



McGill, M., Kehoe, A., Freeman, E. and Brewster, S. (2020) Expanding the bounds of seated virtual workspaces. *ACM Transactions on Computer-Human Interaction*, 27(3), 13.

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

© The Authors 2020. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *ACM Transactions on Computer-Human Interaction*, 27(3), 13.
<http://dx.doi.org/10.1145/3380959>.

<http://eprints.gla.ac.uk/217386/>

Deposited on: 5 June 2020

Expanding The Bounds Of Seated Virtual Workspaces

MARK MCGILL[†], AIDAN KEHOE[‡], EUAN FREEMAN[†], and STEPHEN BREWSTER[†],
[†]University of Glasgow, Scotland and [‡]Logitech Design Lab, Republic of Ireland

Mixed Reality (MR) Augmented and Virtual Reality (AR/VR) headsets can improve upon existing physical multi-display environments by rendering large, ergonomic virtual display spaces whenever and wherever they are needed. However, given the physical and ergonomic limitations of neck movement, users may need assistance to view these display spaces comfortably. Through two studies, we developed new ways of minimising the physical effort and discomfort of viewing such display spaces. We first explored how the mapping between gaze angle and display position could be manipulated, helping users view wider display spaces than currently possible within an acceptable and comfortable range of neck movement. We then compared our implicit control of display position based on head orientation against explicit user control, finding significant benefits in terms of user preference, workload and comfort for implicit control. Our novel techniques create new opportunities for productive work by leveraging MR headsets to create interactive wide virtual workspaces with improved comfort and usability. These workspaces are flexible and can be used on-the-go, e.g., to improve remote working or make better use of commuter journeys.

CCS Concepts: • **Human-centered computing** → **Virtual reality**; *Usability testing*; *Interaction techniques*; • **Computing methodologies** → **Mixed / augmented reality**;

Additional Key Words and Phrases: Virtual Reality; Augmented Reality; Mixed Reality; Workspaces; Productivity; Virtual Desktops; Virtual Displays; Display Space; Displays; Multi-Monitor; Rotational Gain;

ACM Reference Format:

Mark McGill[†], Aidan Kehoe[‡], Euan Freeman[†], and Stephen Brewster[†]. 2020. Expanding The Bounds Of Seated Virtual Workspaces. *ACM Trans. Comput.-Hum. Interact.* 27, 3, Article 13 (May 2020), 43 pages. <https://doi.org/10.1145/3380959>

1 INTRODUCTION

Multi-display workspaces allow us to access more information [20], provide peripheral awareness of information that is not the main focus [33] and increase productivity by enabling more efficient multi-tasking [16]. In effect, they “improve efficiency in ways that are difficult to measure yet can have substantial subjective benefit” [33]. However, arranging physical multi-display environments can be problematic: they are expensive, require large amounts of space and have high energy demands. They have a limited size, shape and orientation. They are not portable, requiring users to resort to the “impoverished” [2] environment of single display laptops and tablets when working on the move. They can also be an ever-present aesthetic blight on workplaces and homes, are difficult to conceal when turned off, and are impossible not to notice when turned on. Use of multi-monitor arrangements has also been suggested to increase the risk of musculoskeletal disorders [79]. However some neck/head movement can protect against such disorders due to greater variation in muscle activity [27].

Authors' address: Mark McGill[†]; Aidan Kehoe[‡]; Euan Freeman[†]; Stephen Brewster[†],
[†]University of Glasgow, School of Computing Science, Glasgow, Scotland, first.last@glasgow.ac.uk,
[‡]Logitech Design Lab, Cork, Republic of Ireland, akehoe@logitech.com.

In contrast, consumer grade Mixed Reality (MR) [74, 99] Virtual and Augmented Reality (VR/AR) headsets have opened up the possibility of rendering content virtually, replacing multiple physical displays with virtual ones. They can render content with depth, at any position, at any size, and can track head orientation, which can be used as a proxy for gaze direction. Wearers of these headsets can view virtual workspaces as and when required, in any position and orientation in the virtual (VR) or real (AR) world. Given such capabilities, we can envisage virtual display spaces that are egocentrically-oriented around the user, providing large virtual canvases upon which to arrange applications and content. Such display spaces could automatically adapt to ergonomic guidelines [121] and adopt beneficial configurations for specific tasks [30, 50]. However, whilst we can escape the limitations of physical monitors, we cannot escape the physical limitations of our bodies; wide virtual workspaces will have at least the same demands on the user’s neck muscles compared to existing multi-display environments, and consequently the same risks of neck strain and developing musculoskeletal disorders.

This paper examines how the physical demands of viewing wide virtual workspaces can be lessened, across two studies. In study 1, we do so by manipulating the mapping between the user’s head rotation and virtual display counter-rotation around the natural pivot point of the user. We map the virtual display space onto a real-world range of motion of half the size, meaning $\pm 120^\circ$ and $\pm 60^\circ$ of virtual space become accessible through $\pm 60^\circ$ and $\pm 30^\circ$ of real-world head rotation respectively. We examine the use of three novel mappings for navigating this space: a constant counter-rotational gain, a central deadzone with dynamically determined gain for peripheral displays, and dynamic gain with deadzones on all virtual displays. In a targeting task ($n=16$) with participants viewing content across this virtual space, our assistive mappings maintained user accuracy whilst lessening the physical discomfort and neck fatigue. The *deadzones on all displays* condition proved particularly effective and was preferred by participants.

In study 2, we then iterated upon the most preferred combination of mapping and workspace size from our first study, enabling users to implicitly control the position of displays in a three display workspace by triggering transitions to the previous/next display when users reached the edge of a display. Using a state-of-the-art VR workspace with interactive desktops and a positionally tracked physical keyboard and trackpad, participants compared this implicit control using gaze to both a control condition where displays were fixed in place, and a condition where users could explicitly shift displays left/right by key press. We found significant benefits in favour of implicit control of display position based on gaze, with positive effects in terms of user preference, workload and comfort. We discuss the implications this research has for the design of virtual workspaces in the future. Our research affirms for the first time the significant productivity benefits to be had through the adoption of VR/AR workspaces that increase users’ capability to view content spatially oriented around them, whilst avoiding negative effects in terms of fatigue and discomfort.

2 LITERATURE REVIEW

“The impoverished environment of the single monitor forces users to make explicit context switches on the introduction of new information, frequently in the form of a new window overlaying the previous one. This severely affects the user’s ability to make comparisons and requires the user to expend valuable mental resources on the minutiae of managing views rather than on the problem at hand” [2].

Expanding The Bounds Of Seated Virtual Workspaces

Much research has been conducted examining the benefits [1] and drawbacks [42] of multi-display environments [29] such as in Figure 1. In general users show a “unanimous preference” for such configurations, as they enable “multi-window and rich information tasks, enhanc(ing) users’ awareness of peripheral applications, and offer(ing) a more ‘immersive experience’ ” [8]. Additional displays are a “partition to physically distribute work across and take advantage of our rich 3D spatial cognition capabilities” [78]. Czerwinski *et al.* [16] found that a larger, wider display space led to increased performance of users when carrying out complex, multi-window productivity tasks. Andrews *et al.* noted that large, high-resolution displays provide “space to think” [2, 3], supporting sensemaking by providing external memory, and allow for more information to be accessed and visually compared at once, more quickly, and mediated by glance-based behaviours. Andrews *et al.* noted that documents could be placed in persistent locations, allowing physical navigation and preventing switches between application level tasks (e.g. reading) and system level view management, and physically arranged into related units or collections. Ball *et al.* [6] noted that physical navigation (moving eyes, head, body) increased user performance and was preferred. Ling *et al.* [65] found that the number of window switches decreased. And Czerwinski *et al.* [15] suggested that multi-display environments allowed users to “engage in more complex multitasking behavior”, whilst numerous papers have noted benefits of having “abundant” display spaces [56] in terms of spatial memory [94] and performance across a variety of productivity tasks [13, 51, 55, 65, 80, 106].

However, larger display spaces do have drawbacks [42]. The necessity of such a space is sometimes questionable, with Endert *et al.* noting that often “the ability to see all of (the workspace) all the time is not needed” [20]; so the physical displays may take up desk space when not needed. There are also cost implications as multiple large displays are expensive,




Fig. 1. A modest example of a wide physical multi-display workspace. It is ergonomically questionable, requiring a large degree of movement to switch from attending to the left/rightmost displays, and effectively permanently partitions the user from the shared office space. Each physical display also has varying characteristics e.g. in terms of pixel density, luminance, degrees of freedom for re-positioning the display etc.

especially if they are not always needed. And [34] noted “often times a user preferred a display that did not match optimal performance”. With respect to multiple separated displays, [89] suggested that “the visual and physical separation between displays requires that users perform visual attention switches between displays”, and there are diminishing returns in adding more displays [75]. The very capability of being able to attend to content peripherally can itself be a distraction [68], with the possibility that dividing attention across displays can be counter-productive [119]. With respect to singular large contiguous display spaces, Jakobsen *et al.* [91] suggested there were potential issues with increased mental demand from having any one application expand to fill the entire display space. Employing a larger display space also creates new design challenges [15, 107] regarding cursor tracking, crossing gaps between physical displays, managing space and layout, failing to leverage the periphery, and helping users comprehend changes on unattended displays [17, 39, 47].

2.0.1 Physical Consequences Of Large Display Spaces. The issue of neck strain has been repeatedly raised, particularly with respect to the use three monitors:

“Because of the width of the three monitors, I frequently tilt my neck to see the side monitors; with two monitors I could see both when looking straight ahead by just turning my eyes, but not with three... If I spend an inordinate amount of time focusing on one of the side monitors my neck becomes a bit stiff. To combat this I try to make myself swivel my chair when I start looking at a side monitor, but it’s second nature for me to turn my head and I don’t realize that I’m not swivelling until my neck starts to bother me.” [75]

Viewing wider display spaces leads to the activation of different muscle groups [102] than viewing a single display. Nimbarde *et al.* noted that “Increased activation of anterior neck muscles caused by asymmetrical, more rotated head-neck postures... may increase the risk of neck musculoskeletal disorders, especially with prolonged computer use” [79]. In effect, prolonged usage has the potential for profoundly negative effects. However, some degree of movement can be beneficial; the movements required for typical dual-monitor usage may provide some protection compared to single display environments, due to the greater variation in muscle activity [27, 103, 123]. User controls for managing the workspace have the potential to minimize physical consequences if they are used, but as noted by Ball and North in a study of a pan and zoom interface for an array of displays, participants preferred not to use the provided controls, instead choosing to physically navigate [5, 100]. In addition, as the desk surface tends to support the peripherals used in daily productivity tasks, there may be a natural inclination to turn the neck, rather than the rest of the body, to keep the hands on peripherals. Whilst Endert *et al.* discussed the potential for keyboard/mouse trays that moved with the rotation of the office chair [20], such a solution is not hugely generalisable to the wide range of peripherals and tasks that occur at, or around, the desk.

2.0.2 Organizing And Managing Display Space. Broadly, there are multiple means by which content on a display can be managed:

Positioning and sizing of content Such as moving application windows around the space, resizing them to fit;

Grouping / layering content From the Windows taskbar for grouping windows under an application shortcut, to structures that link windows or views using “ordering, proximity, and alignment to create structures like piles, clusters, lists, and even heterogeneous interrelated types” [96];

Summary views For example alt-tab overviews in Windows and Exposé on OS X, but alternatives have been explored, see Warr *et al.* [115] for an overview;

Zoomable spaces Similar to summary views, but instead consist of a much larger contiguous display space, with the user then zooming in on the area they wish to focus on [87] e.g. Hutchings *et al.* [44] suggested a scheme whereby part of the display space was dedicated to management, with a “focus region” where tasks were conducted;

Virtual desktops and transitions Referring to the context (*i.e.*, active applications, current layout) of a given display being encapsulated in a virtual display, such that different virtual displays can be transitioned to/from each physical display. Jeuris *et al.* [49] showed that dedicated virtual desktop workspaces allow for faster task resumption and reduced cognitive load, noting there was a “strong argument for supporting goal-oriented dedicated workspaces”.

The usage of these capabilities is highly variable [45, 46], with users exhibiting different space management styles, and utilizing different subsets of the tools available, with multi-display users relying on the taskbar less and window interactions more to switch between windows [45].

2.1 What Constitutes A Virtual Workspace?

VR and AR headsets remove the need for physical displays embedded in the environment; instead, they can render content anywhere, at any depth and size. Moreover, there is nothing that binds a given piece of virtual content to the physical or virtual position it inhabits in the way that a physical monitor is often immovable with respect to user interactions or inputs (with some actuated exceptions [104]). As a consequence, work regarding large and multi-display configurations needs to be reconsidered from the perspective of headsets that enable a display space of infinite possibilities in terms of composition, layout and interactivity.

2.1.1 How Should A Virtual Workspace Present Content? Firstly, there is the question of how productivity content is presented. Should interfaces utilize depth? Should the “window” of the application be free to move around in space, or should it be anchored to some virtual display/container? How should content be shaped and oriented, and should it exist within ego or exocentric space? For some of these questions, there are as-yet no clear answers.

With respect to use of depth, the Google Daydream team [71] suggest that “the further away you get from a purely 2D representation of content, the less efficient it will become at conveying the idea you’re trying to get across” [71]. Use of depth can be problematic because eyes can converge at different depth points, with extra visual information to process, and parallax effects, leading to greater eye fatigue, whilst Tan and Czerwinski [105] suggested that performance may be negatively impacted by mixed display distances.

Entirely 2D workspaces have been noted to be fastest, as there is also a cost in switching between 3D displays [52]. Accordingly, it seems reasonable to suggest that workspace-type content be planar. With respect to the shape and size of these planes, curved planes appear to be preferable, making content more accessible [97], less fatiguing and more legible [81]. Endert *et al.* [20] also state that “the degree of the curvature is somewhat up to personal preference, but we suggest that curving the display to a configuration where the user is equidistant from all areas of the display”. It has been suggested that each virtual piece of content should effectively fit within the comfortable field of view of the headset so that parts are not cut off [72], with task time being optimal if the display size is a maximum of $\frac{3}{4}$ of the headset field of view due to reduced head motion [25].

2.1.2 A Single Continuous Display Space, Or Discrete Displays? Then there is the question of whether each plane is an application window placed with some constraints in space, or a container for application windows, mimicking existing physical displays. Rendering application windows as independent items would have some benefits: *e.g.*, each one could be sized and placed freely, without constraint, with the potential to be anchored based on context *e.g.* Sharma *et al.* aligning virtual ‘insights’ with real objects to create shared virtual workspaces. However, no constraints in terms of edges and boundaries mean that support needs to be designed to help users position, resize and snap windows together whilst avoiding occlusion (*e.g.*, Android widget placement grid). In addition, no compartmentalized displays or application containers mean that for virtual desktop switching, the whole 360° space is effectively a singular virtual workspace.




Fig. 2. *Top* examples of free windows in space, from left: Magic Leap [53], Microsoft Hololens [73], Multiscreens [83]. *Bottom* examples of window containers, from left: Virtual Desktop [31], Bigscreen [10], SPACE [86].

Conversely, having virtual displays/application containers means that the space can be arbitrarily divided, with applications assigned to specific overloaded regions, replicating existing behaviour [44]. Each container can also have its own set of hidden/minimized applications and virtual desktops. Familiar view management behaviours such as the summary views discussed previously, snapping and split-views in Windows [14], and exploiting display edges for layout [113] are possible. As Grudin noted: “very large displays will find significant uses...yet a place will remain for the arbitrary division of space... space with a dedicated purpose, always accessible with a glance” [33]. Hybrid approaches have also been mooted, for example Zhen *et al.* explored workspaces that had both physical and digital content contained within them [63]. Indeed, it would seem reasonable that such hybrid presentations, combining both virtual displays and application containers along with real physical displays and artefacts, would be likely to see adoption. Such environments would still provide access to the familiar features that users rely on to assist in the management and division of their digital workspaces when in Mixed Reality.

2.1.3 Layout of Virtual Workspaces: “Ethereal Planes”. Assuming that the user’s workspace consists of multiple 2D planar elements, referred to by Ens *et al.* as “Ethereal planes”, there is then the layout of these planes to be considered, in terms of what frame of reference and arrangement [22]. With respect to the frame of reference, Ens *et al.* contrasted exocentric and

Expanding The Bounds Of Seated Virtual Workspaces

egocentric layouts. *Exocentric*, or grounded/world-locked [61], refers to content being placed in the world, much as an existing monitor is now. Exocentric content may aid recall by placing elements in a global spatial context. It enables the environment to be augmented, for example overlaying digital content on physical surfaces to provide an element of tactile feedback [118] or creating activity-promoting spaces [85], and allows for collaboration through shared views. However, with exocentric content, discoverability needs to be supported so that the user can find it in the physical world, whilst transitioning content layouts between different physical spaces is problematic e.g. prioritizing visual saliency versus spatial constancy [24].

Conversely, *Egocentric* or display-locked content refers to content that is placed relative to the user, remaining at the same offset regardless of what motion the user undertakes. This means that, irrespective of location, the user can access a spatially familiar layout. However, head-locked egocentric interfaces may provoke simulator sickness when used in-motion, due to visually perceiving fixed content when physically perceiving motion [72].

The frame of reference needs to be considered because both exo- and egocentric frames imply certain layouts: in the former, with content placed in the world, aligned with physical features; in the latter, with content placed around the user in a variety of configurations e.g. vertically or horizontally, two-plus-two [66], or oriented toward the user as in the “personal cockpit” [23, 25], etc. However, as shown in the personal cockpit, egocentric-type layouts are optimal when exocentrically placed, i.e. oriented around a fixed point in world space where the user’s head is. This provides a display space that is effectively oriented and laid out around the user without inducing nausea due to moving with the user. It also retains the ability to be transposed to different environments whilst retaining spatial consistency, and can be adjusted based on need (e.g. height for seated and standing usage [66]). With respect to how content should be oriented when placed around a single point in world space, it has been suggested that displays should be oriented toward/curved around the point to maximize legibility in spherical configurations [25], with the user placed offset to this centre to create a less claustrophobic arrangement [71].

Horizontal arrangements have shown some benefits for background awareness [64] and users appreciate bow shaped screen configurations and symmetry [66]. Su *et al.* [101] suggested that content should not be positioned behind the user. Moreover, the capability to comfortably explore this space via neck/eye movement must be considered. For the horizontal field of view, Hololens documentation [72] recommends that neck rotations of more than 45° off-centre are to be avoided, whilst the Google Daydream team [71] suggest that comfortable eye movement can accommodate $\pm 35^\circ$ on the horizontal, with neck movement leading to a comfortable span of $\pm 60^\circ$. These recommendations are in line with ergonomics research (see [81] for a summary), where Tiley and Dreyfuss suggested a $\pm 45^\circ$ range for easy head rotation, with $\pm 60^\circ$ for maximal head rotation [110]. In spherical configurations of displays, the central display is effectively a central focus area for the most frequently used applications [66], with the left and right sides of the focus area used for supporting content as “shifting attention from the central focus area to the left or right side requires head or body movement” [66]. The more peripheral the content, the less attention is intended to be devoted to it, with applications presenting notifications, messages and status information “mostly placed at the outside of the field of view” [66], with fatigue decreasing as users tend toward the central display zone [81]. Berki also noted some cognitive benefits regarding virtual workspaces, particularly with respect to information density/availability [7].

2.2 Supporting Large, Comfortable VR/AR Display Spaces

VR and AR displays can recreate any large or multi-display workspace, in any environment or context, with all the benefits and drawbacks therein. But they can also potentially go beyond what is possible with physical displays, creating display spaces that can adapt to the current task/environment/user/ergonomics, etc. However, a number of design questions need to be answered, regarding whether the concept of physical display boundaries is retained, how users will be aided in maintaining and populating an abundant display space, and what kinds of layouts are optimal. This paper takes the position that **a.** Virtual displays should be provided; **b.** Existing display management techniques should be retained where possible, without necessitating wholly new headset-specific window management techniques [112] such as world-in-miniature approaches [25] **c.** Spherical or ellipsoid layouts would appear to be preferable based on current research. Finally, we ignore the shared use case [114] and instead concentrate on the optimal workspace for a single user. If we consider the usage of such a configuration during a standard workday, with the user seated at a desk using peripherals resting on it, the negative effects of excessive head and neck movement could impair the usability of such a display space. Whilst content can be placed peripherally across many displays, these displays may not be comfortably reachable, and may cause harm if used everyday for hours at a time.

However, because these display spaces are virtual, and because AR/VR headsets are crudely instrumenting user gaze through headset orientation, they also have the potential to break the physical rules and constraints of physical display spaces. If we consider any spherical or ellipsoid arrangement, the centre of this arrangement is inherently a pivot point, implying that these arrangements can be rotated around this point. Prior research has commonly exploited “rotational gain” where user head movements were dynamically accelerated. For example, this has been used in redirected walking [48, 59, 84, 90, 116] and accelerated panning in VR 360° videos [41]. However, the breadth of ways by which the position of these virtual displays could be manipulated/counter-rotated around the pivot point (e.g. based on head angle, gaze, or explicit user command), and the consequences of performing such counter-rotations (e.g. in terms of workload, physical effort, fatigue, neck discomfort etc.) have yet to be explored.

3 STUDY 1: MANIPULATING THE MAPPING BETWEEN HEAD ROTATION AND WORKSPACE POSITION

Given an egocentrically-oriented virtual display space (a contiguous set of virtual displays positioned around, and facing towards, the user - but fixed in world space), we firstly set out to explore how the mapping between user head rotation and virtual display space counter-rotation could be manipulated. By this, we mean that if a user were to rotate their head left, the virtual displays could, determined by the current head rotation angle of the user, be rotated around the user in the opposite direction. Figure 3 illustrates display counter-rotation - this could require less head rotation to reach a given display, and potentially both expand the reachability of uncomfortably wide virtual workspaces and make smaller widths of workspace more easily accessible.

We defined a virtual display space for this work, based on the findings of our literature review. Our display space consisted of a horizontal arrangement of five virtual displays, each curved such that every point was equidistant from the central pivot point. This display space was positioned so that when the VR user faced forward, the middle of the five displays was aligned with their view; i.e., there was a central display with two peripheral displays

Expanding The Bounds Of Seated Virtual Workspaces



Fig. 3. We manipulate the mapping between head rotation and virtual display position, here seen from above the user, with the arrow indicating the user head rotation and the arc being the virtual displays. *Left*: The user looks straight ahead at the central display of a workspace. *Right*: As the user's head rotates to the left, the virtual displays rotate to the right; i.e., they counter-rotate, with the extent of this counter-rotation determined by the mapping in use.

on each side. Each display was 60° wide: the 'screen' was 50° wide with 5° margin on each side, creating the appearance of five distinct displays, rather than a single wrap-around screen. These virtual displays were placed within a VR living-room (see Figure 4). We used a living-room setting because this would provide similar visual cues to those that would be experienced if using AR headsets (excluding the difference in field of view): i.e., a stable background with moving virtual elements. This also represented the potential worst-case in terms of motion/simulator sickness, as opposed to an empty/featureless VR workspace.

Given this virtual display space, we then defined research questions to examine how counter-rotation of the virtual displays, around the pivot-point of the user, might impact the usability of the peripheral displays; e.g., in terms of physical effort and time taken to transition to a peripheral display. We defined four mappings between a given head rotation angle and its associated virtual display counter-rotation. These affected the way the displays moved when the user moved their head. The first mapping was our control condition, where head rotation occurred in a static environment with no counter-rotation of the virtual displays, as is standard in VR. The second mapping investigated the use of a static counter-rotational gain. As the user rotated their head, the displays would counter-rotate (i.e., rotate in the opposite direction); this effectively results in greater rotation in the virtual environment for a given head movement. This gain could expand the reachable bounds of a display space. The use of static gain raised two research questions:



Fig. 4. The virtual living-room environment with five 60° wide displays positioned around the user.

RQ1.1 How does static counter-rotational gain affect the usability of the peripheral displays?

RQ1.2 How perceivable is counter-rotational gain, when fixating on a particular point on a display?

With **RQ1.1** we consider the impact of gain on usability, when accessing content on the peripheral displays. Gain should allow users to more comfortably reach the furthest extents of the display space with less head rotation, but may come at a cost if the rotation acts against proprioceptive awareness of head angle. By addressing this, we will also measure input performance and characterise the use of static counter-rotational gain in VR.

When a rotational gain is constantly applied, subtle and unintentional natural head movements might be enough to cause perceivable counter-rotation of the display space. Such perceptible movements might cause inaccuracy when fixating on a particular display of interest, or affect comfort and simulator/motion sickness incidence. As such, **RQ1.2** considers the extent of perceptible movement and whether or not it negatively affects usability.

If counter-rotational gain was indeed perceivable when fixating on a display, it could be desirable to mitigate these head movements. To do this, we consider the use of deadzones: areas within the display space where no counter-rotational gain is applied. Deadzones are regions within each display that the user can fixate upon, without perceiving additional counter-rotational movement. The use of deadzones led to our third and fourth mappings.

The third mapping used a deadzone within the central display only. The deadzone was $\pm 12.5^\circ$, meaning no counter-rotational gain was applied when the user's head rotation angle was within this range of the middle of the central display. The peripheral displays were accessible as before using a counter-rotational gain, which was only applied when gaze moved out of the deadzone. The fourth mapping used a deadzone in every virtual display; the width was the same as before. As users rotated their head, they would experience quick sliding transitions at the edges of the displays. The 'sliding' was a result of dynamic between-display gain being experienced at the edges of each display, then no gain being experienced in the deadzone at the center of each display. In both cases, the total size of each deadzone was half of the width of the display (minus margins), lying well within the comfortable range of eye movement if the head was oriented toward the center of the display. A further two research questions arose from these mappings:

RQ1.3 What impact does introducing a deadzone to the central display have on the usability of the peripheral displays?

RQ1.4 What impact does introducing a deadzone to all displays have on the usability of the peripheral displays?

Deadzones were introduced to mitigate undesirable counter-rotation when fixating on a display, but the change in dynamics of the counter-rotational gain may impact usability. As such, **RQ1.3** considers how the use of a deadzone in the central display only affects access to the peripheral displays. If deadzones are indeed beneficial, then it may be desirable to use them on all displays in the virtual environment. **RQ1.4** considers the effect on usability of having a deadzone in every virtual display.

3.1 Independent Variables

Given these research questions, two factors were defined: *Mapping* and *No.OfDisplays*. There were four levels for *Mapping*, the previously-described mappings between head rotation and display counter-rotation:

Expanding The Bounds Of Seated Virtual Workspaces

- 1: No mapping** This was the control condition, with no assistive mapping: neck movement in VR matched the amount of neck movement in reality.
- 2: Constant assistance** This condition corresponded to **RQ1.1**. It had a constant 2x counter-rotational gain.
- 3: Central deadzone** This condition corresponded to **RQ1.3**. It featured a fixed central display, with a $\pm 12.5^\circ$ deadzone. When outside of the deadzone, the mapping enacted a dynamic counter-rotational gain for viewing the peripheral displays.
- 4: All displays have deadzones** This condition corresponded to **RQ1.4**. Every display had a central $\pm 12.5^\circ$ deadzone. Outside of these deadzones, the mapping enacted inter-display counter-rotational gain, resulting in the appearance of rapid transitions occurring in the space between the deadzones.

We implemented these mappings using *Animation Curves* in the Unity3D gaming engine (see Figure 5). This meant that, for every frame rendered in the VR scene, counter-rotation would be applied on the basis of the current head rotation angle.

There were two levels for the *No.OfDisplays* factor, allowing us to evaluate the use of our mappings across two widths of virtual space: three displays and five displays. Since each display was 60° wide (50° with a 10° margin), this meant that without an assistive mapping, the three display condition had an effective virtual range of $\pm 60^\circ$ between the middle of the central display and the middle of the peripheral displays. For the five display condition, the range was $\pm 120^\circ$ to the middle of the left-/right-most displays (see Table 1).

We investigated *No.OfDisplays* because we wanted to understand the efficacy of our mappings across different ranges of physical movement. To evaluate the mappings under the same conditions, we mapped these virtual display space ranges onto a real-world range of motion of half the size, meaning $\pm 120^\circ$ and $\pm 60^\circ$ of virtual space was accessible through $\pm 60^\circ$ and $\pm 30^\circ$ of real-world head rotation, respectively. This meant that each mapping




Fig. 5. Transfer graphs of the mappings of head rotation to display counter-rotation used for each condition and across the *No.OfDisplays* factor (top = 3 displays, bottom = 5 displays). For all, except the “no mapping” conditions, a head rotation of $\pm 30^\circ$ degrees would result in a counter-rotation of the displays by 30° for the three display conditions, and of 60° for the five display conditions.

<i>No. Of Displays</i>	Virtual Span	Physical Head Rotation Required	
		Unassisted	Assistive Mapping
3 Displays	$\pm 60^\circ$	$\pm 60^\circ$	$\pm 30^\circ$
5 Displays	$\pm 120^\circ$	$\pm 120^\circ$	$\pm 60^\circ$

Table 1. Range of head rotation required to reach the center of left/right peripheral displays for each level of *No. Of Displays*. Unassisted refers to the control condition with no counter-rotational gain, and assisted refers to any condition where counter-rotational gain is applied.

would be evaluated within the comfortable range of neck movement (i.e., expanding the immediately accessible workspace) and approaching the maximal neck range (i.e., providing an additional infrequent access space) [110]. See the supplementary video figure for examples of the experimental conditions in motion.

3.2 Task

We designed an experimental task that would realistically stress the user’s viewing of the peripheral displays whilst maintaining their engagement and motivation to perform the task well. Participants watched a nature documentary whilst performing a targeting task (see Figure 6 and the video figure for examples). First, they would dwell for one second on a target placed on the central display. A gaze reticule, based on forward orientation of headset, was used for targeting. Then, they were instructed to look either left or right to the next target. After dwelling on the peripheral target for one second, the participant would continue watching the nature documentary on the display currently in focus, for 8 ± 2 seconds. After this duration, a new target would appear on the current peripheral display; once selected, they would be instructed to move back to the central display, where they would again select a target then wait for 8 ± 2 seconds. This sequence constituted one full trial and always resulted in participants moving from the central display to a peripheral display, then back again.




Fig. 6. Example of targeting task experienced, looking at the transition from central to peripheral display.

Participants had two practice trials for each condition, before then performing ten trials over the course of ~ 5 minutes. The left/right displays were viewed an equal number of times for each condition with a randomized order, and condition order was counter-balanced. Participants ($n=16$, 8 male, average age= 26.9) were recruited from University forums and paid $\pounds 8$ for taking part, with sessions taking ~ 1 hour. The task was built using the Unity3D engine (2017.1) and presented using the Oculus Rift CV1 headset, adjusted so each participant could clearly and comfortably view the content.

3.3 Dependent Variables

We validated the usability of the mappings using both quantitative and qualitative measures. For the quantitative measures, we examined only the central-to-peripheral transitions. If we included the peripheral-to-central return transitions as well, the responses would be confounded by the *Central Deadzone* condition, where targeting would be noticeably different on the central display versus the peripheral displays. The quantitative task metrics were:

Virtual target accuracy (Virtual Accuracy) The angular distance (in degrees) between the head rotation (measured by VR headset orientation) and the virtual position of the target (i.e., where the target is currently located in VR based on the current mapping and head rotational angle).

Real-world target accuracy (Real Accuracy) The angular distance (in degrees) between the head rotation and the effective real-world position of the target (i.e., where the center of the target would be if being looked at directly, based on the current mapping).

Total real-world movement during dwell selection (Dwell Movement) The cumulative head rotation (in degrees) that occurred during a successful dwell selection (i.e., how well users could focus/fixate over time).

Time to peripheral target selection (Time To Selection) Starting when the central display target is selected, the time taken (in seconds) to complete a successful dwell selection of the peripheral target.

Re-targeting duration The time (in seconds) until a successful selection after first looking at the target without dwelling for a selection (i.e., incidence of over/undershooting the target).

These measures would help to understand the quantitative performance impact/inefficiencies a given mapping might have (real-world target accuracy, time to peripheral target selection, re-targeting duration). They also give insight into the stability of the virtual peripheral display when attempting to select a target, both over the course of the duration of the target selection (total real-world movement during dwell) and at the actual moment of selection (virtual target accuracy).

Given the capability for ballistic head rotations to fixate on fixed displays, by dynamically influencing the visual perception of these ballistic head movements through our display mappings we might expect some negative impact in terms of our capability to quickly and accurately fixate on the center of a given display we are transitioning to. As a consequence, we record accuracy metrics regarding accuracy of target selection in VR and in reality. The difference between the virtual target accuracy and the real-world target accuracy is effectively the extent to which the dynamic mapping between head orientation and display position influences the position of the target. Consider the *Constant gain* condition - with every degree offset from the position of the target, in virtual terms this offset would be doubled. So if the user was looking at a target centered at 30°, and their head orientation on the y-axis was 25°, then virtually they would appear to be offset from the target by 10°. However, in reality, they were physically offset from the effective position of the center of the target by only 5°. In effect, virtual accuracy may not reflect the actual physical targeting behaviour of the user, because it is influenced by the mechanics of the mapping being used. For example, if virtual accuracy is maintained given a novel mapping, yet real-world accuracy has increased, this at a minimum suggests that the physical demands of targeting have increased, whether perceptibly or otherwise. As a consequence, both these metrics should be considered when describing the impact on accuracy of a given assistance technique.

For the qualitative measures, the emphasis was on understanding the physical implications of viewing wide virtual display spaces with and without assistive mappings, as well as one additional question examining the perception of display instability during targeting:

Simulator Sickness Questionnaire (SSQ) A measure of simulator/motion sickness used for VR studies [54].

NASA Task Load Index (TLX) A measure of user workload, recording mental demand, physical demand, effort, temporal demand, performance and frustration [35]

Physical discomfort “Please rate your physical discomfort when viewing the left/right displays”, based on the subjective comfort survey from [28] using a 7-point Likert scale ranging from no discomfort to pain.

Comfort viewing peripheral displays “I could view the left/right displays comfortably”, 7-point from strongly disagree to strongly agree. Where physical discomfort concentrated on physical symptoms, this question was to elicit a more general response regarding how comfortable participants were in locating and attending to the peripheral displays.

Neck fatigue “Please rate your neck fatigue” with a 12-point scale used by [82] based on the Borg CR10 scale [12].

Visual discomfort “Please rate your general visual discomfort (e.g. feelings of tiredness, soreness, irritation, watering and/or burning in eyes)”, 7-point from no discomfort to pain, based on [43].

Perceived body movement “I had to turn my body/shoulders to see the left/right displays”, 7-point from strongly disagree to strongly agree.

Perceived display stability “I noticed displays were moving when I stared directly at them”, 7-point from strongly disagree to strongly agree.

User rankings “Please rank the conditions you experienced in order of preference - which would you most prefer to use day to day?”, intended to elicit user preferences regarding the virtual display layouts and mappings experienced.

3.4 Results

For all results, a two-factor repeated measures ANOVA was performed. Where data were non-parametric, an Aligned-Rank Transform [120] was used to allow parametric methods. For effect size, Generalized Eta Squared (η_g^2) is reported (see [18] for interpretation). For *post hoc* contrasts, the *lsmeans* [62] R package was used with Tukey adjustment. The majority of the plots are Violin plots [37], displaying a rotated kernel density plot on either side of a box plot. Kernel density plots are “a variation of a Histogram that uses kernel smoothing to plot values, allowing for smoother distributions by smoothing out the noise. The peaks of a Density Plot help display where values are concentrated over the interval” [109], allowing for density estimation. The box plots feature notches denoting the 95% confidence level [58].

3.4.1 Quantitative Performance. For significance test results for the quantitative measures factors, see Table 2.

Virtual Accuracy. Firstly, regarding the mean virtual target accuracy (i.e. the unassisted angular distance between the head rotation angle and the virtual target position at the point of selection) there was a significant effect on *No .OfDisplays* and *Mapping* with *post hoc* contrasts finding differences between *No assistance* – *All deadzones* ($t=3.31, p<0.01$) and *Central deadzone* – *All deadzones* ($t=3.34, p<0.01$) (see Figure 7). Comparing the assistive conditions, the difference between *Central deadzone* and *All deadzones* suggests that dynamic gain without a deadzone does impair the ability to quickly fixate on a particular display,

Expanding The Bounds Of Seated Virtual Workspaces

Metric	Mapping	No.OfDisplays	Mapping x No.OfDisplays
Virtual accuracy	$F(3,105)=4.92, p<0.01, \eta_g^2=.06$	$F(1,105)=11.62, p<0.01, \eta_g^2=.06$	$F(3,105)=0.31, p=0.82$
Real accuracy	$F(3,105)=50.48, p<0.01, \eta_g^2=.36$	$F(1,105)=23.24, p<0.01, \eta_g^2=.09$	$F(3,105)=0.48, p=0.7$
Dwell movement	$F(3,105)=309.57, p<0.01, \eta_g^2=.85$	$F(1,105)=65.16, p<0.01, \eta_g^2=.26$	$F(3,105)=18.30, p<0.01, \eta_g^2=.30$
Time to selection	$F(3,105)=37.82, p<0.01, \eta_g^2=.31$	$F(1,105)=173.93, p<0.01, \eta_g^2=.41$	$F(3,105)=12.44, p<0.01, \eta_g^2=.14$
Re-targeting duration	$F(3,105)=35.38, p<0.01, \eta_g^2=.43$	$F(1,105)=47.6, p<0.01, \eta_g^2=.07$	$F(3,105)=24.28, p<0.01, \eta_g^2=.22$

Table 2. Statistical testing for main effects on *Mapping* and *No.OfDisplays*, with interaction effects. $p < 0.05$ highlighted

albeit not hugely so (amounting to a difference of a few degrees). This is in contrast to the static *Constant gain* condition which features broadly comparable performance to *All deadzones*.




Fig. 7. Absolute virtual accuracy of dwell target selections, measured as degrees offset from centre of target in virtual space in the y-axis.

The significant difference between the *No assistance* condition, which in effect is entirely a deadzone with no display counter-rotation, and the *All deadzones* condition, also suggests that the difference in accuracy in these cases is a result not of the mapping, but of the head rotation angle. As previously noted (see Table 1), the assistive mappings were created to effectively halve the neck rotation required, meaning that depending on *No.OfDisplays*, the center of the target was effectively at a head rotation of $\pm 30^\circ$ – $\pm 60^\circ$ degrees for the assisted conditions, compared to $\pm 60^\circ$ – $\pm 120^\circ$ unassisted. This suggests that accuracy subtly degrades as the neck angle on the y-axis increases, all things being equal.

Real Accuracy. With respect to the mean accuracy of the target dwell selections in real-world space (i.e. the angular distance between the head rotation angle and the required head rotation to reach the center of the target, as noted in Table 1), there were main effects for both *Mapping* and *No.OfDisplays* with *post hoc* contrasts finding significant differences between all mappings (all $|t| > 2.7, p < 0.05$). As can be seen in Figure 8, accuracy increased

going from *No assistance*, to *All displays having deadzones*, to *Constant gain* and finally having a *Central deadzone*. Accuracy was also significantly better with *3 displays* than *5 displays*.




Fig. 8. Absolute real-world accuracy of dwell target selections, measured as degrees offset from centre of target in real space in the y-axis.

An initial interpretation might conclude that both static and dynamic gain both dramatically increased the targeting accuracy when viewing peripheral displays. However, to fully interpret this result it must also be considered against the virtual accuracy. The conditions with peripheral gain being constantly applied effectively amplified all head movement when looking at the peripheral target. Accordingly, for users to approximate their perceived accuracy in conditions without peripheral gain (denoted by the virtual accuracy as seen by the user) they would have to be approximately twice as accurate in reality, given that being off by 0.9° in reality would result in being off by approximately 2° in VR (given the counter-rotations based on gain). What this does suggest however, is that users can significantly improve on their ability to stabilize their head movement, given the amplified visual feedback they received in the *Constant gain* and *Central deadzone* conditions. Whether this resulted in any significant increase in cognitive or physical workload will be discussed in subsection 3.4.3.

Time To Selection. Mappings in the *3 display* condition were broadly comparable, with significant differences only between *Central deadzone* and $\{Constant gain (t=5.90, p<0.01), All deadzones (t=7.68, p<0.01) \text{ and } No assistance (t=6.43, p<0.01)\}$. There was no improvement seen in the time to reach the target, with mean durations effectively ranging from 1.1 seconds to 1.2 seconds for the conditions that were not significantly different. For *5 displays*, there was a significant *post hoc* effect between *All deadzones* and $\{No assistance (t=4.48, p<0.01), Central deadzone (t=-3.99, p<0.01)\}$, with mean time going from 1.9 seconds with no assistance to 1.5 seconds with assistance.

Re-Targeting Duration. An explanation as to the reason for there being little difference in task time across mappings can be found in Figure 9, specifically looking at the re-targeting duration i.e. the time elapsed between the user gaze reticle first exiting the target bounds (if it did at all) and returning for a successful dwell. For the *3 displays* condition, *post hoc* tests found significant differences between *No assistance* – $\{Central deadzone (t=8.51, p<0.01), All$

Expanding The Bounds Of Seated Virtual Workspaces

deadzones ($t=3.89, p<0.01$), and *Central deadzone* – {*Constant gain* ($t=6.64, p<0.01$), *All deadzones* ($t=4.62, p<0.01$)}. In effect, the conditions that featured variable levels of gain assistance were likely to result in more overshoot behaviour. This behaviour was tempered in the *5 displays* condition however, with significant *post hoc* differences only between *Central deadzone*–{*No assistance*, *Constant gain*, *All deadzones* (all $|t|>3.8, p<0.01$)}. No significant learning effects (i.e. changes as users became more proficient) were found over the 10 trial duration for each mapping.




Fig. 9. Mean duration between first looking at target and starting dwell selection (re-targeting duration).

3.4.2 Stability of Peripheral Displays. This pattern of increased real-world accuracy in the gain-assisted conditions can also be seen in the significant main effects for angular movement during a successful dwell selection (see Figure 10), with the performance of the conditions in the same order as for real-world target accuracy. Pairwise contrasts for the interaction effect found significant differences between all pairs ($|t| > 3.27, p < 0.05$) **except:** *no assistance:3* – {*all deadzones:3/5*}, with the *all deadzone* conditions effectively featuring the same degree of movement as the *3 display* control; between *all deadzone:3/5*; and between *constant gain:3* – {*constant gain:5*, *central deadzone:5*}. As Figure 10 shows, the conditions where gain assistance was present when looking directly at the target featured lower overall angular movement (i.e. were more stable) during the 1 second dwell. The perceived movement question broadly mirrored this result.

3.4.3 Qualitative Measures. For significance test results, see Table 2.

SSQ Perceived Sickness. Regarding simulator sickness, there was a significant main effect on number of displays only, with *5 displays* inducing more nauseogenic symptoms. However, real levels were low, with maximum mean scores of ~ 10 , still well within the limits described as problematic by Kennedy *et al.* [54].

TLX Perceived Workload. There was a significant interaction effect for Overall TLX **Workload**. Pairwise *post hoc* comparisons emphasized the difference between *5 displays* and *3 displays*, with means for *No assistance:5* significantly different from *No assistance:3* ($t=3.67, p<0.01$), *Constant gain:3* ($t=3.81, p<0.01$) and *All deadzones:3* ($t=4.12, p<0.01$). However, there were no significant contrasts within the 3 or 5 display conditions. The TLX **Effort** subscale mirrored these results, but no significant contrasts were found for TLX **Performance**




Fig. 10. *Left*: Angular movement during successful dwell over target. *Right*: Perceived display movement question, ranging from strongly agree (perceiving movement) to strongly disagree (perceiving no movement).

or TLX **Frustration** on their significant interaction effects, and there were no significant effects on **Temporal demand**.

For TLX **Mental demand** there were significant main effects for *No.OfDisplays* and *Mappings*. *Post hoc* tests on *Mappings* found a difference between *Constant gain* and *Central deadzone* ($t=3.6, p<0.01$), which imposed greater demands on the user, we would suggest as a result of the observed oscillation behaviour for this technique. For TLX **Physical**

Metric	Mapping	No.OfDisplays	Mapping x NoOfDisplays
SSQ Simulator sickness	$F(3, 105)=1.96, p=0.13$	$F(1,105)=4.31, p=0.04, \eta_g^2=.01$	$F(3,105)=0.61, p=0.61$
TLX Overall workload	$F(3,105)=6.43, p<0.01, \eta_g^2=.03$	$F(1,105)=29.23, p<0.01, \eta_g^2=.07$	$F(3,105)=7.49, p<0.01, \eta_g^2=.05$
TLX Performance	$F(3,105)=1.61, p=0.19$	$F(1,105)=2.69, p=0.06$	$F(3,105)=2.27, p=0.09$
TLX Effort	$F(3,105)=6.75, p<0.01, \eta_g^2=.05$	$F(1,105)=24.98, p<0.01, \eta_g^2=.11$	$F(3,105)=5.48, p<0.01, \eta_g^2=.05$
TLX Frustration	$F(3,105)=2.57, p=0.06$	$F(1,105)=10.50, p<0.01, \eta_g^2=.02$	$F(3,105)=6.58, p<0.01, \eta_g^2=.05$
TLX Mental demand	$F(3,105)=4.62, p<0.01, \eta_g^2=.01$	$F(1,105)=4.57, p=0.03, \eta_g^2=.01$	$F(3,105)=0.77, p=0.51$
TLX Physical demand	$F(3,105)=9.28, p<0.01, \eta_g^2=.09$	$F(1,105)=64.64, p<0.01, \eta_g^2=.19$	$F(3,105)=7.12, p<0.01, \eta_g^2=.06$
TLX Temporal demand	$F(3,105)=2.01, p=0.12$	$F(1,105)=3.09, p=0.08$	$F(3,105)=2.19, p=0.09$
Physical discomfort viewing periphery	$F(3,105)=9.49, p<0.01, \eta_g^2=.08$	$F(1,105)=21.03, p<0.01, \eta_g^2=.01$	$F(3,105)=0.06, p=0.98$
Comfort viewing periphery	$F(3,105)=9.55, p<0.01, \eta_g^2=.17$	$F(1,105)=33.40, p<0.01, \eta_g^2=.17$	$F(3,105)=3.73, p=0.01, \eta_g^2=.07$
Neck fatigue	$F(3,105)=3.63, p=0.02, \eta_g^2=.04$	$F(1,105)=22.86, p<0.01, \eta_g^2=.08$	$F(3,105)=0.13, p=0.94$
Visual discomfort	$F(3,105)=1.83, p=0.15, \eta_g^2=.01$	$F(1,105)=11.54, p<0.01, \eta_g^2=.02$	$F(3,105)=4.52, p<0.01, \eta_g^2=.02$
Perceived body movement	$F(3,105)=23.87, p<0.01, \eta_g^2=.29$	$F(1,105)=100.21, p<0.01, \eta_g^2=.30$	$F(3,105)=7.77, p<0.01, \eta_g^2=.14$
Perceived display stability	$F(3,105)=25.60, p<0.01, \eta_g^2=.24$	$F(1,105)=4.47, p<0.05, \eta_g^2=.01$	$F(3,105)=2.09, p=0.1, \eta_g^2=.01$
Rankings by preference	$F(3,105)=10.79, p<0.01, \eta_g^2=.19$	$F(1,105)=64.11, p<0.01, \eta_g^2=.35$	$F(3,105)=0.52, p=0.67$

Table 2. Statistical testing for main effects on *Mapping* and *No.OfDisplays*, and interaction effects. $p < 0.05$ highlighted

Expanding The Bounds Of Seated Virtual Workspaces

demand, there was a significant interaction effect, with *post hoc* contrasts finding significant differences between *no assistance:5* – {*No assistance:3*, *Constant gain:3*, *Central deadzone:3*, *All deadzones:3/5* (all $t < -3.13, p < 0.05$)}. This effectively emphasized the significant physical demands imposed by the 120° movements of the non-assisted 5 *display* condition.

Perceived Fatigue. Examining the significant main effect of *Mappings* on neck fatigue, the *no assistance* conditions were statistically more fatiguing than the *All deadzones* conditions ($t=3.05, p=0.02$). There was a significant main effect on *No.OfDisplays* and *Mappings* for discomfort/pain, with *post hoc* tests finding significant differences between *No assistance* and *Constant gain* ($t=4.06, p<0.01$), *Central deadzone* ($t=3.80, p<0.01$) and *All deadzones* ($t=4.40, p<0.01$).

Perceived Physical Discomfort, Comfort In Viewing and Visual Discomfort. Examining perceived physical discomfort (see Figure 11) there were significant main effects on *Mapping* and *No.OfDisplays*, with 5 *displays* being more physically discomfoting than 3 *displays*, and *No assistance* being more physically discomfoting than *Constant gain* ($t=-4.41, p<0.01$) and *All deadzones* ($t=-4.68, p<0.01$).




Fig. 11. *Left:* Perceived physical discomfort when viewing the peripheral displays. *Right:* Perceived comfort viewing the peripheral displays.

Regarding the perceived comfort viewing the peripheral displays showed significant main effects on *No.OfDisplays* (with 5 *displays* being less comfortable) and *Mapping*, with a significant interaction effect observed. *Post hoc* contrasts found significant differences between *no assistance:5* and all other conditions **except** *central deadzone:5* (all $|t| > 3.20, p < 0.05$); and between *central deadzone:5* and *all deadzones:3* ($t=3.18, p=0.39$). For general visual discomfort there was a significant interaction effect but no significant *post hoc* contrasts.

Perceived Body Movement. With a significant interaction effect, participants perceived more body/shoulder movement in the *no assistance:5* condition relative to all other conditions (all $t > 5.00, p < 0.01$), as well as between *Central deadzone:5* and {*Constant gain:3* ($t=3.26, p<0.05$), *All deadzones:3* ($t=3.22, p<0.05$)}.




Fig. 12. Perceived body movement required to view peripheral displays.

Ranked preferences. Examining subjective rankings there were significant main effects on *No.OfDisplays* and *Mapping*, with *post hoc* contrasts on *Mapping* finding significant differences between *No assistance* – {*constant gain* ($t=3.29, p<0.01$), *All deadzones* ($t=5.42, p<0.01$)} and *Central deadzone* – *All deadzones* ($t=3.83, p<0.01$), with *All deadzones* having the lowest mean/median rankings, as can be seen in Figure 13.




Fig. 13. Participant rankings, ordered by mean ranking from top (best) to bottom (worst) condition.

3.5 Interviews

Interviews were loosely guided on the basis of reported rankings, and coded using Initial Coding, where participants’ statements were assigned emergent codes over repeated cycles. These codes were then grouped using a thematic approach and reported based on frequency and interest (see [92]), with representative excerpts quoted.

3.5.1 Five displays were problematic. Whilst this was to be expected given the reasoning behind the choice of the five display workspace (to go beyond the comfortable range of neck movement), it is important to note that, without an assistive mapping, the five display workspace was, as expected, universally considered problematic by participants, predominantly for reasons of comfort. For example “it wasn’t that comfortable, if you didn’t

Expanding The Bounds Of Seated Virtual Workspaces

move your whole body it would actually hurt your neck a little bit” (P4) requiring “too much effort for me to turn all the way round in the seat” (P6).

The three display workspace was much less problematic, with fewer interview comments particularly singling it out, however one participant noted that the $\pm 60^\circ$ angle (utilized for the 3 display, no assistance condition and the assistive mappings on the 5 display workspace) was acceptable only for glance-based usage, but that they “wouldn’t want to work with my head at that angle for more than a couple of seconds” (P2).

3.5.2 Constant assistance was divisive. Providing constant assistance (i.e. some form of gain without peripheral deadzones) was noted by four participants as being “predictable so you could target precisely” (P1), and fluent (P5), “assist(ing) you so you didn’t have to move too far” (P8). However, constant assistance was felt to be either unnecessary or too aggressive in its application by four participants, being “annoying” and “overwhelming” (P3) and “overly sensitive” leading to overshooting (P7).

3.5.3 Deadzone preference based on viewing stability. In comparison to the *Constant assistance* and *Central deadzone* conditions, having deadzones on all the displays was remarked upon by five participants as being preferable because of the viewing stability:

P6: I really disliked the ones that have constant assistance, because its just taking any movement... If I was breathing, coughing, talking, that was affecting the way you were looking at it, it should have remained locked. (I preferred) that the screens would remain rigid, without too much movement from myself.

However, having one central deadzone alone was suggested by three participants to result in more oscillation:

P8: I didn’t like the mix of being still in the middle and then moving to the side, because you overshoot it, so then you would have to overcorrect, it was a bit disorientating.

Also, four participants noted that movement between displays for the *All deadzones* conditions was “too sudden” (P2, P12), noting “the movement didn’t feel natural, it felt like it was something you’d want to control with a keypress” (P2).

3.5.4 Known targets requiring unknown movements. There was also the question of when assistance should begin to be provided. Three participants noted that, with the *3 display* conditions, assistance could be disconcerting because they were transitioning to displays that were within their field of view, and had a known and expected degree of head movement associated with them:

P3: It was easier to have muscle memory. In the real world trying to look here or there I know how far I need to move my head.

3.5.5 Preferences dependent on usage/task/equipment. The issue of what mapping would be appropriate for different tasks was also raised by two participants:

P1: If it was a glance, to look at a twitter feed then look back again, then I can imagine my preference would change based on the rapidity of that movement, either how quickly it takes to go there and come back, and how frequently. I can imagine if i’m predominantly looking at one screen, the deadzone in the centre is good because I can still look around and it’s not going to shift, but I don’t know necessarily which one would be better for frequent flicks.

In addition, two participants noted that the weight of the headset would be likely to contribute to their experience of fatigue and physical discomfort during the task:

P3: I did experience a little bit of a headache and a little bit of tiredness, but I think that was from the heaviness of the headset causing neck pain from leaning forward.

3.6 Implications for Research Questions

RQ1.1 How does static counter-rotational gain affect the usability of the peripheral displays?

Regarding **RQ1.1**, to fully understand the benefits and drawbacks of applying rotational gain for viewing the virtual display space, we firstly need to consider how performance compared across different viewing ranges. Consider the most extreme case: viewing $5 \times 60^\circ$ width displays across a $\pm 120^\circ$ range. With no assistive mapping, this use case was entirely impractical, being well beyond the acceptable/safe viewing angle for head+eye movement, requiring whole body movement, fatiguing the neck and causing discomfort in the process.

However, mapping the wide *5 display* $\pm 120^\circ$ space to a $\pm 60^\circ$ head rotation range for the *Constant gain* condition resulted in body movement, neck fatigue, selection time, physical discomfort and viewing comfort all being brought to levels comparable *3 display, no assistance* condition, with limited effects on workload and targeting accuracy. Those that preferred *Constant gain* felt it to be predictable mapping to use. Considering this condition alone, there is evidence that applying counter-rotational gain in this manner can effectively squeeze more display space into a given range of head rotation.

RQ1.2 How perceivable is counter-rotational gain, when fixating on a particular point on a display?

However, regarding **RQ1.2**, it was noted that subtle head/neck movements would invariably be amplified, and users strongly perceived this motion. Interview responses suggested that for at least five participants the constant movement of the displays based on their head movement was disliked, although there is no evidence that this motion consequently resulted in increased motion sickness. It is important to note, however, that the dislike for the *Constant gain* was not universal.

Interestingly, the perceived movement of the peripheral displays did not affect targeting accuracy, with users seemingly more accurate in real-world degrees in order to maintain accuracy in virtual-world degrees. Whilst this potentially increased effort in targeting did not register on the TLX questionnaire, we would suggest that, over the course of a workday, the potentially subtle increase in physical and mental demand required to stabilize display content may become an impediment. Accordingly, we would suggest this mapping be used only for glance-based activities, or where focus is constantly shifting.

RQ1.3 What impact does introducing a deadzone to the central display have on the usability of the peripheral displays?

The display stabilization problem was, however, anticipated, and **RQ1.3** and **RQ1.4** both examined different mappings that took this issue into account. The intention of **RQ1.3** was to examine the feasibility of a central stable display with a peripheral glance-based region accessible under counter-rotational gain. Whilst this mapping was effective for the *5 display* factor, for *3 displays* this brought about unanticipated problems regarding user expectations when making short leaps in gaze between targets that are within their range of vision. Given a target that was partially visible, users knew the amount of head rotation required to fixate on it. Thus, there was not enough time during the ballistic phase of target acquisition to adjust for the changing magnitude of their movement under assistance, resulting in users spending ~ 0.5 seconds more to fixate on the target due to re-targeting.

Expanding The Bounds Of Seated Virtual Workspaces

Whilst the *Central deadzone* condition did retain many of the benefits of the assistive mappings whilst providing a stable central display (e.g. for on-going work, with peripheral displays acting as more transient information spaces), it is clear that for such a solution to become usable some iteration would have to take place. For example the central deadzone could be made larger, such that the glance-based region is at the very extremity of the user's head rotation, much like edges of a physical display are often used as "infinite" width Fitts law targets. Or the ratio between real/virtual mappings could be altered, with a less aggressive mapping defined, to better allow users to target within the peripheral space. This mapping was also subject to increased targeting demands, as it featured the same increased real-world accuracy to maintain virtual targeting accuracy as with *Constant gain*.

RQ1.4 What impact does introducing a deadzone to all displays have on the usability of the peripheral displays?

In contrast, with respect to **RQ1.4**, having a deadzone on every display appeared to enable users to access the wider display space without impairing targeting, providing a stable view when fixating on any display, with user rankings indicated that this condition was the most preferred of the mappings. Notably, *All deadzones:5* was broadly comparable to *No assistance:3*, meaning that for the *5 display* space this mapping brought performance to a level comparable to the control *3 display* workspace. Given that both the unassisted *3 display* space and the assisted *5 display* space both operated within a range of $\pm 60^\circ$, it is promising to note that similar performance levels can be achieved using one of the mappings evaluated, providing access to a display space of double the size in the process. Similarly, if we compare *No assistance:3* operating over $\pm 60^\circ$ with *All deadzones:3* operating over $\pm 30^\circ$ we again see decreases in terms of body movement, neck fatigue and discomfort, as well as improved performance in selection time. However, the *All deadzones* did still have some notable problems. For example, four participants noted that either the transition movements between displays were too abrupt, or the size of the deadzones was not large enough.

4 STUDY 2: IMPLICIT VERSUS EXPLICIT CONTROL OF WORKSPACE

Study 1 demonstrated a strong potential utility for providing implicit assistance in viewing wide virtual display arrangements by counter-rotating displays based on head rotation, with the *All deadzones:3* Condition being broadly preferred. However, this study had some notable caveats. Firstly, the task was weakly ecologically valid, with no productivity element and limited interaction with our virtual "workspace", instead focussing on tightly controlling the amount of physical movement and examining impact on accuracy in a target selection task. And secondly, the focus was solely on exploring different approaches toward using head orientation to dynamically determine the counter-rotation of the virtual displays. There was no examination of explicit discrete user control of the display positions e.g. through keypress. Whilst the literature makes a persuasive case against user control (e.g. in terms of poor potential adoption, and the significant ergonomic problems that a lack head movement brings), nonetheless user control is the defacto standard for virtual desktop usage currently. For example, Windows 10 supports virtual desktop switching through key press, whilst OS X supports inputs from touchpad/mouse/keyboard for the same actions.

Consequently, we designed a second study to examine two research questions to address these gaps:

RQ 2.1 Do user preferences for implicit counter-rotation of virtual displays hold when performing a productivity task?

RQ 2.2 To what extent is implicit counter-rotation of virtual displays preferable or more performant than giving users explicit mechanisms for selectively managing control through key press?

4.1 Independent Variables

Given these research questions, three conditions were defined:

- 1: Fixed displays** The control condition, equivalent to the *No mapping:3* condition in Study 1, where the virtual displays were fixed in space
- 2: User control** Here, displays could be shifted left/right in 60° increments based on pressing CTRL + Left/Right arrow keys.
- 3: Boundary switching** This condition was a refinement of the *All deadzones:3* condition found to be most preferable/performant in Study 1.

For *Boundary switching*, this time counter-rotations would be event-driven, triggered when a raycast based on headset orientation intersected a boundary margin on the virtual display ($\pm 5\%$), rather than be based on a mapping between head angle and display counter-rotation angle. This small but significant change in design was due to our need for a fully interactive virtual display. In testing with a fully interactive desktop, the size of the deadzones used in the Study 1 *All deadzones* condition was insufficient, with transitions occurring accidentally when viewing close to the edges/corners of the displays. By instead triggering a discrete shift of $\pm 30^\circ$ at the very edge of each display, we could expand the deadzones to encompass nearly the full size of each virtual display, whilst retaining the same effective physical head rotation required to look at each display i.e. a $\pm 30^\circ$ movement being required to move an angular distance of 60° from the middle of the central display to the middle of the left/right displays.

4.2 Virtual Workspace

Again, in the absence of a high resolution, wide field-of-view AR headset, we utilized a VR headset for rendering the virtual workspace (see Figure 14). This time, we used a Samsung Odyssey Microsoft Mixed Reality VR headset [93], as this provided a substantial increase in



Fig. 14. Three display virtual workspace with positionally tracked keyboard and trackpad, with the holiday planning task applications open and assigned to each display (from left: notepad, browser, pdf viewer).

Expanding The Bounds Of Seated Virtual Workspaces

resolution (1440x1600 pixels per eye). Participants were seated this time on a standard office swivel chair, rather than the fixed seating of Study 1, to be more representative of typical desk seating without restriction. Using a Windows 10 machine, we mirrored desktops using [36] at 60Hz and a resolution of 1366 * 768 per desktop, with the scene itself rendered at 90Hz. We arrived at this resolution during pilot testing, as this maximized the resolution whilst resulting in legible text in VR, with the virtual displays positioned at a comfortable distance from the user, with each display being visible within the field of view of the headset. It should be noted however that this resolution is significantly lower than the typical Windows Desktop resolution for modern day monitors, and thus the amount of text/information density of the virtual displays was decreased compared to existing physical monitors. For performance reasons given the increased resolution and desktop mirroring, we were limited to mirroring three desktops within Unity, with the virtual living room backdrop from Study 1 removed.

For interactivity, we used a prototype fully positionally tracked VR keyboard with hand visualization developed by Logitech and Microsoft [67] (see Figure 15), which our previous testing has shown to allow touch typing performance at approximately 80% of non-VR baseline with no significant differences in error rate when selecting single keys. For cursor control, we used an Apple Magic Trackpad attached to the side of the tracked keyboard, with a 3D model of the trackpad viewable in VR, however no hand/finger visualization was possible due to limitations with the Windows drivers of this device. This was chosen as it allowed for cursor control in the blind - using a mouse for example may have led to participants losing the mouse when transitioning between keyboard and mouse, whereas the



Fig. 15. Closeup of (top) the physical Logitech Mixed Reality keyboard and attached Apple Magic Trackpad and (bottom) virtual equivalents in VR with hand visualization on the keyboard.

trackpad was in a known and fixed position relative to the keyboard. This implementation is state of the art, effectively a best-in-class fully functional VR workspace at the time of writing. See the video figure for footage in action.

The trackpad allowed for two-finger scrolling, and we also implemented a cursor relocation gesture. In pilot testing we identified that some participants would lose track of the cursor position in the virtual workspace. Accordingly, we allowed users to relocate the cursor to wherever their gaze currently intercepted any of the virtual displays by performing a three finger tap on the trackpad. We also triggered the Windows mouse locator (an animated reticule which appears over the mouse) after each relocation. For the *User Control* condition, pressing the CTRL key on the virtual keyboard would visualize shortcuts on the left/right arrow keys indicating that the desktops could be shifted left/right, performing in functionally the same way as desktop shifting in Windows or OS X virtual desktop implementations.

4.3 Task

To motivate using the three virtual displays interchangeably and interactively we employed a holiday planning task, where the rightmost display contained a series of questions prompting about aspects of a trip to plan to a given destination (randomized between Antwerp, Budapest and Lucerne, asking participants to find flights, accommodation and specific tourist attractions), the central display providing a Chrome browser instance for search, and the leftmost display providing a Notepad instance for typing answers/notes. There were 6 questions per condition, with each question containing multiple sub-parts requiring separate answers. This was so that participants would be motivated to frequently switch between displays, with participants instructed to answer the questions as quickly and accurately as possible, and to not fixate on lengthy answers for any one question. For the full question sheets see Supplemental Materials.

4.4 Demographics and Dependent Variables

In total, 18 participants (13 reported male, 5 female, mean age 23.14 years, Std.dev=4.66, each paid £10) took part in the study, being recruited from mailing lists and forums. For qualitative measures, we employed select measures from Study 1, recording a subset of the NASA TLX scales (frustration, mental demand, physical demand and performance), neck fatigue, and perceived discomfort in viewing the peripheral displays, again on the Borg CR10 scale.

For the TLX scales specifically, we asked the TLX questions twice at the end of each condition. The first set were labelled to capture the workload of the holiday planning task, with the second set labelled to capture the workload of switching between different virtual displays. In this way, participants would be forced to rate their workload of the underlying task, and their capability viewing the virtual display space, separately. The first set of results were ignored, and the second set then analysed.

For quantitative measures, we recorded all transitions between the three virtual displays, including the duration of each viewing instance and, in the case of *User Control*, whether the transition was triggered by an explicit user key press or enacted through gaze (in which case no display shifting would have occurred). We also recorded all usage of the mouse cursor relocation feature. At completion of the study, interviews were conducted and rankings were captured along with preferences regarding the mouse relocation feature.

4.5 Results

For all results, a repeated measures ANOVA was performed, with all reporting the same as described in Study 1.

4.5.1 Qualitative Results. For significance test results for the qualitative measures, see Table 4.

TLX subscales. There were moderate significant effects in the physical demand and performance subscales. *Post-hoc* tests suggested the *Boundary* condition was less physically demanding than the *Fixed* condition, and perceived as more performant than both *Fixed* and *User control*, with the distribution of the performance results across participants seen in Figure 16.




Fig. 16. TLX Performance subscale for Study 2

Discomfort and fatigue. There was a moderate significant effect with respect to physical discomfort in viewing the peripheral displays. *Post-hoc* differences were found between the *Fixed* condition and both *Boundary* and *User control*, indicating that providing some capability to manipulate the position of the virtual displays was beneficial in enabling access to the peripheral displays. However there were no effects found with regards to neck fatigue.

Rankings. There was a moderate significant effect on Rankings, with *post-hoc* tests showing a significant difference between *Boundary* and *Fixed*, with *Boundary* being preferred by 12/18 users. The distribution of ranks can be seen in Figure 17.




Fig. 17. User rankings of Conditions for Study 2, with conditions ordered by mean ranking from top to bottom, top being best mean ranking.

4.5.2 Usage Behaviour. Viewing. An unexpected result was found when examining the viewing behaviour exhibited across conditions. The *Boundary* approach led to subtly different viewing behaviour, as can be seen both in Table 5 and Figure 18. After filtering for display viewing instances less than 200ms (to rule out potential repeated transitions at boundary edges), there remained a significant but small difference in the quantity of short display viewing instances: i.e., where a user might look at a given display for less than a few seconds before transitioning. *Post-hoc* tests showed differences between *Boundary* compared

Measure	RM ANOVA			95% CI			<i>Post hoc</i> s		
	F(2,34)	<i>p</i>	η_g^2	<i>Fixed</i>	<i>User</i>	<i>Boundary</i>	Fixed-User	Boundary-Fixed	Boundary-User
Ranking	4.87	=0.01	.16	[1.97, 2.70]	[1.74, 2.48]	[1.19, 1.92]	t=0.71, p=0.76	t=-2.48, p<0.05	t=-1.77, p=0.19
TLX Effort	2.15	=0.13	.07	[7.74, 11.71]	[5.91, 9.87]	[4.96, 8.93]	-	-	-
TLX Frustration	0.95	=0.4	.02	[5.49, 9.51]	[4.27, 8.29]	[5.60, 9.62]	-	-	-
TLX Mental Demand	3.00	=0.06	.07	[5.90, 9.77]	[7.79, 11.65]	[5.35, 9.21]	-	-	-
TLX Physical Demand	5.14	=0.01	.12	[8.68, 13.09]	[6.02, 10.43]	[4.80, 9.21]	t=2.13, p=0.09	t=-3.10, p=0.01	t=-0.98, p=0.60
TLX Performance	4.13	<0.05	.10	[9.68, 13.21]	[9.73, 13.27]	[12.23, 15.77]	t=-0.06, p=0.99	t=2.52, p<0.05	t=2.46, p<0.05
Physical Discomfort	7.12	<0.01	.17	[1.70, 2.75]	[0.75, 1.81]	[0.58, 1.64]	t=3.18, p<0.01	t=-3.73, p<0.01	t=-0.56, p=0.84
Neck Fatigue	2.83	=0.07	.10	[2.69, 4.87]	[1.30, 3.48]	[0.97, 3.14]	-	-	-

Table 4. Quantitative results from TLX and additional questionnaires. TLX workload specifically refers to the workload experienced in switching between viewing the different displays. For TLX Overall and Performance, higher is better, for all other TLX results lower is better. For discomfort and fatigue, higher is worse, with fatigue on the Borg CR10 scale (12 point) and discomfort on a 7-point Likert scale ranging from “No Discomfort” to “Pain”.

Measure	RM ANOVA			95% CI			<i>Post hoc</i> s		
	F(2,34)	<i>p</i>	η_g^2	<i>Fixed</i>	<i>User</i>	<i>Boundary</i>	Fixed-User	Boundary-Fixed	Boundary-User
Mean duration viewing displays	6.72	=0.01	.13	[4.73, 6.31]	[5.97, 7.55]	[6.13, 7.71]	t=-2.85, p<0.05	t=3.23, p<0.01	t=0.38, p=0.92
Mean total transitions between displays	6.93	=0.01	.15	[99.51, 119.72]	[84.89, 105.11]	[79.12, 99.33]	t=2.59, p<0.05	t=3.6, p<0.01	t=-1.02, p=0.57

Table 5. Quantitative statistical test results, examining the duration of each viewing instance of a virtual display, and the number of gaze transitions (i.e. where the display the head orientation raycast hits changes) between virtual displays.

Expanding The Bounds Of Seated Virtual Workspaces




Fig. 18. Viewing behaviour averaged across participants, broken down by duration of instances viewing any of the three displays. Top: cumulative viewing by duration spent viewing the display prior to transitioning to another display. As each condition lasted 10 minutes, there was a maximum cumulative viewing of 600 seconds. Bottom: Histogram of viewing instances by duration. Display position assistance refers to whether discrete display shifts occurred at the point of the transition between displays. For the *User* Condition, this was when the key shortcuts for shifting displays left/right were used. For the *Boundary* Condition, this always occurred as participants would cross through the display Boundary when shifting gaze from one display to another.

to both *Fixed* and *User*, with less short instances occurring in *Boundary*. For the *User control* condition, we can see that the key-press triggered transitions constituted $\approx 35\%$ of all transitions between displays (Std.dev=32%, CI[26.49%, 44.39%]), with the rest occurring without assistance (i.e. by head orientation only).

Mouse cursor relocation. Across conditions, the Mouse cursor relocation feature was invoked 4.29 (Std.Dev=4.96) times on average, predominantly being used for inter-monitor jumps. These statistics are highly dichotomous however, with four participants becoming more heavily reliant on it (invoking it 10+ times per condition), and five participants failing to use the feature at all, despite training and a permanent reminder of how to use the feature being rendered on the touchpad. The feature was generally considered by users to be necessary for effective usage of a wide virtual workspace (mean=4.1, Std.dev=2.01, where 0 was strongly disagree and 6 was strongly agree).

4.6 Interviews

Interviews were coded using the same procedure as described in Study 1.

4.6.1 Boundary Assistance Broadly Preferred. As illustrated in the rankings, the *Boundary* condition was most preferred, with those that preferred it noting it was more comfortable and less fatiguing than the *Fixed* displays. Participants picked up on the implicit aspect of control, with descriptions of it being intuitive as it “just worked every time” (P4), “eliminated the extra input (and) was very natural” (P7), “(was) receptive to what you’re doing” (P13) and “more instinctive” (P15).

However, there were some issues which were repeatedly raised, with two participants noting that initial usage of this condition was “surprising” (P2, P5), with others suggesting that there was a period of adaptation as “I had to get used to where those boundaries were, I already had expectations about where the screens were going to be” (P6). The boundary edges also led to occasional false positive transitions for three participants, for example “if I looked at the corner of the screen it would jump to the other one if I didn’t necessarily want it to” (P10), with it being described as “a bit sensitive” (P7). And one participant noted that, when quickly and repeatedly glancing between displays it could be disorientating, being “really weird, it made me feel like I was not looking at the right thing” (P18).

4.6.2 Preferences for User Control. For those that preferred the user control over display position it was found to be “very quick” (P1) and preferable over the *Fixed* workspace as it was less fatiguing. At least three participants noted that they preferred keyboard shortcuts in day-to-day usage “so it came more instinctively to me” (P15), with preferences appearing to be because users would rather retain complete control over when the displays transitioned, if at all. However, five participants noted that having additional shortcuts was problematic, being either “tedious” (P2) or easy to forget (P2, P5, P6). As P13 noted, “I felt it was more of a job to do it than just moving your head around”, with P3 similarly noting “it made things more complicated”.

4.6.3 Fixed Workspace Uncomfortable, Fatiguing. Participants almost unanimously supported the findings from Study 1 that viewing wide workspaces without assistance was physically uncomfortable. The three display workspace in the *Fixed* condition was noted as being: “quite painful with your neck after a while” (P3), “a strain to look at” (P4) and “very uncomfortable” (P10).

4.6.4 Cursor Relocation A Necessity. The utility of the cursor relocation function was noted by ten participants, with the three finger relocation gesture allowing users to “quickly relocate (the cursor)” (P6) making it “easier to navigate” (P7). However, five participants suggested they repeatedly forgot about the existence of the gesture, despite instruction, training, and a permanent notification rendered on the trackpad.

4.7 Implications for Research Questions

RQ 2.1 Do user preferences for implicit counter-rotation of virtual displays hold when performing a productivity task?

Regarding RQ 2.1, these results re-affirm the benefits of implicit counter-rotation of virtual displays. User rankings significantly favoured the refined approach of the *Boundary* condition, where display transitions were triggered when reaching the edges of the virtual displays, with significant benefits seen in perceived performance, discomfort, and overall workload.

RQ 2.2 To what extent is implicit counter-rotation of virtual displays preferable or more performant than giving users explicit mechanisms for selectively managing control through key press?

Regarding RQ 2.2, whilst rankings and subjective performance results for the majority of participants favoured the *Boundary* condition, there was a small subset of participants that preferred having explicit control over manipulating the position of the virtual displays within the workspace. These preferences appear motivated by a number of points. Firstly, keyboard shortcuts are familiar, with pre-existing virtual desktop controls for example often enacted through keyboard shortcuts. Accordingly, such controls are familiar to some users. Secondly,

Expanding The Bounds Of Seated Virtual Workspaces

there were edge cases in the *Boundary* condition which compromised some participants capability to view virtual displays of their choosing. Interviews suggested that false positive transitions were triggered for at least two participants when attempting to fixate on corners of displays, whilst one noted that in attempting rapid and successive transitions between displays some amount of disorientation occurred.

5 GENERAL DISCUSSION

Across two studies, we have built up a significant body of evidence demonstrating that implicit control of virtual display position, where discrete counter-rotations occur based on head orientation/gaze, is preferable and more performant than fixed display configurations, explicit user control, and continuous rotational gain approaches for counter-rotation. The *Boundary switching* condition of Study 2 refined the deadzone-based approach of Study 1, and represents a significant first step in designing for implicit control, allowing users to navigate a 180° wide workspace within 90° of head rotation. This approach retains a healthy amount of head movement (effectively operating in a $\pm 30^\circ$ envelope for viewing three 60° displays), requires no prior knowledge or training, and does not necessitate explicit user input to enact.

However, that is not to say that we have arrived at a solution that is yet ready for day-to-day usage when AR and VR-driven mixed reality virtual workspaces become common place. Our design and implementation were informed by the literature and the scenarios that motivated this work, but this was a formative investigation of techniques to improve the accessibility of virtual displays. Our results should be interpreted with this and some limitations, below, in mind. This work can be used as a starting point to further investigate these techniques, and there are also other areas of work where our techniques might be useful, which we now discuss.

5.1 Caveats and Limitations

5.1.1 Hardware Characteristics. The VR headsets we used had a $\sim 110^\circ$ field of view. Future iterations of VR headsets are likely to have lower weight and a wider field of view. Increased field of view and increased range of motion from lighter headsets may require adapting assistive mappings. An increased field of view could allow users to see more display space at once, so decreased counter-rotational gain may be necessary or deadzones may need to be narrower. The fundamental benefits of our approaches will still be present, however, and could increase the uptake of VR workspaces in the workplace.

5.1.2 Virtual Display Configurations. We used curved virtual displays that were $\pm 50^\circ$ wide with $\pm 5^\circ$ margins on each side. The displays were an appropriate size for this study and accurately represented the size of physical displays used in productivity environments (albeit they lack the same curvature). Our mappings might need to be adapted for different sized displays, but our work provides a starting point. Different amounts of curvature could also be investigated. We evaluated a curvature that was circular (i.e., every point on each display was equidistant from the pivot point). However, as noted by the Google Daydream team [71], the curvature of a virtual plane may impact its perceived comfort (see Figure 19). They found that a relaxed curvature was preferred by users, being less claustrophobic when dealing with large virtual displays. Changing the curvature may require changes to adaptive mapping.

5.1.3 Duration and Context of Usage. Our study sessions lasted for one hour and participants were seated throughout. Practical use in the workplace may change the nature of the




Fig. 19. Design considerations regarding display curvature from [71]. *Left*: Equidistant curvature, which was suggested to feel claustrophobic. *Right*: Relaxed curvature, with the size of the circle increased and the center point placed behind the user, resulting in a lesser degree of curvature. This was suggested to potentially be perceived as more comfortable.

use of virtual display spaces, as users will interact over much longer durations. As such, our designs may need to be adapted to suit working patterns.

5.1.4 Other Forms of Explicit Control. Whilst we have explored an initial comparison of explicit user control of virtual display position against implicit control based on head orientation/gaze, we have only examined one explicit interaction, namely keyboard shortcuts for rotating displays left/right. However, we cannot rule out the possibility that other explicit mechanisms might have elicited different preferences. For example, mid-air gestures, multi-finger swipes on touchpads, or discrete button presses on mice might all be preferred by different subsets of users.

However, as discussed in the literature review, there are persuasive arguments for implicit control e.g. the potential for reduced cognitive load, the physical benefits of retaining some degree of head/neck movement, preferences for physical navigation over using provided controls and the highly variable adoption that explicit management techniques tend to have, given that controls must be learned etc. Answering such questions more fully will require a significant body of further research, however adoption of technology does not necessarily follow the empirically “best” technique on any given metric. As it stands, we may not be able to establish more concretely the utility/application of such mappings until such time as augmented/virtual reality workspaces see significant adoption, providing a platform on which to perform longitudinal deployments and evaluations.

5.2 Future Research

5.2.1 Refinements To Boundary Switching. In Study 1, we relied on a mapping between head orientation and virtual display counter-rotation. However, this approach brought with it natural limitations in the size of the deadzone on the virtual displays. In Study 2, we changed to instead trigger fixed transitions when at the edge boundary of a virtual display. This allowed us to effectively increase the deadzone size to 90% of the virtual display, with a $\pm 5\%$ boundary. However, even with such a boundary, some false positive transitions were reported by users when looking at the far edges of displays. Accordingly, we can envisage refinements where transitions are triggered when exiting a display/entering the nearest

Expanding The Bounds Of Seated Virtual Workspaces

display for example. Moreover, a small subset of users considered the animated transitions to be disorientating. Different transitions might decrease disorientation, for example by instantly repositioning the displays. Support for rapid and successive transitions between displays needs to also be explored, perhaps by allowing the temporary suspension of the movement of displays. Indeed this issue may be intrinsically linked with the point regarding animated transitions.

5.2.2 *The Role of Gaze.* Throughout this paper, we have effectively used head orientation in lieu of gaze, as this is intrinsically captured by AR/VR headsets currently. However, given eye tracking capability, our boundary approaches for example might operate very differently, with users looking at the edge of a display through a combination of neck and eye movement to trigger a transition. What impact this might have on utility, usability, physical fatigue etc. however has yet to be determined.

5.2.3 *Cursor Management in Wide Virtual Display Spaces.* Our initial testing for Study 2 suggested that assistance in managing the position of the mouse cursor over the display space was required. Qualitative and quantitative results from this study suggested that a significant subset of users found the subsequent cursor relocation gesture we implemented both useful and necessary.

Managing cursors over multiple monitors is a topic that gained some traction over a decade ago with the advent of both affordable multi-monitor workspaces and gaze tracking technology, with notable approaches including warping the cursor to gaze position [19, 122]; transitioning the cursor between displays based on gaze [4, 95] or seat orientation [21]; modulating indirect input based on gaze [111]; employing multiple cursors across displays [11, 57] disambiguated by gaze [88]; transitions across displayless spaces [77] with objects [76]; and even more basic user-controlled cursor centering [40].

The advent of virtual workspaces rendered by MR headsets will necessitate that we re-consider and re-contextualise such research, given we effectively have a display space that is not necessarily linearly continuous (dependent on mapping/technique used), and may also be subdivided into discrete “displays” (be they instances of apps or containers/desktops as in this paper). Currently, we use a proxy for gaze in the form of headset orientation, which we employed in an approach similar to MAGIC touch [19] and the work of Ashdown *et al.* [4], which both used gaze - in our implementation user input relocated the cursor based on the current headset orientation. However, future headsets will inevitably incorporate gaze tracking, at which point a virtual workspace rendered by a gaze-tracked MR headset could effectively implement any of the aforementioned cursor management techniques. Consequently, future research should begin to consider which of these techniques might best facilitate general purpose cursor management in wide virtual display spaces, given this new baseline of MR headset and head orientation.

5.2.4 *Alternate Virtual Display Arrangements and Anchor Points.* This work focused on virtual display arrangements that rotate around the vertical axis at the user’s head. Such displays could also be anchored to other body parts or rotate around different axes. For example, content could be placed vertically (rather than horizontally) [23] or spatially around the user [25] (see Figure 20).

Virtual displays could also be anchored on or around parts of the body [32], allowing users to access information easily in VR. If users are immersed in 3D content, then anchoring displays to a visualisation of the wrist (Figure 21), for example, would allow them to access key information by raising their arm, regardless of their orientation in the virtual world.




Fig. 20. Alternate display configurations suggested by Ens *et al.* *left*: from [23], *right*: from [25].

Our techniques could then be adapted and evaluated for these body-mapped displays. For example, counter-rotational gain might help make more content accessible through a shorter range of motion and deadzones may help mitigate display instability due to unintentional arm movements.




Fig. 21. Virtual displays or controls could also wrap around the wrist for easy access in VR. For example, to allow users to access ‘pinned’ information or frequently used commands

5.2.5 Utility of Different Techniques for Implicit/Explicit Control of Display Position. Whilst we have arrived at a condition (*Boundary switching*) that was preferred and most performant for general purpose usage of an egocentric virtual display space, that is not to say that other approaches we have considered here, or different configurations of specific techniques, might not be better suited for specific contexts. For example, let us consider three scenarios that we would speculate might exhibit markedly different results:

Context 1 - Seated passenger on plane journey: Consider the standard passenger using an MR headset on a long haul flight [117]. Plane seating is typically fixed in place, and there may be social acceptability concerns regarding looking significantly left or right, given that other passengers may feel they are being watched. Consequently, the range of acceptable head movement might be considered limited. In such a temporary context, users might shift preferences toward an explicit control mechanism and accept that a lack of head movement will temporarily be preferable, even if physically less comfortable. Or users might choose to map a smaller display space (e.g. 120° of display) to an increasingly compressed range of movement (e.g. 30° of head movement) by utilizing rotational gain.

Expanding The Bounds Of Seated Virtual Workspaces

Context 2 - Transient roomscale VR workspace: An artist in a roomscale virtual editor instantiates a temporary instance of a horizontal workspace containing various tools and applications. Using a high rotational gain, the artist might quickly navigate a large set of potential menus by looking left or right. When they raise their peripheral or hand to interact with the current menu, the display space is locked in place, preventing any additional display movement whilst interaction occurs. The workspace is then dismissed and they continue with their editing of the virtual scene.

Context 3 - Glance-based reference materials: A student configures their custom virtual workspace with 4 virtual displays, each 60° wide: 1 central, 1 left, and 2 right. The central and left displays both use the same size of boundary as in the *Boundary switching* technique, with the central display containing the word processor, and the leftmost display containing a web browser for searching for information. The two rightmost displays however are dedicated to quick, glance based activities, with the first containing instant messaging applications and email, and the furthest display containing media playback controls for music along with various personal desktop widgets. For both of these displays, the boundary sizes are increased to $\pm 25\%$, meaning that each has a central deadzone of $\pm 15^\circ$. As a consequence, glancing at the first right display would require a 30° movement, whilst moving to the rightmost display would require an additional 15° of movement.

In effect, the student has sacrificed some ability to look at the edges of the two rightmost displays for a capability to transition to these displays at-a-glance within a comfortable range of head movement, pushing the least used applications to the furthest away display. Alternatively, the student might change the width and boundary size of the right side displays, instantiating narrower displays sized to suit the applications contained therein, with larger boundaries to facilitate quick transitions within a limited range of head movement.

What these examples illustrate is that the preferences we captured are likely to change based on the context of usage, the task undertaken, and the kind of viewing to be facilitated on the peripheral displays. Bespoke solutions appropriating, re-configuring, or entirely re-imagining techniques discussed in this paper might lead to better user experiences for specific applications and user groups. Moreover, differences in resilience to simulator/motion sickness and visual/physical fatigue might result in power users having their own preferred configurations. However, we suggest that the *Boundary switching* condition provides a strong foundation, being a general purpose means of expanding the bounds of seated workspaces as demonstrated across our two studies. Our exploration of this space will also help others in creating solutions for specific application areas, with our work representing an initial building block in facilitating a range of new VR/AR interfaces.

5.2.6 Other Benefits of Mixed Reality Workspaces. Finally, this paper has focused predominantly on the underlying mechanics of presenting and accessing virtual displays, in ways that are both efficient and ergonomic. However, assuming the productive worker is wearing a MR headset, we could envisage a multitude of other beneficial virtual changes that could be made to the general working environment. We might for example augment peripherals or the desk surface to enhance functionality and usability [9], create hybrid physical-virtual interfaces [63], communicate through embodied telepresence [26, 70] or alter our physical or virtual environment to encourage well-being [98] and mental health. On the latter point, consider the VR worker that uses their headset to block out reality and its surrounding distractions, much as workers might currently use noise cancelling headphones. Instead, they might alter their virtual context, from the shared virtual meeting room during a conference

call, to a relaxing beach landscape or 360 degree real-time imagery of a remote location (e.g. their home, or a public park), much as Big Screen VR [10] currently facilitates different virtual environments to view content within. Or consider the AR worker whose headset renders virtual partitions to block out distractions [60] and modifies the surrounding physical environment to better match their own particular tastes and interests. Their headset might reflect their own cherished relationships by dynamically rendering images of loved ones in visually salient places throughout their day. Windows to other places or worlds might be rendered in AR in environments where there are no windows (e.g. the cubicle), allowing workers to look out onto a pleasant or familiar environment. Cultural shifts toward working at home may eventually negate the need for such interventions. But if cubicle-based, shared, open-plan offices remain prevalent in the future, MR headsets could be used to improve other aspects of the working experience such as well-being, going beyond functional improvements.

6 CONCLUSIONS

Multi-display workspaces can be beneficial in productivity environments but have several drawbacks. VR/AR headsets are an emerging technology that can be used to create virtual multi-display workspaces, allowing users to multi-task and access more information without the need for physical displays, especially when out of the office or in transit [38, 69, 117]. The use of virtual displays also allows us to manipulate the appearance and behaviour of workspaces in ways that are not possible in the physical world, potentially reducing discomfort and fatigue.

Across two studies, this paper explored how we can better facilitate navigation of wide, egocentric, horizontally oriented collections of virtual displays by manipulating the position of these displays implicitly based on head orientation. We firstly examined how the mapping between head orientation and display position could be manipulated for virtual display spaces oriented around a user. Display counter-rotation based on a dynamic mapping to head angle allows users to access a wider range of display space with less head movement. We evaluated two widths of virtual display space ($\pm 60^\circ$ and $\pm 120^\circ$, for three and five displays, respectively) with three assistive mappings that halved the physical range-of-motion required to view the display space. We found that using and manipulating this mapping helped VR users access wider display spaces whilst minimizing neck fatigue and discomfort. In a follow-up, we then iterated upon the most preferred mapping from our first study, allowing for users to implicitly control the position of displays in a three display workspace by triggering transitions to the nearest display as users looked at the edges of the displays. Compared to both a control condition where displays were fixed in place, and a condition where users could explicitly shift displays left/right by key press, we again found significant benefits for our approach in terms of user preference, workload and comfort.

Implicit control of virtual display position provides a new justification for the adoption of VR/AR headsets for productivity in the future. Doing so can reduce the physical requirements of using a given size of display space, and expand the available display space given a physical range to work within, taking into account physical capabilities, ergonomics, and the current environment and tasks. Moreover, such control could be applied to any VR or AR headset with rotational tracking, and could potentially improve productivity when seated at a desk or in physically restricted environments such as planes or cars [38, 69, 117]. Given headsets such as the Varjo VR-2 [108] (which combines a high-resolution center panel for focused detailed work with a lower resolution panel for peripheral vision) and the Microsoft HoloLens 2 [73] (with resolution catching up with consumer VR headsets) it is becoming increasingly feasible to conduct text-heavy work in both VR and AR. This research can

Expanding The Bounds Of Seated Virtual Workspaces

help Mixed Reality headset users to view wider display spaces than are currently possible within acceptable ranges of neck/head movement, and demonstrates a novel use of VR/AR headsets for productivity, creating possibilities for new working environments and practices.

7 ACKNOWLEDGEMENTS

This research was funded in-part by an EPSRC IAA project (EP/R511705/1) and a joint IAA project funded by the EPSRC (303740) and ESRC (77563/1). This research also received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement #835197: ViAJeRo).

REFERENCES

- [1] James A. Anderson, Jennifer Hill, Paul Parkin, and Autumn Garisson. 2007. Productivity, screens, and aspect ratios; A comparison of single, traditional aspect, dual, traditional aspect and single, widescreen aspect computer. *Illinois: NEC Display Solutions* (2007).
- [2] Christopher Andrews, Alex Endert, and Chris North. 2010. Space to Think: Large, High-Resolution Displays for Sensemaking. In *Proceedings of the 28th international conference on Human factors in computing systems - CHI '10*. ACM Press, New York, New York, USA, 55. <https://doi.org/10.1145/1753326.1753336>
- [3] Christopher Andrews and Chris North. 2012. Analyst's Workspace: An embodied sensemaking environment for large, high-resolution displays. In *2012 IEEE Conference on Visual Analytics Science and Technology (VAST)*. IEEE, 123–131. <https://doi.org/10.1109/VAST.2012.6400559>
- [4] Mark Ashdown, Kenji Oka, and Yoichi Sato. 2005. Combining Head Tracking and Mouse Input for a GUI on Multiple Monitors. In *CHI '05 Extended Abstracts on Human Factors in Computing Systems (CHI EA '05)*. ACM, New York, NY, USA, 1188–1191. <https://doi.org/10.1145/1056808.1056873>
- [5] Robert Ball and Chris North. 2005. Effects of tiled high-resolution display on basic visualization and navigation tasks. In *CHI '05 extended abstracts on Human factors in computing systems - CHI '05*. ACM Press, New York, New York, USA, 1196. <https://doi.org/10.1145/1056808.1056875>
- [6] Robert Ball, Chris North, and Doug A Bowman. 2007. Move to Improve : Promoting Physical Navigation to Increase User Performance with Large Displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI'07*. 191–200. <https://doi.org/10.1145/1240624.1240656>
- [7] Borbála Berki. 2019. Desktop VR as a Virtual Workspace: a Cognitive Aspect. *Acta Polytechnica Hungarica* 16, 2 (2019). <https://doi.org/10.12700/APH.16.2.2019.2.13>
- [8] Xiaojun Bi and Ravin Balakrishnan. 2009. Comparing Usage of a Large High-Resolution Display to Single or Dual Desktop Displays for Daily Work. In *CHI '09 Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1005–1014. <https://doi.org/10.1145/1518701.1518855>
- [9] Xiaojun Bi, Tovi Grossman, Justin Matejka, and George Fitzmaurice. 2011. Magic desk: bringing multi-touch surfaces into desktop work. In *Proceedings of the 2011 annual conference on Human factors in computing systems - CHI '11*. 2511. <https://doi.org/10.1145/1978942.1979309>
- [10] Bigscreen Inc. 2017. Bigscreen Beta. *steampowered.com* (2017). <http://store.steampowered.com/app/457550/Bigscreen>
- [11] Renaud Blanch and Michaël Ortega. 2009. Rake Cursor: Improving Pointing Performance with Concurrent Input Channels. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09)*. ACM, New York, NY, USA, 1415–1418. <https://doi.org/10.1145/1518701.1518914>
- [12] Gunnar Borg. 1998. Borg's perceived exertion and pain scales. *Human Kinetics* (1998), 104 vii. <https://doi.org/10.1097/00005768-199809000-00018>
- [13] J. Carlton Collins. 2011. Increase Productivity With Multiple Monitors. (2011). <http://www.journalofaccountancy.com/issues/2011/feb/20103260.html>
- [14] Mark Coppock. 2017. How to split your screen and manage your work in Windows 10. (2017). digitaltrends.com/computing/how-to-split-your-screen-in-windows-10/
- [15] Mary Czerwinski, George Robertson, Brian Meyers, Greg Smith, Daniel Robbins, and Desney Tan. 2006. Large display research overview. *CHI '06 extended abstracts* (2006), 69. <https://doi.org/10.1145/1125451.1125471>
- [16] Mary Czerwinski, Greg Smith, Tim Regan, Brian Meyers, George Robertson, and Gary Starkweather. 2003. Toward Characterizing the Productivity Benefits of Very Large Displays. *Interact* 3 (2003), 9–16. <http://research.microsoft.com/en-us/um/redmond/groups/cue/publications/>

interact2003-productivitylargedisplays.pdf

- [17] Jakub Dostal, Per Ola Kristensson, and Aaron Quigley. 2013. Subtle gaze-dependent techniques for visualising display changes in multi-display environments. *Proceedings of the 2013 international conference on Intelligent user interfaces - IUI '13* (2013), 137. <https://doi.org/10.1145/2449396.2449416> arXiv:arXiv:1011.1669v3
- [18] Steve Draper. 2018. Effect size. (2018). <http://www.psy.gla.ac.uk/~steve/best/effect.html>
- [19] Heiko Drewes and Albrecht Schmidt. 2009. The MAGIC Touch: Combining MAGIC-Pointing with a Touch-Sensitive Mouse. In *Human-Computer Interaction – INTERACT 2009*. Springer Berlin Heidelberg, Berlin, Heidelberg, 415–428. https://doi.org/10.1007/978-3-642-03658-3_46
- [20] Alex Endert, Lauren Bradel, Jessica Zeitz, Christopher Andrews, Chris North, and Virginia Tech. 2012. Designing Large High-Resolution Display Workspaces. *AVI '12 Proceedings of the International Working Conference on Advanced Visual Interfaces* (2012), 58–65. <https://doi.org/10.1145/2254556.2254570>
- [21] Alex Endert, Patrick Fiaux, Haeyong Chung, Michael Stewart, Christopher Andrews, and Chris North. 2011. ChairMouse: Leveraging Natural Chair Rotation for Cursor Navigation on Large, High-resolution Displays. In *CHI '11 Extended Abstracts on Human Factors in Computing Systems (CHI EA '11)*. ACM, New York, NY, USA, 571–580. <https://doi.org/10.1145/1979742.1979628>
- [22] Barrett Ens, Juan David Hincapié-Ramos, and Pourang Irani. 2014. Ethereal Planes: A Design Framework for 2D Information Space in 3D Mixed Reality Environments. *Proceedings of the 2nd ACM Symposium on Spatial User Interaction* (2014), 2–12. <https://doi.org/10.1145/2659766.2659769>
- [23] Barrett Ens and Pourang Irani. 2017. Spatial Analytic Interfaces: Spatial User Interfaces for In Situ Visual Analytics. *IEEE Computer Graphics and Applications* 37, 2 (mar 2017), 66–79. <https://doi.org/10.1109/MCG.2016.38>
- [24] Barrett Ens, Eyal Ofek, Neil Bruce, and Pourang Irani. 2016. Shared Façades: Surface-Embedded Layout Management for Ad Hoc Collaboration Using Head-Worn Displays. In *Collaboration Meets Interactive Spaces*. Springer International Publishing, Cham, 153–176. https://doi.org/10.1007/978-3-319-45853-3_8
- [25] Barrett M. Ens, Rory Finnegan, and Pourang P. Irani. 2014. The personal cockpit: A Spatial Interface for Effective Task Switching on Head-Worn Displays. *Proceedings of the 32nd annual ACM conference on Human factors in computing systems - CHI '14* August 2016 (2014), 3171–3180. <https://doi.org/10.1145/2556288.2557058>
- [26] Sean Fanello, Sergio Orts-escolano Christoph Rhemann, Mingsong Dou, Vladimir Tankovich, Charles Loop, and Philip Chou. 2016. Holoportation : Virtual 3D Teleportation in Real-time. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology (UIST '16)*. 741–754. <https://doi.org/10.1145/2984511.2984517>
- [27] Amanda M. Farias Zuniga and Julie N. Côté. 2017. Effects of Dual Monitor Computer Work Versus Laptop Work on Cervical Muscular and Proprioceptive Characteristics of Males and Females. *Human Factors* 59, 4 (2017), 546–563. <https://doi.org/10.1177/0018720816684690>
- [28] Hilary L. Gallagher, Erin Caldwell, and Christopher B. Albery. 2008. *Neck muscle fatigue resulting from prolonged wear of weighted helmets*. Technical Report. 1–27 pages. <http://www.dtic.mil/docs/citations/ADA491626>
- [29] Fernando Garcia-Sanjuan, Javier Jaen, and Vicente Nacher. 2016. Toward a General Conceptualization of Multi-Display Environments. *Frontiers in ICT* 3 (sep 2016), 20. <https://doi.org/10.3389/fict.2016.00020>
- [30] Thomas Geymayer, Manuela Waldner, Alexander Lex, and Dieter Schmalstieg. 2017. How Sensemaking Tools Influence Display Space Usage. *EuroVis Workshop on Visual Analytics* (apr 2017). <https://doi.org/10.2312/eurova.20171112> arXiv:1704.03949
- [31] Guy Godin. 2017. Virtual Desktop. (2017). store.steampowered.com/app/382110/Virtual
- [32] Jens Grubert, Matthias Heinisch, Aaron Quigley, and Dieter Schmalstieg. 2015. MultiFi : Multi-Fidelity Interaction with Displays On and Around the Body. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI'15* (2015), 3933–3942. <https://doi.org/10.1145/2702123.2702331>
- [33] Jonathan Grudin. 2001. Partitioning Digital Worlds: Focal and Peripheral Awareness in Multiple Monitor Use. *Proceedings of the SIGCHI conference on Human factors in computing systems - CHI '01* 3 (2001), 458–465. <https://doi.org/10.1145/365024.365312>
- [34] Peter A. Hancock, Ben D. Sawyer, and Shawn c. Stafford. 2015. The effects of display size on performance. *Ergonomics* 58, 3 (mar 2015), 337–354. <https://doi.org/10.1080/00140139.2014.973914>
- [35] Sandra G. Hart and Lowell E Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. In *Human mental workload*. <http://humanfactors.arc.nasa.gov/>

Expanding The Bounds Of Seated Virtual Workspaces

groups/TLX/downloads/NASA-TLXChapter.pdf

- [36] Hecomi. 2019. Desktop Duplication API implementation for Unity. (2019). <https://github.com/hecomi/uDesktopDuplication>
- [37] Jerry L. Hintze and Ray D. Nelson. 1998. Violin plots: a box plot-density trace synergism. *The American Statistician* (1998). <http://amstat.tandfonline.com/doi/abs/10.1080/00031305.1998.10480559>
- [38] Philipp Hock, Sebastian Benedikter, Jan Gugenheimer, and Enrico Rukzio. 2017. CarVR: Enabling In-Car Virtual Reality Entertainment. (2017). <https://doi.org/10.1145/3025453.3025665>
- [39] Raphael Hoffmann, Patrick Baudisch, and Daniel S Weld. 2008. Evaluating visual cues for window switching on large screens. In *Proceedings of ACM CHI 2008 Conference on Human Factors in Computing Systems*, Vol. 1. 929–938. <https://doi.org/10.1145/1357054.1357199>
- [40] Nic Hollinworth and Faustina Hwang. 2011. Cursor Relocation Techniques to Help Older Adults Find 'Lost' Cursors. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. ACM, New York, NY, USA, 863–866. <https://doi.org/10.1145/1978942.1979068>
- [41] Seokjun Hong and Gerard J Kim. 2016. Accelerated viewpoint panning with rotational gain in 360 degree videos. In *Proceedings of the 22nd ACM Conference on Virtual Reality Software and Technology - VRST '16*. 303–304. <https://doi.org/10.1145/2993369.2996309>
- [42] Cory House. 2017. Why I Stopped Using Multiple Monitors. (2017). <https://hackernoon.com/why-i-stopped-using-multiple-monitors-bfd87efa2e5b>
- [43] Peter Alan Howarth and Howell Owen Istance. 2007. The association between visual discomfort and the use of visual display units. *Behaviour & Information Technology* 4, 2 (apr 2007), 131–149. <https://doi.org/10.1080/01449298508901794>
- [44] Dugald R. Hutchings and Mary Czerwinski. 2004. Exploring the use and affordances of multiple display environments. In *Workshop on Ubiquitous Display Environments at UbiComp*. 1–6. <http://www.academia.edu/download/43905386/Exploring>
- [45] Dugald R. Hutchings, Greg Smith, Brian Meyers, Mary Czerwinski, and George Robertson. 2004. Display Space Usage and Window Management Operation Comparisons between Single Monitor and Multiple Monitor Users. *Proceedings of the Working Conference on Advanced Visual Interfaces (AVI'04)* (2004), 32–39. <https://doi.org/10.1145/989863.989867>
- [46] Dugald R. Hutchings and John Stasko. 2004. Revisiting Display Space Management: Understanding Current Practice to Inform Next-generation Design. In *Proceedings of the Conference on Graphics, Visualization & HCI (GI'04)*, Vol. 62. 127–134. <http://dl.acm.org/citation.cfm?id=1006058.1006074>
- [47] Dugald R. Hutchings and John Stasko. 2007. Consistency, Multiple Monitors, and Multiple Windows. In *Proceedings of the SIGCHI conference on Human factors in computing systems - CHI '07*. 211. <https://doi.org/10.1145/1240624.1240658>
- [48] P. M. Jaekl, M. R. Jenkin, and Laurence R. Harris. 2005. Perceiving a stable world during active rotational and translational head movements. *Experimental Brain Research* 163, 3 (jun 2005), 388–399. <https://doi.org/10.1007/s00221-004-2191-8>
- [49] Steven Jeuris and Jakob E. Bardram. 2016. Dedicated workspaces: Faster resumption times and reduced cognitive load in sequential multitasking. *Computers in Human Behavior* 62 (sep 2016), 404–414. <https://doi.org/10.1016/j.chb.2016.03.059>
- [50] James C. Johnston, Eric Ruthruff, and M.-C. Lien. 2014. Visual Information Processing From Multiple Displays. *Human Factors: The Journal of the Human Factors and Ergonomics Society* 57, 2 (mar 2014), 276–297. <https://doi.org/10.1177/0018720814545974>
- [51] Youn-ah Kang and John Stasko. 2008. *Lightweight task/application performance using single versus multiple monitors: a comparative study*. 17–24 pages. <https://doi.org/10.1145/1056808.1056875>
- [52] Abhijit Karnik, Tovi Grossman, and Sriram Subramanian. 2013. Comparison of User Performance in Mixed 2D-3D Multi-Display Environments. Springer, Berlin, Heidelberg, 260–277. https://doi.org/10.1007/978-3-642-40483-2_18
- [53] Kevin Kelly. 2016. The Untold Story of Magic Leap, the World's Most Secretive Startup. (2016). <https://www.wired.com/?p=1999666>
- [54] Robert S. Kennedy, Norman E. Lane, Kevin S. Berbaum, and Michael G. Lienthal. 1993. Simulator Sickness Questionnaire: An Enhanced Method for Quantifying Simulator Sickness. *The International Journal of Aviation Psychology* 3, 3 (jul 1993), 203–220. https://doi.org/10.1207/s15327108ijap0303_3
- [55] Hermann Klinke, Christoph Krieger, and Sebastian Pickl. 2014. Analysis of the influence of screen size and resolution on work efficiency. (jan 2014). <https://doi.org/10.18419/OPUS-3489>
- [56] Søren Knudsen, Mikkel Rønne Jakobsen, and Kasper Hornbæk. 2012. An Exploratory Study of How Abundant Display Space May Support Data Analysis. In *Proceedings of the 7th Nordic Conference*

- on *Human-Computer Interaction Making Sense Through Design - NordiCHI '12*. 558–567. <https://doi.org/10.1145/2399016.2399102>
- [57] Masatomo Kobayashi and Takeo Igarashi. 2008. Ninja Cursors: Using Multiple Cursors to Assist Target Acquisition on Large Screens. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*. ACM, New York, NY, USA, 949–958. <https://doi.org/10.1145/1357054.1357201>
- [58] Martin Krzywinski and Naomi Altman. 2014. Points of Significance: Visualizing samples with box plots. *Nature Methods* 11, 2 (jan 2014), 119–120. <https://doi.org/10.1038/nmeth.2813>
- [59] Scott A. Kuhl, Sarah H. Creem-Regehr, and William B. Thompson. 2008. Recalibration of rotational locomotion in immersive virtual environments. *ACM Transactions on Applied Perception* 5, 3 (aug 2008), 1–11. <https://doi.org/10.1145/1402236.1402241>
- [60] Hyelip Lee, Seungwoo Je, Rachel Kim, Himanshu Verma, Hamed Alavi, and Andrea Bianchi. 2019. Partitioning open-plan workspaces via augmented reality. *Personal and Ubiquitous Computing* (02 Oct 2019). <https://doi.org/10.1007/s00779-019-01306-0>
- [61] Hae Jin Lee. 2017. Taking the Leap from 2D to 3D. (2017). <https://medium.com/microsoft-design/taking-the-leap-from-2d-to-3d-4199ccc777a1>
- [62] Russell V. Lenth. 2016. Least-squares means: the R package lsmeans. *jstatsoft.org* (2016). <https://www.jstatsoft.org/article/view/v069i01/v69i01.pdf>
- [63] Zhen Li, Michelle Annett, Ken Hinckley, Karan Singh, and Daniel Wigdor. 2019. 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)*. ACM, New York, NY, USA, Article 687, 14 pages. <https://doi.org/10.1145/3290605.3300917>
- [64] David Lindlbauer, Klemen Liliija, Robert Walter, and Jörg Müller. 2016. Influence of Display Transparency on Background Awareness and Task Performance. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 1705–1716. <https://doi.org/10.1145/2858036.2858453>
- [65] Chen Ling, Alex Stegman, Chintan Barhbaya, and Randa Shehab. 2017. Are Two Better Than One? A Comparison Between Single- and Dual-Monitor Work Stations in Productivity and User’s Windows Management Style. *International Journal of Human-Computer Interaction* 33, 3 (mar 2017), 190–198. <https://doi.org/10.1080/10447318.2016.1231392>
- [66] Lars Lischke, Sven Mayer, Katrin Wolf, Niels Henze, Harald Reiterer, and Albrecht Schmidt. 2016. Screen arrangements and interaction areas for large display work places. *Proceedings of the 5th ACM International Symposium on Pervasive Displays - PerDis '16* (2016), 228–234. <https://doi.org/10.1145/2914920.2915027>
- [67] Logitech, Inc. 2019. Microsoft and Logitech MR keyboard SDK. (2019). https://github.com/Logitech/labs_mrkeyboard_sdk
- [68] Gloria Mark, Daniela Gudith, and Ulrich Klocke. 2008. The Cost of Interrupted Work : More Speed and Stress. *CHI 2008: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2008), 107–110. <https://doi.org/10.1145/1357054.1357072>
- [69] Mark McGill, Alexander Ng, and Stephen A. Brewster. 2017. How Visual Motion Cues Can Influence Sickness For In-Car VR. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '17)*. ACM, New York, NY, USA, 469–469. <https://doi.org/10.1145/3027063.3049790>
- [70] Mark McGill, John H. Williamson, and Stephen Brewster. 2016. Examining The Role of Smart TVs and VR HMDs in Synchronous At-a-Distance Media Consumption. *ACM TOCHI* (2016). <https://doi.org/10.1145/2983530>
- [71] Chris McKenzie and Adam Glazier. 2017. Designing screen interfaces for VR. (2017). <https://www.youtube.com/watch?v=ES9jArHRFHQ>
- [72] Microsoft. 2017. Mixed Reality Design. (2017). <https://developer.microsoft.com/en-us/windows/mixed-reality/design>
- [73] Microsoft. 2019. HoloLens 2. (2019). <https://www.microsoft.com/en-gb/hololens>
- [74] Paul Milgram and Fumio Kishino. 1994. A Taxonomy of Mixed Reality Visual Displays. *IEICE Transactions on Information and Systems* 77, 12 (1994), 1321—1329. <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.102.4646>
- [75] Scott Mitchell. 2010. Two Monitors Are Better Than One, But Are Three Better Than Two? (2010). <http://scottonwriting.net/sowblog/archive/2008/07/10/163293.aspx>
- [76] Miguel A Nacenta, Carl Gutwin, Dzmitry Aliakseyeu, and Sriram Subramanian. 2009. There and back again: Cross-display object movement in multi-display environments. *Human-Computer Interaction*

Expanding The Bounds Of Seated Virtual Workspaces

24, 1-2 (2009), 170–229. <https://doi.org/10.1080/07370020902819882>

- [77] Miguel A. Nacenta, Regan L. Mandryk, and Carl Gutwin. 2008. Targeting Across Displayless Space. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*. ACM, New York, NY, USA, 777–786. <https://doi.org/10.1145/1357054.1357178>
- [78] Timothy Neate, Matt Jones, and Michael Evans. 2017. Cross-device media: a review of second screening and multi-device television. *Personal and Ubiquitous Computing* 21, 2 (apr 2017), 391–405. <https://doi.org/10.1007/s00779-017-1016-2>
- [79] Ashish D. Nimbarte, Rabab T. Alabdulmohsen, Steve E. Guffey, and John R. Etherton. 2013. The Impact of Use of Dual Monitor Screens on 3D Head-Neck Posture and Activity of Neck Muscles. *IIE Transactions on Occupational Ergonomics and Human Factors* 1, 1 (jan 2013), 38–49. <https://doi.org/10.1080/21577323.2012.667051>
- [80] Justin W Owens, Jennifer Teves, Bobby Nguyen, Amanda Smith, Mandy C Phelps, and Barbara S Chaparro. 2012. Examination of Dual vs. Single Monitor Use during Common Office Tasks. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 56, 1 (2012), 1506–1510. <https://doi.org/10.1177/1071181312561299>
- [81] Sungryul Park, Donghee Choi, Jihhyeon Yi, Seungbae Songil Lee, Ja Eun Lee, Byeonghwa Choi, Seungbae Songil Lee, and Gyouhyung Kyung. 2017. Effects of display curvature, display zone, and task duration on legibility and visual fatigue during visual search task. *Applied Ergonomics* 60 (apr 2017), 183–193. <https://doi.org/10.1016/j.apergo.2016.11.012>
- [82] Gunnel Peterson, David Nilsson, Johan Trygg, Deborah Falla, Åsa Dederig, Thorne Wallman, and Anneli Peolsson. 2015. Novel insights into the interplay between ventral neck muscles in individuals with whiplash-associated disorders. *Scientific Reports* 5 (2015), 15289. <https://doi.org/10.1038/srep15289>
- [83] Phyllida. 2017. Multiscreens VR. (2017). store.steampowered.com/app/512400/Multiscreens/
- [84] Herbert L. Pick Jr, John J. Rieser, Douglas Wagner, and Anne E. Garing. 1999. The recalibration of rotational locomotion. *Journal of Experimental Psychology: Human Perception and Performance* 25, 5 (1999), 1179–1188. <https://doi.org/10.1037/0096-1523.25.5.1179>
- [85] Kathrin Probst, David Lindlbauer, Florian Perteneder, Michael Haller, Bernhard Schwartz, and Andreas Schrempf. 2013. Exploring the Use of Distributed Multiple Monitors Within an Activity-Promoting Sit-and-Stand Office Workspace. *Human-Computer Interaction - INTERACT 2013* (2013), 476–493. <https://link.springer.com/chapter/10.1007/978-3-642-40477-1>
- [86] Pygmal Technologies. 2017. SPACE. (2017). <https://getspace.io/>
- [87] Roman Rädle, Hans-Christian Jetter, Simon Butscher, and Harald Reiterer. 2013. The effect of egocentric body movements on users' navigation performance and spatial memory in zoomable user interfaces. *Proceedings of the 2013 ACM international conference on Interactive tabletops and surfaces - ITS '13* (2013), 23–32. <https://doi.org/10.1145/2512349.2512811>
- [88] Kari-Jouko Rähkä and Oleg Špakov. 2009. Disambiguating Ninja Cursors with Eye Gaze. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09)*. ACM, New York, NY, USA, 1411–1414. <https://doi.org/10.1145/1518701.1518913>
- [89] Umar Rashid, Miguel A Nacenta, and Aaron Quigley. 2012. Factors influencing visual attention switch in multi-display user interfaces. In *Proceedings of the 2012 International Symposium on Pervasive Displays - PerDis '12*. 1–6. <https://doi.org/10.1145/2307798.2307799>
- [90] Sharif Razzaque, Zachariah Kohn, and Mary C. Whitton. 2001. Redirected Walking. In *Eurographics 2001 - Short Presentations*. Eurographics Association. <https://doi.org/10.2312/egs.20011036>
- [91] Mikkel Rønne Jakobsen and Kasper Hornbæk. 2011. Sizing up visualizations. In *Proceedings of the 2011 annual conference on Human factors in computing systems - CHI '11*. ACM Press, New York, New York, USA, 1451. <https://doi.org/10.1145/1978942.1979156>
- [92] Johnny Saldaña. 2015. *The coding manual for qualitative researchers*. Sage.
- [93] Samsung.com. 2018. Samsung Odyssey Mixed Reality HMD. (2018). <https://www.samsung.com/us/computing/hmd/windows-mixed-reality/xe800zaa-hc1us-xe800zaa-hc1us/>
- [94] Joey Scarr, Andy Cockburn, and Carl Gutwin. 2012. Supporting and Exploiting Spatial Memory in User Interfaces. *Foundations and Trends in Human-Computer Interaction* 6, 1 (2012), 1–84. <https://doi.org/10.1561/1100000046>
- [95] Marcos Serrano, Barrett Ens, Xing-Dong Yang, and Pourang Irani. 2015. Gluey: Developing a Head-Worn Display Interface to Unify the Interaction Experience in Distributed Display Environments. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '15)*. ACM, New York, NY, USA, 161–171. <https://doi.org/10.1145/2785830.2785838>

- [96] Frank M Shipman, Catherine C Marshall, and Thomas P Moran. 1995. Finding and Using Implicit Structure in Human-organized Spatial Layouts of Information. In *Proceedings of the International Conference on Human Factors in Computing Systems (CHI'95)*. 346–353. <https://doi.org/10.1145/223904.223949>
- [97] Lauren Shupp, Christopher Andrews, Margaret Dickey-Kurdziolek, Beth Yost, and Chris North. 2009. Shaping the Display of the Future: The Effects of Display Size and Curvature on User Performance and Insights. *Human-Computer Interaction* 24, 1-2 (apr 2009), 230–272. <https://doi.org/10.1080/07370020902739429>
- [98] Kate Sparks, Brian Faragher, and Cary L Cooper. 2001. Well-being and occupational health in the 21st century workplace. *Journal of occupational and organizational psychology* 74, 4 (2001), 489–509.
- [99] Maximilian Speicher, Brian D. Hall, and Michael Nebeling. 2019. What is Mixed Reality?. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*. ACM Press, New York, New York, USA, 1–15. <https://doi.org/10.1145/3290605.3300767>
- [100] Matthew J Stainer, Kenneth C Scott-Brown, Benjamin W Tatler, and Kenneth C Scott-Brown. 2017. On the Factors Causing Processing Difficulty of Multiple-Scene Displays. *i-Perception* 8, 2 (2017). <https://doi.org/10.1177/2041669516689572>
- [101] Ramona E Su and Brian P Bailey. 2005. Put Them Where? Towards Guidelines for Positioning Large Displays in Interactive Workspaces. *LNCS* 3585 (2005), 337–349. https://doi.org/10.1007/11555261_29
- [102] Grace P.Y. Szeto, Carey C.Y. Chan, Simon K.M. Chan, H.Y. Lai, and Emily P.Y. Lau. 2014. The effects of using a single display screen versus dual screens on neck-shoulder muscle activity during computer tasks. *International Journal of Industrial Ergonomics* 44, 3 (may 2014), 460–465. <https://doi.org/10.1016/j.ergon.2014.01.003>
- [103] Grace P.Y. Szeto, Carey C.Y. Chan, Simon K.M. Chan, H.Y. Lai, and Emily P.Y. Lau. 2014. The effects of using a single display screen versus dual screens on neck-shoulder muscle activity during computer tasks. *International Journal of Industrial Ergonomics* 44, 3 (may 2014), 460–465. <https://doi.org/10.1016/j.ergon.2014.01.003>
- [104] Kazuki Takashima, Takafumi Oyama, Yusuke Asari, Ehud Sharlin, Saul Greenberg, and Yoshifumi Kitamura. 2016. Study and Design of a Shape-Shifting Wall Display. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems (DIS '16)*. ACM, New York, NY, USA, 796–806. <https://doi.org/10.1145/2901790.2901892>
- [105] Desney Tan and Mary Czerwinski. 2003. Effects of Visual Separation and Physical Discontinuities when Distributing Information across Multiple Displays. *Proceedings Interact'03* (2003), 252–259. <http://research.microsoft.com/en-us/um/people/desney/publications/interact2003short-visualseparation.pdf>
- [106] Desney Tan, Darren Gergle, Peter Scupelli, and Randy Pausch. 2006. Physically large displays improve performance on spatial tasks. *ACM Transactions on Computer-Human Interaction* 13, 1 (2006), 71–99. <https://doi.org/10.1145/1143518.1143521>
- [107] Tao Ni, Greg S. Schmidt, Oliver G. Staadt, Mark A. Livingston, Robert Ball, and Richard May. 2006. A Survey of Large High-Resolution Display Technologies, Techniques, and Applications. In *IEEE Virtual Reality Conference (VR 2006)*. IEEE, 223–236. <https://doi.org/10.1109/VR.2006.20>
- [108] Varjo Technologies. 2019. Varjo VR-2. (2019). <https://varjo.com/products/vr-2/>
- [109] The Data Visualisation Catalogue. 2019. Density Plot. (2019). https://datavizcatalogue.com/methods/density_plot.html
- [110] Alvin R. Tilley and Henry Dreyfuss Associates. 2002. *The measure of man and woman: human factors in design*. John Wiley & Sons.
- [111] Simon Voelker, Andrii Matvienko, Johannes Schöning, and Jan Borchers. 2015. Combining Direct and Indirect Touch Input for Interactive Workspaces Using Gaze Input. In *Proceedings of the 3rd ACM Symposium on Spatial User Interaction (SUI '15)*. ACM, New York, NY, USA, 79–88. <https://doi.org/10.1145/2788940.2788949>
- [112] James Walker. 2017. *Improving everyday computing tasks with head-mounted displays*. Ph.D. Dissertation. Michigan Technological University. <http://digitalcommons.mtu.edu/etdr/388>
- [113] Jim Wallace, Vicki Ha, Ryder Ziola, and Kori Inkpen. 2006. Swordfish: user tailored workspaces in multi-display environments. (2006), 1487 pages. <https://doi.org/10.1145/1125451.1125724>
- [114] James R Wallace, Nancy Iskander, and Edward Lank. 2016. Creating Your Bubble: Personal Space On and Around Large Public Displays. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (2016), 2087–2092. <https://doi.org/10.1145/2858036.2858118>

Expanding The Bounds Of Seated Virtual Workspaces

- [115] Andrew Warr, Ed H Chi, Helen Harris, Alexander Kuscher, Jenn Chen, Robert Flack, and Nicholas Jitkoff. 2016. Window Shopping: A Study of Desktop Window Switching. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 3335–3338. <https://doi.org/10.1145/2858036.2858526>
- [116] Betsy Williams, Gayathri Narasimham, Bjoern Rump, Timothy P. McNamara, Thomas H. Carr, John Rieser, and Bobby Bodenheimer. 2007. Exploring large virtual environments with an HMD when physical space is limited. In *Proceedings of the 4th symposium on Applied perception in graphics and visualization - APGV '07*. ACM Press, New York, New York, USA, 41. <https://doi.org/10.1145/1272582.1272590>
- [117] Julie R. Williamson, Mark McGill, and Khari Outram. 2019. PlaneVR: Social Acceptability of Virtual Reality for Aeroplane Passengers. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. ACM, New York, NY, USA, Article 80, 14 pages. <https://doi.org/10.1145/3290605.3300310>
- [118] Hanno Wirtz, Jan R uth, Martin Serror, J  Agila Bitsch Link, and Klaus Wehrle. 2014. Opportunistic interaction in the challenged internet of things. *Proceedings of the 9th ACM MobiCom workshop on Challenged networks - CHANTS '14* (2014), 7–12. <https://doi.org/10.1145/2645672.2645679>
- [119] Viktoria Witka. 2016. *Perceptual Issues in Multi-Display Environments*. Technical Report. <http://www.eislab.fim.uni-passau.de/files/publications/2016/mmd>
- [120] Jacob O. Wobbrock, Leah Findlater, Darren Gergle, and James J. Higgins. 2011. The aligned rank transform for nonparametric factorial analyses using only anova procedures. In *Proceedings of the 2011 annual conference on Human factors in computing systems - CHI '11*. ACM Press, New York, New York, USA, 143. <https://doi.org/10.1145/1978942.1978963>
- [121] Hoi Chi Woo, Peter. White, and Christopher W. K. Lai. 2016. Ergonomics standards and guidelines for computer workstation design and the impact on users' health - a review. *Ergonomics* 59, 3 (mar 2016), 464–475. <https://doi.org/10.1080/00140139.2015.1076528>
- [122] Shumin Zhai, Carlos Morimoto, and Steven Ihde. 1999. Manual and Gaze Input Cascaded (MAGIC) Pointing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '99)*. ACM, New York, NY, USA, 246–253. <https://doi.org/10.1145/302979.303053>
- [123] Amanda M.F. Zuniga. 2015. *Muscular, visual and proprioceptive outcomes of computer work with one versus two computer monitors*. Ph.D. Dissertation. http://digitool.library.mcgill.ca/webclient/StreamGate?folder_id=0&dvs=1571847232431-956