# Understanding the Co-Dependence of Cursor Warping and Boundary Switching in XR Workspaces

**Yuzheng Chen, MRes**

School of Computing and Communications

Lancaster University

A thesis submitted for the degree of

*Master by Research*

February, 2026

**Understanding the Co-Dependence of Cursor Warping and Boundary Switching in XR Workspaces**

Yuzheng Chen, MRes.

School of Computing and Communications, Lancaster University

A thesis submitted for the degree of *Master by Research.* February, 2026.

# Abstract

**Extended reality (XR) enables flexible, multi-display workspaces without the need for physical monitors, yet extensive pointer travel between displays can compromise both ergonomics and productivity. To address this challenge, we evaluated five cursor warping techniques and their interaction with Boundary Switching—a counter-rotation method that shifts displays opposite to head yaw to reduce neck strain. The techniques initialize the pointer at different locations when moving between displays: (1) at the display's center, (2) at the last-known cursor position on the target display, (3) at the same relative position as on the previous display, (4) at the user's first gaze fixation within the target display, and (5) projected into world space as is typical for desktop environments. In a within-subjects study (N=20), participants selected cued targets across five horizontally arranged virtual displays. Without Boundary Switching, center, last-known, and relative-position warping proved reliable defaults. With Boundary Switching enabled, gaze-based warping significantly outperformed all alternatives, while world-space warping consistently underperformed. Our findings highlight how display motion changes pointer initialization strategies and inform design guidelines for ergonomic and efficient multi-display interfaces in immersive workspaces.**

# Declaration

I declare that the work presented in this thesis is, to the best of my knowledge and belief, original and my own work. The material has not been submitted, either in whole or in part, for a degree at this, or any other university. This thesis does not exceed the maximum permitted word length of 35,000 words including appendices and footnotes, but excluding the bibliography. A rough estimate of the word count is: 13730

Yuzheng Chen

# Publications

The following publication has been produced during the course of this research:

**Yuzheng Chen**, Haopeng Wang, Hock Siang Lee, Florian Weidner, Jinghui Hu, and Hans Gellersen (2025). "The Interaction of Cursor Warping and View Manipulation in Multi-Display XR Environments". Submitted to CHI 2026, under review

# Contents

# List of Tables

# List of Figures

# Chapter 1

# Introduction

Multi-display workspaces have become a central element of modern knowledge work, as they provide expanded screen real estate for multitasking, parallel document viewing, and cross-application coordination. Early adoption was driven by the increasing availability of affordable flat-panel monitors, with research consistently showing that additional displays improve productivity by reducing window-switching costs and externalizing short-term memory across applications (Czerwinski et al., n.d.; Ball and North, 2005). These setups are now common in domains such as software development, financial trading, data analysis, and design, where users must monitor and manipulate multiple information streams simultaneously (Cetin, Stuerzlinger, and Dill, 2018). With the advent of virtual and augmented reality, multi-display workspaces are no longer constrained by physical hardware: XR systems allow users to instantiate arbitrarily many displays, arrange them flexibly in 2D or 3D space, and scale them at negligible cost (Pavanatto, North, et al., 2021). Beyond scalability, XR environments offer unique benefits such as immersive focus through reduced real-world distractions (Ruvimova et al., 2020), enhanced privacy by shielding content from bystanders (Grubert et al., 2018), and opportunities for portable "offices" that replicate multi-monitor setups without bulky equipment. Consequently, XR multi-display environments are gaining recognition in the literature as a viable approach to emulate, and potentially exceed, the productivity enhancements associated with traditional physical multi-monitor configurations.

As XR systems increasingly aim to support everyday productivity tasks, an important design goal is to preserve familiar desktop interaction workflows within immersive environments. This work is motivated by emerging XR productivity scenarios in which users rely on traditional desktop input devices, such as a mouse and keyboard, while working within immersive multi-display environments. A representative example is the use of head-mounted displays such as Apple Vision Pro in conjunction with a connected laptop or desktop computer, where applications and windows from the host system are extended into an XR workspace. In such configurations, users interact with multiple virtual displays using conventional pointer input, creating a strong demand for efficient cursor control across screens.

The primary target users of this work are knowledge workers who perform prolonged desktop-style tasks in XR, including software development, data analysis, document editing, and other multi-window workflows that require frequent attention shifts across displays.

While XR enables flexible placement and manipulation of virtual displays, it also introduces new challenges for maintaining consistent and predictable cursor behavior across screens. At the core of this challenge is the lack of principled support for cursor transitions across virtual displays when the workspace itself can move relative to the user. One central issue is cursor management. Traditional continuous mouse movement techniques from physical multi-monitor setups do not scale well to XR, where users face discontinuities at display borders, cursor loss, and long travel across large gaps (Pavanatto, Davari, et al., 2023). Poorly chosen cursor warping further slows target acquisition, while mismatches between mouse motion and cursor trajectories can disrupt users' mental models of the workspace (Baudisch et al., 2004a). These challenges are amplified in XR because displays can be freely arranged and reconfigured at any time. Focusing on head movement, McGill et al. proposed *counter-rotational gain* to reduce head rotation by rotating the workspace opposite to head yaw (Mcgill et al., 2020a), but its effects on precise cursor interaction remain unexplored and may even worsen spatial inconsistencies.

For desktop environments, Benko et al. (Benko and Feiner, 2007) introduced Multi-Monitor Mouse (M3), which varied both the trigger for switching displays and the cursor's initialization strategy. While this approach reduced the need for long cursor movements, it was never evaluated in XR environments, it relied solely on head direction (which does not always accurately reflect user intent (Stiefelhagen and Zhu, 2002)), and M3's interaction with Boundary Switching remains unexplored but poses a challenge due to the disruption of the coupling between hand motor space and counter-rotation.

Building on these gaps, my study addresses how cursor placement and counter-rotation can be co-dependent in XR multi-display workspaces. In this thesis, co-dependence refers to interaction mechanisms whose effectiveness cannot be meaningfully evaluated in isolation, because user performance emerges from their combined behavior rather than from either mechanism alone.

To systematically examine this design space and disentangle how cursor strategies and counter-rotation interact, I pose the following research questions:

**RQ1:** Which cursor $\times$ counter-rotation combinations are optimal?

**RQ2:** Is advanced counter-rotation beneficial compared to traditional cursor behaviour?

**RQ3:** Where should the cursor be initialised when shifting attention to another display?

**RQ4:** Is counter-rotational gain (Boundary-Switching) beneficial for multi-display cursor interaction?

To answer these research questions and optimize mouse-based interaction in XR, I investigate five cursor warping techniques designed for VR-based multi-display workspaces. These techniques define how the cursor is positioned when a user moves their attention from one screen to another. Figure 1.1 outlines the behaviour of these cursor warping techniques. I examine four commonly used approaches: WorldSpace,

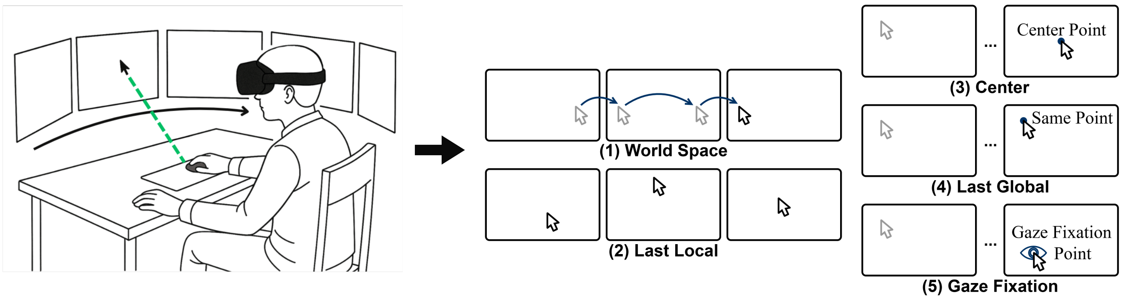


Figure 1.1: Illustration of cursor warping strategies for XR multi-display workspaces. Left: a user in VR interacting with multiple virtual monitors. Right: (1) WorldSpace: the cursor moves continuously across adjacent screens and keeps its current position when switching. (2) LastLocal: each display maintains its own cursor, and warping activates the target display by restoring its cursor (i.e., the last remembered position on that display). (3) Center: the cursor always reappears at the center of the new display. (4) LastGlobal: the cursor appears at the same relative position it was in on the previous display. (5) GazeFixation: the cursor appears at the user's gaze landing point (first fixation) on the target display.

where the cursor moves continuously across adjacent screens; Center, where the cursor reappears at the center of the new display; LastLocal and LastGlobal warping, where the cursor is restored to its last position on that screen or to the same relative position as on the previous screen. In addition, I propose a novel GazeFixation technique (inspired by MAGIC (Zhai, Morimoto, and Ihde, 1999) and gaze supported mouse (Flegel, Pick, and Mentler, 2021)), which places the cursor at the user's initial gaze point on the new screen. This technique is intended to reduce cursor travel distance and better align with user attention at the moment of transition.

In addition to cursor transitions, I examine if cursor warping and counter-rotation interact and how they jointly affect user performance. I specifically focus on *counter-rotational gain*, a technique that decouples head motion from scene rotation so that a small head turn results in a larger change in viewing angle (Mcgill et al., 2020a).

To explore this interaction, I compare all cursor transition techniques under both *fixed-view* and *counter-rotational gain* conditions.

To empirically evaluate these interactions under controlled yet ecologically plausible conditions, Through a controlled user study (N=20), simulating regular use of multi-display VR environments, I evaluate performance and user experience. I analyze a range of metrics, including selection time, cursor movement distance, overshoot behavior, cursor dynamics (speed and acceleration), and subjective measures (satisfaction, ease of use, and motion sickness).

My results show how cursor transition techniques and counter-rotation strategies interact, highlighting which combinations provide the best balance of efficiency, accuracy, and comfort. My study shows that cursor warping strategies such as Center and LastLocal consistently reduced selection times and pointer travel without adding errors or workload, making them well-suited for XR desktops. Critically, their effectiveness depended on counter-rotation control: Fixed View favored Center and LastLocal, while GazeFixation best offset the costs of counter-rotated displays, underscoring the need for co-dependence between cursor and counter-rotation control. These findings demonstrate that cursor behaviour and view manipulation form a coupled interaction system, in which changes to one component directly reshape the usability of the other.

Based on this investigation, this thesis makes the following empirical and design-oriented contributions:

1. I show that advanced cursor warping techniques outperform the conventional WorldSpace approach (LastLocal, Center, and GazeFixation perform best).

2. I demonstrate that a cursor warping's effectiveness depends on counter-rotation (FixedView: LastLocal and Center. HeadGain: GazeFixation).

3. I derive design implications for XR desktop systems, including the need for automatic cursor repositioning, the pairing of counter-rotation with gaze-based warping, and the avoidance of persistent world-space cursors.

Taken together, my findings show that seemingly low-level design choices, how and where the cursor appears when switching displays, have system-wide consequences for efficiency, ergonomics, and usability in XR.

The remainder of this thesis is structured as follows. Chapter 2 reviews prior work on multi-display interaction in both physical and XR contexts. Chapter 3 describes my methodology, including experimental design, apparatus, and measures. Chapter 4 presents the results of my controlled user study. Chapter 5 discusses the findings in light of my research questions and outlines implications for XR workspace design. Finally, Chapter 6 concludes with a summary of contributions and directions for future research.

# Chapter 2

# Literature Review

This chapter situates my work within prior research on multi-display interaction. I first review studies of physical multi-monitor setups, then examine how these ideas extend to XR workspaces, and finally consider counter-rotation techniques. This synthesis provides the necessary context for my investigation of cursor warping strategies in XR.

## 2.1 Physical Multi-Monitor Workspaces

Human–computer interaction (HCI) research has long examined why people use multiple monitors and how they affect productivity. In general, adding more displays provides users with more screen space to spread out their work. For example, a person can keep a reference document on one screen and actively work on another without constantly opening, closing, or minimizing windows (Saleem and Weiler, 2018). An example of such a physical multi-monitor workspace is shown in Figure 2.1. This additional screen real estate reduces what HCI researchers call *coordination cost*—the time and effort spent finding and arranging application windows—and supports multitasking by making multiple tasks visible at once (Colvin, Tobler, and Anderson, 2004). Surveys and experiments consistently report that people prefer multi-monitor setups and feel more productive when using them

(Miller, Stenmark, and Ittersum, 2020; Ling et al., 2017). In short, multiple displays often improve efficiency and align with users' reported preferences and perceived productivity (Miller, Stenmark, and Ittersum, 2020; Ling et al., 2017).

Figure 2.1: Example of a physical multi-monitor workspace with five displays arranged horizontally.

However, having several screens also introduces new interaction challenges. When the user moves the mouse cursor from one monitor to another, they may need to travel a long distance or physically turn their head farther than with a single screen. Screens can have physical frames or *bezels* between them, and resolutions or sizes may not match perfectly. These discontinuities can cause the cursor to vanish or "get lost" at a border, a phenomenon often called *cursor loss* (Baudisch et al., 2004b; Waldner, Kruijff, and Schmalstieg, 2010). Reacquiring a lost cursor takes extra effort: in practice, users may overshoot a screen or repeatedly jiggle the mouse trying to see the pointer, which hurts efficiency. Such cross-screen friction—long cursor travel, frequent pointer disappearance, and head/eye movement—can diminish the benefits of multiple monitors (Baudisch et al., 2004b; Nacenta, Mandryk, and Gutwin, 2008).

To mitigate these issues, researchers have proposed techniques to shorten or automate the transition between screens. One common approach is *cursor warping*, which automatically moves the cursor instantly to a target screen once a switching condition is met, rather than requiring the user to physically drag it. For example, a system might teleport the mouse pointer to the opposite side of the next screen when the user presses a special key. Other methods include using multiple independent cursors (one per display) or leveraging head or eye tracking to hand over control to the screen where the user is looking. Studies indicate that these approaches can improve user satisfaction and speed by reducing the awkward travel between screens, even if they introduce an explicit switching step (Benko and Feiner, 2005; Ashdown, Oka, and Sato, 2005).

### 2.1.1 Switching Triggers and Cursor Warping.

Prior work has explored different ways to decide *when* to switch the cursor and *how* to position it on the new screen. For example, the Multi-Monitor Mouse (M3) system by Benko et al. (2005) systematically experimented with various *triggers* and *warping strategies*. The trigger could be a keyboard shortcut, a mouse button, or a head orientation (using head tracking). When a trigger fired, the cursor would "warp" to the next display according to one of several rules. In one strategy, the cursor appeared at the center of the destination screen; in another, it appeared in the same relative position as on the previous screen; and in a third ("frame-memory"), the cursor jumped to the exact position where it was last seen on that screen. Follow-up work refined these ideas, for example combining head orientation with a button press to avoid unintended activations (the so-called "Midas touch" problem) (Benko and Feiner, 2007). Remembering the last cursor location on each display was found to provide the fastest and most comfortable performance (Benko and Feiner, 2007).

Other researchers have explored related techniques. For instance, some systems use explicit key bindings or gestures to jump the cursor across screens (Baudisch et al., 2004b; Nacenta, Mandryk, and Gutwin, 2008), while others rely on gaze direction or eye tracking to select the target screen and warp the cursor there (Flegel, Pick, and Mentler, 2021). All these methods share the goal of reducing unnecessary mouse movement by placing the cursor closer to the intended target in the new display. Controlled experiments show that users often prefer these intelligent warping techniques and complete cross-screen tasks more quickly with them (Benko and Feiner, 2005; Benko and Feiner, 2007).

### 2.1.2 Performance Costs of Crossing Displays.

Even aside from pointer movement, using multiple monitors can add overhead to tasks. Studies have observed that when a user's focus shifts to a non-primary or distant screen, task completion times increase and more actions (e.g., window switches) are required (Lin and Y.-s. Chen, 2023; Ball and North, 2005). Physically, users tend to swivel their chair, lean forward, or stretch to reach peripheral monitors, which adds time and effort. Compared to a single display, multi-monitor setups induce more head and body movements, especially if the off-screen target is far away (Rashid, Nacenta, and Quigley, 2012). For example, a user might take longer to locate a document on a second monitor or to drag a window from one screen to another. These effects motivate designing interactions that minimize travel distance and reorientation. Ideally, the system should reduce the need to physically turn or hunt for the cursor, thereby lowering the cognitive and physical costs of working across monitors (Lin and Y.-s. Chen, 2023).

Taken together, prior work on physical multi-monitor workspaces establishes cursor transition as a central bottleneck: cross-display performance depends not only on how users select a target screen, but critically on where the cursor is re-initialised when attention shifts. This thesis builds on these insights, but examines them in XR where displays can be spatially distributed and may move relative to the user.

## 2.2 XR Workspaces as Virtual Multi-Display Environments

Extended Reality (XR) technologies, which include virtual reality (VR) and augmented reality (AR), bring multi-monitor concepts into fully virtual spaces (Medeiros et al., 2022; Ng et al., 2021). In an XR workspace, users can have many virtual "monitors" arranged around them, even in a small physical area. This allows powerful multi-screen setups to be used anywhere, which is especially useful for working on the go or in public settings where privacy is a concern (Biener et al., 2024; Mcgill et al., 2020b). In principle, these virtual displays can provide the same advantages as real monitors: more total screen area, dedicated windows for different tasks, and quick visual access to multiple information sources simultaneously (Medeiros et al., 2022).

At the same time, XR headsets have hardware limits that affect usability. Most current VR/AR headsets have narrower fields of view and lower angular resolution than physical monitors. In practice, this means that virtual screens can appear smaller or blurrier, and users must turn their head more often to scan across them. In particular, looking at the periphery of a wide virtual screen requires more physical head rotation than with a large physical monitor, which can be tiring (Pavanatto, Davari, et al., 2023). Studies also find that when many virtual screens are available, users spend more time searching visually, since it can be harder to know where a window is placed in 3D space (Zhan et al., 2020). These constraints make cross-display cursor behaviour more consequential in XR than on the desktop: wider spatial layouts increase traversal distances, peripheral viewing increases search effort, and dynamic window placement makes cursor re-entry less predictable. As a result, cursor transition techniques that work well on fixed physical monitors may behave differently when applied to virtual displays distributed in 3D space.

To support traditional desktop-style work in XR, toolkits have been developed that bring standard applications and mouse/keyboard input into the virtual environment. For example, WindowMirror is an open-source toolkit that mirrors desktop windows in VR and suggests arranging them on cylindrical surfaces centered around the user (Bovo et al., 2024). By rendering windows as floating rectangles in space, such tools allow users to use their normal desktop applications through an XR display. One challenge, however, is how to manage the cursor when windows can exist anywhere in 3D space. If a window is tilted or out of view, it can be unclear how the flat mouse cursor maps onto it. Some research addresses this by projecting the 2D cursor onto a "virtual screen" that is always perpendicular to the user's gaze, ensuring consistent interaction (Spittle et al., 2023; A. Hill and Johnson, 2008). In other words, instead of letting the pointer float freely in 3D, the system treats it as if it lies on a visible plane so the user can see and control it normally.

### 2.2.1 Canvas vs. Subdivided Layouts.

In XR, there are different ways to arrange virtual windows. A *continuous canvas* approach gives the user a very large undivided workspace (like one huge whiteboard) where windows can be placed anywhere. This flexibility lets users spread out content without fixed boundaries, potentially reducing the need for extreme turning if the workspace is arranged ergonomically (Pavanatto, Davari, et al., 2023). In contrast, a *subdivided layout* simulates multiple separate screens (for example, an array of virtual monitors or a spherical grid). These distinct slots or "regions" can help organize tasks (e.g., coding on one screen, reference on another) and allow snapping windows into place for quick management. Comparative studies suggest a trade-off: while a continuous canvas allows creative use of space and can improve large-scale orientation, the structured multi-screen arrangement often makes it easier to keep track of multiple apps and switch between them (Pavanatto, Lu, et al., 2025; Pavanatto, Davari, et al., 2023). As a result, XR systems often adopt familiar multi-monitor metaphors to reduce cognitive load in organizing many applications simultaneously (Pavanatto, Lu, et al., 2025). Given this trade-off, my work focuses on subdivided, multi-display–style XR workspaces, where cursor transitions occur frequently and cursor placement policies can have immediate effects on efficiency and comfort.

## 2.3 Counter-Rotation in XR

When users wear an HMD and look around a virtual environment, the system typically lets their head movement control the view with a 1:1 mapping. This direct, head-tracked view usually feels natural and helps users maintain spatial orientation and presence (Seibert and Shafer, 2018; Ragan et al., 2017). However, when the virtual environment is very large or the content of interest is spread widely, constantly turning the head or body can become tiring. To address this, researchers have explored methods for modifying how the view responds to head motion.

One idea is to use *rotation gains*, where the virtual scene rotates by a larger or smaller angle than the user's head actually moves. For example, a gain of 2× would rotate the virtual world twice as much, so a small head turn covers more of the scene. Moderate amplification gains can reduce the physical effort of scanning a wide area, at least for stationary or seated users (Ragan et al., 2017). However, if the gain is too high, users can become disoriented or lose track of where they are looking, which harms spatial memory and may increase motion sickness (Freitag, Weyers, and Kuhlen, 2016; Ragan et al., 2017).

Another class of techniques uses *discrete* view changes instead of continuous rotation. A well-known example in VR is *snap turning*, where the user's view instantaneously jumps by a fixed angle (e.g., 30° or 45°) when they press a button. Studies indicate that snap rotations can reduce simulator sickness without significantly affecting task performance, as long as the increments are not too large (Farmani and Teather, 2020). Other research has explored using hand-held controllers or mid-air gestures to pan the view, or providing mini-maps and guided navigation for large virtual worlds (Tara and Teng, 2017; Yang et al., 2018). Collaborative techniques (sharing another user's perspective) and unusual methods like oscillating the camera to mimic walking motion have also been tested, with mixed results depending on context (Jeong, Choi, and Lee, 2025; Tserenchimed and Kim, 2024; Moullec et al., 2024).

A related concept is *counter-rotation* (or *rotational assistance*), where the virtual environment automatically rotates slightly opposite to the user's head movements. In effect, the system tries to keep the content of interest in front of the user by "pulling" the world toward them as they look around. Early work found that applying a small counter-rotation (gain less than 1) can improve ergonomic comfort: users need to turn their body less and experience less strain (Mcgill et al., 2020b). More recent work has also tested this on the vertical axis (pitch) to ease looking up and down (O'Hagan et al., 2025). These methods generally maintain presence while reducing the amount of physical motion required.

Not all effects of counter-rotation are beneficial, however. For tasks requiring precise pointing or spatial learning, the results are mixed. Some studies report that moderate counter-rotation does not change performance in simple pointing tasks, but it can impair memory of the spatial layout of objects (Freitag, Weyers, and Kuhlen, 2016). Additionally, if the counter-rotation is not perfectly smooth, small jitter can make it harder to click accurately on a moving cursor (Batmaz et al., 2021). In contrast, very subtle gains (below the user's notice threshold) can often be applied without users realizing and without harming task accuracy (Wang et al., 2023). Overall, more research is needed to understand how these view-control techniques interact with precise input. Critically, most evaluations of counter-rotation focus on ergonomic coverage and coarse navigation, and rarely examine how view manipulation alters the moment-to-moment mapping between hand movement, cursor placement, and target acquisition. This leaves an open question that directly motivates this thesis: when the workspace itself moves during attention shifts, which cursor warping strategies remain effective, and which break down?

## 2.4 Relation to This Work

My work bridges these areas by investigating how to hand off cursor control among multiple virtual displays in XR and how that interacts with counter-rotation techniques. Prior studies suggest that explicit cursor warping and remembering the cursor's last position tend to speed up multi-display interaction and are preferred by users (Baudisch et al., 2004b; Benko and Feiner, 2005; Benko and Feiner, 2007; Nacenta, Mandryk, and Gutwin, 2008; Waldner, Kruijff, and Schmalstieg, 2010). Building on this evidence, I implement and compare five different warping strategies in a virtual workspace—for example, warping the cursor relative to the world, to the center of the screen, or to its last known location on the target display. I then measure how these techniques perform both with a normal 1:1 view and with a counter-rotational gain applied. This approach lets us directly test how view manipulation (counter-rotation) affects the experience of switching cursors between virtual monitors, a topic that has not been extensively studied. In doing so, I aim to identify methods that enable fast, accurate switching in VR while keeping users comfortable and well-oriented (Benko and Feiner, 2007; Baudisch et al., 2004b).

# Chapter 3

# Methodology

This chapter details the methodological approach used to evaluate cursor warping in XR multi-display workspaces. I first motivate the problem and pose four research questions (RQ1–RQ4), then delimit the design spaces for display layout, display rotation control, and cursor warping. I describe the resulting $5 \times 2$ within-subjects experiment (cursor warping technique $\times$ counter-rotation), including task, procedure, apparatus, and participant sample, followed by the performance and subjective measures used for analysis. These components provide a controlled basis for assessing main effects and interactions between cursor placement and counter-rotation.

## 3.1 Motivation and Research Questions

Building on prior work in XR multi-display workspaces, I sought to examine how pointer warping techniques interact with counter-rotation strategies to influence user performance and experience. In traditional multi-monitor desktop setups, the cursor travels to an adjacent display by crossing a shared boundary (often appearing at a corresponding edge position). XR environments, however, introduce additional complexity: virtual displays can be arranged in 3D space (not strictly edge-aligned) and context switches may be triggered by non-mouse inputs like gaze or head movement. This freedom allows the system to reposition the cursor arbitrarily (e.g., “warping” it to a new location) when the user shifts focus to another display. Such flexibility could reduce cursor travel and prevent getting “lost,” but might also disrupt the user’s spatial frame of reference. At the same time, XR permits novel counter-rotation manipulations; for example, rotating the entire virtual workspace opposite to the user’s head turn (to reduce neck strain) could inadvertently affect pointer alignment. Given these possibilities, I identified a gap in understanding how cursor placement and dynamic counter-rotation jointly impact cross-display interactions.

Mcgill et al., 2020a demonstrated that introducing a counter-rotational gain (rotating the virtual displays opposite to head movement) can reduce neck strain and expand the effective field of view in seated VR. However, their study did not address how this technique might interact with different pointer behaviors. By moving the display space relative to the user, such counter-rotation could exacerbate issues like inconsistent spatial mapping or inefficient cursor travel if the cursor is not appropriately managed, potentially degrading selection performance. To explore these issues, I closely replicated the Study 1 setup of Mcgill et al., 2020a in my user study and extended it to test multiple cursor warping methods under both 1:1 and Boundary Switching strategies.

## 3.2 Design space for layouts

XR supports both 1D ribbons and 2D grids of virtual displays (Mcgill et al., 2020a; O’Hagan et al., 2025). As illustrated in Figure 3.1, to keep rotation well-posed and avoid vertical–horizontal coupling, I fix a 1D horizontal arc of five displays for all experiments and defer 2D layouts to future work. This choice keeps cross-display geometry and path lengths well defined, reduces confounds from vertical head motion, and supports systematic sampling of spans and directions under the same rotational policies.

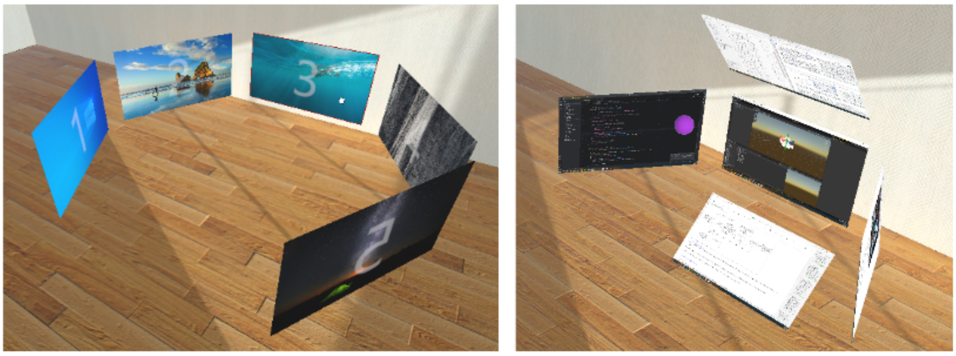

Figure 3.1: Layout rationale. Left: 1D horizontal arc used in the study. Right: representative 2D layouts (explored conceptually, not tested).

## 3.3 Design space for displays rotation control

I implemented five rotation strategies, organized by trigger modality (head vs. gaze), plus a non-rotating baseline. Only two were taken forward to formal testing; all five are reported here to fully document the exploration path and motivate future studies (Figure 3.2). The strategies can be seen as spanning a continuum from strict spatial fidelity (baseline) to increasingly assisted reorientation driven by either head or gaze input

### 3.3.1 Baseline

*1:1.* Head yaw maps 1:1 to the workspace; no gain and no snapping. The display array remains static with respect to the real world: turning the head by $\theta$ rotates the view by $\theta$, and the displays do not reorient. This establishes a natural control mapping and preserves spatial reference. The baseline condition is directly analogous to a physical multi-monitor setup, where the user must rotate their head (and sometimes torso) to look from one display to another. It offers high predictability and spatial fidelity, but at the cost of potentially large physical rotations for extreme angles.

### 3.3.2 Head-triggered strategies

*Boundary Switching (Boundary-Switching, head-triggered).* When the head "cursor" crosses a display boundary, the array rotates opposite to the head turn to center the next display; this is a discrete, boundary-triggered counter-rotation with an effective gain of about 2:1 (Mcgill et al., 2020a). In practice, the user initiates the transition with a moderate head turn, and the system completes the motion by snapping the next display into central view. This reduces the physical yaw needed to bring the target display into view, while keeping orientation changes legible and predictable at display boundaries. The sharp, binary trigger avoids the ambiguity and jitter of continuous coupling, making the technique easy to learn.

*Continuous head-gain with per-display deadzones.* The head forward-vector yaw is continuously mapped to counter-rotation with a central deadzone for each display (e.g., $\pm 12.5°$). Outside deadzones, inter-display gain accelerates transitions so the array appears to "slide" between displays, producing rapid inter-display motion in the space between deadzones. This follows the mapping described by McGill et al. for conditions where all displays define central deadzones with amplified motion outside them (Mcgill et al., 2020a).

### 3.3.3 Gaze-triggered strategies

*Boundary Switching, gaze-triggered.* The same boundary-crossing logic as Boundary Switching, but using gaze (the head→gaze vector) as the trigger. The intention is to align automatic reorientation with attentional focus, potentially reducing instances where the head briefly crosses a boundary without a genuine switch of task focus. While conceptually attractive, this method risked unintended rotations from quick glances or saccades, and thus was retained as exploratory.

*Continuous gaze-gain with deadzones.* This variant replaces the head forward-vector with the head-gaze vector while retaining per-display deadzones. It explores whether attention-driven continuous assistance can reduce unnecessary head effort while keeping the display containing the fixation preferentially centered.

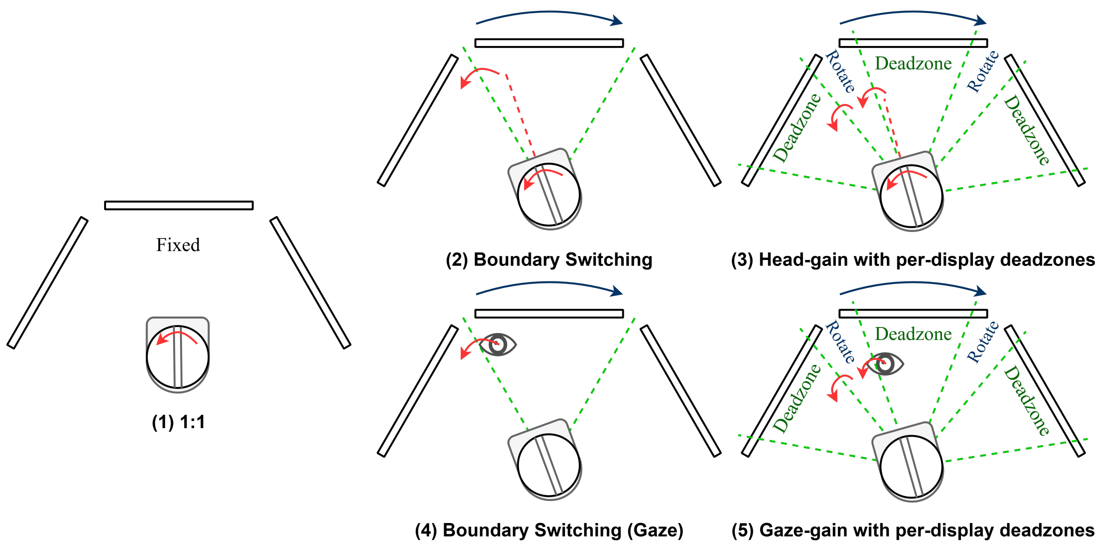


Figure 3.2: Displays rotation control methods. (1): 1:1. (2), (3): Boundary Switching (head-based) and continuous head-gain with deadzones. (4), (5): boundary switching and continuous gain with deadzones (gaze-based).

### 3.3.4 Experimental design

In the user study I only compared 1:1 and Boundary Switching. The main reason was to keep the study duration manageable while focusing on the behavior of cursor warping techniques; a broader comparison of display rotation methods is left for future work. I chose Boundary Switching as the assisted condition since boundary switching has been established in prior work by both McGill and O'Hagan (Mcgill et al., 2020a; O'Hagan et al., 2025).

## 3.4 Design space for cursor warping techniques

I compare five widely applicable methods that determine where the pointer appears on the destination display when attention shifts (Figure 1.1):

1. WorldSpace: The cursor behaves like in a traditional multi-monitor setup: it moves continuously across adjacent screen boundaries and keeps its current global position when the user looks at a new display (i.e., no warping is applied).

2. LastLocal: Each display maintains its own cursor, and when the user switches to a display, its cursor is restored at the last remembered position on that display (frame-memory Benko and Feiner, 2007).

3. Center: When switching displays, the cursor always reappears at the center of the newly focused display, providing a predictable starting point (screen-center Benko and Feiner, 2005).

4. LastGlobal: The cursor appears at the same relative position on the new display as it was on the previous display. For instance, if the cursor was located near the bottom-left of one screen, it will warp to the bottom-left of the target screen, analogous to frame-relative warping Benko and Feiner, 2005.

5. GazeFixation: When the user shifts to a new display, the cursor is immediately warped to the gaze landing point, i.e., the first stable fixation detected on that display, leveraging eye-tracking to align with the user's visual attention, similar to Flegel, Pick, and Mentler, 2021.

### 3.4.1 Design alternatives for gaze-driven warping

I considered three strategies before adopting GazeFixation: real-time gaze warping (jump at every sample), prediction-based warping (jump to a predicted landing point), and fixation-based warping (warp once upon fixation). Real-time updates systematically biased placement near the entry boundary of the destination display and were susceptible to micro-saccade jitter. Prediction reduced latency but introduced noticeable placement errors requiring corrective mouse motion. Hybrids that switched to real-time after a manual nudge created trigger ambiguity (participants could not tell whether the jump was caused by fixation confirmation or the real-time override). These observations motivated a fixation-only design.

### 3.4.2 First-gaze warping via I-VT fixation detection

I detect fixations using an I-VT scheme in head-centric angular space. For each frame $t$, I compute the inter-sample visual angle $\theta_t$ between consecutive gaze direction vectors and the angular velocity $v_t = \theta_t/\Delta t$. A sample is labeled as fixation when $v_t <$ `velocityThreshold`. I store the onset time $t_0$ of the most recent below-threshold sample and confirm fixation once $(t - t_0) \geq$ `minFixationDuration`. Upon confirmation, the cursor warps to the gaze landing point (not the first entry point). I tuned a short confirmation interval of approximately 60 ms using conservative motion thresholds, including an acceleration threshold of less than $60°/s^2$, which yielded accurate landings with tolerable delay in pilots.

### 3.4.3 Experimental design

All five cursor warping techniques were included, as each represents a unique strategy worth direct comparison.

## 3.5 Study Design

Using the above techniques, I designed a controlled within-subjects experiment to evaluate their performance and to answer the research questions. The study employed a 5 × 2 factorial design: **cursor warping technique** (five levels: WorldSpace, LastLocal, Center, LastGlobal, GazeFixation) × **rotation mode** (two levels: 1:1 vs. Boundary Switching). This yielded 10 distinct conditions combining pointer and display counter-rotation strategies. Every participant experienced all conditions (within-subjects); condition order was counterbalanced to mitigate learning effects. In practice, I blocked the study by rotation mode: half of the participants tried all five cursor techniques under the Fixed (no-rotation) condition first, then all five under Boundary Switching, while the other half did the opposite. Within each block, the order of the five cursor warping conditions was balanced using a Latin square. This approach controlled for potential order biases and allowed us to analyze both the main effects of each factor and their interaction.

### 3.5.1 Task

Participants performed a multi-display target selection task that simulated moving a cursor between virtual monitors. They sat at a virtual desk and used a physical mouse to control a cursor within the VR workspace. On each trial, the goal was to click a circular "button" target that could appear on any of the five displays.

At the beginning of each block of trials, the first target appeared on the center display (Display 3) at a random location. Thereafter, each trial proceeded as follows (illustrated in Figure 3.3): the participant would click the current target, then a visual cue would immediately indicate the next target's display number and an approximate location on that display. This cue was presented as a highlighted rectangle on the target display (Figure 3.3b), designed to minimize visual search time by making the upcoming target's location obvious. Once the participant shifted their attention to the indicated display, the cursor was re-warping according to the active warping technique (e.g., jumping to the center, last position, etc., depending on condition). The next target then became visible (Figure 3.3c), and the participant would acquire it by moving the cursor (if needed) and clicking. A new trial would then begin with another cue for the following display.

The targets were white rectangular buttons (approximately the size of a standard desktop icon, 74×83 pixels in virtual size). Target positions on each display were randomized, constrained to lie within the central 80% of the display area (i.e., with at least a 10% margin from the edges) to avoid extremely peripheral placements. Participants received immediate visual feedback upon clicking a target (the button would disappear), and the trial ended once the target was successfully selected.

### 3.5.2 Procedure

After giving informed consent, participants completed a brief tutorial to familiarize themselves with the virtual multi-display setup and the different cursor warping

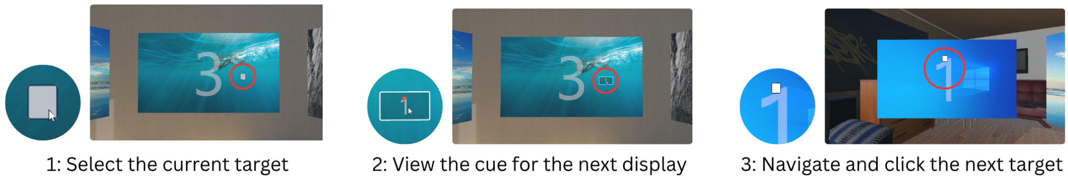


(a) Initial target is selected on Display 3.

(b) After selection, a cue highlights Display 1 and approximate target area.

(c) Cursor warps to Display 1 (here to center, in Center) and the target appears for selection.

Figure 3.3: Example trial sequence.

behaviors. They practiced each of the five cursor techniques (with both counter-rotation modes) in a guided training scenario ( 10 trials per technique). The experimenter ensured that participants understood how each warping method worked (e.g., noting that the cursor would jump to screen centers in Center, or to gaze points in GazeFixation). Participants were allowed to repeat training trials until they felt comfortable.

For the main trials, participants completed all 10 conditions (5 warping techniques $\times$ 2 counter-rotation modes). The order of conditions was as described above (counter-rotation block order counterbalanced between subjects, and warping technique order counterbalanced within each block). Each condition consisted of 40 selection trials. To ensure a thorough and unbiased sample of display transitions, I generated sequences using a balanced tournament graph traversal. In each block of 20 trials, every ordered pair of distinct displays (e.g., $1 \rightarrow 3$, $3 \rightarrow 1$, $2 \rightarrow 5$, etc.) occurred exactly once as a consecutive trial transition (Figure 3.4 a). Two such sequences (each covering all unique transitions) were executed back-to-back to make the 40 trials per condition. This ensured that participants experienced an even distribution of jumps between all pairs of displays, eliminating bias toward any particular screen or direction.

Participants were instructed to perform the task as quickly and accurately as possible. They remained seated throughout. Short breaks were provided between

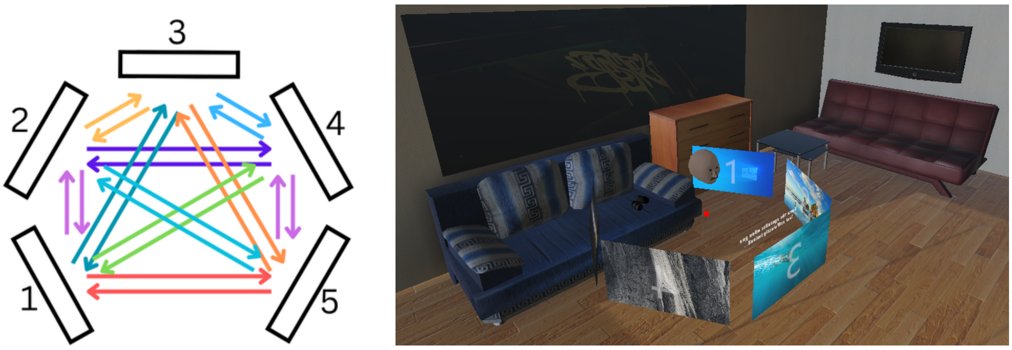


(a) (b)

Figure 3.4: Left: Example of the Tournament Graph traversal method used to balance all possible display-to-display transitions; Right: Virtual workspace setup replicating McGill et al.'s Study 1 Mcgill et al., 2020b, showing the five horizontally aligned displays arranged in a cylindrical configuration around the user.

conditions (after each set of 40 trials) to prevent fatigue; additional rest was offered upon request. After finishing the five warping conditions under a given counter-rotation mode (e.g., after all 1:1 trials), I asked participants to rank or choose their preferred cursor warping method for that viewing mode and to briefly explain their choice. They could type their explanation on a keyboard or speak it aloud for the experimenter to record. A final questionnaire at the end gathered overall strategy preferences and additional comments. Each session lasted approximately 60 minutes. Participants received £10 as compensation. The study was approved by the institutional ethics board.

### 3.5.3 Apparatus

The study was conducted using a Meta Quest Pro VR headset (per-eye resolution 1800×1920, 106° horizontal field of view), which provided head-tracked stereo rendering and integrated eye tracking. The virtual workspace was rendered in Unity 2021.3 and resembled the setup from McGill et al.'s XR desktop studyMcgill et al., 2020a. Five virtual monitors (each 16:9 aspect ratio) were arranged side-by-side in a semi-circular formation around the user (Figure 3.4 b). The displays were positioned at roughly 1.2 m virtual distance from the user, with their centers aligned at the user's eye height. Each display spanned about 60° of the horizontal field, with a slight bezel gap between them (5° separation), so that the five-display array covered a wide panorama. The virtual monitors were placed within a static virtual living-room environment to provide a stable frame of reference and naturalistic setting.

Participants used a standard USB optical mouse (Lenovo model SM8823) on a desk to control the cursor. The mouse input was captured through the host PC (Windows 10) via Oculus Link. I calibrated the cursor movement to match typical desktop behavior: moving the physical mouse by a certain distance produced a proportional cursor movement on the virtual displays equivalent to the same motion on a real 1080p monitor. I did not apply any custom acceleration or gain beyond the default Windows settings, so the cursor's speed and feel were familiar to users (ensuring that muscle memory from regular PC use would transfer to the XR setting).

### 3.5.4 Participants

I recruited 20 participants (12 female, 8 male), aged 22–34 ($M = 27.5$, $SD = 4.2$). All were right-handed, able-bodied, and had normal or corrected-to-normal vision. Participants had diverse computer backgrounds but generally limited XR experience. On a 0–3 self-rating scale of VR familiarity (0 = "never used", 1 = "once or twice", 2 = "few times", 3 = "regularly"), the mean was 1.2 ($SD = 0.83$). Experience with eye-tracking or gaze-based interaction was even lower ($M = 0.8$, $SD = 0.70$ on the same scale). Thus, for most, the study represented one of their first exposures to a VR desktop and gaze-assisted cursor techniques. All participants completed the study without noticeable simulator sickness (confirmed via post-study questioning and low FMS scores).

## 3.6 Measures

I collected both objective performance metrics and subjective ratings for each condition.

*Performance metrics.*

- Completion Time: per-trial mean completion time, measured from the end of the previous trial to the end of the current trial.
- Target Acquisition Time: per-trial mean time from the end of cursor warping (cursor appears) to selecting the target.
- Head Movement: mean accumulated head-rotation angle (deg). warping point to the target center.
- Mouse Movement: relative accumulated mouse-movement value recorded by the computer (relative unit).
- Error Rate: per-trial mean error rate, where each additional (unnecessary) click is counted as one error.

*Subjective measures.*

- Raw NASA Task Load Index (TLX) Byers and S. G. Hill, 1989: a measure of subjective workload across mental demand, physical demand, effort, temporal demand, performance, and frustration.
- Satisfaction: "Overall, how satisfied were you with your experience using the system?" (7-point Likert scale).
- Single Ease Question (SEQ) Sauro and Dumas, 2009: "How easy was it to use the system?" (7-point Likert scale).
- Fast Motion Sickness Scale (FMS)Keshavarz and Hecht, 2011: "How sick do you feel right now?" (21-point scale).

After completing all conditions of a counter-rotation strategy (e.g., 1:1), participants also indicated their preferred cursor warping technique ("Which cursor warping method do you prefer the most?") and provided a brief open-ended explanation of their choice. Participants could answer this explanation either in writing (via an optional open-ended text box) or verbally, in which case I recorded their response as notes.

# Chapter 4

# Results

This chapter presents the findings of the within-subjects experiment. I first report the outcomes for the full five-display configuration, followed by the subset analysis of the three central displays. The results are organized around the key performance and subjective measures, with figures and tables summarizing the main effects and interactions. Key post-hoc findings are summarized in the main text; full post-hoc comparisons (Bonferroni-adjusted) are provided in Appendix A.1 and Appendix A.2.

## 4.1 Data Analysis Approach

In this study, I investigate the relationship between counter-rotation techniques and cursor warping techniques via a selection task. Across 20 participants, the design comprised 10 conditions $\times$ 20 trials $\times$ 2 repetitions, yielding 8,000 planned observations. All dependent measures except Target Acquisition Time included the full 8,000 trials. For Target Acquisition Time, the last trial of each condition for each participant was affected by a logging problem ($10 \times 20 = 200$ trials) and was discarded, leaving 7,800 valid trials for analysis.

Unless stated otherwise, I analyze two datasets: the full *five-display* configuration and the *three-center-display* subset. Assumption checks focused on the repeated-measures ANOVA model: I inspected Q–Q plots of residuals and used Shapiro–Wilk tests on cell-wise values (Shapiro and Wilk, 1965).

For variables that were approximately normal, I ran two-way repeated-measures ANOVAs (Type III). When sphericity was violated (Mauchly's test), I applied Greenhouse–Geisser corrections and report GG-corrected degrees of freedom, $\epsilon_{GG}$, and $p$-values.If normality was violated, I used the aligned rank transform (ART) (Wobbrock et al., 2011). Post-hoc tests were conducted only when the corresponding omnibus effect (main or interaction) was significant. After repeated-measures ANOVA, follow-up analyses employed paired $t$-tests with Bonferroni adjustment. After ART, I obtained estimated marginal means and pairwise contrasts via `emmeans` with Bonferroni adjustment (Lenth, 2025; Dunn, 1961). Omnibus effects are reported as generalized $\eta^2$ (ges) for repeated-measures ANOVAs and as partial $\eta^2$ for ART-based tests (Cohen, 1973). For RM-ANOVAs I report generalized $\eta^2$ (ges), which is recommended for repeated-measures designs with more than one factor (Olejnik and Algina, 2003). For ART analyses, I report partial $\eta^2$, which is the conventional effect size reported for nonparametric ANOVA models.

## 4.2 Five Displays

### 4.2.1 Completion Time

This measure is the per-trial mean completion time, measured from the end of the previous trial to the end of the current trial.

A repeated-measures ANOVA (Greenhouse–Geisser corrected) revealed significant main effects of cursor warping, $F(2.93, 55.72) = 6.54$, $p < .001$, $ges$=.057, and counter-rotation, $F(1, 19) = 22.29$, $p < .001$, $ges$=.072, as well as a significant interaction, $F(3, 56.99) = 7.56$, $p < .001$, $ges$=.045; see Figure 4.1.

To unpack the interaction, simple-effects analyses showed: within 1:1, Center was faster than GazeFixation. Within Boundary Switching, all alternatives outperformed WorldSpace. Considering counter-rotation within each cursor-warping level, 1:1 outperformed Boundary Switching for Center, LastLocal, and WorldSpace,

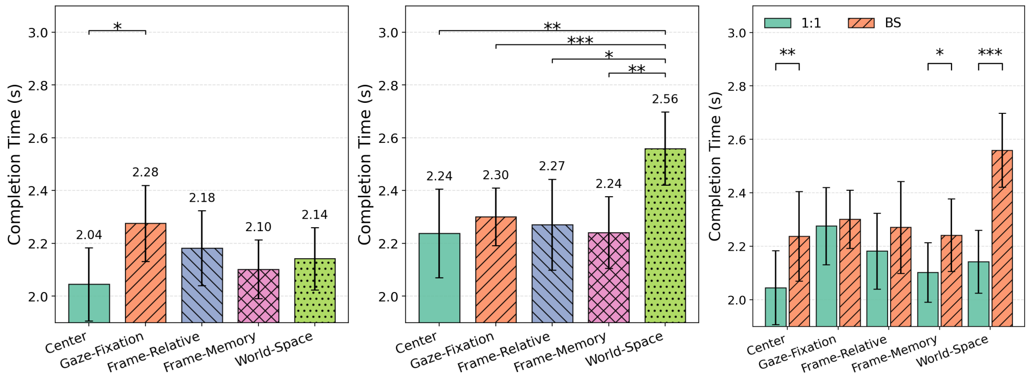


(a) 1:1 (b) Boundary Switching (c) Merged Interaction

Figure 4.1: 5D Completion Time (same y-axis limits across panels; error bars are within-subject SE, Cousineau–Morey).

with no reliable differences for GazeFixation or LastGlobal. Collapsed across factors, Center and LastLocal outperformed WorldSpace, and overall 1:1 was faster than Boundary Switching. Bonferroni-corrected post-hoc comparisons supporting these simple-effects findings are reported in Appendix A.1.

Overall, completion time exhibited a clear interaction: WorldSpace was comparatively slow (especially under Boundary Switching), whereas 1:1 yielded shorter times, particularly for Center, LastLocal, and WorldSpace.

### 4.2.2 Target Acquisition Time

This measure is the per-trial mean time from the end of cursor warping (cursor appears) to selecting the target.

A repeated-measures ANOVA (Greenhouse–Geisser corrected) indicated significant main effects of cursor warping, $F(3.11, 59.09) = 11.06$, $p < .001$, $ges$=.117, and counter-rotation, $F(1, 19) = 13.84$, $p = .001$, $ges$=.062, together with a strong interaction, $F(2.33, 44.19) = 24.79$, $p < .001$, $ges$=.202; see Figure 4.2.

To explore the interaction, simple-effects analyses showed: within 1:1, Center was faster than GazeFixation. Within Boundary Switching, GazeFixation outperformed all other cursor-warping methods, while Center, LastGlobal, and

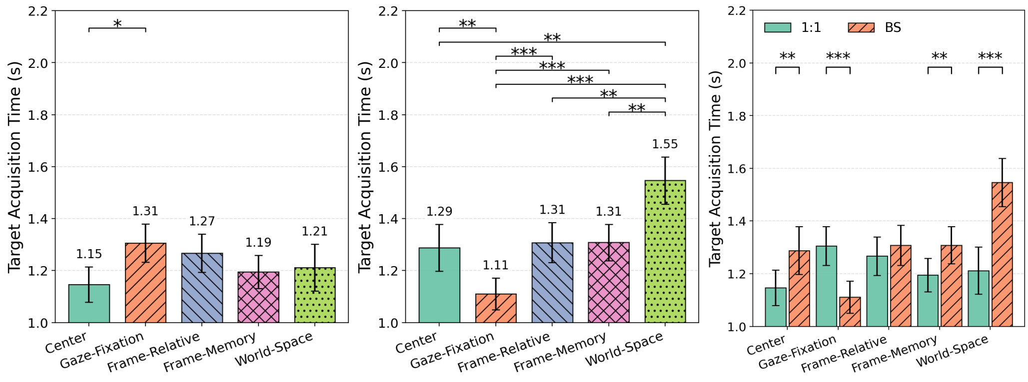


(a) 1:1 (b) Boundary Switching (c) Merged Interaction

Figure 4.2: 5D Target Acquisition Time (same y-axis limits across panels; error bars are within-subject SE, Cousineau–Morey).

LastLocal each outperformed WorldSpace. Considering counter-rotation within each cursor-warping level, 1:1 was faster for Center, LastLocal, and WorldSpace, whereas Boundary Switching was faster for GazeFixation; LastGlobal showed no reliable difference. Collapsed across factors, Center, GazeFixation, and LastLocal outperformed WorldSpace, GazeFixation also outperformed LastGlobal, and overall 1:1 was faster than Boundary Switching. Bonferroni-corrected post-hoc comparisons supporting these simple-effects findings are reported in Appendix A.1.

Overall, target acquisition exhibited a pronounced interaction: GazeFixation excelled under Boundary Switching, whereas 1:1 favored Center and LastLocal and mitigated the cost of WorldSpace.

### 4.2.3 Head Movement

This measure is the mean accumulated head-rotation angle (deg).

A repeated-measures ANOVA (Greenhouse–Geisser corrected) revealed significant main effects of cursor warping, $F(2.92, 55.40) = 18.34$, $p < .001$, $ges$=.201, and counter-rotation, $F(1, 19) = 139.07$, $p < .001$, $ges$=.303, as well as a significant interaction, $F(2.86, 54.35) = 8.87$, $p < .001$, $ges$=.087; see Figure 4.3.

Decomposing the interaction showed that within 1:1, GazeFixation reduced

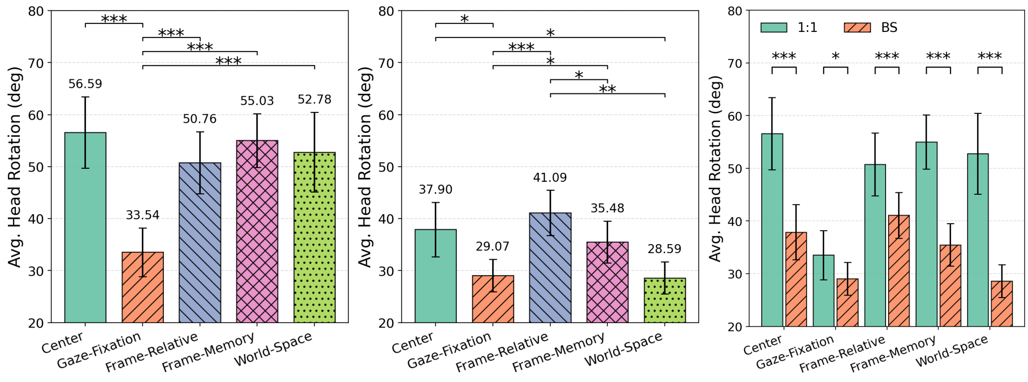


(a) 1:1 (b) Boundary Switching (c) Merged Interaction

Figure 4.3: 5D Head Movement (same y-axis limits across panels; error bars are within-subject SE, Cousineau–Morey).

head movement relative to all other cursor-warping levels. Within Boundary Switching, Center exceeded GazeFixation and WorldSpace; GazeFixation was lower than LastGlobal and LastLocal; and LastGlobal exceeded both LastLocal and WorldSpace. Considering counter-rotation within each cursor-warping level, 1:1 produced more head movement than Boundary Switching for all five cursor-warping conditions.

Collapsed across factors, post-hoc tests confirmed that Center exceeded GazeFixation, and GazeFixation was lower than LastGlobal, LastLocal, and WorldSpace. Collapsed across cursor warping, 1:1 produced more head movement than Boundary Switching.

Overall, GazeFixation consistently minimized head movement, whereas 1:1 induced substantially more head rotation than Boundary Switching across cursor-warping levels.

### 4.2.4 Mouse Movement

This measure is the relative accumulated mouse-movement value recorded by the computer (relative unit). Because normality was violated, I analyzed this measure using the ART.

A repeated-measures ANOVA on aligned ranks revealed a robust main effect of cursor warping, $F(4, 76) = 59.52$, $p < .001$, *partial* $\eta^2$ =.758, no main effect of counter-rotation, $F(1, 19) = 0.13$, $p = .727$, *partial* $\eta^2$ =.007, and a significant interaction, $F(4, 76) = 3.19$, $p = .018$, *partial* $\eta^2$ =.144; see Figure 4.4.

To explore the interaction, simple-effects tests did not reveal reliable pairwise differences either among cursor-warping levels within each counter-rotation mode or between counter-rotation modes within each cursor-warping level. Collapsed across counter-rotation, post-hoc comparisons showed that WorldSpace required more mouse movement than all other cursor-warping methods, and GazeFixation required less movement than LastLocal; no other pairwise differences were reliable. Bonferroni-corrected post-hoc comparisons supporting these simple-effects findings are reported in Appendix A.1.

Overall, mouse-path length was primarily driven by cursor warping: WorldSpace consistently induced the most movement, whereas GazeFixation was comparatively economical. No consistent modulation by counter-rotation was observed.

### 4.2.5 Error Rate

This measure is the per-trial mean error rate, where each additional (unnecessary) click is counted as one error. Because normality was violated, I analyzed this measure using the ART.

A repeated-measures ANOVA on aligned ranks revealed no significant effects of cursor warping, $F(4, 76) = 1.66$, $p = .168$, *partial* $\eta^2$ =.080, counter-rotation, $F(1, 19) = 0.41$, $p = .530$, *partial* $\eta^2$ =.021, or their interaction, $F(4, 76) = 1.10$, $p = .365$, *partial* $\eta^2$ =.054. Consistent with these omnibus results, post-hoc comparisons did not reveal any reliable pairwise differences.

Overall, error rates were comparable across cursor-warping and counter-rotation conditions.

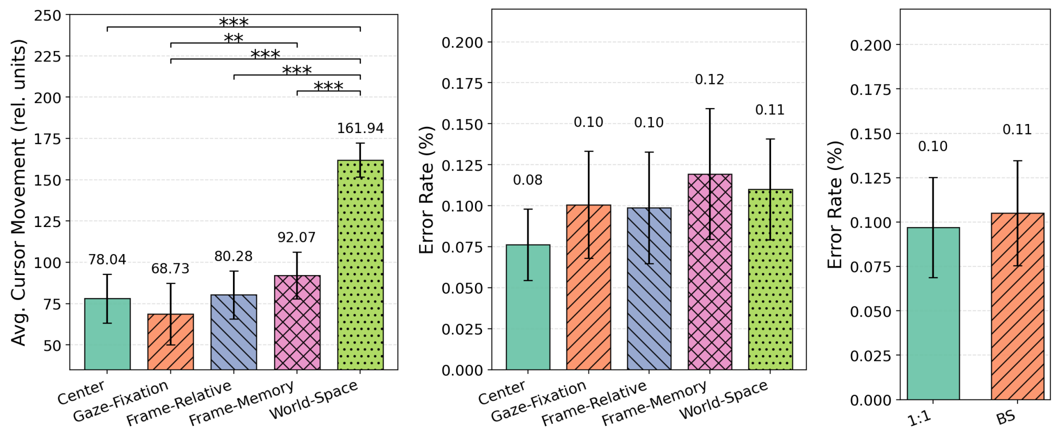


(a) (b) (c)

Figure 4.4: (a) Mouse Movement (Main Effect of CursorWarping); (b) Error Rate (Main Effect of CursorWarping); (c) Error Rate (Main Effect of Counter-rotation). All panels share the same y-axis limits for comparability; error bars show within-subject SE (Cousineau–Morey).

## 4.3 Three Displays Subset

This three-display analysis is a subset of the five-display dataset, not a separate experiment. It was conducted to examine whether the main interaction patterns generalize when considering only the three most relevant display conditions.

### 4.3.1 Completion Time

A repeated-measures ANOVA (Greenhouse–Geisser corrected) revealed significant main effects of cursor warping, $F(3.12, 59.32) = 5.56$, $p = .002$, $ges$=.061, and counter-rotation, $F(1, 19) = 48.16$, $p < .001$, $ges$=.146, as well as a significant interaction, $F(3.12, 59.24) = 5.93$, $p = .001$, $ges$=.045; see Figure 4.5.

To unpack the interaction, within 1:1, Center was faster than GazeFixation, and GazeFixation was slower than both LastGlobal and LastLocal. Within Boundary Switching, no significant pairwise differences among cursor-warping levels were found. Considering counter-rotation within each cursor-warping level, 1:1 was

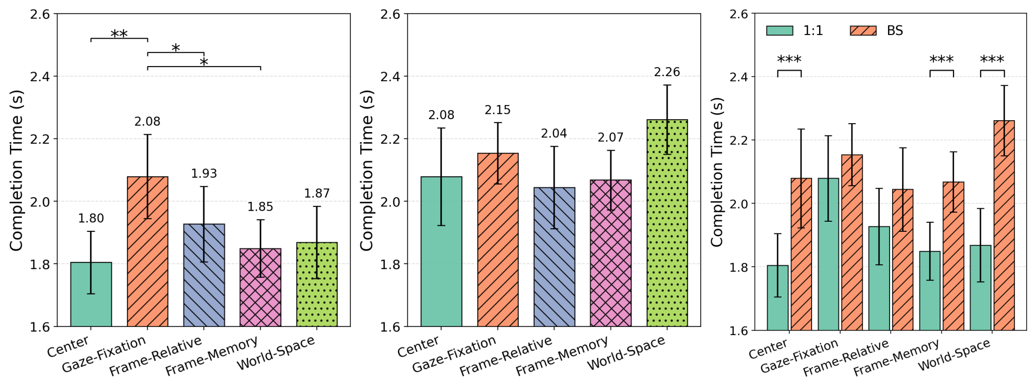


(a) 1:1 (b) Boundary Switching (c) Merged Interaction

Figure 4.5: 3D Completion Time (same y-axis limits across panels; error bars are within-subject SE, Cousineau–Morey).

faster than Boundary Switching for Center, LastLocal, and WorldSpace, with no significant differences for GazeFixation or LastGlobal. Collapsed across counter-rotation, GazeFixation was slower than both LastGlobal and LastLocal; collapsed across cursor warping, 1:1 was faster than Boundary Switching. Bonferroni-corrected post-hoc comparisons supporting these simple-effects findings are reported in Appendix A.1.

Overall, this three-display subset again shows that 1:1 benefited Center, LastLocal, and WorldSpace, whereas GazeFixation was specifically disadvantaged under 1:1 but not under Boundary Switching.

### 4.3.2 Target Acquisition Time

A repeated-measures ANOVA (Greenhouse–Geisser corrected) revealed significant main effects of cursor warping, $F(3.34, 63.48) = 3.55$, $p = .016$, $ges$=.049, and counter-rotation, $F(1, 19) = 57.62$, $p < .001$, $ges$=.177, as well as a strong interaction, $F(2.60, 49.44) = 17.32$, $p < .001$, $ges$=.206; see Figure 4.6.

To unpack the interaction, within 1:1, GazeFixation was slower than Center, LastGlobal, LastLocal, and WorldSpace (all $p \leq .043$). Within Boundary Switching, GazeFixation was faster than Center, LastLocal, and WorldSpace (all $p \leq .028$), and

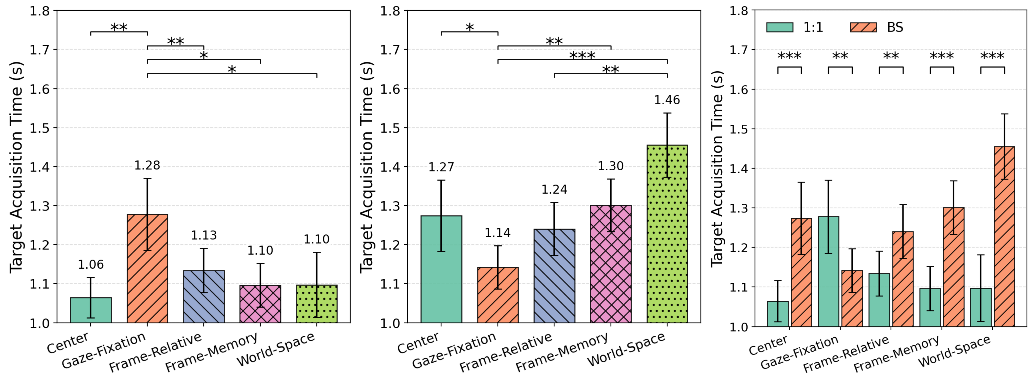


(a) 1:1 (b) Boundary Switching (c) Merged Interaction

Figure 4.6: 3D Target Acquisition Time (same y-axis limits across panels; error bars are within-subject SE, Cousineau–Morey).

LastGlobal was faster than WorldSpace ($p = .006$). Considering counter-rotation within each cursor-warping level, 1:1 was faster than Boundary Switching for Center, LastGlobal, LastLocal, and WorldSpace (all $p \leq .001$), whereas Boundary Switching was faster for GazeFixation ($p = .008$). Bonferroni-corrected post-hoc comparisons supporting these simple-effects findings are reported in Appendix A.1.

Collapsed across counter-rotation, no significant pairwise differences among cursor-warping levels were found; collapsed across cursor warping, 1:1 was faster than Boundary Switching ($p < .001$).

Overall, this three-display subset again shows a pronounced interaction: GazeFixation excelled under Boundary Switching, whereas 1:1 favored Center and LastLocal and reduced the cost of WorldSpace.

### 4.3.3 Head Movement

A repeated-measures ANOVA (Greenhouse–Geisser corrected) revealed significant main effects of cursor warping, $F(2.95, 56.04) = 20.27$, $p < .001$, $ges$=.231, and counter-rotation, $F(1, 19) = 113.27$, $p < .001$, $ges$=.312, as well as a significant interaction, $F(3.13, 59.40) = 8.60$, $p < .001$, $ges$=.094; see Figure 4.7.

Decomposing the interaction, within 1:1, GazeFixation reduced head movement

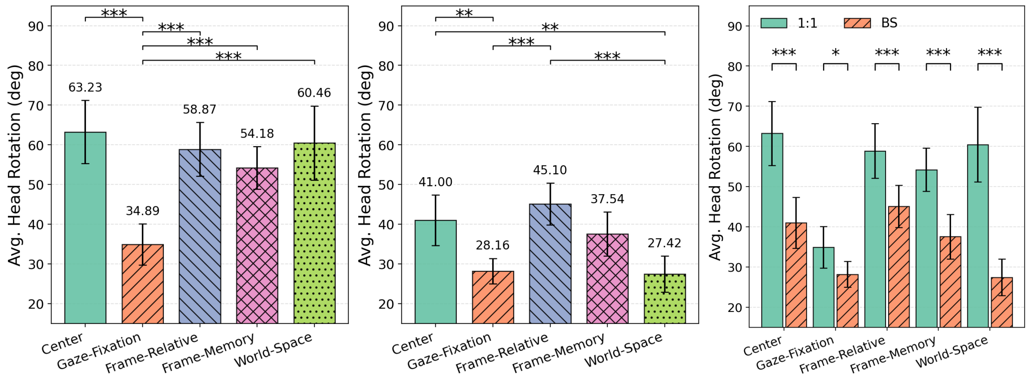


(a) 1:1 (b) Boundary Switching (c) Merged Interaction

Figure 4.7: 3D Head Movement (same y-axis limits across panels; error bars are within-subject SE, Cousineau–Morey).

relative to all other cursor-warping levels (all $p < .001$). Within Boundary Switching, Center exceeded GazeFixation and WorldSpace, GazeFixation was lower than LastGlobal, and LastGlobal exceeded WorldSpace (all $p \leq .007$); no other differences were significant. Considering counter-rotation within each cursor-warping level, 1:1 produced more head movement than Boundary Switching for Center, GazeFixation, LastGlobal, LastLocal, and WorldSpace (all $p \leq .016$). Bonferroni-corrected post-hoc comparisons supporting these simple-effects findings are reported in Appendix A.1.

Collapsed across counter-rotation, post-hoc tests showed that GazeFixation yielded less head movement than Center, LastGlobal, LastLocal, and WorldSpace (all $p \leq .004$), and LastGlobal exceeded WorldSpace ($p = .002$). Collapsed across cursor warping, 1:1 produced more head movement than Boundary Switching ($p < .001$).

Overall, GazeFixation consistently minimized head movement, whereas Boundary Switching reduced motion relative to 1:1 across cursor-warping levels.

### 4.3.4 Mouse Movement

As normality was violated, I analyzed this measure using ART. The analysis revealed a strong main effect of cursor warping, $F(4, 76) = 56.50$, $p < .001$, *partial* $\eta^2$=.748, no main effect of counter-rotation, $F(1, 19) = 1.77$, $p = .199$, *partial* $\eta^2$=.085, and a significant interaction, $F(4, 76) = 3.39$, $p = .013$, *partial* $\eta^2$=.151 (see Figure 4.8).

Collapsed across counter-rotation, post-hoc comparisons showed that WorldSpace required more mouse movement than all other cursor-warping levels (all $p < .001$), and Center and GazeFixation required less movement than LastLocal (both $p \leq .035$). No significant differences were found between counter-rotation modes overall. Given the significant interaction, I examined simple effects. However, post-hoc comparisons within each simple effect did not reveal consistent significant differences. I therefore summarize the omnibus interaction pattern only.

Overall, mouse-path length was primarily governed by cursor warping: WorldSpace demanded the most movement, while Center and GazeFixation minimized it. counter-rotation did not reliably modulate this pattern in the three-display subset.

### 4.3.5 Errors

As normality was violated, I analyzed this measure using ART. The analysis revealed no significant main effects of cursor warping, $F(4, 76) = 1.61$, $p = .180$, *partial* $\eta^2$=.078, or counter-rotation, $F(1, 19) = 0.25$, $p = .625$, *partial* $\eta^2$=.013, and no significant interaction, $F(4, 76) = 1.20$, $p = .316$, *partial* $\eta^2$=.060. Accordingly, no post-hoc comparisons are interpreted.

Overall, error rates were comparable across cursor-warping and counter-rotation conditions in the three-display subset, consistent with the five-display analysis.

## 4.4 Subjective Measures

All subjective measures were analyzed with Friedman tests (within-subjects; block = participant; factor = technique). For the overall analysis, scores were averaged

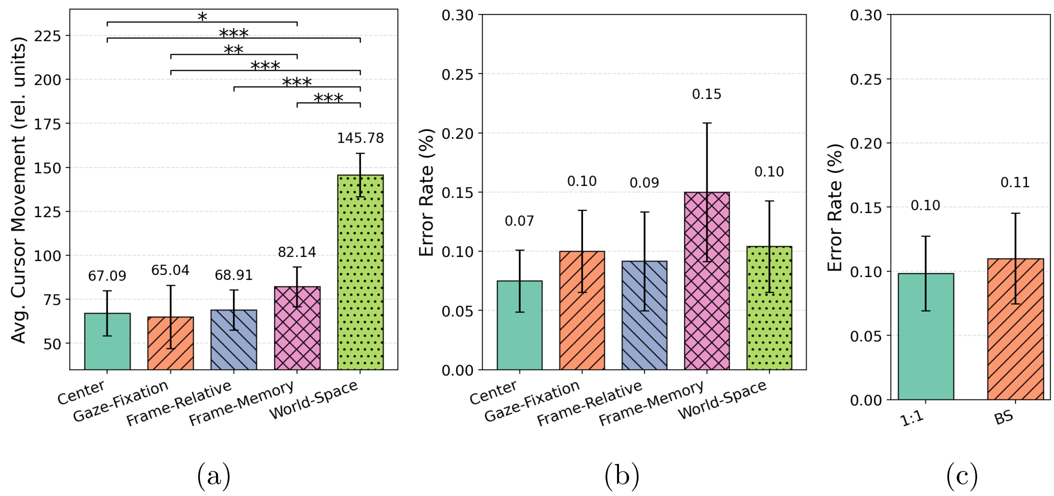


Figure 4.8: (a) Mouse Movement (Main Effect of CursorWarping); (b) Error Rate (Main Effect of CursorWarping); (c) Error Rate (Main Effect of Counter-rotation), all in the 3D condition. All panels share the same y-axis limits for comparability; error bars show within-subject SE (Cousineau–Morey).

at the participant×technique level across counter-rotation modes to satisfy the one-observation-per-block-and-condition requirement. I report Kendall's $W$ as an effect size (small ≈ .10, medium ≈ .30, large ≈ .50). When the omnibus test was significant, I ran paired Wilcoxon signed-rank tests with Bonferroni adjustment. Higher values indicate more *workload* for NASA-TLX subscales (mental, physical, temporal, effort, frustration, performance[1]), more *sickness* for the 21-point scale, and higher *satisfaction/SEQ* on 7-point scales. Medians illustrate direction.

For visualization, I show per-technique medians with 95% CIs; see Figure 4.9 a–c. Significant Bonferroni-adjusted pairs are annotated in the panels.

There was a significant effect of technique on overall *mental demand*, $\chi^2(4) = 13.69$, $p = .008$, $W = .171$ (small). Bonferroni-adjusted pairwise tests showed WorldSpace > Center ($p_{\text{adj}} = .033$); medians were 30.00 vs. 13.75, respectively (see Figure 4.9a).

[1]Higher "Performance" reflects greater workload on the NASA-TLX dimension, not objective accuracy.

*Overall satisfaction* also varied by technique, $\chi^2(4) = 12.25$, $p = .016$, $W = .153$ (small), but no Bonferroni-adjusted pair reached significance; medians suggested Center/GazeFixation/Last-global (all 5.00) > WorldSpace/Last-local (both 4.50) (see Figure 4.9c).

*Physical demand* showed a similar pattern, $\chi^2(4) = 12.12$, $p = .016$, $W = .151$ (small), with no Bonferroni-significant pairs; medians indicated higher load for *Last-global* (40.00) and WorldSpace (35.00) than GazeFixation (22.50) (see Figure 4.9b).

*SEQ* did not show a significant omnibus effect overall, $\chi^2(4) = 7.41$, $p = .116$, $W = .093$ (trivial).

Other measures were not significant overall (all $p \geq .054$), including *NASA-TLX total* ($\chi^2(4) = 9.31$, $p = .054$, $W = .116$), *frustration* ($\chi^2(4) = 9.12$, $p = .058$, $W = .114$), *temporal demand* ($\chi^2(4) = 8.55$, $p = .073$, $W = .107$), *effort* ($\chi^2(4) = 6.62$, $p = .157$, $W = .083$), and *performance* ($\chi^2(4) = 7.51$, $p = .111$, $W = .094$). *Motion sickness* was also not significant overall, $\chi^2(4) = 2.98$, $p = .562$, $W = .037$ (trivial).[2]

### 4.4.1 1:1

Technique significantly affected *mental demand*, $\chi^2(4) = 15.92$, $p = .003$, $W = .199$ (small). Bonferroni tests showed Center < Last-local ($p_{\text{adj}} = .006$); medians: 12.50 vs. 20.00. *Physical demand* was also significant, $\chi^2(4) = 10.32$, $p = .035$, $W = .129$ (small). Bonferroni tests indicated Last-global > GazeFixation ($p_{\text{adj}} = .012$); medians: 50.00 vs. 25.00. Overall satisfaction reached significance, $\chi^2(4) = 9.63$, $p = .047$, $W = .120$ (small), but no Bonferroni-significant pairs; medians again favored Center/GazeFixation/Last-global (5.00) over WorldSpace/Last-local (4.50). All other scales were not significant in this counter-rotation.

[2] Although one Bonferroni-adjusted pair for Performance (WorldSpace vs. Last-local) reached nominal significance ($p_{\text{adj}} = .040$), the omnibus test for Performance was not significant; I therefore do not interpret this comparison.

### 4.4.2 Boundary Switching

Technique significantly affected *SEQ*, $\chi^2(4) = 20.32$, $p < .001$, $W = .254$ (small). Bonferroni tests found WorldSpace < Center ($p_{\text{adj}} = .036$) and WorldSpace < GazeFixation ($p_{\text{adj}} = .049$); medians: 4.00 vs. 5.00/6.00. *Frustration* and *temporal demand* also had significant omnibus effects ($\chi^2(4) = 13.96$, $p = .007$, $W = .174$; $\chi^2(4) = 10.76$, $p = .029$, $W = .134$; both small), but no Bonferroni-significant pairs; medians suggested lower frustration and temporal demand for GazeFixation/Last-global compared to WorldSpace, yet these trends should be interpreted cautiously. Other scales were not significant in Boundary Switching.

Across techniques, WorldSpace tended to increase mental and physical workload and depress SEQ (under Boundary Switching). GazeFixation consistently ranked among the best (lower workload, especially physical, and higher SEQ), while *Last-global* incurred higher physical demand. Center often delivered lower mental demand with good satisfaction. Effects were small (all $W \approx .11$–$.25$), indicating modest but reliable differences.

A few scales showed omnibus $p$-values slightly above threshold, suggesting trends that did not survive Bonferroni correction. For example, overall *NASA-TLX total* approached significance ($\chi^2(4) = 9.31$, $p = .054$, $W = .116$), with higher medians for WorldSpace (32.5) than GazeFixation (22.5). Similarly, *frustration* ($p = .058$) and *temporal demand* ($p = .073$) hinted at higher workload for WorldSpace relative to other techniques. These effects should be interpreted as exploratory and not confirmatory.

I also asked participants which cursor-warping technique they preferred in each counter-rotation. In 1:1, Center was chosen most frequently (8/20, 40%), followed by *Last-global* (6/20, 30%), with GazeFixation and WorldSpace tied (3/20, 15% each). In Boundary Switching, GazeFixation led the preferences (9/20, 45%), followed by *Last-global* (4/20, 20%) and Center (4/20, 20%), with WorldSpace (2/20, 10%) and *Last-local* (1/20, 5%) less preferred. These descriptive preferences align with my inferential results: GazeFixation tends to yield lower workload and higher

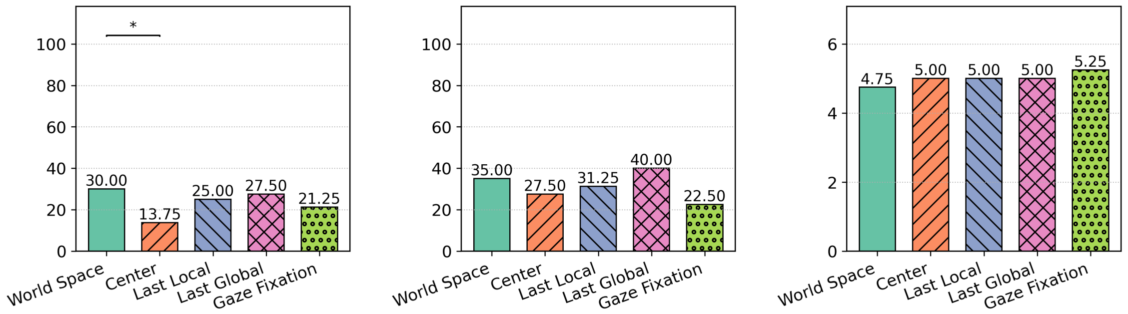


(a) Mental demand (overall) (b) Physical demand (overall) (c) Satisfaction / SEQ (overall)

Figure 4.9: Subjective measures by technique (overall across counter-rotation modes). Bars show medians; Bonferroni-significant pairs are annotated.

SEQ (especially under Boundary Switching), whereas WorldSpace is generally less favored.

## 4.5 Qualitative Feedback

Participants' comments reflected a mix of practical reasoning and personal preference. One recurring theme was the value of predictability. Several participants emphasized that knowing where the cursor would appear reduced effort: for example, one explained that Center was preferable because "you know where the starting point is; don't have to search" (P8, 1:1), while another highlighted that LastGlobal was "the quickest in my opinion" (P11, 1:1).

By contrast, GazeFixation provoked stronger reactions. Some described it as effortless, "Don't have to find the cursor. GazeFixation > Global > World space > Center & Local" (P7, Boundary Switching), while others found that its effectiveness varied with counter-rotation, noting that "the Boundary Switching turns very fast which can cause the eye tracking to be not as efficient" (P11, Boundary Switching). This divide suggests that gaze coupling can be compelling when it aligns with attentional focus, but is also sensitive to system dynamics.

Familiarity also shaped preferences. A few participants favored WorldSpace simply because it felt "the most familiar" (P7, 1:1), while others leaned toward Center as a "neutral" and intuitive option (P6, 1:1). These comments indicate that subjective comfort sometimes outweighed observed efficiency.

Boundary Switching itself was described in ambivalent terms. Some welcomed the reduction in manual navigation, "easier to navigate" (P8, Boundary Switching), whereas others experienced slower operation or even discomfort: "The 'Sea' wallpaper in Boundary Switching makes me motion sickness." (P15). Thus, while Boundary Switching could relieve physical effort, it occasionally introduced perceptual strain.

# Chapter 5

# Discussion

This section interprets the empirical results in light of my research questions and prior work. I first synthesize the main effects of cursor warping and counter-rotation and explain their interaction, then address each RQ (RQ1–RQ4) to identify effective pairings of view and cursor control. I conclude with design implications for XR desktops, broader applicability beyond XR, and limitations that bound the scope of my findings.

## 5.1 Overview

The study examined how counter-rotation techniques (1:1 vs. Boundary Switching) and cursor warping techniques (Center, LastLocal, LastGlobal, GazeFixation, WorldSpace) affect user performance in a selection task across five displays (5D) and three displays (3D).

My findings show that the way the cursor is initialized on a new display and how the user's view is controlled (with or without counter-rotational gain) both significantly influence performance and experience in multi-display XR workspaces. In general, I found that the baseline approach, akin to a standard multi-monitor pointer (continuous "crossing the boundary" without repositioning), was the least efficient. Participants took longer to complete trials, required more time to acquire targets, and moved the mouse much further under this baseline than with any cursor warping technique. In contrast, warping the cursor to a defined location on the target display yielded clear benefits: it sped up task completion and reduced unnecessary cursor travel. Notably, two simple warping strategies: placing the cursor at the screen's center, or at its last position on that screen, provided consistently good performance across conditions. A gaze-triggered warping strategy showed exceptional speed in certain situations, but its advantages depended strongly on the counter-rotation method.

Meanwhile, introducing a Boundary Switching counter-rotation (which subtly rotates the virtual workspace opposite to head turns) created an ergonomic benefit (less physical movement), but it also introduced a small performance penalty in speed for most warping techniques. Importantly, I observed a strong interaction: the effectiveness of a given cursor warping method often changed under Boundary Switching versus a fixed 1:1 counter-rotation. In the following, I discuss these outcomes in light of my research questions and outline implications for the design of XR multi-display interfaces.

## 5.2 RQ1: Is a screen-based cursor beneficial compared to traditional cursor placement?

Treating the cursor as bound to virtual screens rather than a continuous workspace proved advantageous. Across conditions, the conventional WorldSpace cursor (analogous to a standard multi-monitor pointer) consistently underperformed, particularly with Boundary Switching enabled. This echoes prior findings from physical multi-monitor setups, where continuous cross-screen motion incurs extra travel and *cursor loss* when monitors are separated (Benko and Feiner, 2005; Benko and Feiner, 2007; Baudisch et al., 2004b). By contrast, screen-based cursor warping, such as LastLocal, which resembles memory-based strategies shown to improve targeting times by up to 30% (Benko and Feiner, 2005), reduced acquisition time without compromising accuracy or comfort. Even in immersive XR, the drawbacks of a continuously moving cursor mirrored those in desktop environments: WorldSpace required longer movements for distant targets but provided no measurable benefit.

Importantly, screen-based warping did not increase error rates or workload. Participants adapted quickly to the cursor "jumping" to a relevant position rather than dragging it manually across displays. This confirms that the advantages of multi-monitor cursor assistance (Ashdown, Oka, and Sato, 2005; Benko and Feiner, 2007) extend to XR: reducing pointer travel and eliminating display-border friction improves efficiency without cost to accuracy.

## 5.3 RQ2: Where should the cursor be initialized when shifting attention to another display?

My results show that three strategies consistently supported efficient cross-display interaction: Center, LastLocal, and LastGlobal. All three significantly outperformed WorldSpace in speed, though each offered distinct advantages. Center provided high predictability, which participants valued (e.g., P8: "you know where the starting point is; don't have to search"). LastLocal reduced travel when participants resumed work on the same screen, while LastGlobal preserved frame-relative transitions, which some found natural (P20: "When I move to the new display, I still remember where I put the cursor at the previous display").

These findings extend prior multi-monitor studies. Benko et al. (Benko and Feiner, 2005; Benko and Feiner, 2007) reported that restoring the pointer to its last position on the target screen yielded the greatest benefits in physical setups. I similarly observed that LastLocal was among the fastest techniques, particularly under 1:1 where displays did not move. However, as earlier work cautioned (Benko and Feiner, 2007), such memory-based placement can impose cognitive demands: participants in my study noted that resuming from a blank screen felt less "obvious" than a center start, and NASA-TLX scores confirmed higher mental demand for

LastLocal than Center under 1:1. Thus, predictability can outweigh minimal travel distance when users re-engage with a display.

A hybrid approach appears promising: LastLocal can minimize movement when prior context is clear, while Center serves as a reliable fallback when uncertainty is high (e.g., after long gaps or window rearrangements). Indeed, Center performed robustly across conditions—it never lagged significantly behind alternatives in 1:1 and even outperformed GazeFixation in some cases. Participants also ranked it highly for ease of use, underscoring the value of a simple, well-known baseline: snapping to screen center offers a predictable "neutral" starting point, even if not always optimal in distance.

## 5.4 RQ3: Is counter-rotational gain (McGill's Boundary Switching) beneficial for multi-display cursor interaction?

The impact of Boundary Switching depended on the metric: it substantially reduced physical effort but introduced a modest penalty in pointing speed. Enabling Boundary Switching (Mcgill et al., 2020b) rotated the virtual workspace counter to head yaw, reducing the head movement needed to view the five-display panorama. From an ergonomic perspective this worked as intended, yet performance data revealed slower pointing unless Boundary Switching was paired with an appropriate cursor technique (RQ4).

Overall, trial completion times were slower with Boundary Switching than with 1:1 (selection times increased by about 5–10%, and ANOVAs showed significant main effects favoring 1:1). Why would a technique designed to reduce effort hinder selection speed? One likely factor is the disruption of spatial reference frames. With 1:1, head motion directly maps to gaze direction, yielding stable proprioceptive and visual cues. Under Boundary Switching, the counter-rotating environment requires recalibration, which may slow hand–cursor coordination. Prior work similarly observed that altered head rotations can impair spatial orientation and memory of target locations (Ragan et al., 2017; Freitag, Weyers, and Kuhlen, 2016). In my study, participants completed all selections without higher error rates or frustration, but slightly longer times suggest subtle adaptation, such as extra corrective movements or momentary hesitation.

The observed penalty was not universal. It mainly appeared when cursor techniques were not optimized for the moving view. Thus, Boundary Switching can be highly effective for multi-display interaction if interfaces accommodate it. Consistent with prior reports of reduced fatigue and workload (O'Hagan et al., 2025), Boundary Switching improves ergonomics and may support sustained comfort. However, naive use of Boundary Switching in XR desktops risks slower fine-grained pointing.

## 5.5 RQ4: Which cursor × counter-rotation combinations are optimal?

The strongest pairings were (i) Boundary Switching with GazeFixation and (ii) 1:1 with Center or LastLocal. These combinations capitalized on the strengths of both display control and cursor placement while limiting drawbacks. By contrast, the weakest outcome occurred when Boundary Switching was combined with WorldSpace, the baseline cursor technique that involves no warping. I discuss each notable pairing below.

### 5.5.1 Boundary Switching + GazeFixation: A winning synergy

Boundary Switching paired with GazeFixation formed the most effective combination. Under Boundary Switching, warping the cursor to the first fixation on the new display significantly outperformed all other warping methods in acquisition time (Figure 4.2). This pairing effectively offset Boundary Switching 's performance penalty: participants could rotate their view with reduced physical effort and have the cursor appear directly at the attended target, avoiding search or reorientation. Several participants highlighted this benefit (e.g., P7: "Don't have to find the cursor"), and selection times were often as fast as the best conditions across the study.

This result aligns with prior research showing that gaze naturally signals user intent and can streamline cross-display coordination (Nacenta, Mandryk, and Gutwin, 2008; Flegel, Pick, and Mentler, 2021). My contribution is to demonstrate, in an XR context with counter-rotating views, that gaze-based warping compensates for spatial disorientation. By aligning the cursor with the user's visual reference frame, the system maintains continuity even when the environment shifts. This echoes findings on mixed-initiative cursor control, where head movement supports coarse positioning and gaze supports fine targeting (Ashdown, Oka, and Sato, 2005; Zhai, Morimoto, and Ihde, 1999).

Practically, my data suggest that XR systems employing dynamic counter-rotation modes should integrate gaze-based cursor placement. With eye tracking now common in VR/AR headsets, this combination offers both ergonomic benefits and top performance: Boundary Switching + GazeFixation yielded the fastest selections and was rated most satisfying under Boundary Switching (45% of participants selected it as their preferred method).

### 5.5.2 1:1 + Center/LastLocal: Reliable defaults

Under 1:1, Center and LastLocal were the most efficient warping strategies. Both produced faster completion times than alternatives, with significant pairwise advantages over GazeFixation and large gains over WorldSpace. The two strategies addressed different priorities: Center maximized predictability, whereas LastLocal minimized pointer travel when users resumed work on the same screen.

In 1:1, Center yielded faster selections, was the most predictable option, and received the most first-choice votes (8/20, 40%). Center is therefore a strong default when the counter-rotation is static.

LastLocal achieved comparable efficiency, outperforming WorldSpace and several warp baselines on time, yet it received few first-choice votes in 1:1. As discussed in RQ2, the likely reason is predictability: when the prior cursor location was uncertain, participants needed a brief search. My task design (randomized displays and abstract targets) provided few semantic anchors, which may understate LastLocal 's value for real desktops where last locations often coincide with meaningful UI elements (e.g., text fields, search boxes, recently clicked controls). Nevertheless, long intervals between visits or many active displays can still undermine memory for last positions, reducing subjective acceptance.

LastGlobal did not show clear speed gains in 1:1 but attracted 6/20 (30%) first-choice votes, suggesting that some users perceived it as more predictable or comfortable than LastLocal.

In sum, for static-view XR desktops, 1:1 combined with Center or LastLocal are practical defaults, and LastGlobal is a viable alternative for users who prefer frame-relative continuity. Exposing these options and allowing context-aware selection can balance speed, predictability, and effort.

### 5.5.3 The worst-case pairing: Boundary Switching + WorldSpace

It is worth highlighting the combination that consistently yielded poor results: using Boundary Switching while keeping the cursor in a traditional WorldSpace mode (no warping). This pairing was effectively a "naïve" implementation of an XR multi-display workspace – the system moves the world for the user's convenience, but the cursor remains anchored in the global environment, forcing the user to realign it manually for each new view. Across my measures, Boundary Switching + WorldSpace was either the slowest or among the slowest conditions. For example, in the five-display analysis, every other warping method outperformed

WorldSpace under Boundary Switching on completion time, and target acquisition times for Boundary Switching + WorldSpace were dramatically worse than for Boundary Switching + any warping (in fact, simple-effects tests showed WorldSpace was significantly slower than all other techniques within the Boundary Switching condition). Users also found this combination cumbersome: only 10% chose WorldSpace as their favorite under Boundary Switching, and many participants' comments alluded to the frustration of "losing track of the cursor" or feeling that it lagged behind their gaze when the view rotated.

The underlying issue is that Boundary Switching accentuates the drawbacks of the unassisted cursor. When the system rotates the whole workspace to help the user, the cursor doesn't magically stay under the user's focus – it stays where it was in the global coordinate system. So a participant might turn their head 30° to a new display (with Boundary Switching perhaps only requiring a 15° physical turn thanks to amplification), but now the cursor is still back on the previous screen or floating somewhere off-target. The user must then engage in a catch-up behavior: either a large mouse movement to "carry" the cursor over, or an extra visual search to find that tiny white dot in a now-shifting scene. Prior research on multi-display interaction noted that when visual space and device space diverge, it severely impacts performance(Benko and Feiner, 2007), and this is exactly what happens here: Boundary Switching causes the visual frames to shift without a corresponding device-space pointer update. My findings make a strong case that designers should avoid using counter-rotation movement alone without coupling it to cursor repositioning. Simply put, enabling something like Boundary Switching without any cursor warping is asking for trouble: it is likely to slow users down and increase cognitive load, essentially negating some benefits of the counter-rotation assistance. This is a crucial insight for XR interface developers inspired by solutions like McGill's(Mcgill et al., 2020b). If you implement a counter-rotational workspace, be sure to also implement an intelligent cursor strategy (such as gaze warping or at least centering on the active display). Otherwise, as my data showed, users get the worst of both worlds: an “easier” view that actually feels slower to interact with.

## 5.6 3 Displays vs 5 Displays

Throughout my analysis, I examined both the full five-display scenario and a subset of three center displays. Importantly, I found that all the main patterns of results held true in both cases. Note, the three-display subset was not a separate condition but rather a way to confirm that my findings were not driven solely by the extreme peripheral monitors. It appears they do generalize: the interactions between cursor warping and counter-rotation remained significant and qualitatively similar when focusing only on the three central screens.

## 5.7 Implications for XR Desktop

The implications discussed in this section primarily target XR systems designed to support desktop-style productivity workflows. These include scenarios in which users rely on traditional input devices such as a mouse and keyboard while performing knowledge work across multiple virtual displays. Representative examples include headset-based extensions of conventional operating systems, such as using an XR device to augment a connected laptop or desktop computer with additional virtual monitors. My findings provide concrete implications for the design of next-generation XR desktop environments, such as virtual multi-monitor office setups. A key implication is that smart cursor warping should be implemented as a built-in feature of XR operating systems. For instance, an XR OS could offer a setting such as “Automatically reposition cursor when focusing on another window,” with configurable options including Center, LastLocal, and GazeFixation. Since all evaluated warping techniques demonstrated either neutral or positive effects compared to the WorldSpace baseline, enabling one by default could improve user performance immediately. Omitting a screen-based cursor concept in XR would risk repeating a well-documented shortcoming of early physical multi-monitor systems—an issue already addressed in prior research (Benko and Feiner, 2007; Waldner, Kruijff, and Schmalstieg, 2010).

At a broader design level, my results suggest that viewpoint control and pointer control should be treated as interdependent components rather than independent features. Developing mechanisms such as virtual display rotation (Boundary Switching) in isolation from pointer logic, or vice versa, overlooks the strong coupling between them. I therefore recommend that XR OS designers and developers treat these elements as co-dependent components; for example, when implementing Boundary Switching, an appropriate cursor management policy should be provided in tandem.

## 5.8 Implications Beyond XR

Although my study was conducted in VR, the patterns likely extend to other multi-display contexts with head or eye tracking. For physical multi-monitor setups, prior work cautions that pushing a WorldSpace pointer across screens is suboptimal (Benko and Feiner, 2007). My results provide concrete defaults for tracked desktops: adopt screen-based warping by default; prefer Center or LastLocal on re-entry to a display; and, when eye tracking is available, use GazeFixation to reduce search and reorientation overhead.

Augmented Reality presents a related opportunity. Because many AR devices support eye tracking, GazeFixation can align the pointer with the user's visual reference frame. Memory cues analogous to LastLocal may further aid resumption. Lystbæk et al.'s *Spatial Gaze Markers* help users return to prior points of interest by leaving transient traces at gaze endpoints (Lystbæk et al., 2024). This principle suggests AR desktops could retain per-display last positions (or a small shortlist) to accelerate re-engagement, complementing gaze-based placement.

Large, high-resolution single displays pose similar challenges due to long pointer travel distances (Pavanatto, Grubert, and Bowman, 2025). Treating regions of a powerwall as virtual monitors enables the same toolkit: partition the surface, apply Center or LastLocal on region switches, and, with eye tracking, warp toward distant regions using GazeFixation. In all cases, matching dynamic view or attention shifts with proactive, screen-based cursor placement preserves continuity and reduces effort.

## 5.9 Limitations

My study examined discrete target selection across multiple virtual screens, but several limitations remain. First, the task was a simplified point-and-click with clearly marked targets. Realistic workflows involve compound operations such as window management, text editing, or drag-and-drop across displays. These activities may reveal different dynamics, for instance when sustained cursor movement across boundaries is required. Prior multi-monitor research suggests that explicit triggers or visual cues can assist such transitions (Benko and Feiner, 2005; Nacenta, Mandryk, and Gutwin, 2008; Waldner, Kruijff, and Schmalstieg, 2010). Complex activities may also amplify coordination costs of switching, as observed in earlier work on multi-monitor behavior (Ball and North, 2005; Rashid, Nacenta, and Quigley, 2012). Future studies should test cursor warping in richer scenarios to assess whether the observed benefits extend beyond isolated selections.

Second, I tested one layout and gain setting: five displays in a horizontal row with a single counter-rotation magnitude for Boundary Switching. XR environments, however, support diverse geometries such as grids, curved panoramas, or vertical stacks. Prior work shows that display arrangements influence pointing behavior (Pavanatto, Davari, et al., 2023; Pavanatto, Lu, et al., 2025), suggesting that preferences for cursor placement could shift in non-horizontal or more immersive setups. Likewise, my counter-rotation gain was moderate, consistent with earlier ergonomic findings (Mcgill et al., 2020b). Larger gains (e.g., 4:1) can impair orientation and performance (Ragan et al., 2017; Freitag, Weyers, and Kuhlen, 2016), potentially increasing pointing time unless paired with appropriate warping. Conversely, weaker gains or alternative view control (e.g., vertical-axis assistance (O'Hagan et al., 2025)) may better balance ergonomics and speed. Broader exploration of view manipulations and layouts is needed to establish robustness.

Third, the accuracy and reliability of eye tracking constrained GazeFixation. My implementation generally worked well but was occasionally offset by calibration drift or latency. As VR/AR eye trackers improve, gaze-initialized warping should become more precise and less error-prone. Prior work indicates that gaze-assisted pointing strongly depends on tracker quality and calibration (Flegel, Pick, and Mentler, 2021). Designers may meanwhile adopt hybrid strategies, such as confirming gaze targets with a click or dwell, to mitigate occasional errors.

Fourth, my participant sample and exposure duration limit generality. Most participants had little XR experience, and preferences were likely influenced by familiarity with desktop metaphors. This helps explain why WorldSpace, despite slower performance, was perceived as "familiar," whereas GazeFixation was not immediately preferred. Longer-term use may shift perceptions: with practice, users may adapt to gaze-based warping and value its efficiency. My study captured only initial performance and impressions. Longitudinal deployments could reveal how users integrate warping into daily workflows, whether preferences evolve (e.g., greater acceptance of GazeFixation), and whether long-term strategies emerge. Expert or frequent XR users may also behave differently from novices. Broader sampling and extended exposure are therefore necessary to inform lasting design guidelines for XR multi-display pointing.

Finally, although I designed a broader set of displays rotation techniques, the experiment formally compared only 1:1 and Boundary Switching. This establishes a baseline–versus–boundary-switching contrast that is widely discussed in the literature (Mcgill et al., 2020a), but it does not exhaust the space of boundary-triggered or gain-based implementations. Different boundary-switching triggers (e.g., gaze-triggered) or continuous gain mappings with deadzones could interact with mouse-based cursor control in ways that alter both performance and preference. Notably, in my mouse-driven task, 1:1 outperformed Boundary Switching across speed measures even though 1:1 incurred greater head movement—an outcome that contrasts with prior reports favoring boundary switching for ergonomics and coverage (Mcgill et al., 2020a). Future work should therefore re-evaluate multiple boundary-switching variants and gain schedules under mouse input, and extend comparisons to 2D layouts, to determine whether the "best" rotation assistance in prior work remains optimal once cursor warping and mouse dynamics are jointly considered.

## 5.10 Expectations and Generalizations

Regarding population and setting generalizability, I expect the underlying human factors to extend beyond VR. My results can apply broadly across AR, CAVE systems, and multi-monitor configurations with head tracking. However, context-specific calibration remains necessary. Display size, rotation gain magnitude, and gaze tracking precision establish boundary conditions. I anticipate core findings (e.g., gaze warp efficacy with dynamic viewports) will generalize provided moderate gain values and adequate tracking accuracy are maintained.

Future work should explore adaptive and intelligent enhancements. Generalizing cursor warping to predictive placement based on learned user workflows offers promising directions. Systems could anticipate cursor targets following common action sequences, extending beyond static rules toward dynamic, personalized interaction. While such approaches align with intelligent interface trends and suggest potential gains, they introduce unpredictability risks.

I also consider generalization to broader XR display concepts. While my study examined clearly delineated virtual monitors, XR environments increasingly feature borderless floating windows and 3D content. A natural extension treats any region of interest or UI cluster as a logical display for pointer management. For instance, floating tool palettes in 3D design applications could employ magnetic cursor behavior, snapping focus as users direct attention toward them. The fundamental question, "Where should the cursor transition when user attention shifts?", remains relevant as XR UIs become more spatially distributed. Future interfaces for coding, design, and analysis will benefit from treating pointer transitions as central design considerations rather than afterthoughts. My research provides foundational guidelines: define logical zones and implement intelligent pointer transfer mechanisms.

In conclusion, this study validates established solutions within XR contexts while revealing critical interaction effects. My findings demonstrate that XR workspaces can achieve both enhanced ergonomics and improved efficiency through thoughtful design. By addressing the interplay between head and pointer movement, I establish groundwork for immersive desktops that enhance productivity without compromising comfort. I encourage researchers and designers to extend these findings across diverse scenarios, continuing to refine the integration of view and cursor control in spatial computing environments.

# Chapter 6

# Conclusions

This paper systematically examined how cursor warping techniques and counter-rotation interact in immersive multi-display XR. My study showed that cursor warping on each target display clearly outperforms the conventional WorldSpace approach. Screen-based strategies such as Center, LastLocal, and LastGlobal consistently reduced selection times and pointer travel without increasing errors or workload. These results indicate that cursor warping is well suited for XR desktops and avoids the inefficiencies often observed in multi-display workspaces.

The effectiveness of a given cursor warping strategy depended strongly on counter-rotation. With Fixed View, Center and LastLocal provided efficient and predictable defaults that minimized both search and mental effort. With Boundary Switching, by contrast, GazeFixation was the only strategy that fully offset the orientation costs of counter-rotated displays, enabling rapid and reliable selection. The combination of Boundary Switching with WorldSpace yielded the weakest outcomes, reinforcing that cursor control and counter-rotation exhibit strong co-dependence, rather than functioning as independent mechanisms.

From these findings we draw three main implications. First, XR desktops should enable automatic cursor repositioning by default, using strategies such as Center or LastLocal to minimize unnecessary search. Second, if systems apply counter-rotation for ergonomic benefits, they should couple it with gaze-based warping (GazeFixation) to preserve both comfort and speed. Third, designers should avoid WorldSpace cursors in XR multi-display environments, as they consistently underperformed. Together, these guidelines provide a concrete design space for pointer initialization in immersive desktops.

Beyond these immediate implications, the study contributes a broader methodological advance by combining controlled evaluation of cursor warping with manipulations of display orientation. This approach demonstrates how input techniques must be evaluated not in isolation but in concert with viewport and workspace dynamics. While my evaluation focused on a horizontal five-display arrangement and a single pointing task, the principles extend to larger or irregular layouts, hybrid XR-desktop setups, and tasks involving mixed input modalities such as touch or pen.

Future research should examine adaptive systems that automatically choose among cursor warping strategies based on user context, display motion, or task demands. Larger-scale longitudinal studies are also needed to capture how users internalize warping rules over time and whether preferences shift with expertise. In addition, exploring how cursor warping interacts with emerging XR workspace features, such as predictive window placement or multimodal selection. These could further refine the ergonomics and productivity of immersive environments.

In summary, my results provide a roadmap for XR workspace design. By aligning counter-rotation with adaptive cursor warping, systems can support both efficiency and ergonomics, helping XR desktops evolve from experimental prototypes into practical, everyday alternatives to physical multi-monitor setups.

# Appendix A

# Results

## A.1 Additional Tables

Table A.1: Critical post-hoc comparisons for the **5-displays** subset.

| DV | Contrast | $p_{\text{adj}}$ |
|---|---|---|
| Completion Time | LastLocal vs WorldSpace | 0.0030 |
| Completion Time | Center vs WorldSpace | 0.024 |
| Completion Time | 1:1 vs Boundary Switching | $1.5 \times 10^{-4}$ |
| Completion Time | Center vs GazeFixation | 0.010 |
| Completion Time | GazeFixation vs WorldSpace | $9.2 \times 10^{-4}$ |
| Completion Time | LastLocal vs WorldSpace | 0.0020 |
| Completion Time | Center vs WorldSpace | 0.0060 |
| Completion Time | LastGlobal vs WorldSpace | 0.039 |
| Completion Time | 1:1 vs Boundary Switching | 0.0070 |
| Completion Time | 1:1 vs Boundary Switching | 0.010 |
| Completion Time | 1:1 vs Boundary Switching | $6.7 \times 10^{-6}$ |
| Head Movement | GazeFixation vs LastGlobal | $1.6 \times 10^{-6}$ |

*Continued on next page*

| DV | Contrast | $p_{\text{adj}}$ |
|---|---|---|
| Head Movement | GazeFixation vs LastLocal | $3.4 \times 10^{-6}$ |
| Head Movement | Center vs GazeFixation | $1.2 \times 10^{-4}$ |
| Head Movement | GazeFixation vs WorldSpace | 0.0020 |
| Head Movement | 1:1 vs Boundary Switching | $3.5 \times 10^{-10}$ |
| Head Movement | GazeFixation vs LastGlobal | $3.5 \times 10^{-6}$ |
| Head Movement | GazeFixation vs LastLocal | $1.2 \times 10^{-5}$ |
| Head Movement | Center vs GazeFixation | $2.8 \times 10^{-4}$ |
| Head Movement | GazeFixation vs WorldSpace | $6.8 \times 10^{-4}$ |
| Head Movement | GazeFixation vs LastGlobal | $9.0 \times 10^{-6}$ |
| Head Movement | LastGlobal vs WorldSpace | 0.0010 |
| Head Movement | Center vs GazeFixation | 0.010 |
| Head Movement | GazeFixation vs LastLocal | 0.014 |
| Head Movement | Center vs WorldSpace | 0.029 |
| Head Movement | LastGlobal vs LastLocal | 0.045 |
| Head Movement | 1:1 vs Boundary Switching | $1.4 \times 10^{-5}$ |
| Head Movement | 1:1 vs Boundary Switching | 0.048 |
| Head Movement | 1:1 vs Boundary Switching | $2.6 \times 10^{-4}$ |
| Head Movement | 1:1 vs Boundary Switching | $2.9 \times 10^{-8}$ |
| Head Movement | 1:1 vs Boundary Switching | $2.3 \times 10^{-6}$ |
| Mouse Movement | GazeFixation vs WorldSpace | $5.4 \times 10^{-21}$ |
| Mouse Movement | Center vs WorldSpace | $1.0 \times 10^{-18}$ |
| Mouse Movement | LastGlobal vs WorldSpace | $1.5 \times 10^{-17}$ |
| Mouse Movement | LastLocal vs WorldSpace | $5.7 \times 10^{-14}$ |
| Mouse Movement | GazeFixation vs LastLocal | 0.0025 |
| Target Acquisition | GazeFixation vs WorldSpace | $6.9 \times 10^{-5}$ |
| Target Acquisition | Center vs WorldSpace | 0.0030 |

*Continued on next page*

| DV | Contrast | $p_{\text{adj}}$ |
|---|---|---|
| Target Acquisition | LastLocal vs WorldSpace | 0.0050 |
| Target Acquisition | GazeFixation vs LastGlobal | 0.031 |
| Target Acquisition | 1:1 vs Boundary Switching | 0.0010 |
| Target Acquisition | Center vs GazeFixation | 0.023 |
| Target Acquisition | GazeFixation vs WorldSpace | $3.8 \times 10^{-8}$ |
| Target Acquisition | GazeFixation vs LastGlobal | $2.1 \times 10^{-5}$ |
| Target Acquisition | GazeFixation vs LastLocal | $2.8 \times 10^{-5}$ |
| Target Acquisition | Center vs GazeFixation | 0.0010 |
| Target Acquisition | Center vs WorldSpace | 0.0010 |
| Target Acquisition | LastLocal vs WorldSpace | 0.0020 |
| Target Acquisition | LastGlobal vs WorldSpace | 0.0040 |
| Target Acquisition | 1:1 vs Boundary Switching | 0.0020 |
| Target Acquisition | 1:1 vs Boundary Switching | $1.4 \times 10^{-5}$ |
| Target Acquisition | 1:1 vs Boundary Switching | 0.0050 |
| Target Acquisition | 1:1 vs Boundary Switching | $2.2 \times 10^{-5}$ |

Table A.2: Critical post-hoc comparisons for the **3-displays** subset.

| DV | Contrast | $p_{\text{adj}}$ |
|---|---|---|
| Completion Time | GazeFixation vs LastLocal | 0.0060 |
| Completion Time | GazeFixation vs LastGlobal | 0.015 |
| Completion Time | 1:1 vs Boundary Switching | $1.3 \times 10^{-6}$ |
| Completion Time | Center vs GazeFixation | 0.0010 |
| Completion Time | GazeFixation vs LastGlobal | 0.011 |
| Completion Time | GazeFixation vs LastLocal | 0.012 |
| Completion Time | 1:1 vs Boundary Switching | $8.6 \times 10^{-5}$ |

*Continued on next page*

| DV | Contrast | $p_{\text{adj}}$ |
|---|---|---|
| Completion Time | 1:1 vs Boundary Switching | $3.5 \times 10^{-4}$ |
| Completion Time | 1:1 vs Boundary Switching | $3.0 \times 10^{-6}$ |
| Head Movement | GazeFixation vs LastGlobal | $1.1 \times 10^{-6}$ |
| Head Movement | Center vs GazeFixation | $1.3 \times 10^{-4}$ |
| Head Movement | GazeFixation vs LastLocal | $2.1 \times 10^{-4}$ |
| Head Movement | LastGlobal vs WorldSpace | 0.0020 |
| Head Movement | GazeFixation vs WorldSpace | 0.0040 |
| Head Movement | 1:1 vs Boundary Switching | $1.9 \times 10^{-9}$ |
| Head Movement | GazeFixation vs LastGlobal | $3.2 \times 10^{-5}$ |
| Head Movement | Center vs GazeFixation | $1.5 \times 10^{-4}$ |
| Head Movement | GazeFixation vs LastLocal | $2.4 \times 10^{-4}$ |
| Head Movement | GazeFixation vs WorldSpace | $3.3 \times 10^{-4}$ |
| Head Movement | GazeFixation vs LastGlobal | $5.2 \times 10^{-6}$ |
| Head Movement | LastGlobal vs WorldSpace | $7.0 \times 10^{-5}$ |
| Head Movement | Center vs GazeFixation | 0.0060 |
| Head Movement | Center vs WorldSpace | 0.0070 |
| Head Movement | 1:1 vs Boundary Switching | $9.4 \times 10^{-6}$ |
| Head Movement | 1:1 vs Boundary Switching | 0.016 |
| Head Movement | 1:1 vs Boundary Switching | $1.5 \times 10^{-4}$ |
| Head Movement | 1:1 vs Boundary Switching | $8.3 \times 10^{-5}$ |
| Head Movement | 1:1 vs Boundary Switching | $6.2 \times 10^{-7}$ |
| Mouse Movement | GazeFixation vs WorldSpace | $1.0 \times 10^{-19}$ |
| Mouse Movement | Center vs WorldSpace | $1.0 \times 10^{-18}$ |
| Mouse Movement | LastGlobal vs WorldSpace | $1.1 \times 10^{-17}$ |
| Mouse Movement | LastLocal vs WorldSpace | $5.7 \times 10^{-14}$ |
| Mouse Movement | GazeFixation vs LastLocal | 0.0025 |

*Continued on next page*

| DV | Contrast | $p_{\text{adj}}$ |
|---|---|---|
| Mouse Movement | Center vs LastLocal | 0.035 |
| Target Acquisition | 1:1 vs Boundary Switching | $3.6 \times 10^{-7}$ |
| Target Acquisition | GazeFixation vs LastGlobal | 0.0020 |
| Target Acquisition | Center vs GazeFixation | 0.0030 |
| Target Acquisition | GazeFixation vs LastLocal | 0.019 |
| Target Acquisition | GazeFixation vs WorldSpace | 0.043 |
| Target Acquisition | GazeFixation vs WorldSpace | $7.6 \times 10^{-6}$ |
| Target Acquisition | GazeFixation vs LastLocal | 0.0010 |
| Target Acquisition | LastGlobal vs WorldSpace | 0.0060 |
| Target Acquisition | Center vs GazeFixation | 0.028 |
| Target Acquisition | 1:1 vs Boundary Switching | $2.5 \times 10^{-5}$ |
| Target Acquisition | 1:1 vs Boundary Switching | 0.0080 |
| Target Acquisition | 1:1 vs Boundary Switching | 0.0010 |
| Target Acquisition | 1:1 vs Boundary Switching | $2.0 \times 10^{-4}$ |
| Target Acquisition | 1:1 vs Boundary Switching | $5.4 \times 10^{-6}$ |

# References

Ashdown, Mark, Kenji Oka, and Yoichi Sato (Apr. 2005). “Combining Head Tracking and Mouse Input for a GUI on Multiple Monitors”. In: *CHI ’05 Extended Abstracts on Human Factors in Computing Systems*. Portland OR USA: ACM, pp. 1188–1191. ISBN: 978-1-59593-002-6. DOI: 10.1145/1056808.1056873. (Visited on 07/27/2025).

Ball, Robert and Chris North (2005). “Analysis of User Behavior on High-Resolution Tiled Displays”. In: *Human-Computer Interaction - INTERACT 2005*. Ed. by Maria Francesca Costabile and Fabio Paternò. Vol. 3585. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 350–363. ISBN: 978-3-540-28943-2 978-3-540-31722-7. DOI: 10.1007/11555261_30. (Visited on 07/27/2025).

Batmaz, Anil Ufuk et al. (2021). “No Jitter Please: Effects of Rotational and Positional Jitter on 3D Mid-Air Interaction”. In: *Proceedings of the Future Technologies Conference (FTC) 2020, Volume 2*. Ed. by Kohei Arai, Supriya Kapoor, and Rahul Bhatia. Vol. 1289. Cham: Springer International Publishing, pp. 792–808. ISBN: 978-3-030-63088-1 978-3-030-63089-8. DOI: 10.1007/978-3-030-63089-8_52. (Visited on 07/27/2025).

Baudisch, Patrick et al. (Apr. 2004a). “Mouse Ether: Accelerating the Acquisition of Targets across Multi-Monitor Displays”. In: *CHI ’04 Extended Abstracts on Human Factors in Computing Systems*. CHI EA ’04. New York, NY, USA: Association for Computing Machinery, pp. 1379–1382. ISBN: 978-1-58113-703-3. DOI: 10.1145/985921.986069. (Visited on 02/28/2025).

— (Apr. 2004b). “Mouse Ether: Accelerating the Acquisition of Targets across Multi-Monitor Displays”. In: *CHI ’04 Extended Abstracts on Human Factors in Computing Systems*. Vienna Austria: ACM, pp. 1379–1382. ISBN: 978-1-58113-703-3. DOI: 10.1145/985921.986069. (Visited on 07/27/2025).

Benko, Hrvoje and Steven Feiner (Apr. 2005). “Multi-Monitor Mouse”. In: *CHI ’05 Extended Abstracts on Human Factors in Computing Systems*. Portland OR USA: ACM, pp. 1208–1211. ISBN: 978-1-59593-002-6. DOI: 10.1145/1056808.1056878. (Visited on 07/27/2025).

— (2007). “Pointer Warping in Heterogeneous Multi-Monitor Environments”. In: *Proceedings of Graphics Interface 2007 on - GI ’07*. Montreal, Canada: ACM Press, p. 111. ISBN: 978-1-56881-337-0. DOI: 10.1145/1268517.1268537. (Visited on 07/27/2025).

Biener, Verena et al. (Mar. 2024). “Working with XR in Public: Effects on Users and Bystanders”. In: *2024 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*. Orlando, FL, USA: IEEE, pp. 779–780. ISBN: 979-8-3503-7449-0. DOI: 10.1109/VRW62533.2024.00186. (Visited on 07/27/2025).

Bovo, Riccardo et al. (Mar. 2024). “WindowMirror: An Opensource Toolkit to Bring Interactive Multi-Window Views into XR”. In: *2024 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*. Orlando, FL, USA: IEEE, pp. 436–438. ISBN: 979-8-3503-7449-0. DOI: 10.1109/VRW62533.2024.00085. (Visited on 07/27/2025).

Byers, James C and Susan G Hill (1989). “Traditional and raw task load index (TLX) correlations: Are paired comparisons necessary”. In: pp. 481, 485.

Cetin, Gokhan, Wolfgang Stuerzlinger, and John Dill (Oct. 2018). “Visual Analytics on Large Displays: Exploring User Spatialization and How Size and Resolution Affect Task Performance”. In: *2018 International Symposium on Big Data Visual and Immersive Analytics (BDVA)*, pp. 1–10. DOI: 10.1109/BDVA.2018.8534027. (Visited on 02/28/2025).

Chen, Yuzheng et al. (2025). "The Interaction of Cursor Warping and View Manipulation in Multi-Display XR Environments". Submitted to CHI 2026, under review.

Cohen, Jacob (Apr. 1973). "Eta-Squared and Partial Eta-Squared in Fixed Factor Anova Designs". In: *Educational and Psychological Measurement* 33.1, pp. 107–112. ISSN: 0013-1644. DOI: 10.1177/001316447303300111. (Visited on 09/01/2025).

Colvin, Janet, Nancy Tobler, and James A Anderson (2004). "Productivity and multi-screen computer displays". In: *Rocky Mountain Communication Review* 2.1, pp. 31–53.

Czerwinski, Mary et al. (n.d.). "Toward Characterizing the Productivity Benefits of Very Large Displays". In: ().

Dunn, Olive Jean (Mar. 1961). "Multiple Comparisons among Means". In: *Journal of the American Statistical Association* 56.293, pp. 52–64. ISSN: 0162-1459. DOI: 10.1080/01621459.1961.10482090. (Visited on 09/08/2025).

Farmani, Yasin and Robert J. Teather (Dec. 2020). "Evaluating Discrete Viewpoint Control to Reduce Cybersickness in Virtual Reality". In: *Virtual Reality* 24.4, pp. 645–664. ISSN: 1359-4338, 1434-9957. DOI: 10.1007/s10055-020-00425-x. (Visited on 07/27/2025).

Flegel, Nadine, Christian Pick, and Tilo Mentler (2021). "A Gaze-Supported Mouse Interaction Design Concept for State-of-the-Art Control Rooms". In: *Human Interaction, Emerging Technologies and Future Applications IV*. Ed. by Tareq Ahram, Redha Taiar, and Fabienne Groff. Vol. 1378. Cham: Springer International Publishing, pp. 208–216. ISBN: 978-3-030-73270-7 978-3-030-74009-2. DOI: 10.1007/978-3-030-74009-2_26. (Visited on 07/27/2025).

Freitag, Sebastian, Benjamin Weyers, and Torsten W. Kuhlen (Apr. 2016). "Examining Rotation Gain in CAVE-like Virtual Environments". In: *IEEE Transactions on Visualization and Computer Graphics* 22.4, pp. 1462–1471. ISSN: 1077-2626. DOI: 10.1109/TVCG.2016.2518298. (Visited on 07/27/2025).

Grubert, Jens et al. (Nov. 2018). “The Office of the Future: Virtual, Portable, and Global”. In: *IEEE Computer Graphics and Applications* 38.6, pp. 125–133. ISSN: 1558-1756. DOI: 10.1109/MCG.2018.2875609. (Visited on 02/28/2025).

Hill, Alex and Andrew Johnson (Mar. 2008). “Withindows: A Framework for Transitional Desktop and Immersive User Interfaces”. In: *2008 IEEE Symposium on 3D User Interfaces*. Reno, NE: IEEE, pp. 3–10. ISBN: 978-1-4244-2047-6. DOI: 10.1109/3DUI.2008.4476584. (Visited on 07/27/2025).

Jeong, Jaehoon, Haegyeom Choi, and Donghun Lee (Feb. 2025). “Multi-Mode Hand Gesture-Based VR Locomotion Technique for Intuitive Telemanipulation Viewpoint Control in Tightly Arranged Logistic Environments”. In: *Sensors* 25.4, p. 1181. ISSN: 1424-8220. DOI: 10.3390/s25041181. (Visited on 07/27/2025).

Keshavarz, Behrang and Heiko Hecht (Aug. 2011). “Validating an Efficient Method to Quantify Motion Sickness”. In: *Human Factors* 53.4, pp. 415–426. ISSN: 0018-7208. DOI: 10.1177/0018720811403736.

Lenth, Russell V. (2025). *Emmeans: Estimated Marginal Means, Aka Least-Squares Means*. Manual.

Lin, Cheng-Jhe and Yu-shao Chen (2023). “The Effects of Dual Display Configuration and Users’ Experience on Performance and Preference in Complex Spatial-Verbal Tasks”. In: *HCI International 2023 Posters*. Ed. by Constantine Stephanidis et al. Vol. 1832. Cham: Springer Nature Switzerland, pp. 351–357. ISBN: 978-3-031-35988-0 978-3-031-35989-7. DOI: 10.1007/978-3-031-35989-7_45. (Visited on 07/27/2025).

Ling, Chen et al. (Mar. 2017). “Are Two Better Than One? A Comparison Between Single- and Dual-Monitor Work Stations in Productivity and User’s Windows Management Style”. In: *International Journal of Human–Computer Interaction* 33.3, pp. 190–198. ISSN: 1044-7318, 1532-7590. DOI: 10.1080/10447318.2016.1231392. (Visited on 07/27/2025).

Lystbæk, Mathias N. et al. (May 2024). “Spatial Gaze Markers: Supporting Effective Task Switching in Augmented Reality”. In: *Proceedings of the 2024*

*CHI Conference on Human Factors in Computing Systems*. CHI '24. New York, NY, USA: Association for Computing Machinery, pp. 1–11. ISBN: 979-8-4007-0330-0. DOI: 10.1145/3613904.3642811. (Visited on 09/11/2025).

Mcgill, Mark et al. (May 2020a). "Expanding the Bounds of Seated Virtual Workspaces". In: *ACM Trans. Comput.-Hum. Interact.* 27.3, 13:1–13:40. ISSN: 1073-0516. DOI: 10.1145/3380959. (Visited on 02/28/2025).

— (June 2020b). "Expanding the Bounds of Seated Virtual Workspaces". In: *ACM Transactions on Computer-Human Interaction* 27.3, pp. 1–40. ISSN: 1073-0516, 1557-7325. DOI: 10.1145/3380959. (Visited on 07/27/2025).

Medeiros, Daniel et al. (Nov. 2022). "From Shielding to Avoidance: Passenger Augmented Reality and the Layout of Virtual Displays for Productivity in Shared Transit". In: *IEEE Transactions on Visualization and Computer Graphics* 28.11, pp. 3640–3650. ISSN: 1077-2626, 1941-0506, 2160-9306. DOI: 10.1109/TVCG.2022.3203002. (Visited on 07/27/2025).

Miller, Robert A., Cheryl K. Stenmark, and Kyle Van Ittersum (Dec. 2020). "Dual Computer Displays Reduce Extraneous Cognitive Load". In: *Journal of Computer Assisted Learning* 36.6, pp. 890–897. ISSN: 0266-4909, 1365-2729. DOI: 10.1111/jcal.12442. (Visited on 07/27/2025).

Moullec, Yann et al. (2024). "To Use or Not to Use Viewpoint Oscillations When Walking in VR ? State of the Art and Perspectives". In: *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–19. ISSN: 1077-2626, 1941-0506, 2160-9306. DOI: 10.1109/TVCG.2024.3436858. (Visited on 07/27/2025).

Nacenta, Miguel A., Regan L. Mandryk, and Carl Gutwin (Apr. 2008). "Targeting across Displayless Space". In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Florence Italy: ACM, pp. 777–786. ISBN: 978-1-60558-011-1. DOI: 10.1145/1357054.1357178. (Visited on 07/27/2025).

Ng, Alexander et al. (Oct. 2021). "The Passenger Experience of Mixed Reality Virtual Display Layouts in Airplane Environments". In: *2021 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. Bari, Italy: IEEE, pp. 265–274. ISBN: 978-1-6654-0158-6. DOI: 10.1109/ISMAR52148.2021.00042. (Visited on 07/27/2025).

O'Hagan, Joseph et al. (Apr. 2025). "Improving Ergonomic Viewing of Spatial XR Workspaces Through 2D Rotational Assistance". In: *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems.* Yokohama Japan: ACM, pp. 1–8. ISBN: 979-8-4007-1395-8. DOI: 10.1145/3706599.3719920. (Visited on 07/27/2025).

Olejnik, Stephen and James Algina (2003). "Generalized Eta and Omega Squared Statistics: Measures of Effect Size for Some Common Research Designs". In: *Psychological Methods* 8.4, pp. 434–447. ISSN: 1939-1463. DOI: 10.1037/1082-989X.8.4.434.

Pavanatto, Leonardo, Shakiba Davari, et al. (Dec. 2023). "Virtual Monitors vs. Physical Monitors: An Empirical Comparison for Productivity Work". In: *Frontiers in Virtual Reality* 4, p. 1215820. ISSN: 2673-4192. DOI: 10.3389/frvir.2023.1215820. (Visited on 07/27/2025).

Pavanatto, Leonardo, Jens Grubert, and Doug A. Bowman (Mar. 2025). "Spatial Bar: Exploring Window Switching Techniques for Large Virtual Displays". In: *2025 IEEE Conference Virtual Reality and 3D User Interfaces (VR).* Saint Malo, France: IEEE, pp. 186–194. ISBN: 979-8-3315-3645-9. DOI: 10.1109/VR59515.2025.00043. (Visited on 07/27/2025).

Pavanatto, Leonardo, Feiyu Lu, et al. (Mar. 2025). "Multiple Monitors or Single Canvas? Evaluating Window Management and Layout Strategies on Virtual Displays". In: *IEEE Transactions on Visualization and Computer Graphics* 31.3, pp. 1713–1730. ISSN: 1077-2626, 1941-0506, 2160-9306. DOI: 10.1109/TVCG.2024.3368930. (Visited on 07/27/2025).

Pavanatto, Leonardo, Chris North, et al. (Mar. 2021). "Do We Still Need Physical Monitors? An Evaluation of the Usability of AR Virtual Monitors for Productivity Work". In: *2021 IEEE Virtual Reality and 3D User Interfaces (VR)*, pp. 759–767. DOI: 10.1109/VR50410.2021.00103. (Visited on 02/28/2025).

Ragan, Eric D. et al. (Aug. 2017). "Amplified Head Rotation in Virtual Reality and the Effects on 3D Search, Training Transfer, and Spatial Orientation". In: *IEEE Transactions on Visualization and Computer Graphics* 23.8, pp. 1880–1895. ISSN: 1077-2626. DOI: 10.1109/TVCG.2016.2601607. (Visited on 07/27/2025).

Rashid, Umar, Miguel A. Nacenta, and Aaron Quigley (June 2012). "Factors Influencing Visual Attention Switch in Multi-Display User Interfaces: A Survey". In: *Proceedings of the 2012 International Symposium on Pervasive Displays.* Porto Portugal: ACM, pp. 1–6. ISBN: 978-1-4503-1414-5. DOI: 10.1145/2307798.2307799. (Visited on 07/27/2025).

Ruvimova, Anastasia et al. (Apr. 2020). ""Transport Me Away": Fostering Flow in Open Offices through Virtual Reality". In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems.* CHI '20. New York, NY, USA: Association for Computing Machinery, pp. 1–14. ISBN: 978-1-4503-6708-0. DOI: 10.1145/3313831.3376724. (Visited on 02/28/2025).

Saleem, Jason J. and Dustin T. Weiler (Sept. 2018). "Performance, Workload, and Usability in a Multiscreen, Multi-Device, Information-Rich Environment". In: *PeerJ Computer Science* 4, e162. ISSN: 2376-5992. DOI: 10.7717/peerj-cs.162. (Visited on 09/30/2025).

Sauro, Jeff and Joseph S. Dumas (Apr. 2009). "Comparison of Three One-Question, Post-Task Usability Questionnaires". In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems.* CHI '09. New York, NY, USA: Association for Computing Machinery, pp. 1599–1608. ISBN: 978-1-60558-246-7. DOI: 10.1145/1518701.1518946. (Visited on 09/08/2025).

Seibert, Jonmichael and Daniel M. Shafer (Mar. 2018). "Control Mapping in Virtual Reality: Effects on Spatial Presence and Controller Naturalness". In: *Virtual Reality* 22.1, pp. 79–88. ISSN: 1359-4338, 1434-9957. DOI: 10.1007/s10055-017-0316-1. (Visited on 07/27/2025).

Shapiro, S. S. and M. B. Wilk (1965). “An Analysis of Variance Test for Normality (Complete Samples)”. In: *Biometrika* 52.3/4, pp. 591–611. ISSN: 0006-3444. DOI: 10.2307/2333709. JSTOR: 2333709. (Visited on 09/01/2025).

Spittle, Becky et al. (Sept. 2023). “A Review of Interaction Techniques for Immersive Environments”. In: *IEEE Transactions on Visualization and Computer Graphics* 29.9, pp. 3900–3921. ISSN: 1077-2626, 1941-0506, 2160-9306. DOI: 10.1109/TVCG.2022.3174805. (Visited on 07/27/2025).

Stiefelhagen, Rainer and Jie Zhu (Apr. 2002). “Head Orientation and Gaze Direction in Meetings”. In: *CHI ’02 Extended Abstracts on Human Factors in Computing Systems*. CHI EA ’02. New York, NY, USA: Association for Computing Machinery, pp. 858–859. ISBN: 978-1-58113-454-4. DOI: 10.1145/506443.506634. (Visited on 09/12/2025).

Tara, Rayi Yanu and Wei-Chung Teng (2017). “A Suitability Evaluation of Controlling 3D Map Viewpoint by Gamepad Orientation for Remote Navigation”. In: *IEEE Access* 5, pp. 10686–10693. ISSN: 2169-3536. DOI: 10.1109/ACCESS.2017.2682275. (Visited on 07/27/2025).

Tserenchimed, Tuvshintulga and Hyungki Kim (June 2024). “Viewpoint-Sharing Method with Reduced Motion Sickness in Object-Based VR/AR Collaborative Virtual Environment”. In: *Virtual Reality* 28.3, p. 122. ISSN: 1434-9957. DOI: 10.1007/s10055-024-01005-z. (Visited on 07/27/2025).

Waldner, Manuela, Ernst Kruijff, and Dieter Schmalstieg (Oct. 2010). “Bridging Gaps with Pointer Warping in Multi-Display Environments”. In: *Proceedings of the 6th Nordic Conference on Human-Computer Interaction: Extending Boundaries*. Reykjavik Iceland: ACM, pp. 813–816. ISBN: 978-1-60558-934-3. DOI: 10.1145/1868914.1869036. (Visited on 07/27/2025).

Wang, Chen et al. (July 2023). “On Rotation Gains Within and Beyond Perceptual Limitations for Seated VR”. In: *IEEE Transactions on Visualization and Computer Graphics* 29.7, pp. 3380–3391. ISSN: 1077-2626, 1941-0506, 2160-9306. DOI: 10.1109/TVCG.2022.3159799. (Visited on 07/27/2025).

Wobbrock, Jacob O. et al. (May 2011). "The Aligned Rank Transform for Nonparametric Factorial Analyses Using Only Anova Procedures". In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '11. New York, NY, USA: Association for Computing Machinery, pp. 143–146. ISBN: 978-1-4503-0228-9. DOI: 10.1145/1978942.1978963. (Visited on 09/08/2025).

Yang, Yalong et al. (June 2018). "Maps and Globes in Virtual Reality". In: *Computer Graphics Forum* 37.3, pp. 427–438. ISSN: 0167-7055, 1467-8659. DOI: 10.1111/cgf.13431. (Visited on 07/27/2025).

Zhai, Shumin, 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 the CHI Is the Limit - CHI '99*. Pittsburgh, Pennsylvania, United States: ACM Press, pp. 246–253. ISBN: 978-0-201-48559-2. DOI: 10.1145/302979.303053. (Visited on 09/11/2025).

Zhan, Tao et al. (Aug. 2020). "Augmented Reality and Virtual Reality Displays: Perspectives and Challenges". In: *iScience* 23.8, p. 101397. ISSN: 25890042. DOI: 10.1016/j.isci.2020.101397. (Visited on 07/27/2025).