse, and touchpad can be hard to use with peripheral portions of the display space that are distant from the input devices, forcing additional head rotation that could possibly result in neck pain [21]. In setups with many displays organized horizontally in a cylinder around the user, it was shown that amplified head rotation (i.e., performing a virtual movement opposite to the physical head movement) could reduce the amount of head rotation required to access peripheral displays [30].

Previous work also reveals some of the issues of virtual monitors. Context switching between physical and virtual environments and focal distance switching between displays reduce task performance and increase visual fatigue [18]. Having multiple depth layers, such as combining an HWD with a smartwatch, can induce more errors when interacting [12]. Social acceptance and monitor placement are also shown to play important roles in the use of virtual monitors in public places like airplanes [33], and in the layout distribution of content across multiple shared-transit modalities [32].

On one hand, these displays provide us the opportunity to investigate a truly seamless display, with uniformity across the space. Since

those displays are highly configurable, we need more empirical information on how user experience is impacted by design choices. On the other hand, since virtual displays have potential for much larger screen spaces, we need to specifically understand the trade-offs of such choices, which will enable us to design systems that better manage that space without generating window management overhead. Therefore, this study focuses on a direct comparison between conditions with and without boundaries, and with and without transitions between physical and virtual displays.

3 CANVAS DISPLAY CONCEPT

In this section, we present one possible design of a canvas display. Conceptually, we define a canvas display as an implementation that reduces or eliminates the number of virtual display's boundaries or divisions. Although it's possible to produce a (nearly perfect) canvas display physically, this is usually not practical, but virtual displays can do this because they do not rely on physical display hardware and do not need to represent every pixel all at once. In this paper we investigate a Canvas Display that eliminates boundaries in the X and Y axes, focusing on sub-divisions between multi-monitor setups.

Our version of *canvas display* is a single, unified, virtual display that is curved over a cylinder around the user and can have any size. The cylindrical organization implies that it possesses normal curvatures of zero (vertical) and $1/r$ (horizontal), where r is the radius of the cylinder—or, in our case, how far away the display is from the user's initial position. Any size implies that it can have a maximum width that surrounds the user's entire field of regard, being as large as a full cylinder ($2 * \pi * r$), and a technically infinite maximum height—although extremes will not be distinguishable for the user due to distortions. Thus we argue that the viewing angle between the content normal and the user will define how large the height can be. We expect it to be smaller than the field of regard.

More interestingly, single and unified imply that it represents one display and does not contain sub-divisions. This is in contrast with multi-monitor setups, the go-to solution for people requiring large screen real estate. A unified display can present multiple benefits, such as a non-fragmented workspace (everything is together) and better space utilization (people will often avoid placing content crossing boundaries). It also avoids displays with heterogeneous characteristics (colors, resolutions, aspect ratios) and the need to set up and manage separate devices. Examples of Canvas display can be seen in Fig. 1 (c,d).

4 USER STUDY

We conducted a user study to investigate the user experience of various configurations of displays with varying levels of screen sub-divisions and virtuality. We wanted to understand the best ways to make use of virtual displays for productivity work, which brings the question of using a canvas display or a sub-divided one. Namely, we were interested in understanding how does the presence or absence of boundaries and snapping behavior would affect window management, layout strategies, and user experience. Existing implementations of virtual displays focus on replicating a multi-monitor form factor, consisting of multiple surfaces with fixed width and height [6, 15, 30, 32, 33]. We further examine how virtuality impacts those results by also testing both concepts coupled with a central physical monitor. We think that there are some important benefits of using a central physical monitor (resolution, brightness, opacity). Therefore, in the case when users choose to extend a physical monitor with virtual displays, should the extension be similar (multi-monitor setup), or does the flexibility of canvas help?

This study aims to understand how virtual display characteristics impact window organization and access. We believe characteristics such as field of view and head movement will affect the most suitable organization approach. For instance, physical multi-monitor sub-divisions may be beneficial to facilitate organization and window placement. However, we must consider elements such as reduced field of view (FOV), resolution, and opacity in virtual displays, as those could provide difficulty in placing and accessing windows efficiently.

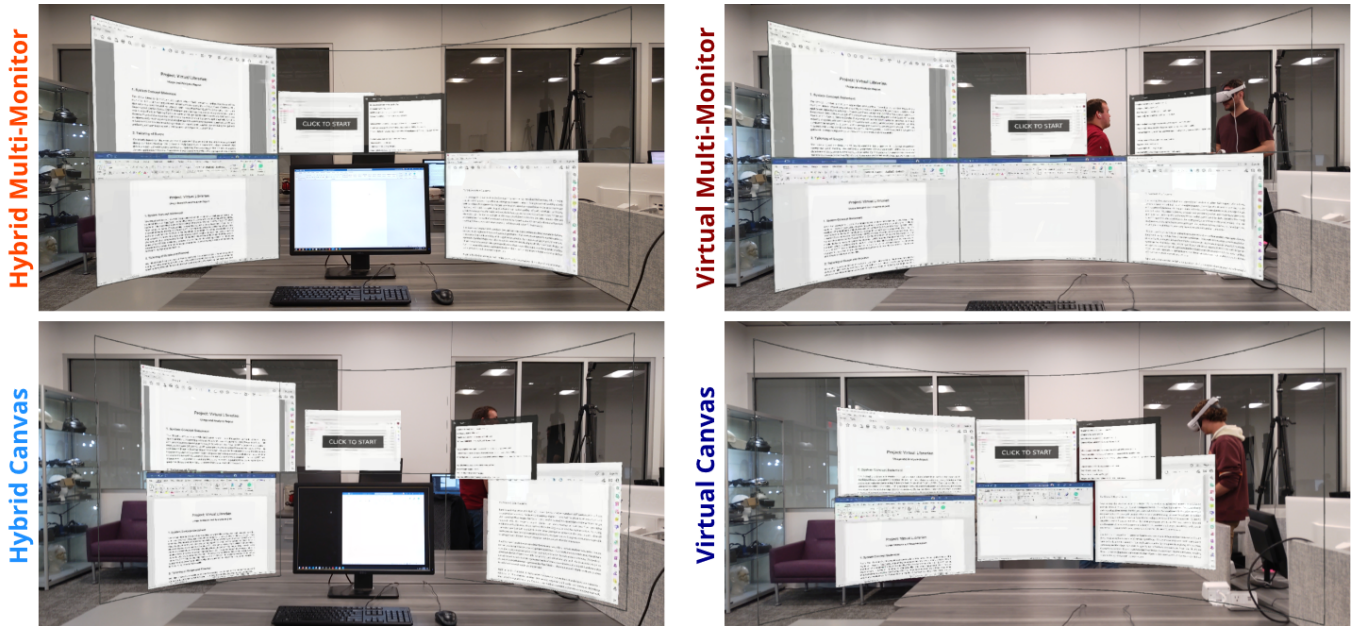


Fig. 2. Conditions in this study viewed from a HoloLens 2: Hybrid Multi-Monitor, Virtual Multi-Monitor; Hybrid Canvas; and Virtual Canvas.

4.1 Conditions

Our study included four virtual display configurations, arising from two independent variables with two levels each. First, we varied the *style* of the display setup (multi-monitor or Canvas). The **multi-monitor** setup, seen in Fig. 2 (top), consisted of six monitors placed in a three (horizontal) by two (vertical) grid. The monitors had a one-centimeter boundary between them. In the multi-monitor conditions, in addition to regular window placement, we allowed participants to use a snapping feature called Fancy Zones provided by Microsoft Power Tools¹. While the user was moving a window, they could hold down the Shift key to see the monitor divided horizontally into two equal portions. Then, they could move the window over the space and release the key for it to quickly attach to that space. The **Canvas** setup, seen in Fig. 2 (bottom), used the Canvas display concept, providing a single unified space without any boundaries. Participants could not use the snapping feature in these conditions, as we wanted them to have more freedom during placement, but they could move and resize windows freely.

Second, we varied the *virtuality* of the display setup (purely virtual or hybrid). The **virtual** conditions, seen in Fig. 2 (right), used only virtual displays, while the **hybrid** conditions, seen in Fig. 2 (left), combined virtual displays (either Canvas or Multi-Monitor) with a single physical monitor in the bottom center of the display area. We combined those by placing the virtual displays at the same visual angle as the physical monitor from the user’s perspective. The combination between physical and virtual was not seamless due to depth differences that were perceptible to users.

In all conditions, the displays were world-fixed, being placed one meter away from the user during calibration, and the total screen space available was the same (equivalent to 6 monitors with 1920x1200 pixels) – even though the canvas display could theoretically not have outer boundaries, we decided to maintain the same size as the Multi-Monitor to allow for a fair comparison, and focus on the sub-divisions aspect.

The surfaces were curved along a cylinder of one meter radius, with a width of 2.267m and a height of 0.925m, leading to an angular width of 129.88 degrees, or 0.46 degrees per pixel. Also, across all conditions, the background of the virtual displays was transparent. Given the large size of the screens, we didn’t want them to present a barrier to visual contact with the real world—one of the points of using AR is to allow

users to see and interact with the physical world. We further asked participants not to minimize windows during any of the conditions, as we aimed to understand how users would organize the workspace, not how they would vary which windows were displayed.

4.2 Implementation and Apparatus

A Microsoft HoloLens 2 was used as the HWD in all conditions. It has a field of view of 43 degrees horizontally and 29 degrees vertically. It has a resolution of 2048x1080, with a 3:2 aspect ratio. Head tracking is done with four visible light cameras, while eye tracking uses two IR cameras. Participants wore the HoloLens 2 during all experimental conditions. The device was used both for rendering virtual displays when necessary and for obtaining metrics from the usage of the displays, such as head rotation and eye-tracking.

The physical monitor used in the study was a DELL U2412Mc, with 24 inches along the diagonal, native resolution of 1920x1200, 16:10 widescreen aspect ratio, 16.78 million colors, 2M:1 high dynamic contrast ratio, positioned at 70cm from the user. Both brightness and contrast were set to 75%, while color profile was set to standard. Even though this model of monitor has an adjustable height and angle, we locked those in place using tape, to ensure they would not be moved or rotated between participants. In all conditions, we used this monitor to calibrate the position of the virtual displays, by asking the user to pan a virtual “ghost” of the screen over the physical one using their index finger. Calibration was completed once the models visually matched.

Virtual displays were rendered at the same physical size of 24” and same resolution of 1920x1200. However, we changed the scale of the entire system to 150% (including physical monitor), a magnification that has been shown in the past to match a readable resolution on HoloLens [36]. The Canvas conditions simply merged the six monitors together, keeping the same resolution and screen space - thus, resulting in the elimination of a gap of about two centimeters that exist between the screen in the Multi-Monitor conditions.

We used a full version of the Windows 10 operating system as the interaction environment. Our system contains two main components: a back-end WPF .NET Framework 4.7 application that was created on Visual Studio 2022, using a proprietary solution for virtual display creation, and the Windows Graphics Capture API²; and a front-end application that was designed on the Unity Engine version 2019.2.21f1

¹<https://learn.microsoft.com/en-us/windows/powertoys/>

²<https://learn.microsoft.com/en-us/uwp/api/windows.graphics.capture>

³, using the Mixed Reality Toolkit (MRTK) version 2.3.0⁴.

There are multiple means in which we can create virtual displays within Windows 10. That includes plugging in external monitors and taking them out of the user's viewpoint [36], using HDMI dummy dongles that fake a monitor connection, and using custom drivers or services that simulate a monitor through software. We used a proprietary implementation of a service that creates such monitors. While our solution is not available publicly, the monitors used in this study can be replicated through any of the options described above. Once monitors exist, the application uses unmanaged code from Win32 and Windows Graphics Capture APIs to enumerate and share the render buffer of each monitor through Direct3D⁵ and SharpDX⁶.

The Unity scene contained surfaces that were warped in a circumference around the camera. During calibration they were positioned at 1m from the camera and at the same height of the physical monitor (in Virtual conditions we removed the monitor after calibration). These surfaces were placed at world-coordinates, allowing participants to perform small head movements while the monitors would stay registered in place. Each of these surfaces rendered an external texture of the monitor captured obtained from the back-end application. It again uses Direct3D and SharpDX to access the shared textures from the graphics card's output buffer, achieving real-time display in Unity. We set the wallpaper of Windows to a solid black color, ensuring that the background would not be rendered in HoloLens.

MRTK managed HoloLens integration, including spatial, hand, and eye tracking. We used the Holographic Remoting Player⁷ for mirroring the Windows 10 monitors on the HoloLens. This application streams content from a computer to a HoloLens in real-time, through a tethered connection. The HoloLens sends the information obtained from its sensors (such as head and eye tracking) back to the PC, which uses them to make all necessary computations. The result is sent back to the HoloLens and displayed to the user. For input, we used a standard DELL keyboard, model KB2012-B, and Logitech mouse, model M-U0026, in all conditions; the mouse could be moved naturally across all virtual displays, using Windows built-in multi-monitor alignment. We did not use any boundary compensation (i.e., the cursor jumped directly from one display to another when crossing a boundary).

The experiment was run on a PC with an Intel i7-8700K CPU, 16GB of 3200MHz DDR4 DRAM, a Samsung SSD, a GeForce GTX 1070 8GB GPU. The connection with the HoloLens was managed through a 16FT USB A-C cable with up to 5Gbps transfer rate.

4.3 Experimental Design

Our mixed-design study had two independent variables, display style (within subjects) and virtuality (between subjects). The presentation order was counterbalanced within each group. We recruited 40 participants from the general population that fit the following inclusion criteria: (1) were at least 18 years old, (2) had normal vision (corrected or uncorrected), (3) were proficient with the English language, and (4) used a computer with Windows 10 for daily work.

Our objective dependent variables included how many windows were moved and resized, how large those movements were, and the spread of window placement on the screen. We also gathered data about how users perceived and rated multiple user experience metrics. We captured screenshots every ten seconds, to understand user strategies visually. We further obtained performance and accuracy measured on the task, but we do not analyze these results in this paper.

4.4 Hypotheses

Our hypotheses regarding style and virtuality of virtual displays were as follows:

³<https://unity.com/>

⁴<https://github.com/microsoft/MixedRealityToolkit-Unity>

⁵<https://learn.microsoft.com/en-us/windows/win32/direct3d>

⁶<http://sharpdx.org/>

⁷<https://learn.microsoft.com/en-us/windows/mixed-reality/develop/native/holographic-remoting-player>

H1. Canvas will lead to a more optimized utilization of the space than Multi-Monitor.

As participants have less affordances to organize the space, we believe that they will organize the layout in a way that better fits their content, rather than conforming their content to preexisting spaces. We believe that the final result will utilize the space better, with less space being wasted, and more content density. To test this hypothesis, we will analyze how optimized the layout was at the end of the task (Sect. 5.1).

H2. Canvas will lead to less head rotation and cursor movement than Multi-Monitor.

We believe that the affordances provided by boundaries and the snapping feature in the Multi-Monitor conditions will compel participants to use areas of the display that they would not in less constrained scenarios. Conversely, we expect Canvas participants to naturally avoid using the extremes of the display, and thus reduce both the frequency and amplitude of head rotations and cursor movements. We will analyze how much head movement (Sect. 5.4) and how much mouse movement there was (Sect. 5.5).

H3. Canvas will lead to participants performing more window management operations than Multi-Monitor.

With Canvas, space is less structured and snapping is unavailable, which we expect will make participants spend more effort in placing and resizing windows. On the other hand, a structured multi-monitor setup with snapping behavior should provide an easy way for participants to quickly find a place for each window. We will analyze how much window management was completed during the task (Sect. 5.2).

H4. Canvas will support more diverse window placement strategies than Multi-Monitor.

While boundaries and snapping can be seen as affordances that simplify window placement, we suggest they are also rigid structures that enforce certain patterns of window management. We believe that by removing boundaries, Canvas will support a greater of window placement strategies. We will examine what strategies are used (Sect. 5.6), and how much window management varied within each group (Sect. 5.2).

H5. Canvas will have a worse overall user experience than Multi-Monitor.

We believe that the unconstrained nature of Canvas will lead to participants having to do more interaction and think more about the windows and layout, and that this will lead to lower perceived performance, perceived accuracy, ease of use, comfort, confidence, satisfaction, preference, and mental workload. We will evaluate subjective ratings provided by participants (Sect. 5.8).

H6. Hybrid conditions will lead to participants taking more time organizing windows than Virtual conditions.

We believe that the asymmetric nature of the hybrid conditions will lead to participants rearranging windows more frequently, to take advantage of the high-resolution properties of the physical monitor. The switch between virtual and physical should also affect this, as participants will have to switch focus between physical and virtual contexts. We will compare the total time moving windows between conditions (Sect. 5.3).

H7. Multi-monitor will be most preferred in Hybrid conditions, while Canvas will be most preferred in Virtual conditions.

We think that most participants in the Hybrid group will prefer Multi-Monitor because of the symmetry with the physical monitor, as they will have a single mental model both cognitively and mechanically. On the other hand, we think that most participants in the Virtual group will prefer Canvas, as monitor sub-division may be perceived as artificial and rather limiting of their freedom. We will compare preferences provided by participants (Sect. 5.7).

4.5 Experimental Task

The task performed by participants was a derivation of the one proposed by Jeuris et al. [23]. It consisted of three sub-tasks, which can be seen in Fig. 3. We focused on simulating multitasking, during which the user must simultaneously manage (move and resize) windows while alternating between heterogeneous tasks. Participants always had 6 windows opened at all times and had to decide how to organize them depending on the active tasks. Tasks were designed in a way that they

could not be completed in the available time. We used these tasks merely as proxies for multi-window complex tasks, and did not look into performance and accuracy measures. The task is a combination of sequential and parallel tasks. In parallel, we have two workflows:

- **Primary task: Productivity.** Diverse sub-tasks will be executed sequentially *twice*. Since each was executed twice in a fixed cycle, users had to resume a task from where they stopped before:
 1. **Copying task:** The user must copy the text of a PDF document in this template needs to be replaced by specific student information from an image. This subtask requires 3 windows (Image, Word, PDF).
 2. **Comparing task:** The user must read text in a document and compare against a copy. Whenever a word is missing, the user must double click the word in the original, and press highlight. This subtask requires 2 windows (Word, PDF).
- **Secondary task: Monitoring.** The user needs to keep track of an email inbox. If a new email comes in, the user must stop what they are doing and click over it immediately.

We implemented the monitoring task with a Unity desktop app that shows screenshots of an email inbox. Initially it shows an empty inbox, then after a few seconds it displays a new message, and then it waits for the user to click over it. The time for a message to arrive was randomly chosen between 30, 40, 60, 80, or 90. Since each task took 120 seconds, those numbers represent varying levels of beginning, middle, and end of task. Once the new email arrived, the participant had up to 30s to click over it, otherwise it disappeared.

4.6 Procedure

The university’s Institutional Review Board approved the study. The study took place face-to-face at our laboratory in a single session of 90 minutes. We recruited participants through mailing lists and asked them to complete a screening questionnaire for our inclusion criteria discussed in the study design section. They scheduled a session time and received a digital copy of the consent form. Since this study was conducted during the COVID pandemic, before and after each session we thoroughly disinfected the table, keyboard, mouse, pen, and the AR display. During the study, the participant and investigator wore face masks and stayed at least two meters apart.

Upon arrival, we greeted the participant at our laboratory. They signed the consent form and answered a background questionnaire on a tablet. Next, they received general instructions and completed the standard HoloLens 2 eye calibration procedure. The participant practiced each of the three sub-tasks on a large Microsoft Surface Hub display for up to two minutes. This allowed them to understand the task without compromising our measures of each condition.

We calibrated the virtual coordinate system with the physical monitor position to match eye-tracking data and render virtual displays in the same position across participants. We provided the participants free time to explore the condition before the main task began. The investigator would ask them to move a Windows Explorer window across the display, ensuring they could visualize the entire space and that there were no issues with connectivity. In multi-monitor conditions, we also taught participants how to snap windows to the spaces, which they could use at their discretion. We asked participants to avoid placing windows across boundaries or between virtual and physical screens, as those situations can represent a loss in performance and real-world users usually avoid them.

The participant flipped the headset up between each condition and answered a condition questionnaire on the physical monitor. The questionnaire asked participants to provide ratings for statements based on their perception and asked them to write down the pros and cons of the condition they just tried. Once all conditions were completed, the participant answered a final questionnaire, where they could select their

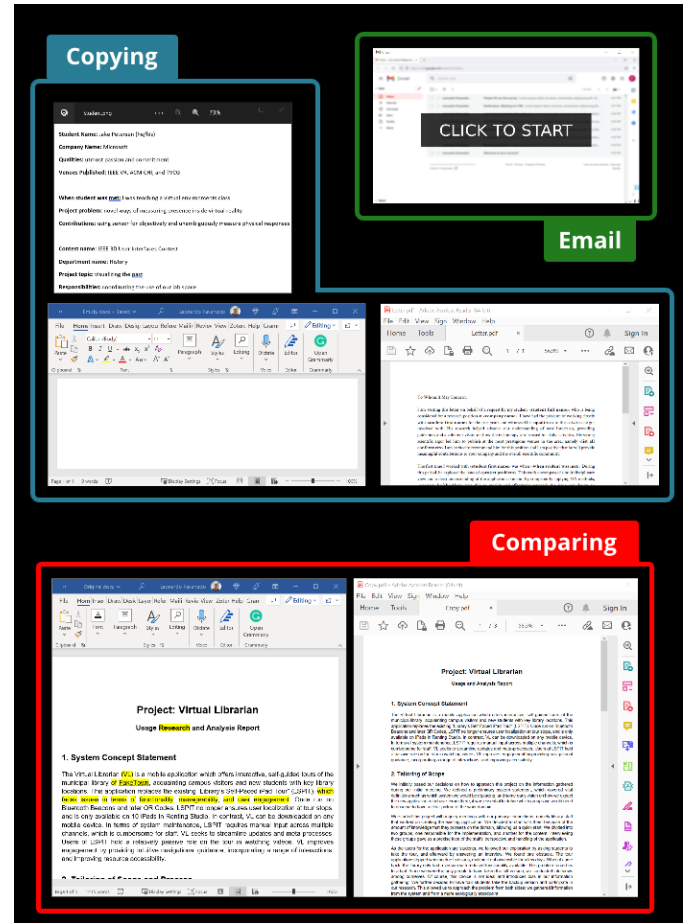


Fig. 3. Three tasks that participants completed. Copying and Comparing executed twice sequentially; Email executed in parallel.

preferred condition and express their feelings about all the conditions they used. We completed the study with a ten-minute semi-structured interview conducted at a distance of 2.5 meters between the participant and the investigator.

4.7 Participants

Forty participants (aged 19 to 32, 14 female) from the campus population took part in the experiment in individual sessions of around 90 minutes. Eleven participants were graduate students, and 29 were undergraduate students. All participants used a computer daily for work, with 35 people reporting using a computer for at least four hours on a typical weekday, with 10 using it for more than 8 hours. Almost all participants (38) reported at least intermediate experience with the Windows operating system. The majority of participants (24) had little to no experience with AR.

5 RESULTS

We collected our results from multiple sources. A Qualtrics survey recorded the questionnaires, including qualitative and subjective quantitative measures. From Unity, we obtained a frame-by-frame log of all events during the sessions, such as time, frame time, head orientation angle, and cursor 2D coordinates. In our back-end server, we obtained a log of all window management. Finally, we recorded audio files with the responses given by participants during the semi-structured interviews. These were transcribed by Office Online, with manual verification and fixes completed by the authors.

We exported Qualtrics and Unity outputs to “.csv” formats, which were ideal for further processing through Python scripts. We performed the statistical analysis using the JMP Pro 16 software. We used an α




Fig. 4. Layout characteristics at the end of the task. (a) Total pixels in all windows, (b) total pixels overlapping among windows, and (c) percentage of the screen filled by windows. Error bars represent 95% confidence intervals.

level of 0.05 in all significance tests. In the results figures, significantly different pairs are marked with * when $p \leq .05$, ** when $p \leq .01$, and *** when $p \leq .001$.

We verified normality through Shapiro-Wilk tests and normal quantile plot inspections for all the cases before deciding whether to apply two-way mixed-design factorial analysis of variance (ANOVA), non-parametric tests, or apply a transformation before using ANOVA. We further performed pairwise comparisons using Tukey HSD when appropriate. Our two factors were the display style (Canvas or Multi-Monitor) within subjects and the virtuality (Virtual or Hybrid) between subjects.

5.1 Layout at the End of Task

We recreated a snapshot of the windows' positioning at the end of the task from our log of window manipulations, and we examined the overall characteristics of the layout.

Regarding the **total pixels for all windows**, we found a main effect of display style ($F_{1,1} = 18.58, p < 0.001$). Canvas ($M = 7,063,533.0, SD = 1,685,574.1$) led to smaller windows than Multi-Monitor ($M = 8,651,483.8, SD = 2,346,664.2$), and this difference was especially present in hybrid systems. While we didn't find main effect for virtuality ($p = 0.69$), we found an interaction effect ($F_{1,1} = 5.35, p = 0.0264$). Post-hoc tests ($p < 0.001$) show that Hybrid Multi-Monitor had the largest window sizes ($M = 9,202,624.2, SD = 2,657,789.6$), while Hybrid Canvas ($M = 6,728,641.5, SD = 1,684,241.7$) the smallest. These results are shown in Fig. 4 (a).

Regarding the **total pixels overlapping among windows**, we found a main effect of display style ($F_{1,1} = 14.86, p < 0.001$). Canvas ($M = 1,596,434.7, SD = 1,303,605.2$) leads to more overlapping of windows than Multi-Monitor ($M = 716,847.34, SD = 1,137,224.3$). We didn't find a main effect of virtuality ($p = 0.19$) or a significant interaction effect ($p = 0.225$). These results are shown in Fig. 4 (b).

Regarding the **percentage of screen space utilized**, we found a main effect of display style ($F_{1,1} = 45.68, p < 0.001$). Canvas ($M = 0.395, SD = 0.09$) leads to less screen space utilization than Multi-Monitor ($M = 0.574, SD = 0.160$). We didn't find a main effect of virtuality ($p = 0.16$) or a significant interaction effect ($p = 0.12$). These differences can be observed in Fig. 4 (c).

While **total pixels for all windows** and **percentage of screen space utilized** may look like similar metrics, the first refers to the total size of windows, independent of whether there was overlapping happening. The later refers specifically to how many of the pixels on the display were filled with one or more windows.

We also measured the spread of windows both horizontally and vertically, which we defined as the difference between the maximum and the minimum pixel coordinates that contain a window. Regarding **spread in the horizontal axis**, we found main effects for both virtuality ($F_{1,1} = 5.71, p = 0.022$) and display style ($F_{1,1} = 10.30, p = 0.003$). Hybrid ($M = 5130.00, SD = 1090.46$) led to a larger hori-

zontal spread than Virtual ($M = 4396.85, SD = 1245.91$), and Multi-Monitor ($M = 5074.46, SD = 1308.95$) led to a larger spread than Canvas ($M = 4443.59, SD = 1049.99$). We didn't find a significant interaction effect ($p = 0.74$).

Regarding **spread in the vertical axis**, we found main effects for both virtuality ($F_{1,1} = 6.41, p = 0.0157$) and display style ($F_{1,1} = 32.82, p < 0.001$). Hybrid ($M = 2424.42, SD = 389.74$) led to a larger spread than Virtual ($M = 2181.22, SD = 513.41$), and Multi-Monitor ($M = 2534.85, SD = 410.05$) led to a larger spread than Canvas ($M = 2064.56, SD = 408.81$). We didn't find a significant interaction effect ($p = 0.74$).

5.2 Window Management During the Task

While the layout at the end of the task can show us some organization and management strategy tendencies, we need more information to understand how the task evolved in each configuration. We obtained a log of all window-related operations for each session. From those, we counted the operations a user performed (moving or resizing), accumulated the amplitude of those operations (number of pixels moved or resized), and created a distribution of the amplitude of operations (such as moving and resizing across horizontal and vertical axes).

A move operation is defined as changing top left coordinates of the window without changes in dimensions. A resizing operation considers changes in the dimensions, independent of whether those also moved the window (which can happen during a snap operation).

Regarding the **total number of move operations**, we found a main effect for style ($F_{1,1} = 29.73, p < 0.001$), with Canvas ($M = 28.18, SD = 1.76$) and Multi-Monitor ($M = 18.33, SD = 1.76$). Canvas led to a larger number of moves than Multi-Monitor (1.53x higher). Regarding **total number of move operation that ended with a snap**, we found a main effect on virtuality ($F_{1,1} = 5.09, p = 0.03$), with Hybrid ($M = 2.87, SD = 1.76$) and Virtual ($M = 4.97, SD = 1.76$). Virtual led to a larger number of moving operations that ended on a snap than Hybrid (1.73x larger). Overall, in Multi-Monitor, 42% of the window movements ended in snaps.

Regarding the **total number of resize operations**, we found a marginally significant main effect for style ($F_{1,1} = 3.93, p = 0.054$), with Canvas ($M = 11.87, SD = 1.76$) and Multi-Monitor ($M = 14.15, SD = 1.76$). Multi-Monitor tends to have a higher number of resizes than Canvas (1.2x higher). Regarding **total number of resize operation that ended with a snap**, we found a main effect on virtuality ($F_{1,1} = 4.24, p = 0.047$), with Hybrid ($M = 5.55, SD = 1.76$) and Virtual ($M = 3.85, SD = 1.76$). Hybrid led to a larger number of resize operations that ended on a snap than virtual (1.44x larger). Overall, in Multi-Monitor, 65.5% of window resizes ended on snaps.

To explore the magnitude of window management operations, we created distributions of the amplitudes of all move and resize operations. The distributions for move operations can be seen in Fig. 5. Unsurprisingly, for all conditions, small movements were always predominant. It

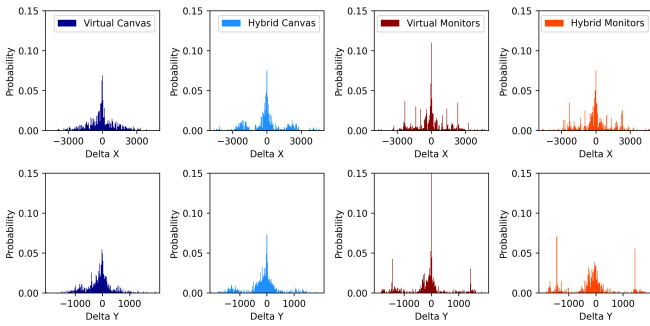


Fig. 5. Distribution of X and Y movements in the four conditions.

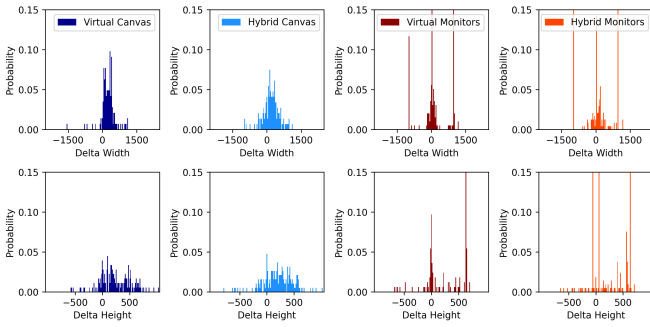


Fig. 6. Distribution of Width and Height resizes in the four conditions.

is interesting, however, how the shape of the distribution changes across conditions. In Virtual Canvas, we have the most similar to a normal distribution, with a smoother shape. In Hybrid Canvas, on the other hand, we observe two extra peaks that are symmetric to each other. These seem to represent the operations of moving a window from the central physical monitor to the virtual Canvas or vice-versa, since we instructed participants not to place windows across the boundary. In both Virtual Monitors and Hybrid Monitors, we see somewhat similar shapes. For movements in X, we observe that there are not two, but multiple extra peaks. This happens because of the snapping behavior, which provides two snapping locations across each monitor—thus, six side-by-side spaces. For Y movements, since there were no vertical sub-divisions within each monitor, we see two extra peaks representing snapping to the monitor above or below the original location of the window.

We conducted the same analysis for resize operations. The distributions for resizes in the Width (X) and Height (Y) axes, for our four conditions, can be seen in Fig. 6. Again, Virtual Canvas produced a result similar to a normal distribution, although the actual changes in resizing are considerably smaller than movement. For resizes, the Hybrid Canvas distribution has a similar shape to Virtual Canvas, indicating that although separate screen spaces will disrupt movement, participants were still able to exert freedom in resize. We can also see the effects of the snapping on both Virtual Monitors and Hybrid Monitors.

5.3 Time in Organization

We analyzed the total amount of time that participants were conducting window management (move or resize) operations. We calculated the time between a grab and a release for each operation. While it could be argued that we are not considering sense-making time where users are trying to plan their next steps, this is an objective variable that will show us if there is a difference in the amount of time doing actual window management. We again conducted a two-way ANOVA, and found no main effect of either virtuality ($p = 0.87$) or display style ($p = 0.74$), and no significant interaction ($p = 0.13$). Although not

significant, we found that Hybrid Canvas ($M = 77.47, SD = 25.06$) had the most window management time, followed by Virtual Monitors ($M = 75.15, SD = 21.33$), then Virtual Canvas ($M = 70.1, SD = 20.11$) and Hybrid Monitors ($M = 69.68, SD = 22.79$).

5.4 Head Movement

We started this analysis by removing outliers, through the use of a moving median window filter of size 60 across the frame data from each user. We then analyzed the total amount of head movement by summing the absolute differences between each resulting frame. We again conducted a two-way ANOVA on the results. There was a significant main effect of display style on Pitch ($F_{1,1} = 18.60, p < 0.001$) and Roll ($F_{1,1} = 12.35, p = 0.001$), but no effect on Yaw. Canvas ($M = 2582.85, SD = 1139.74$) resulted in significantly less pitch rotation than Multi-Monitor ($M = 3226.86, SD = 1242.58$). Canvas ($M = 669.13, SD = 340.57$) also resulted in significantly less roll than Multi-Monitor ($M = 880.49, SD = 537.91$).

5.5 Cursor Movement

In this analysis we also removed outliers, through the use of a moving median window filter of size 60 across the frame data from each user. We then analyzed the total amount of cursor movement by summing the absolute differences between each resulting frame. We again conducted a two-way ANOVA on the results. There were no significant effects for horizontal movement. We found a significant main effect of display style on vertical movement ($F_{1,1} = 4.54, p = 0.039$), with Canvas ($M = 67, 199.05, SD = 17659.94$) resulting in significantly less vertical cursor movement than Multi-Monitor ($M = 74, 785.44, SD = 18486.03$).

5.6 Strategies

Given our mixed design, twenty participants completed each condition. We analyzed participant strategy by observing the screenshots of each trial, and looking into elements such as overlapping, spread, active task, and then identified the possible strategies used. Some participants used multiple or hybrid strategies, but we report the most predominant for each participant. For Virtual Canvas, some strategies used by participants included: placing windows in the center with some visibility of occluded windows for quickly bringing them to the front (9 participants), placing active task windows in the center of the Canvas and stacking other windows on the side (5 participants), placing all the windows in an organized mosaic without overlapping (4 participants), and placing them on areas of the screen grouped by task (2 participants). Those examples can be seen in Fig. 7.

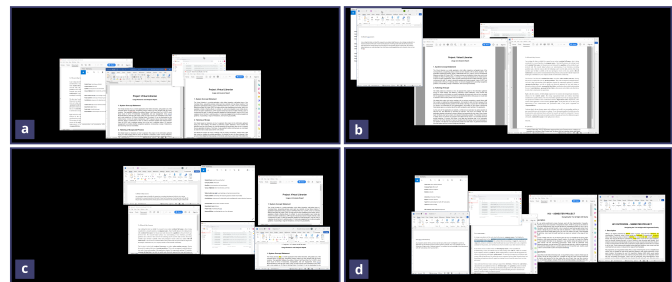


Fig. 7. Strategies used in Virtual Canvas: (a) partial occlusion of unused windows, and (b) stacking of unused windows on the side, (c) mosaic, (d) divide the screen by task.

In the Virtual Multi-Monitor conditions, strategies were less varied: participants snapped windows to an entire monitor or half of the monitor space (17 participants), or they ignored the snapping and place windows such that monitors were divided vertically (3 participants). Also, 2 participants opted for not using the top row of monitors at all, as they judged they would need to rotate their heads too much. Those examples can be seen in Fig. 8.




Fig. 8. Strategies used in Virtual Monitors: (a) snapping windows to either a whole or half monitor, (b) ignoring the snapping suggestions and trying to divide monitors vertically.

In the Hybrid Canvas condition, participants tended to try to place as many windows inside the physical monitor as they could, and then place other windows on the Canvas, as close to the monitor as possible (10 participants), or they treated the physical and virtual space in a more equivalent way (10 participants). In Virtual Canvas, almost all participants used the physical and virtual space equivalently (19 participants). Some examples can be seen in Fig. 9.

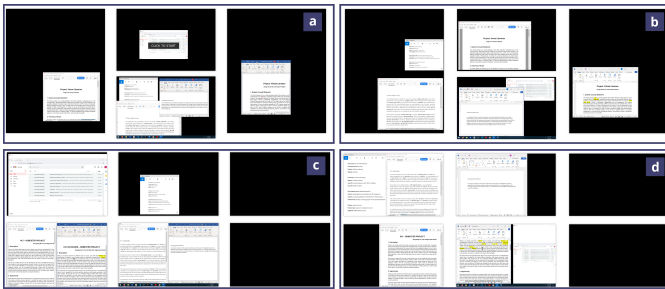


Fig. 9. Strategies in Hybrid conditions: (a, b) Canvas: placing many windows on the physical monitor and placing others on the Canvas as close as possible to the monitor, (c, d) Multi-Monitor: placing windows similarly between physical and virtual parts.

5.7 User Preference

We asked participants at the end of their session which condition they preferred. In the Hybrid group, eleven out of twenty people chose the Multi-Monitor setup as their favorite (55%). Reasons given by those participants included similarity with the systems they usually worked on, frustration about not being able to split space in Canvas, and the snapping feature that made it easier to move windows between screens without requiring precision. One participant mentioned using the main monitor for work, and leaving the virtual display only for looking at information not currently in use. Reasons given by people who chose Canvas include more freedom to interact with, tighter control on window size (not being too large or too small), easier to group windows together and more closely, not being confined to preset monitors, and complains about having to press a button to trigger snapping in Multi-Monitor.

In the Virtual group, fourteen out of twenty people chose the Canvas setup as their favorite (70%). Reasons given by those participants included being able to resize windows to whatever they needed, being able to place them tighter together with some overlap, more freedom in terms of window management, being able to center windows, not getting distracted by the boundaries, having to move their heads less, and not being constrained to the size of a single screen in the Multi-Monitor setup. One participant mentioned that they liked the snapping feature of virtual monitors, but it wasn't versatile enough for how they wanted to organize their screen. Another one mentioned that their preference was due to the limited FOV of the HoloLens, making it easier to access things closer together. Considering both groups, Canvas was slightly preferred over Multi-Monitor, with twenty-three out of forty of participants choosing it (57%).

5.8 Subjective Ratings

We asked participants to subjectively rate many different usability factors, helping us characterize each configuration. Since our data has two independent variables, we opted to use an Aligned Rank Transform (ART) [46] to fix normality before using ANOVA. Unlike Friedman's non-parametric tests, this approach allows us to obtain both interaction and main effects for our data, which is essential for this study. We further used ART-C [13] coupled with One-Way ANOVA and Each Pair Student's t to obtain pairwise comparisons.

5.8.1 Speed

In *"I felt that it took me too long to finish the task"*, there was no significant main effect for either virtuality ($p = 0.73$) or display style ($p = 0.16$), but there was a significant interaction effect ($F_{1,1} = 5.62, p = 0.023$). Pairwise analysis showed marginal significance ($p = 0.053$) between Virtual Multi-Monitor ($M = 4.55, SD = 1.39$) and Virtual Canvas ($M = 3.55, SD = 1.82$).

Other statements, such as *"I thought that I couldn't find the window that I needed quickly"* and *"I believe it took too long for me to switch attention between windows"* did not result in any significant differences. This could indicate that the perception of taking longer was not related to finding windows or switching attention.

5.8.2 Accuracy

In *"I felt that I delivered a quality result on the task."*, there was a significant effect of display style ($F_{1,1} = 4.80, p = 0.034$), with Canvas ($M = 5.12, SD = 1.30$) resulting in higher scores than Multi-Monitor ($M = 4.67, SD = 1.36$).

5.8.3 Comfort

In *"I found that I could see anything on the monitors at a glance"*, there was no significant main effect of either virtuality ($p = 0.25$) or display style ($p = 0.11$), but there was an interaction effect ($F_{1,1} = 11.03, p = 0.002$). Pairwise analysis showed significance between Virtual Multi-Monitor ($M = 3.7, SD = 1.78$) and both Hybrid Multi-Monitor ($p = 0.013, M = 5.05, SD = 1.76$) and Virtual Canvas ($p = 0.024, M = 4.55, SD = 1.95$). These differences could indicate the issues of field of view, with virtual multi-monitors having windows more spread out, and less peripheral vision.

In *"I found the screen space to be too small."*, there was a significant effect of display style ($F_{1,1} = 6.32, p = 0.016$), with Canvas ($M = 2.32, SD = 1.16$) having lower scores than Multi-Monitor ($M = 3.1, SD = 1.83$).

There was no significance for *"I was able to read text information displayed on the monitor(s)"* and *"I think the text on the virtual monitors should be bigger."*, indicating that participants didn't find important differences in readability across the configurations.

5.8.4 Other Aspects of User Experience

For *"I would trust this condition to do serious work"*, there was a significant interaction effect ($F_{1,1} = 4.71, p = 0.036$). Pairwise analysis showed marginal significance ($p = 0.095$) between Virtual Multi-Monitor ($M = 4.35, SD = 1.50$) and Virtual Canvas ($M = 5.2, SD = 1.40$).

For *"I found that the system behaved exactly how I expected."*, there was a marginally significant effect of display style ($F_{1,1} = 3.29, p = 0.077$), with Canvas ($M = 5.65, SD = 1.14$) having higher scores than Multi-Monitor ($M = 5.22, SD = 1.48$).

We didn't find differences on the following statements: *"I think it was easy to organize the windows in the screen,"* *"I felt that it was hard to go back and forth between my various windows,"* *"I found it intuitive to operate the system using this condition,"* *"I found the system behaved similarly to my previous experiences with computers,"* *"I think that I was able to focus on the tasks,"* *"I found it easy to keep track of my email inbox,"* *"I believe that I would use this system for my daily work instead of the one I currently use,"* *"I believe that this system is useful,"* and *"I believe that the task I completed was complicated."*

5.9 Qualitative Analysis

Using the transcribed interview data, we marked common topics in each interview using Taguette [37], from which we extracted the most recurring themes. In this section we report the most common themes identified in our interviews, some common to all conditions, and others specific to a given one.

5.9.1 Virtual Displays in AR

Participants commented on some effects of viewing the content in AR. Some participants mentioned that the virtual readability was worse than the physical monitor because of resolution, color fidelity, or background light/reflections, but all participants agreed they could read the text. P1 added, "I would say the only thing that was a little bit difficult was the Gmail [window]. I sometimes was unsure if it (the text) was bolded or not, and because there is no audio notification there was no extra reinforcement" (bolded text meant a new email had arrived).

They also reported on the positive aspect of seeing a large number of windows side-by-side. As P15 mentioned, "For the AR, it was cool that I was able to just look over instead of having to do keyboard shortcuts to switch it." P26 said, "I wasn't as distracted as I usually am ... I had the opportunity to open as many work documents as I can. That kind of made me focus a little better." Some participants discussed the similarity of our conditions and traditional multi-monitor setups. P3 said, "I think it has the same benefits as physical multi-monitor ... for example, for this task, you're able to see information that you need clearly and on a larger screen, instead of having everything on the same screen." P17 went further, "I wouldn't call it worse than a physical multi-monitor setup at all. I would definitely prefer an augmented reality setup to a single monitor setup. I would consider it for this type of work absolutely." P34 said, "I honestly don't think that [AR] made a difference at all. I think doing work on an augmented reality screen doesn't hinder what you can do." P13 summarized these ideas, "The system itself works basically almost exactly how Windows would work. There isn't really that much of a difference. I'd say it's more just the limitations of the AR headset."

Given the size of the displays, and our choice to render a transparent background, we were curious about distraction caused by seeing the lab background behind their virtual displays. Multiple participants mentioned they were not distracted by the existing environment. P3 said: "I didn't really mind. I didn't really notice any of the things in the background just because the display itself was clear enough that I was able to focus on that instead of being distracted by what was behind it, and also the task itself kind of led me to focus more." P32 added, "I think there's not too much influence, except for the lights." P31 commented on the light influence, "sometimes when I did prioritize the top, the light glare would come in. If it could recognize there's light, maybe darken or increase the screen brightness for that part. Because sometimes if I'm looking at text and there's light coming in, the text could get a little washed out."

5.9.2 Device Field of View

As found in previous studies, participants reported more head rotation than they would usually expect due to the small FOV of the HoloLens. As P12 explained, "I found that I couldn't really use peripheral vision because of the goggles. That kind of caused me to have to turn my head, which slowed me down." P28 said, "I guess just the way the glasses work is you have to really turn your head to be able to see that part of it, which makes sense. For some reason I was expecting to be able to see everything, all at once, but you have to kind of look around for it to appear in the vision." And that led to the choice of strategy by some participants. For instance, P40 said, "I wanted to focus on moving my head as little as possible. On Canvas that just meant sort of pushing everything as close to the center as I could. But on multi-monitor I only use the bottom three monitors essentially."

Some participants also lost the cursor a few times during the task, without any clear distinction between the conditions. The cost of recovering the cursor after it had been lost seemed to be higher due to the small FOV. P5 explained it, "I did [lose the cursor] a couple of times because when I'm looking I don't see the whole entire area. If I

put the mouse too much to the right, I have to look around to see it." As P7 said, "there was a couple of times that if I got distracted with something else, I would forget where I left it." P14 added, "I don't think it was because of one condition or the other. I just think I forgot where I left it and took a little while to find it again." Then, as P26 said, "I'd have to move it a couple of times until I see it."

Participants talked about some strategies they used to mitigate the reduced peripheral vision. P40 avoided missing the cursor, "I essentially had my windows in a circle. So I did not have to spend time looking for the cursor in Canvas because they were so much closer together." P7 improved their view of the email client windows, "I actually found with the Canvas condition that I liked having my Gmail [window] on my physical monitor as I would catch the e-mail coming in my peripheral vision while I was working." P24 agreed, "I think having the physical monitor is good for peripheral vision because with the AR goggles peripheral vision just simply does not exist."

5.9.3 Keyboard Usage

An important aspect of virtual display system is ensuring that the keyboard is fully visible for users, as that can be an obstacle for typing. All participants confirmed the keyboard was visible during the experiment, from which 11 explicitly stated they were able to touch type and thus didn't need to look at the keyboard. P23 preferred to look at the keyboard through a gap under the glasses, "I guess in the headset there's a gap under here [the lenses]. I just looked through there rather than looking through the glasses." On the other hand, P24 mentioned that the keyboard was much farther away than usual with our larger screen space, "Whenever I did look down on the keyboard, it was a lot more distance than I had anticipated."

Some participants commented on the windows not being directly over the keyboard, making them further away from looking at keys. P16 stated, "I felt like I was usually looking at the document while I was typing, which was on the right or the left, so I would have to kind of look back so I could see it, but it was kind of sluggish to see it (the keyboard)." P28 actively try to avoid looking at the keyboard, "I usually type looking at the keyboard a lot. But for this task I was trying to not do that just because it was hard to read the top one and then type into the bottom." This is a known negative effect of large displays, the distance between the keyboard and the content. But since it is easier to achieve such systems with AR, it is important to note it.

5.9.4 Canvas Conditions

With the Canvas conditions, participants reported more freedom to decide how to organize their windows. P9 stated, "I think removing all the boundaries made things a bit easier ... let's say you want to place windows somewhere in between the boundaries ... it would just [be] like really choppy and weird." P11 said, "when I did the monitor one, it felt my setup was kind of constrained to the size of one monitor. But when I had the Canvas I felt that I can make it as big as I want (window) and I can just put it right there and another one (window) right next to it." P13 tied their preference to the FOV, "I like Canvas better because I could display the windows how I wanted to, but that was mostly due to the fact that I couldn't see the whole screen." P26 added, "I felt I've got a chance to overlap windows I didn't care about, while I was working on the windows that I actually did care about"

Centrality of the windows was also discussed. P8 stated, "I think the multi-monitor has potential, but the way that it was organized didn't allow me to take advantage of the centrality, the central space on the screen and so I think that because the Canvas let me organize things the way I wanted to, I was able to use that center space as kind of my focus area." P9 mentioned that, "I just felt on Canvas I was able to find my things a lot more easily", because "they were more closely together and since the multi-monitors would probably push the windows a little too far away from me or and the fact that in the Canvas I was able to resize appropriately so that I could see better." P36 added, "In the multi-monitor, because of the way it snaps, you kind of can't have as much overlap, so I had to look further to see the other screens. When I had the Canvas, it was easier to get everything closer together, in a way." P14 mentioned, "The Canvas one I didn't really have much of

a structure going on. I kind of just picked an area so I didn't have to keep looking back and forth. It didn't have the limitation of the area being too small, but I kind of just didn't organize it really at all, 'cause there wasn't any structure to follow."

The lack of structure made it harder for some participants to organize their windows. P28 said, "It was just kind of slower to have to resize them (windows) and everything manually," P33 mentioned, "I needed same size of windows at the same time, so it was difficult to resize in the Canvas condition and it was easier to do in the multi-monitor one." P39 discussed having to think more actively on Canvas, "In Canvas I had to think about which would be the best place to put things in order for me to do this task. I didn't have to think about that with multi-monitor."

5.9.5 Multi-Monitor Conditions

Regarding the Multi-Monitor conditions, participants reported it being more structured, which helped them organize their windows. P3 said, "I think the borders helped at least mentally for organizing my space." P5 added, "I think the boundaries did help me figure out how to put the windows around each other so I could put them in a more like square-ish design, which was good." P14 mentioned about familiarity, "I felt that the structure where it had the lines, felt more natural. It felt more like I was looking at actual computer screens like I normally do, so being able to organize things on there felt a lot better and easier." P14 also felt "the boundaries made it so that I could pick a place to put things and remember it was there." P16 talked about the structure changing their behavior, "having boundaries definitely made me want to fill up the boundaries; having the boundaries made me want to have windows snap to filling up an entire window rather than kind of having them float around in variable sizes." P22 used the structure to split the tasks, "Multi-Monitor made it easier to organize the different monitors that you had. I would organize one of the monitors on one side to be one of the tasks and another one to be another task. And then in the center, I could just do whatever else I needed."

On the other hand, participants mentioned downsides introduced by having more structure. P40 said, "That limitation [boundaries] forced me to spread out the windows more than I would have liked since my field of view could only be one monitor at a time, essentially." P31 went further, "I feel like for the HoloLens, your restrictions are how much you can actually just see. And having more restrictions inside of that is a lot more difficult. My center (rest position looking forward) would probably be in the middle of the intersection of the boundaries where I would actually move them around more, if they were all free." P6 talked about limiting their choices, "When I did not have a boundary, I could just like put them wherever I want and it was much more intuitive for me to keep them wherever I wanted. But because of the boundary, it was kind of restricted for me and I'm ... OK, I'll keep this here."

One specific issue in the multi-monitor was the insufficient size of each monitor. As P1 explains, "When I was splitting a screen between windows in the multi-monitor, I had to, very often, scroll horizontally in order to be able to see the full screen, and I think that was the biggest limiting factor." P40 extended, "And also with the multi-monitors the snapping regions are like vertically divided. But since the windows for the PDF and word files are wider than half of a screen. This way if I snap it to one-half vertical section, I won't be able to read the entire page without scrolling." P16 said, "I had to use up more space to actually like read a whole document so there were times when I had to use the whole space even if I didn't really want to."

The top monitors were avoided by some participants. P22 added, "I thought that the top ones were pretty unwieldy, just because you have to actually look up a little bit. Versus the other ones to the sides were just a quick head turn left or right. I thought the bottom three monitors were basically the most useful ones than the top three monitors since I didn't end up using those at all." P9 mentioned that "The bottom half was just in front of my face. Up top, I have to like lift my head a bit and that takes a bit more effort."

5.9.6 Snapping

Participants commented very positively on the snapping feature that was present in the multi-monitor condition. P12 said, "I found it easier to like organize the screens because of the snapping feature and I didn't have to manually resize the screens every time, whereas it would automatically be resized when I snapped it into the screen in the multi-monitor." P21 added, "I think the snapping was easier to set up. That was probably the only benefit of the multi-monitor was being able to snap them into place. And just have it easier set up time compared to Canvas." P35 mentioned how easy snapping was to use: "I thought it was pretty intuitive, it was pretty quick to pick up." P38 added, "I did yeah, because I found like manually resizing was taking more time than the snapping. Snapping was easy and it saved some time." P3 said, "I didn't really have to think about organizing my windows and resizing things just because the snapping feature allowed me to do it quickly."

There were also comments about the snapping feature being too constraining at times. P8 said, "I just didn't like the way it was out laid out because it didn't allow me to take advantage of this idea that I want the things I'm focused on to be front—and the big keyword for me is center." P9 added, "You have a fixed size on each window and that lack of freedom kind of hinders the way I can see/read the windows there." P13 mentioned, "It's just that, sometimes you need to have more than two halves of the screen, which in this context you probably do." P15 said, "It would be nice to kind of be able to split the screen in different ways, so like I know we split it into left and right, maybe even being able to split it up and down" P21 said, "Snapping it on fixed sizes that weren't exactly what I needed. So that was a bit of a problem. It made my work on the task slower 'cause I had to like, you know, drag the [scroll] bar to move the view over to see what was said." P36 talked about wasted space: "I think the snapping is more of an issue when you're trying to move multiple things ... for example the student image, you don't necessarily need that to be really big ... it didn't need to be half the screen, but the way it was set up, it had snapped anyway, so it took up more space than it was needed."

5.10 Discussion

In our first five hypotheses, we focused on the differences between display style. We first hypothesized that utilization of space would be more optimized in Canvas (**H1**). Our results support H1. We found that Canvas windows sizes were smaller, that there was more overlapping of windows, and that they utilized a smaller percentage of the available space, with a smaller spread of windows in both horizontal and vertical. Those results indicated that Canvas layouts reflected a greater degree of optimization than the Multi-Monitor conditions. Of course, this was influenced by the snapping behavior present in the Multi-Monitor conditions. While we believe that this pairing is not a confound, but a feature frequently associated with this approach, we do not imply that our results here are generalizable to all multi-monitor setups.

Our second hypothesis was that there would be less head rotation and cursor movement in Canvas (**H2**). This hypothesis was only partially supported. We didn't find significant effects on horizontal movements, but we did on vertical ones. Coupled with other metrics on spread, we can see that participants in Canvas didn't place windows as high as they did in Multi-Monitor, which could indicate that the existence of the boundaries forced users to use go higher to be able to use the space. The lack of significance on horizontal movements was unexpected. Considering that the objective spread measure was significant, we believe that although windows occupied more space in Multi-Monitor, those effects did not fully translate to head movement due to the white space at left and right sides of the windows, which users do not need to look at or interact with. Therefore, the evidence suggests that the less-optimized placement of windows on Multi-Monitor will lead to some degree of larger cursor movement and head rotations.

Hypothesis three stated that participants would perform more window management in Canvas than Multi-Monitor (**H3**). Our evidence supports this hypothesis. We found that Canvas led to a larger number of grab/release actions, and that 52% of the releases on Multi-Monitor displays were snaps (which had pre-defined sizes and locations). We also found that Canvas had more move operations than Multi-Monitor,

and Multi-Monitor had more resize operations than Canvas. Once again, this indicates that in Canvas participants performed more optimizations on their windows. However, it also shows that (1) some of the movements on Canvas became resizes in Multi-Monitor, because of the snapping behavior, and (2) considering that 65.5% of resizes were snaps, Multi-Monitor users still had to perform adjustments in the window sizes.

This becomes clearer when we take into account the interview comments, where multiple participants complained that in the Multi-Monitor conditions they couldn't snap two windows side-by-side without them being a little bit smaller than the width of the text. The distributions of the move and resize operations also show how Multi-Monitor setups reduced size changes that were not at the snapping locations, but maintained a similar percentage of small adjustments. This shows one of the limitations of partitioning display space, as the partitions may not be optimal for the content one wants to display. In general, Canvas does indicate a higher need for layout optimizations, but with the caveat that characteristics of the Multi-Monitor setup and the windows being used can affect how much effort users need to put into window arrangement.

Our fourth hypothesis stated that more diverse window placement strategies would be supported by Canvas than Multi-Monitor (**H4**). Our results partially support this hypothesis. Virtual Canvas seems to indeed lead to more diverse organizations, while the asymmetry in Hybrid Canvas seems to introduce barriers that make users try to focus their use on the physical monitor. Qualitative feedback also suggests that participants felt Canvas led to less structure and more freedom, while Multi-Monitor and snapping were pushing them to conform to a certain style of organization, with windows more spread out and with more symmetry.

The fifth hypothesis expected a worse user experience from Canvas than Multi-Monitor (**H5**). We did not find evidence to support this hypothesis. Virtual Multi-Monitor was rated 28% higher in terms of the perceived task completion time than Virtual Canvas. Canvas was perceived as the condition where participants delivered a higher quality result on the task (9% higher). Participants rated Virtual Multi-Monitor as worse than both Hybrid Multi-Monitor and Virtual Canvas regarding being able to see anything on the monitors at a glance. The Hybrid Multi-Monitor result is likely more related to the physical monitor enhancing resolution and FOV. It was also more likely for Multi-Monitor users to perceive the amount of screen space to be too small than it was for Canvas users. While Multi-Monitor scored marginally higher on ease of window organization, participants were more likely to trust Virtual Canvas than Virtual Multi-Monitor for performing serious work, and they felt Canvas behaved more similarly to how they expected. Finally, most participants preferred Canvas (57% overall, 70% in Virtual, and 45% in Hybrid). Overall, our evidence suggests that Canvas delivered a higher level of user experience, contrary to our hypothesis.

For the next two hypotheses, we focused on the differences affected by virtuality. In our sixth hypothesis we expected participants in Hybrid conditions to take longer organizing windows than Virtual (**H6**). The results do not support this hypothesis. When looking at the time taken to organize windows, we didn't find any significant effects. There was a trend for Hybrid Multi-Monitor to actually perform the fastest, while Hybrid Canvas the slowest. While this trend was not significant, we found significant pair-wise interaction effects on number of operations that followed a similar pattern. Therefore, it seems that the pairing between the conditions seems to have important implications. While Hybrid Multi-Monitor preserved symmetry in interaction (although not visual properties), Hybrid Canvas actually led to different mental models that may have been responsible for participants performing more and longer operations.

Our final hypothesis was that the pairing of window style and virtuality would actually change which condition would be most preferred (**H7**). Our evidence supports this hypothesis. We found that 55% of participants in the Hybrid group preferred Multi-Monitor and that 70% of participants in the Virtual group preferred Canvas. Their comments revealed that while Multi-Monitor's snapping feature and sub-division can help structure the space in both Virtual and Hybrid settings, Hybrid

specifically enhances similarity with the systems they usually work with, and opens up the possibility to use the virtual display for windows that require glancing while working primarily on the physical monitor. From the Virtual group, we found that Canvas allowed them to place windows more tightly together, use partial overlap of windows, enjoy more freedom, and move their heads less. Virtual Canvas users were able to center the windows in their sweet spot without being limited by a fixed monitor size or being distracted by the boundaries, which was not possible in any of the other conditions.

These results from H5, H6, and H7 are particularly interesting when we contrast with the literature. A prior study [36] found that their version of a hybrid multi-monitor condition was better than virtual multi-monitor, which is consistent with our results. The conclusions in that paper were that a physical monitor still has significant benefits over a virtual monitor, and so it's a better choice to extend a physical monitor rather than replace it with virtual monitors. However, our study shows that a purely virtual display can have some important benefits for complex multi-window tasks when we remove the boundaries to create a Canvas. Therefore, with these new results, our recommendation is more nuanced: if you are going to extend a physical monitor, use a virtual multi-monitor setup, but if a physical monitor is not available or not needed, use a virtual Canvas.

6 CONCLUSIONS AND FUTURE WORK

In this paper, we proposed the concept of a Virtual Canvas achieved through AR technology for eliminating boundaries between monitors, and performed a user study to characterize how this design affects user experience on a complex multi-window task. Our results show that eliminating boundaries between multi-monitor setups can lead to better utilization of the space, but at the cost of more interaction. On the other hand, a conventional multi-monitor display can take advantage of snapping features and structure of the divisions to place windows quicker and, initially, with less interaction. Still, the sizes of those windows will be constrained, leading to less efficient utilization and more head movement. In some cases, it may also force users to perform small interactions to fix the size of windows.

In the future, we plan to do a follow up study between our Virtual Canvas and a physical canvas display, that could be achieved with high-resolution projection. Each may have their own advantages, such as the curved nature of virtual, or the high resolution nature of physical displays. We further plan to extend our design of unbounded displays to include depth, and explore approaches to dealing with occlusion when stacking content. The main rationale for pursuing this idea is the advantages presented by placing windows closer together to reduce the effects of the small FOV. A third topic we plan to investigate is how to optimize window management in both Canvas and Multi-Monitor approaches by using novel characteristics of virtual displays, such as dynamically changing the sizes of multi-monitors, or designing intelligent constraints for simplifying window placement on Canvas.

ACKNOWLEDGMENTS

We would like to thank all the participants for their time in participating in the studies.

REFERENCES

- [1] C. Andrews, A. Endert, and C. North. Space to Think: Large High-Resolution Displays for Sensemaking. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '10, pp. 55–64. Association for Computing Machinery, New York, NY, USA, 2010. event-place: Atlanta, Georgia, USA. doi: 10.1145/1753326.1753336
- [2] R. Ball and C. North. Effects of Tiled High-Resolution Display on Basic Visualization and Navigation Tasks. In *CHI '05 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '05, pp. 1196–1199. Association for Computing Machinery, New York, NY, USA, 2005. event-place: Portland, OR, USA. doi: 10.1145/1056808.1056875
- [3] R. Ball, C. North, and D. A. Bowman. 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, pp. 191–200. Association for Computing Machinery,

- New York, NY, USA, 2007. event-place: San Jose, California, USA. doi: 10.1145/1240624.1240656
- [4] M. Bellgardt, S. Pick, D. Zielasko, T. Vierjahn, B. Weyers, and T. W. Kuhlen. Utilizing immersive virtual reality in everydaywork. In *2017 IEEE 3rd Workshop on Everyday Virtual Reality (WEVR)*, pp. 1–4, 2017. doi: 10.1109/WEVR.2017.7957708
 - [5] X. Bi and R. Balakrishnan. 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, pp. 1005–1014. Association for Computing Machinery, New York, NY, USA, 2009. event-place: Boston, MA, USA. doi: 10.1145/1518701.1518855
 - [6] V. Biener, S. Kalamkar, N. Nouri, E. Ofek, M. Pahud, J. J. Dudley, J. Hu, P. O. Kristensson, M. Weerasinghe, K. C. Pucihar, M. Kljun, S. Streuber, and J. Grubert. Quantifying the Effects of Working in VR for One Week. *IEEE Transactions on Visualization and Computer Graphics*, 28(11):3810–3820, 2022. doi: 10.1109/TVCG.2022.3203103
 - [7] V. Biener, D. Schneider, T. Gesslein, A. Otte, B. Kuth, P. O. Kristensson, E. Ofek, M. Pahud, and J. Grubert. Breaking the Screen: Interaction Across Touchscreen Boundaries in Virtual Reality for Mobile Knowledge Workers. *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–1, 2020. doi: 10.1109/TVCG.2020.3023567
 - [8] S. Butscher, S. Hubenschmid, J. Müller, J. Fuchs, and H. Reiterer. Clusters, Trends, and Outliers: How Immersive Technologies Can Facilitate the Collaborative Analysis of Multidimensional Data. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, pp. 1–12. Association for Computing Machinery, New York, NY, USA, 2018. event-place: Montreal QC, Canada. doi: 10.1145/3173574.3173664
 - [9] G. Cetin, W. Stuerzlinger, and J. Dill. Visual Analytics on Large Displays: Exploring User Spatialization and How Size and Resolution Affect Task Performance. In *2018 International Symposium on Big Data Visual and Immersive Analytics (BDVA)*, pp. 1–10, 2018. doi: 10.1109/BDVA.2018.8534027
 - [10] M. Czerwinski, G. Smith, T. Regan, B. Meyers, G. G. Robertson, and G. K. Starkweather. Toward characterizing the productivity benefits of very large displays. In *Interact*, vol. 3, pp. 9–16, 2003.
 - [11] S. Davari, F. Lu, and D. A. Bowman. Occlusion Management Techniques for Everyday Glanceable AR Interfaces. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*, pp. 324–330, 2020. doi: 10.1109/VRW50115.2020.00072
 - [12] A. Eiberger, P. O. Kristensson, S. Mayr, M. Kranz, and J. Grubert. Effects of Depth Layer Switching between an Optical See-Through Head-Mounted Display and a Body-Proximate Display. In *Symposium on Spatial User Interaction*, SUI '19. Association for Computing Machinery, New York, NY, USA, 2019. event-place: New Orleans, LA, USA. doi: 10.1145/3357251.3357588
 - [13] L. A. Elkin, M. Kay, J. J. Higgins, and J. O. Wobbrock. An Aligned Rank Transform Procedure for Multifactor Contrast Tests. In *The 34th Annual ACM Symposium on User Interface Software and Technology*, UIST '21, pp. 754–768. Association for Computing Machinery, New York, NY, USA, 2021. event-place: Virtual Event, USA. doi: 10.1145/3472749.3474784
 - [14] A. Endert, L. Bradel, J. Zeitz, C. Andrews, and C. North. Designing Large High-Resolution Display Workspaces. In *Proceedings of the International Working Conference on Advanced Visual Interfaces*, AVI '12, pp. 58–65. Association for Computing Machinery, New York, NY, USA, 2012. event-place: Capri Island, Italy. doi: 10.1145/2254556.2254570
 - [15] Engadget. Facebook's Infinite Office is a virtual office space for the WFH crowd, 2022.
 - [16] B. M. Ens, R. Finnegan, and P. P. Irani. The personal cockpit: a spatial interface for effective task switching on head-worn displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 3171–3180, 2014.
 - [17] S. Feiner, B. MacIntyre, M. Haupt, and E. Solomon. Windows on the World: 2D Windows for 3D Augmented Reality. In *Proceedings of the 6th Annual ACM Symposium on User Interface Software and Technology*, UIST '93, pp. 145–155. Association for Computing Machinery, New York, NY, USA, 1993. event-place: Atlanta, Georgia, USA. doi: 10.1145/168642.168657
 - [18] J. L. Gabbard, D. G. Mehra, and J. E. Swan. Effects of AR Display Context Switching and Focal Distance Switching on Human Performance. *IEEE Transactions on Visualization and Computer Graphics*, 25(6):2228–2241, 2019. doi: 10.1109/TVCG.2018.2832633
 - [19] F. Garcia-Sanjuan, J. Jaen, and V. Nacher. Toward a General Conceptualization of Multi-Display Environments. *Frontiers in ICT*, 3, Sept. 2016. doi: 10.3389/fict.2016.00020
 - [20] J. Grubert, T. Langlotz, S. Zollmann, and H. Regenbrecht. Towards Pervasive Augmented Reality: Context-Awareness in Augmented Reality. *IEEE Transactions on Visualization and Computer Graphics*, 23(6):1706–1724, 2017. doi: 10.1109/TVCG.2016.2543720
 - [21] 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(6):125–133, 2018.
 - [22] J. Grudin. Partitioning Digital Worlds: Focal and Peripheral Awareness in Multiple Monitor Use. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '01, pp. 458–465. Association for Computing Machinery, New York, NY, USA, 2001. event-place: Seattle, Washington, USA. doi: 10.1145/365024.365312
 - [23] S. Jeuris, P. Tell, S. Houben, and J. E. Bardram. The Hidden Cost of Window Management. *arXiv:1810.04673 [cs]*, Oct. 2018. arXiv: 1810.04673.
 - [24] D. Kobayashi, N. Kirshenbaum, R. S. Tabalba, R. Theriot, and J. Leigh. Translating The Benefits Of Wide-band Display Environments Into An XR Space. In *Symposium on Spatial User Interaction*, pp. 1–11, 2021.
 - [25] K.-D. Le, T. Q. Tran, K. Chlasta, K. Krejtz, M. Fjeld, and A. Kunz. VXSLate: Exploring combination of head movements and mobile touch for large virtual display interaction. In *Designing Interactive Systems Conference 2021*, pp. 283–297, 2021.
 - [26] J. H. Lee, S.-G. An, Y. Kim, and S.-H. Bae. Projective windows: bringing windows in space to the fingertip. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pp. 1–8, 2018.
 - [27] Z. Li, M. Annett, K. Hinckley, K. Singh, and D. Wigdor. HoloDoc: Enabling Mixed Reality Workspaces That Harness Physical and Digital Content. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI '19, pp. 1–14. Association for Computing Machinery, New York, NY, USA, 2019. event-place: Glasgow, Scotland Uk. doi: 10.1145/3290605.3300917
 - [28] F. Lu, S. Davari, L. Lisle, Y. Li, and D. A. Bowman. Glanceable AR: Evaluating Information Access Methods for Head-Worn Augmented Reality. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 930–939, 2020. doi: 10.1109/VR46266.2020.00113
 - [29] T. Mahmood, E. Butler, N. Davis, J. Huang, and A. Lu. Building Multiple Coordinated Spaces for Effective Immersive Analytics through Distributed Cognition. In *2018 International Symposium on Big Data Visual and Immersive Analytics (BDVA)*, pp. 1–11, 2018. doi: 10.1109/BDVA.2018.8533893
 - [30] M. McGill, A. Kehoe, E. Freeman, and S. Brewster. Expanding the Bounds of Seated Virtual Workspaces. *ACM Trans. Comput.-Hum. Interact.*, 27(3), May 2020. Place: New York, NY, USA Publisher: Association for Computing Machinery. doi: 10.1145/3380959
 - [31] M. J. McGuffin. Augmented Reality Knowledge Work: Towards a Research Agenda. In *Workshop on Immersive Analytics: Interaction Design and Prototyping for Immersive Analytics*, p. 11, 2019.
 - [32] D. Medeiros, M. McGill, A. Ng, R. McDermid, N. Pantidi, J. Williamson, and S. Brewster. From Shielding to Avoidance: Passenger Augmented Reality and the Layout of Virtual Displays for Productivity in Shared Transit. *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–11, 2022. doi: 10.1109/TVCG.2022.3203002
 - [33] A. Ng, D. Medeiros, M. McGill, J. Williamson, and S. Brewster. The Passenger Experience of Mixed Reality Virtual Display Layouts in Airplane Environments. In *2021 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pp. 265–274, 2021. doi: 10.1109/ISMAR52148.2021.00042
 - [34] T. Ni, D. A. Bowman, and J. Chen. Increased Display Size and Resolution Improve Task Performance in Information-Rich Virtual Environments. In *Proceedings of Graphics Interface 2006*, GI '06, pp. 139–146. Canadian Information Processing Society, CAN, 2006. event-place: Quebec, Canada.
 - [35] E. Ofek, J. Grubert, M. Pahud, M. Phillips, and P. O. Kristensson. Towards a practical virtual office for mobile knowledge workers. *arXiv preprint arXiv:2009.02947*, 2020.
 - [36] 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
 - [37] R. Rampin and V. Rampin. Taguette: open-source qualitative data analysis. *Journal of Open Source Software*, 6(68):3522, 2021. Publisher: The Open Journal. doi: 10.21105/joss.03522

- [38] R. Raskar, G. Welch, M. Cutts, A. Lake, L. Stesin, and H. Fuchs. The Office of the Future: A Unified Approach to Image-Based Modeling and Spatially Immersive Displays. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '98*, pp. 179–188. Association for Computing Machinery, New York, NY, USA, 1998. doi: 10.1145/280814.280861
- [39] P. Reipschläger, T. Flemisch, and R. Dachsel. Personal Augmented Reality for Information Visualization on Large Interactive Displays. *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–1, 2020. doi: 10.1109/TVCG.2020.3030460
- [40] G. Robertson, M. Czerwinski, P. Baudisch, B. Meyers, D. Robbins, G. Smith, and D. Tan. The large-display user experience. *IEEE Computer Graphics and Applications*, 25(4):44–51, 2005. doi: 10.1109/MCG.2005.88
- [41] A. Ruvimova, J. Kim, T. Fritz, M. Hancock, and D. C. Shepherd. "Transport Me Away": Fostering Flow in Open Offices through Virtual Reality. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI '20*, pp. 1–14. Association for Computing Machinery, New York, NY, USA, 2020. event-place: Honolulu, HI, USA. doi: 10.1145/3313831.3376724
- [42] M. Stone. Color and brightness appearance issues in tiled displays. *IEEE Computer Graphics and Applications*, 21(5):58–66, 2001. doi: 10.1109/38.946632
- [43] D. S. Tan and M. Czerwinski. Effects of visual separation and physical discontinuities when distributing information across multiple displays. In *Proc. Interact*, vol. 3, pp. 252–255, 2003.
- [44] Tao Ni, G. Schmidt, O. Staadt, M. Livingston, R. Ball, and R. May. A Survey of Large High-Resolution Display Technologies, Techniques, and Applications. In *IEEE Virtual Reality Conference (VR 2006)*, pp. 223–236. IEEE, Alexandria, VA, USA, 2006. doi: 10.1109/VR.2006.20
- [45] J. R. Wallace, D. Vogel, and E. Lank. Effect of Bezel Presence and Width on Visual Search. In *Proceedings of The International Symposium on Pervasive Displays, PerDis '14*, pp. 118–123. Association for Computing Machinery, New York, NY, USA, 2014. event-place: Copenhagen, Denmark. doi: 10.1145/2611009.2611019
- [46] J. O. Wobbrock, L. Findlater, D. Gergle, and J. J. Higgins. The aligned rank transform for nonparametric factorial analyses using only anova procedures. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 143–146. ACM, Vancouver BC Canada, May 2011. 893 citations (Crossref) [2023-03-13]. doi: 10.1145/1978942.1978963



Leonardo Pavanatto is a Ph.D. candidate at Virginia Tech, a member of the 3D Interaction Group and the Center for Human-Computer Interaction. His expertise area is VR/AR within the scope of 3D User Interfaces. Some specific research interests include how to use augmented reality for improving productivity on real-world tasks.



Feiyu Lu obtained his Ph.D. from Virginia Tech in May 2023. His research interests lie broadly in the intersections of AR/VR, 3DUI, and HCI. His Ph.D. work focuses on enabling lightweight and unobtrusive information display and interactions on AR HWDs to support a variety of everyday tasks.



Chris North is a Professor of Computer Science at Virginia Tech, and Associate Director of the Sanghani Center for AI and Data Analytics. His research examines the use of large high-resolution and immersive display spaces for visual analytics.



Doug A. Bowman is the Frank J. Maher Professor of Computer Science and Director of the Center for Human-Computer Interaction at Virginia Tech. He is the principal investigator of the 3D Interaction Group, focusing on the topics of three-dimensional user interfaces, VR/AR user experience, and the benefits of immersion in virtual environments.