FlexiCAVE: A Dynamically Configurable High-Resolution Display Facility

ZAINAB AAMIR, Stony Brook University, USA
SAEED BOORBOOR, University of Illinois Chicago, USA
AHAMED SHOAIB, Stony Brook University, USA
ARIE KAUFMAN, Stony Brook University, USA

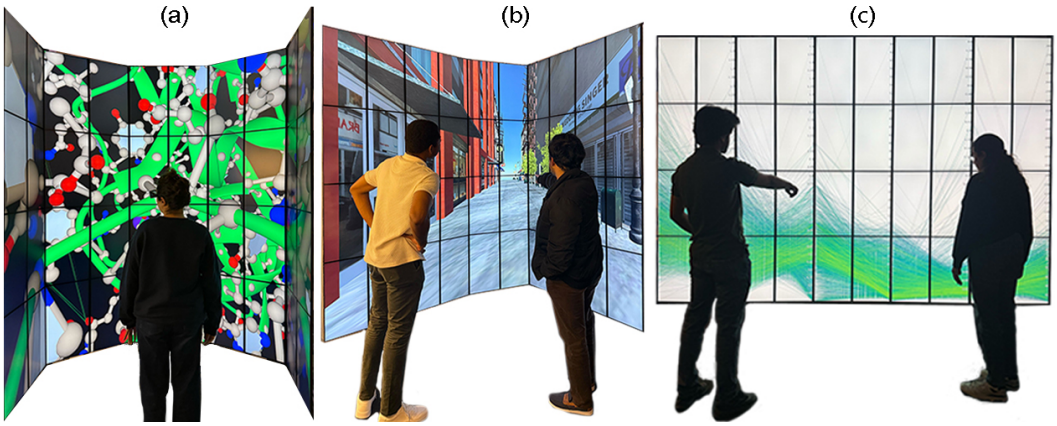
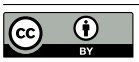


Fig. 1. The FlexiCAVE display facility can be dynamically transformed into various display configurations by adjusting the hinged display columns into various configurations, such as: (a) curved, (b) two-sided, or (c) flat.

Large high-resolution displays (LHRDs) are increasingly central to visualization and visual analytics, offering expansive visual space, simultaneous support for context and detail, and intuitive physical navigation. While advances in resolution and form factor have broadened their use, a persistent challenge remains: most LHRDs adopt static geometries, and even reconfigurable systems typically treat layout adjustment as a setup mechanism rather than as an interaction technique during active analysis. To address this, we introduce the FlexiCAVE, a dynamically reconfigurable LHRD environment featuring rotatable display columns that enable real-time adjustment of horizontal curvature. The system comprises forty high-pixel-density panels arranged across ten columns, providing active stereo visualization at approximately 83 million pixels. Each column rotates inward up to 90°, allowing the display to transform between flat and highly curved layouts during active use. A custom rendering engine synchronizes view perspectives with physical reconfiguration in real time. Through controlled studies and expert evaluations centered on 3D navigation, we found that dynamic reconfiguration reduced spatial judgment error by 51% compared to static flat layouts and by 26% compared to static curved layouts. It also enhanced users' sense of direction, with participants reporting that the physical display structure provided directional anchoring absent in static systems. We further explored

Authors' Contact Information: Zainab Aamir, Computer Science, Stony Brook University, Stony Brook, New York, USA, zaamir@cs.stonybrook.edu; Saeed Boorboor, Computer Science, University of Illinois Chicago, Chicago, Illinois, USA, boorboor@uic.edu; Ahamed Shoaib, Computer Science, Stony Brook University, Stony Brook, New York, USA, aahamedshoai@cs.stonybrook.edu; Arie Kaufman, Computer Science Dept, Stony Brook University, Stony Brook, New York, USA, ari@cs.stonybrook.edu.



This work is licensed under a Creative Commons Attribution 4.0 International License.
© 2026 Copyright held by the owner/author(s).
ACM 2573-0142/2026/11-ARTISS003
https://doi.org/10.1145/3859315

PIVoT (Physical Interaction to Virtual Transformation), mapping physical display rotation to visualization updates. Our findings reveal that physical reconfiguration is most effective when display geometry provides spatially meaningful correspondence to data operations. We derive design principles for reconfigurable display environments and discuss implications for next-generation large-scale visualization systems.

CCS Concepts: • **Human-centered computing → Displays and imagers**.

Additional Key Words and Phrases: Large High-Resolution Displays, Reconfigurable Displays, Interaction

ACM Reference Format:
Zainab Aamir, Saeed Boorboor, Ahamed Shoaib, and Arie Kaufman. 2026. FlexiCAVE: A Dynamically Configurable High-Resolution Display Facility. *Proc. ACM Hum.-Comput. Interact.* 10, 8, Article ISS003 (November 2026), 25 pages. https://doi.org/10.1145/3859315

1 Introduction

Large high-resolution displays (LHRDs) are considered powerful instruments to explore vast data sizes, offering unique perceptual and cognitive benefits for visualization and visual analytics [3, 6, 8, 46, 61]. Compared to traditional desktop settings, where a display screen occupies only a limited field of view (FoV) of the user, LHRDs provide expansive screen real estate, enabling users to engage with larger volumes of data at once. Moreover, they facilitate intuitive data exploration through physical navigation, where a novel view of the data is obtained simply by looking at different sections of the display. As such, it helps maintain information within the user's field of regard (FoR) and reduces the need for virtual navigation [6], while simultaneously providing details and context.

Over the past three decades, since the introduction of the first generation CAVE [21] and PowerWall [59], the proliferation of inexpensive GPU power and high-pixel-density displays has driven the evolution of LHRDs from mega- to giga-pixel resolutions [46], and with novel display configurations, spanning from planar and rectangular setups to more immersive, dome-like [27] and cylindrical [12, 26]. Parallel to these developments has been an ongoing question of the effects of display form factor and configuration on user perception and performance. Specifically, there have been studies comparing *flat* versus *curved* LHRDs, as well as different display curvatures across tasks such as visual search [34, 36, 49, 52], data visualization [35], and interaction [11, 53]. Taken together, these studies suggest that display geometry matters, but not in a uniform way: curved environments can increase immersion and support some search, comparison, and interaction tasks [12, 26, 49], whereas flatter layouts can remain advantageous for overview tasks, collaboration, and lower-density visualizations [35, 46, 49]. Interestingly, some studies have indicated that a significant difference in user performance cannot be conclusively identified between flat and varying curvature display configurations [36].

Rather than indicating a universally optimal form factor, this body of literature suggests that the most appropriate display geometry depends on the task, the spatial structure of the data, and the kind of navigation required. Important, however, is the realization that display geometry is treated as a static choice made before analysis begins. That is to say, researchers compare fixed flat and curved configurations, or evaluate several predetermined curvatures, but the geometry itself is rarely made available as part of interaction. This has led us to formulate an important question: **Can display geometry become an embodied interaction resource that users manipulate during analysis, particularly for tasks that require maintaining orientation while navigating 3D content?** We investigate this question through the lens of 3D navigation, where the relationship between body movement, display shape, and spatial understanding is especially direct, and where the effects of reconfiguration can be examined most clearly.

To explore this space, we designed and built *FlexiCAVE*, a dynamically reconfigurable LHRD that treats horizontal curvature not as a fixed installation parameter, but as an interactive property of

the environment. FlexiCAVE consists of forty high-pixel-density displays arranged into ten hinged columns, with each column composed of monitors mounted in a rotated (portrait) orientation. Each pair of columns can be rotated inward by up to 90°, allowing the overall display curvature to be reconfigured in real time. In its flat configuration, the system spans 3.13 m × 2.14m with a total resolution of 10800 × 7680 pixels (approximately 83 million pixels) and supports active stereoscopic rendering. Rather than serving as a single fixed layout, FlexiCAVE supports a range of configurations, including flat, curved, and *U*-shaped arrangements, effectively functioning as a reconfigurable immersive environment. A custom rendering architecture senses physical layout changes and keeps the virtual scene synchronized with the current display geometry. This allows users to reshape the display while interacting with data, enabling us to study when physical reconfiguration helps and when its benefits are limited. Figure 1 illustrates example configurations, and Figure 4 provides an overview of the system's mechanical and display components.

We further introduce *PIVoT* (Physical Interaction to Virtual Transformation), an interaction approach that maps physical display reconfiguration to corresponding changes in the visualization. Through this interface, we examine if (1) reconfigurable geometry supports spatial judgment and orientation during 3D navigation, and (2) coupling geometry changes to visualization updates is effective when the display shape and data structure are meaningfully aligned. To that end, we combine a controlled study centered on 3D navigation with application scenarios that help illustrate both the promise and the limits of this approach.

In summary, this paper makes the following contributions:

- We present *FlexiCAVE*, a dynamically reconfigurable LHRD with rotatable display columns that support real-time horizontal curvature adjustment, together with a rendering architecture that synchronizes physical reconfiguration with visualization updates.
- Through a controlled study focused on 3D navigation and application scenarios with domain experts, we show how dynamic reconfiguration can improve spatial judgment and orientation relative to fixed flat and fixed curved baselines, and we synthesize implications for the design of reconfigurable LHRD interactions.
- We introduce *PIVoT*, an interaction approach that maps physical display rotation to spatially aware visualization updates, and we characterize when physical reconfiguration is an effective interaction mechanism.

2 Background and Related Work

2.1 Evolution of LHRDs

LHRDs have evolved along two primary dimensions: visual acuity and immersion. The Cave Automatic Virtual Environment (CAVE) [21, 37], developed at the Electronic Visualization Lab (EVL), pioneered immersive projection-based visualization but was limited to 4 Megapixels (MP). The Immersive Cabin [47] and the PowerWall [59] doubled this resolution to 8MP, addressing the need to visualize high-resolution scientific simulations. Subsequent systems have pushed the boundaries of both resolution and form factor [22]. Enhanced stereoscopic projection systems include the six-sided Cornea [22] at KAUST and the five-sided AixCAVE [33] at RWTH Aachen University, both exceeding 100 MPs. Alternative geometries emerged with the pentagonal StarCAVE [23] and the three-story spherical AlloSphere [27].

The advancements in LCD technology in the early 2000s resulted in the emergence of tiled-LHRD facilities as a more feasible solution for large-scale visualization. These facilities, often spanning extensive areas, enabled a balance between visual acuity and immersion. Among the first to construct the next-generation LHRD was EVL with its CAVE2 [26], using LCD technology with 72 displays arranged in a cylindrical configuration. The Stony Brook University Reality Deck [46] achieved 1.6 Gigapixels (GP) across 416 displays powered by 104 GPUs. For a more immersive

setting, they have also introduced the Silo [12], a 0.6GP stereoscopic cylindrical facility. More recently, LED-based systems [38] have reduced space and maintenance requirements. Despite this diversity in resolution and form factor, a fundamental design question persists: how should these displays be geometrically configured to best support varied visualization tasks?

2.2 LHRDs and Display Geometry

LHRDs enable interaction with data at scales that fundamentally change how users engage with visual information [2, 7, 60]. Across form factors ranging from flat powerwalls to immersive CAVEs, these systems support visual search, comparison, navigation, and collaborative sensemaking by affording physical navigation and leveraging spatial memory [3, 7, 10]. Beyond scale, display geometry itself shapes how users interact with data [2, 30, 31, 50]. Prior work shows that curved displays can improve immersion and depth perception [12], while flat displays often better support overview and comparison tasks [35]. At human scale, these tradeoffs become especially important: LHRDs are not simply larger screens, but physical view spaces that users move within and around [9]. Shupp et al. [49] further show that the spatial placement of information and its relationship to the user directly shape interpretation and engagement. Nederlander et al. further demonstrate that physical extent and pixel density shape embodied navigation, with users physically moving to reason across scales in astronomical survey imagery [41]. Geometry is therefore not merely a construction decision, but a rich interaction resource. However, existing systems often leave this resource fixed.

2.3 Display Form Factors and User Perception

Beyond the structural evolution of LHRDs, previous works have examined how display geometry affects user performance, consistently revealing task-dependent tradeoffs. Curved displays tend to support spatial and navigational tasks: Shupp et al. [50] found that large curved displays enhance spatial performance and reduce physical navigation effort during multiscale geo-spatial search, and Kyung et al. [34] showed that moderate curvature optimizes accuracy and speed in visual search tasks, while flat and extreme curvatures may increase fatigue and response time. Physical navigation studies further show that locomotion combined with visual overview outperforms either alone for spatial memory [32], though movement does not uniformly improve performance without appropriate task motivation [31]. Flat layouts, by contrast, have shown to be more effective for structured comparison tasks. Liu et al. [35] demonstrated that flat configurations are more efficient for visualizing small multiples of smaller datasets, though semi-circular arrangements become preferable as the dataset size grows. Their subsequent study [36] found that full-circle layouts impaired spatial memory compared to flat arrangements due to disorientation from excessive head rotation. Curved displays have also been shown to enhance depth perception and navigation, outperforming flat powerwalls in visual exploration tasks [12]. Effects extend beyond navigation and search: Wei et al. [57] showed that curvature significantly influences reading legibility and comprehension in virtual reality (VR), while Ullah et al. [53] and Bihani et al. [11] demonstrated that pointing performance varies with both curvature and user position, and Sultana et al. [52] showed that display field-of-view and type significantly affect text readability and visual search on large VR displays. These findings establish that display geometry meaningfully shapes task performance, but in ways that are task- and context-dependent, the configuration that supports spatial navigation may hinder overview tasks, and vice versa [34, 35, 50]. This body of work points to a fundamental limitation of existing LHRD design: geometry is fixed at construction time, which leaves open the question of how to support users when task demands shift within a session or when different users have distinct preferences. FlexiCAVE addresses this by enabling dynamic physical reconfiguration during active use.

2.4 Reconfigurable Display Spaces and Interaction

A limited number of systems have explored adaptive large display environments. NASA Hyperwall [48] enables non-planar configurations through individual screen adjustment in pitch, yaw and translation of its 49-screen array, supporting flexible layouts for visualization and analysis. Dataspace [17] integrates 15 high-resolution displays mounted on 7-degree-of-freedom robotic arms with VR and augmented reality (AR) components to support collaborative immersive analytics in data-centric workspaces. These systems demonstrate the technical feasibility of adaptive display setups and the potential of flexible layouts for complex analytical workflows. However, these systems treat display adaptation primarily as a configuration or setup mechanism rather than as an interaction method. Geometry is adjusted before or between tasks, not continuously manipulated during active exploration.

FlexiCAVE addresses this gap directly: it is the first system to treat dynamic physical reconfiguration as a dedicated interaction technique, investigating how manipulation of display geometry can support navigation, spatial reasoning, and visualization in real time.

3 The FlexiCAVE

Belkacem et al. [9] envisioned LHRDs as "*a physical view space that is at least the size of the human body and exhibits a significantly higher resolution than a conventional display.*" In addition to scale, Shupp et al. [49] emphasized that the positioning of displayed information within the environment and its relationship to the user influence how the users interpret and engage with data. These insights have prompted the visualization and human-computer interaction communities to make significant contributions with respect to LHRD resolution and display form factors.

Beyond building a generic large-display installation, designing FlexiCAVE was driven by the goal of supporting a novel interaction setting for LHRDs, in which users can establish a relationship between displays and 3D scene navigation while preserving depth cues, and inspecting content at a human scale without sacrificing visual fidelity. Grounded in prior LHRD research on perceptual fidelity, task-dependent effects of display curvature, and the role of embodied navigation in immersive environments, these considerations led us to three system-level design requirements:

D1 – Perceptual Fidelity at Human Scale. A defining characteristic of LHRDs is that they support physical navigation as a mode of interaction: users move closer to inspect detail and step back for overview, leveraging their own bodies as part of the analytical process. As Belkacem et al. [9] note, "users are free to move physically…to move closer to get more details or farther out to get an overview…Physical navigation has the advantage that users can explore data naturally as they would inspect objects in the real world."

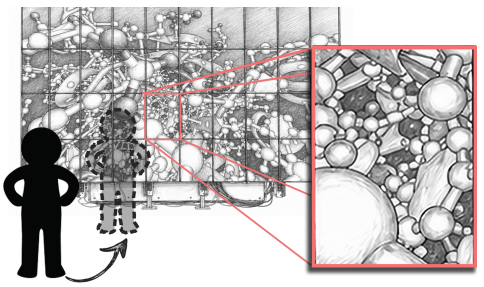


Fig. 2. D1 Illustrated: Moving between far and near viewing distances; inset shows fine detail.

This physical relationship between the user and the display is meaningful only if the display maintains perceptual fidelity across the full range of viewing distances it affords. As users approach the surface, individual pixels become discernible, and insufficient pixel density can compromise the perceptual continuity required for close-range inspection. Prior work has established that LHRDs should combine large physical extent with pixel density substantially exceeding that of conventional displays, so that users can transition fluidly between overview and detail without perceptual discontinuity [2, 26, 46]. Practitioner experience corroborates this. Sinaei et al. [51] have

reported that variations in viewing distance are a persistent challenge in real LHRD deployments. Specifically, practitioners struggle to design visual encodings that remain legible across the full range of distances users adopt, underscoring that perceptual fidelity at the human scale is not just a desirable property but a prerequisite for effective use. Recent work on the resolution limits of the human visual system suggests that human foveal achromatic resolution is approximately 94 pixels per degree (ppd), substantially higher than the long-accepted 60 ppd based on 20/20 acuity standards [4]. That is to say, users standing near a wall-sized display are well within the range where pixel density has meaningful perceptual consequences. This requirement was especially relevant for FlexiCAVE because our target interaction setting places users in close proximity to the display. Fig. 2 illustrates this concept.

D2 - Stereoscopic Spatial Cues. Spatial navigation in 3D environments requires users to interpret relative depth, volumetric extent, and structural arrangement while maintaining orientation. These are precisely the settings in which binocular disparity provides meaningful perceptual benefit. Prior work has shown that stereoscopic viewing improves understanding of 3D structure and can enhance navigation performance relative to monocular viewing [56].

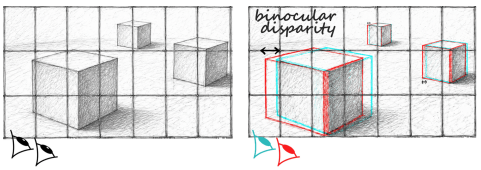


Fig. 3. D2 Illustrated: Without stereo (left), depth is ambiguous. With stereo (right), binocular disparity reveals the relative depth of the scene.

Evidence from large-display and immersive visualization settings further suggests that these benefits extend beyond small desktop conditions to more complex 3D scenes and wall-scale environments [12, 18, 44]. A broader survey of stereoscopic display studies reinforces this trend, reporting benefits most consistently for depth-related spatial tasks, particularly in near-field interaction settings [39]. The value of stereo is also well established in visualization more generally. Prior work argues that stereoscopic viewing helps make depth a more reliable perceptual channel for 3D visualization, addressing limitations that arise when spatial structure must be inferred from monocular displays alone [58]. Fig. 3 illustrates this concept. Studies that isolate stereoscopy as an experimental factor further show that removing stereo can substantially degrade 3D perception and interaction performance [25]. Additional results from large-scale immersive systems likewise indicate that stereo strengthens depth perception and spatial reasoning in complex visual scenes [16, 44].

D3 - Physically Reconfigurable Geometry. As established in Section 2, the literature on display geometry shows that no single configuration is optimal across tasks and users, curved layouts support navigation and immersion while flat layouts better serve overview and comparison [34, 36, 49]. Practitioners building and deploying LHRD systems consistently report having to adapt their workflows to fixed physical environments, with display relocation and reconfiguration identified as persistent infrastructure challenges [51]. FlexiCAVE addresses this and treats physical reconfiguration itself as a design requirement. The display layout can be reconfigured in real time to support varying horizontal curvatures and interaction needs (see Fig. 1). This is a core contribution of the system, and we investigate this as an interaction paradigm in Sections 4 and 5.

FlexiCAVE addresses these design requirements by comprising ten vertical columns of high-resolution displays (**D1**), arranged such that each adjacent column pair is mounted on a pivot, allowing inward rotation and continuous adjustment of horizontal curvature (**D3**), as shown in (Fig. 4a). This structure enables the display geometry to be reshaped during use rather than fixed at design time. Each column integrates four portrait-oriented 1920×1080 displays, providing sufficient resolution for close-range inspection while maintaining a continuous, human-scale surface. To

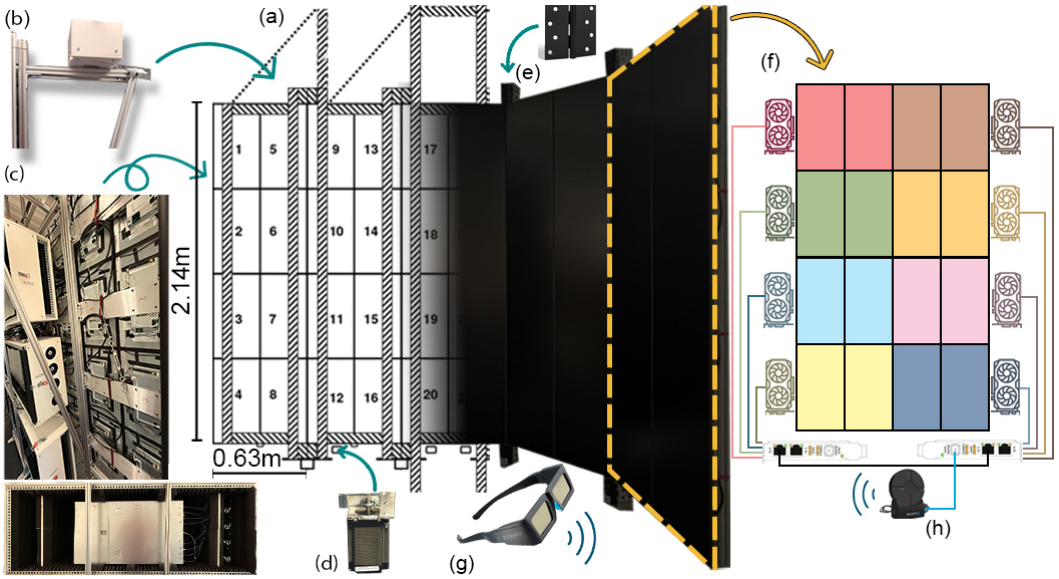


Fig. 4. (a) Partial 2D diagram of the FlexiCAVE showing frames, components, and a sample configuration. (b) Counterweight for stabilizing suspended column-pairs. (c) Rear view with mounted displays and custom noise-canceling server cabinets (bottom inset shows interior of cabinet). (d) Hinge-mounted angle encoders and motors. (e) Hinges for rotating column-pairs. (f) GPU configuration diagram per node. (g) Active stereo glasses. (h) RF emitter connected to the GPU timing server.

support spatial perception during immersive interaction, the displays are configured for active stereoscopic rendering (**D2**).

3.1 Hardware Setup

When building an LHRD facility, each component, from the displays to the visualization cluster, influences user experience and its overall performance. This section outlines several hardware and engineering decisions we made during the construction of the FlexiCAVE. At full load, the FlexiCAVE draws about 9.84 kW in power. Detailed hardware specifications and cost information are provided in the Supplementary Materials.

Display Selection. Arguably, the most essential component of any visualization environment is its display. In addition to the design goals, we considered the following criteria informed by prior LHRD work [22, 26, 46]: image quality, resolution, bezel size, and stereo support. Projector-based, Micro-LED, and OLED panels were ruled out due to recalibration overhead, cost, and stereo compatibility constraints. Among LCD options, the key challenge was balancing bezel size with active stereo support. We opted for FHD IPS panels with a 240 Hz refresh rate and 1 ms response time, which allowed us to operate an active stereo system without compromising pixel resolution. Custom mounts further reduced the bottom bezel from 6mm to 4mm. More details are provided in the Supplementary Materials.

Display Layout. Guided by design goal (**D3**) and the constraints of available physical space, FlexiCAVE is constructed from ten vertical display columns mounted on a custom lightweight aluminum frame (Fig. 4a). The two central columns are fixed to the ceiling and floor, while the remaining columns are organized into four pairs that are suspended from the central frame. Each column pair is connected by a hinge that allows inward rotation from a flat configuration (0°) to a

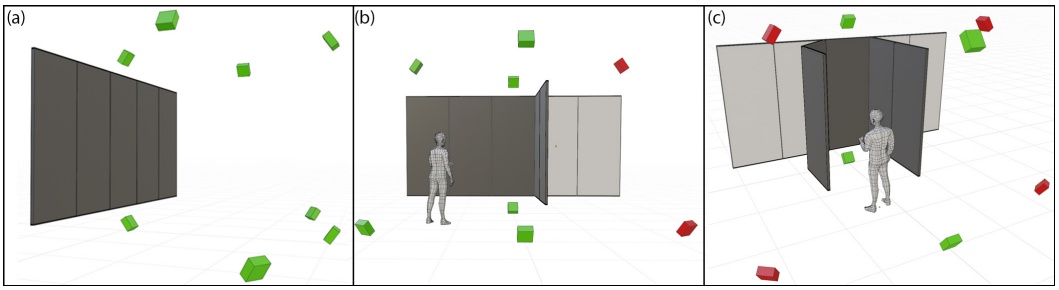


Fig. 5. 3D rendering of optimal IR camera placement in the FlexiCAVE. Green cuboids indicate IR cameras with full or partial marker visibility, and red cuboids indicate cameras without marker visibility. Black quads represent the FlexiCAVE display layout, and gray quads show the reference flat layout. (a) All markers are visible in the flat layout. (b) An example illustrating one side fully folded, obstructing the side cameras. (c) An example illustrating both side cameras are obstructed.

folded configuration (90°), enabling continuous adjustment of horizontal curvature. This design allows the display surface to transition between a flat layout spanning 3.13 m in width and more enclosed, curved configurations with an approximate radius of 1 m. Fig. 1 illustrates several example configurations. To support dynamic reconfiguration during interaction, an angle encoder is mounted on each hinge (Fig. 4d), continuously reporting column rotation to the visualization system so that rendered views can be updated in real time. Counterweights are attached to the rotatable column pairs to ensure stability and balanced motion during both manual and motorized reconfiguration (Fig. 4b).

Motorized Reconfiguration. In addition to manual manipulation, FlexiCAVE supports motorized reconfiguration of the display geometry. Each rotatable column pair is actuated by a stepper motor with an integrated encoder, coupled to the hinge mechanism through a reduction gearbox. This setup provides sufficient torque and positional accuracy to smoothly rotate the display columns while maintaining precise control over curvature. Motorized actuation complements manual interaction by enabling repeatable, fine-grained, and synchronized layout transitions. This allows users to return to predefined configurations, specify exact angles, or perform gradual geometry changes during interaction. To support real-time use, we implemented a lightweight web-based control interface, served to a handheld tablet, that accepts both directional controls and numerical angle input. The web app requires no software installation on the control device, allowing physical reconfiguration from any browser-enabled tablet or laptop. Changes are immediately reflected in both the physical display and the rendered visualization. Motorized actuation complements rather than replaces manual reconfiguration, together supporting a spectrum of interaction styles from embodied manipulation to programmatic control.

Tracking. Embodied interaction in LHRDs relies on accurate tracking of user position and input devices, particularly to support off-axis stereoscopic rendering and spatial navigation [5, 40]. For the FlexiCAVE, we adopted a hybrid tracking setup combining HTC SteamVR base stations [54] for wand-based interactions with OptiTrack Flex 13 infrared (IR) cameras [42] for head tracking. Although both systems employ IR, interference is avoided by synchronizing them through an OptiHub device [43], ensuring seamless operation.

A key challenge in deploying the IR cameras was determining an optimal layout. Since recalibrating extrinsic parameters is both labor-intensive and computationally costly, cameras could not be mounted on the rotatable FlexiCAVE frame. Instead, we installed eight IR cameras along the ceiling and floor of the room, ensuring that the headset's IR markers remain visible to at least three cameras

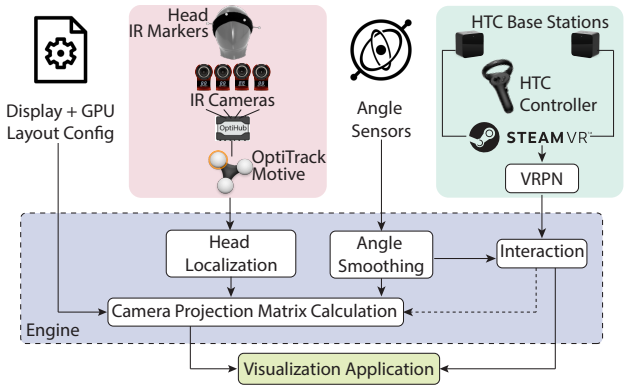


Fig. 6. Our Engine system design for the multi-node and multi-GPU FlexiCAVE setup. Using display and system configurations, along with user and controller tracking, Engine updates camera views and visualizations, abstracting them from the application design.

in any configuration. Following recommended best practices, four HTC base stations were mounted overhead on the ceiling to provide consistent coverage for controller tracking (Fig. 5). Together, this setup provides stable head pose estimation and controller input throughout interaction, supporting layout-aware rendering and embodied exploration. Supplementary Materials describe the infrared camera layout, calibration process, and validation steps used to ensure robust tracking across all display configurations.

3.2 Visualization System

FlexiCAVE is powered by a four-node visualization cluster. We use the term *panel* to refer to a single physical display unit. Two nodes drive the suspended display columns (end-nodes), each supporting 16 panels (4 columns). One node drives the central two fixed columns (center-node), supporting 8 panels. The remaining node serves as the head-node, coordinating rendering and system management.

One node in the synchronization network is designated as the timing server for the active stereo system and is connected to an RF emitter (Fig. 4h) to synchronize left- and right-eye framebuffer swaps across all GPUs and stereo glasses [55]. As RF-based synchronization does not require line-of-sight, the setup supports multiple simultaneous users, allowing any number of glasses to remain in sync.

Prior work has reported reliability issues with fiber-optic connections for high-framerate stereoscopic rendering, where GPU drivers may fail to interpret coherent timing signals [12]. To avoid this limitation, we use high-quality copper display cables for all GPU-to-panel connections. The visualization nodes are placed adjacent to the FlexiCAVE and housed in noise-attenuating enclosures to minimize cable length, signal degradation, and acoustic interference (Fig. 4c). Displays are mounted in portrait orientation, as the native polarization of the panels in landscape mode is orthogonal to the polarization filters used by the RF stereo glasses.

3.3 Software Architecture

FlexiCAVE is supported by a custom rendering and synchronization framework, referred to as the *Engine*, designed to accommodate dynamic physical reconfiguration of the display geometry in real time. The Engine follows a replicated execution model [24], where identical application instances run across multiple nodes while a designated master instance coordinates global state. As shown in Fig. 6, the Engine integrates system configuration parameters, real-time display

geometry updates, and 6 DoF tracking input to maintain consistent rendering across all displays. The master instance, running on the head node, synchronizes camera parameters, interaction state, and inter-node communication, ensuring coherent multi-display rendering during both static use and active reconfiguration. System configuration files specify the physical dimensions and placement of display columns, per-node viewport assignments, GPU bindings, and initial virtual camera parameters. This abstraction allows applications to remain independent of the underlying display topology while enabling the Engine to adapt rendering behavior as the physical layout changes. User head pose and controller input are incorporated into the rendering loop via real-time tracking streams. At each frame, the Engine updates view matrices based on the current display configuration, head pose, and interaction state, supporting off-axis stereoscopic rendering and spatially consistent interaction across configurations. Head position can be treated as fixed or dynamically tracked, depending on application needs.

To support smooth physical reconfiguration, encoder-reported hinge angles are filtered to suppress sensor noise and unintended micro-movements. Angle updates are validated over a short temporal window of five frames (83 ms at 60 fps), and only sustained angle changes exceeding 1.0° are propagated to the rendering pipeline. Once accepted, layout changes are interpolated at a fixed rate to balance responsiveness with visual stability. Across all tested configurations, the system maintained a stable 60 fps rendering rate. Encoder-to-GPU latency averaged approximately 15 ms, with an end-to-end reconfiguration latency of roughly 100 ms when filtering was enabled.

By abstracting display synchronization, tracking integration, and layout-aware rendering, the Engine allows application developers to focus on visualization and interaction design rather than system-level concerns. At the same time, encoder states are exposed as programmable triggers, enabling applications to directly map physical display motion to data transformations, as demonstrated in Section 4 and 5 .

4 Dynamic Reconfiguration for 3D Navigation

Virtual navigation in LHRDs typically relies on controller-based camera rotation or automatic flythrough to reveal hidden regions of 3D scenes. However, these approaches decouple navigation from the physical display structure, treating the screen as a static window into a virtual world. This can limit spatial awareness, as users must mentally track their orientation relative to the scene without stable physical reference points [3, 8]. FlexiCAVE introduces a different paradigm: users physically reconfigure display geometry during navigation, creating a direct coupling between the display physical structure and their viewpoint. This approach treats the display itself as a spatial reference frame that users can actively manipulate, potentially reinforcing embodied spatial reasoning [29].

We investigate this hypothesis through two complementary studies. First, we conduct a controlled experiment measuring angle estimation accuracy across static and dynamically reconfigurable display layouts. Second, a qualitative evaluation with domain experts performing authentic navigation tasks in an urban flood simulation assesses whether these benefits transfer to real-world analytic contexts. Together, these studies address the question posed in Section 1: Can display geometry become an embodied interaction resource that users manipulate during analysis, particularly for tasks that require maintaining orientation while navigating 3D content?

All studies in this paper were approved by the Institutional Review Board of the Office of Research Compliance at Stony Brook University (1173920_MODCR005). Informed consent was obtained from all participants prior to their involvement in the studies. For all qualitative evaluations, involving expert sessions, audio or video recordings were transcribed and analyzed following inductive thematic analysis approach of Braun and Clarke [15] to identify common themes and patterns within the data. Initial codes were generated by the first author, bottom-up from participants' own

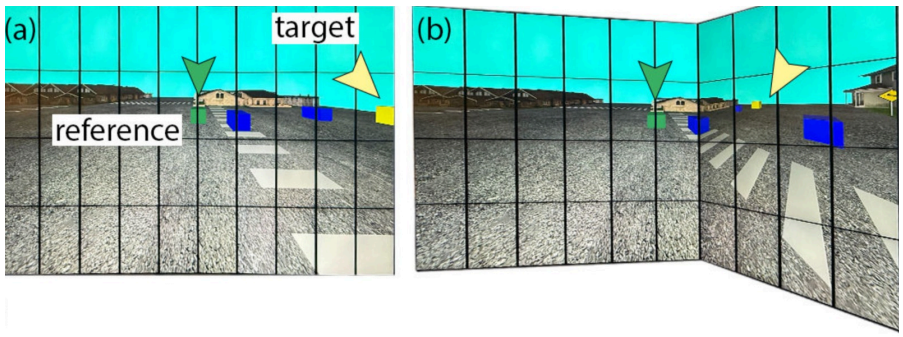


Fig. 7. Examples from the angle perception study for (a) the flat layout and (b) a user-set dynamic layout.

language. The remaining authors reviewed the initial codebook and corresponding data excerpts, and disagreements were resolved through discussion. We describe the specific themes yielded by each evaluation in the respective sections.

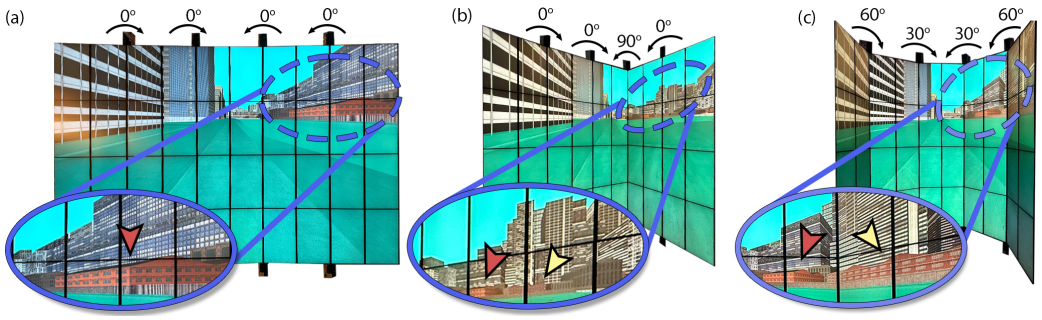


Fig. 8. Visualizing an urban flood simulation in the FlexiCAVE. Scenes shown in (a) flat, (b) *L*-shaped, and (c) curved layouts. Top arrows indicate column-pair rotation angles; bottom insets highlight view changes (e.g., building position/orientation) caused by layout reconfiguration. Views update in real time with FlexiCAVE adjustments.

4.1 Controlled Angle Estimation Study

We conducted a within-subjects experiment comparing spatial angle estimation across three display configurations: static flat, static curved, and dynamically reconfigurable (dynamic). Our primary hypothesis was that allowing users to physically reconfigure the display during navigation would reduce angle estimation error compared to static configurations.

Participants and Apparatus. We recruited 18 participants (12 male, 6 female; age range 23–37 years) through university mailing lists. Participants received a $15 gift card as compensation. All participants reported normal or corrected-to-normal vision and no physical disabilities that would impair their ability to manipulate display columns. None had prior experience with the FlexiCAVE.

The study used FlexiCAVE in three configurations: (1) *Flat*, where all columns remained at an angle of 0°; (2) *Curved*, where the angle between the column pairs was 45°; and (3) *Dynamic*, where participants could physically rotate the columns. In the static conditions (Flat and Curved), navigation was achieved by yaw rotation using a handheld controller. In the Dynamic condition, navigation was performed exclusively through physical reconfiguration of the columns, without controller input. To isolate embodied interaction, motorized reconfiguration was disabled, participants adjusted the display manually in the Dynamic condition.

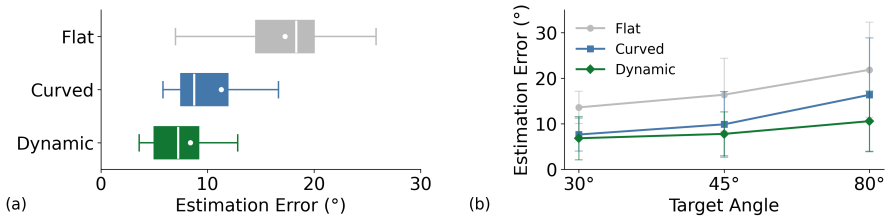


Fig. 9. Angle estimation error across display configurations. (a) Overall distribution of absolute error for Flat, Curved, and Dynamic layouts; the white bar and dot represent median and mean, respectively. (b) Mean absolute error (±SD) broken down by target angle (30°, 45°, 80°)

Task and Procedure. Our task design consisted of a reference cube placed at a fixed distance along the camera initial look-at direction, and a target cube positioned at yaw offsets of 30°, 45°, or 80° at a distance of 6m from the user's initial position. Participants were asked to estimate the angle between the reference and target cubes (shown in Figs. 7a and b). Participants were free to view the scene and perform navigation using either controller-based camera rotation or physical reconfiguration of the display, depending on the active condition. For each user, we recorded the angle estimation error, defined as the absolute difference between the participant's reported estimate and the true yaw angle.

Prior to the study, participants completed a brief orientation session in which they were introduced to the FlexiCAVE facility and its reconfiguration mechanism. They observed a demonstration of column movement, explored a simple 3D scene to familiarize themselves with navigation, and were shown examples of different angles using visual references. Each participant then completed one practice trial without feedback to acclimate to the task.

To control for order effects, a balanced Latin square design was employed. Each participant completed 18 trials: 3 conditions (Flat, Curved, Dynamic) × 3 angles (30°, 45°, 80°) × 2 repetitions, with angle order randomized within each condition block. Participants were free to take breaks at any point during the session and were explicitly asked if they needed rest between condition blocks. The entire session lasted approximately 30 minutes.

Results. Fig. 9a shows the distribution of angle estimation error across the three configurations. A repeated-measures ANOVA revealed a significant main effect of layout on angle estimation error ($F(2, 34) = 17.90$, $p < 0.001$, $\eta_p^2 = 0.51$), indicating a large effect of display configuration on spatial judgment accuracy. Pairwise comparisons using paired t-tests indicated that Dynamic layouts significantly reduced error relative to Flat ($t(17) = 5.80$, $p < 0.001$, $d_z = 1.37$), and Curved also outperformed Flat ($t(17) = 3.53$, $p = 0.003$, $d_z = 0.83$). The difference between Curved and Dynamic also reached significance ($t(17) = 2.25$, $p = 0.038$, $d_z = 0.53$), though with a smaller effect size. Mean absolute errors were lowest for Dynamic ($M = 8.40°$, $SD = 6.71°$), followed by Curved ($M = 11.30°$, $SD = 10.23°$), and highest for Flat ($M = 17.29°$, $SD = 9.49°$). Median errors showed a similar pattern: Dynamic (5.0°), Curved (10.0°), and Flat (15.0°). These results indicate that physical reconfiguration reduced angle estimation error by 51% compared to static flat layouts, with curved layouts providing intermediate improvement (35% reduction). Notably, Dynamic outperformed Curved by 26%, suggesting that the ability to actively manipulate display geometry provides benefits beyond static curvature alone.

To examine whether the benefit of dynamic reconfiguration varied across target angles, we evaluate mean absolute angle estimation error for each combination of display configuration and target yaw offset in Table 1. A two-way repeated-measures ANOVA confirmed significant main effects of both display configuration ($F(2, 34) = 17.90$, $p < .001$) and target angle ($F(2, 34) = 15.04$,

$p < .001$), but no significant Condition × Angle interaction ($F(4, 68) = 1.07$, $p = .380$). The absence

Table 1. Mean absolute angle estimation error (°) by display configuration and target angle. Values are M (SD).

	Target Angle		
Configuration	**30°**	**45°**	**80°**
Flat	13.6 (3.5)	16.4 (8.0)	21.9 (10.5)
Curved	7.7 (3.6)	9.9 (7.2)	16.4 (12.5)
Dynamic	6.8 (4.7)	7.8 (4.8)	10.6 (6.6)

of a significant interaction indicates that the advantage of dynamic reconfiguration over static layouts was consistent in magnitude across all three angular offsets rather than being specific to wider angles. The per-angle trajectories are further illustrated in Fig. 9b which makes two descriptive patterns visible. First, error in the Flat condition increased steeply with angle (13.6°, 16.4°, 21.9° for 30°, 45°, and 80°, respectively), whereas Dynamic error grew more modestly (6.8°, 7.8°, 10.6°), suggesting that the physical alignment afforded by reconfiguration partially attenuates the perceptual cost of larger offsets. Second, Curved and Dynamic performance converged at the narrowest offset: at 30° the two configurations differed by less than 1° on average and the pairwise comparison was non-significant ($t(17) = 0.88$, $p = .393$, $d_z = 0.21$), whereas at 80° the gap widened to 5.8° and approached significance ($t(17) = 1.92$, $p = .072$, $d_z = 0.45$). This pattern is consistent with the idea that a fixed curved arrangement provides a sufficient spatial scaffold for narrow angles, but that the ability to dynamically align display columns with the target direction offers an additional benefit as the yaw offset increases.

4.2 Qualitative Evaluation with Domain Experts

To assess whether the findings transfer to authentic domain tasks, we conducted qualitative evaluation with domain experts navigating an urban flood simulation, Submerse [14] in FlexiCAVE (Fig. 8).

Participants and Apparatus. We invited 5 domain experts (3 male, 2 female, age range 27-33 years) with professional experience in emergency management and flood analysis. All participants had previously used Submerse in static flat or curved LHRD facilities during prior research collaborations, providing a basis for comparative assessment. Participants received no compensation, as the evaluation was conducted during a scheduled research workshop. Participants were asked to perform an exploratory flood analysis task using a scenario in an unfamiliar urban area. The task was intentionally open-ended to reflect authentic emergency management workflows: participants explored flooding patterns, identified critical infrastructure at risk, assessed evacuation routes, and discussed flood progression across the neighborhood. This task required extensive spatial navigation to inspect flooding from multiple vantage points and maintain orientation across a large geographic area.

Task and Procedure. Each participant received an individual 15-minute orientation to FlexiCAVE, including demonstration of manual column reconfiguration and practice navigating a simple 3D scene. All participants were then together presented with the flooding scenario and asked to conduct their analysis while thinking aloud. A handheld 6DoF controller enabled virtual camera translation and rotation. Participants were free to reconfigure display columns, through the web-interface or manually, at any time during the task. The session lasted approximately 90 minutes. We encouraged all participants to verbalize their navigation strategies, observations, and

comparisons to prior experiences with static LHRDs. Following the task, we conducted individual brief semi-structured interviews focusing on perceived differences from prior static LHRD experiences and the role of physical reconfiguration in their analysis process.

Findings. Our inductive analysis yielded three overarching themes. We report each theme below, supported by representative participant quotes and session observations.

Spatial Anchoring Through Display Geometry. Participants consistently reported that physical display structure provided spatial anchoring that was absent in their prior experiences with static LHRDs. One participant explicitly stated that "the position of the columns gave me a clearer sense of direction." This observation aligns with our controlled angle estimation findings, suggesting that visible display geometry reinforces users' internal spatial reference frame. Participants contrasted this with their experiences in static facilities, where virtual camera rotation frequently caused disorientation. Multiple participants described navigation in static displays as feeling passive, like "being a passenger within a vessel observing the flooding scenario." In contrast, FlexiCAVE navigation felt more intentional and active. While participants acknowledged a brief learning curve in understanding how column adjustments affected the viewpoint, once familiar with the interaction, they felt a markedly greater sense of control and navigation intentionality.

Embodied Navigation Intentionality. Participants deliberately used physical reconfiguration to align display surfaces with specific urban features, treating column rotation as an intentional navigation act rather than incidental adjustment. For example, bending a side column by 90° revealed perpendicular streets and cross-block flooding patterns (Fig. 8b), while incremental column rotation provided intuitive access to different view directions (Fig. 8c). Participants noted that the physicality of FlexiCAVE helped them maintain orientation and return to prior viewpoints. Unlike in flat or curved LHRDs, where rotating the virtual camera often caused loss of directional awareness (described as losing "*north*"), the visible arrangement of the displays provided a reliable anchor for spatial navigation. This embodied engagement also strengthened memory of their analysis: one participant stated "I can remember the flooding at that intersection because I bent the right column to look at it," suggesting that physical reconfiguration may create stronger spatial memory associations than purely virtual navigation.

Collaborative Reconfiguration. Collaborative behaviors emerge naturally during the sessions. This theme was identified through thematic analysis and further supported by visual observations of participant interactions. We observed two instances where participants rotated the columns to share viewpoints with colleagues, and one instance where a participant adjusted their viewing position to observe a teammate's perspective. When participants worked in pairs, a natural division of labor emerged: one participant manipulated the display columns while the other navigated with the controller and interpreted flood progression. Participants reported this collaboration felt intuitive and mirrored real-world emergency operations center workflows where responsibilities are distributed among team members.

4.3 Discussion

The two studies address the same underlying question from complementary directions. The controlled angle estimation study demonstrates a measurable perceptual benefit: allowing users to physically reconfigure display geometry reduced angle estimation error by 51% compared to static flat layouts, with a large overall effect ($\eta_p^2 = 0.51$),and a strong Dynamic vs Flat comparison($d_z = 1.37$). The domain expert evaluation extends these findings to authentic analytic contexts. Participants spontaneously reproduced the same behavior in a real task, deliberately aligning columns with features of interest rather than relying on controller-based rotation. The *Spatial Anchoring* theme directly echoes why Dynamic outperformed even static Curved: it is not the presence of a curved

surface but the ability to actively align that surface with a target direction that scaffolds spatial judgment. The *Embodied Navigation Intentionality* theme further explains the mechanism: treating column rotation as a deliberate act may reinforce spatial memory in ways that purely virtual navigation cannot. Notably, the emergence of *Collaborative Reconfiguration* reveals something the controlled study could not capture: manipulable display geometry naturally becomes a shared spatial resource in multi-user settings.

Together, these findings align with theories of embodied cognition [29],by externalizing spatial reference into the physical environment, FlexiCAVE reduces the cognitive load of maintaining orientation through mental imagery alone. We discuss these findings and their design implications in more detail in Section 6.

5 PIVoT: Mapping Physical to Virtual Interaction

Section 4 established that reconfigurable display geometry can scaffold spatial navigation by coupling physical structure with the user's viewpoint. Building on this, we explore whether the same principle extends to visualization updates beyond 3D navigation. In particular, we examine whether reconfigurable display geometry can help in other forms of spatial exploration, and whether this benefit extends to other visualization tasks when spatial correspondence is absent.

Jacob et al. [29] highlight *environmental awareness and skills* as central to effective interaction, referring to people's ability to sense, navigate, and manipulate their physical surroundings as part of human–computer interaction. Yet in most LHRD systems, the display layout functions as a fixed window into a virtual scene, limiting how spatial structure can be used as an interaction resource beyond navigation. To explore the boundaries of this principle, we introduce PIVoT (Physical Interaction to Virtual Transformation), an interaction paradigm that maps physical reconfiguration of the display directly to real-time visualization updates, treating the display itself as an input device. This approach positions physical display geometry as a stand-alone interaction technique.

We illustrate the PIVoT design space through two exploratory applications. The first, volume visualization, shows a case where a physical reconfiguration maps naturally to a spatial data operation. The second, multivariate information visualization, serves as a boundary condition exploration: it examines whether the approach generalizes when display geometry no longer maps meaningfully onto the underlying data operation, and in doing so, sharpens our understanding of when and why physical reconfiguration works.

5.1 Volume Rendering Cutting Plane

Volume rendering is a common technique in immersive visualization and LHRD environments for exploring complex scientific data [13, 28, 45, 46]. A central interaction in volume exploration is the use of cutting planes to reveal internal structures and spatial relationships within the data. While existing systems support plane manipulation through controllers or graphical widgets, these interactions typically treat slicing as an abstract operation decoupled from the physical structure of the display. This separation limits opportunities to leverage embodied spatial cues in large display environments.

With PIVoT, the FlexiCAVE rotating column pairs serve as tangible proxies for cutting planes. Each hinge angle directly defines a plane's orientation, establishing a one-to-one mapping between physical motion and virtual slicing. The volume is centered on the middle columns, and planes aligned with column offsets are re-sliced at every sensor update. Orthographic projection is used for plane-aligned views, producing clean radial slices without distortion (Figs. 10b,c). This design turns the FlexiCAVE display itself into an embodied input device for radial slicing. Further details for slice computation can be found in the Supplementary Material.

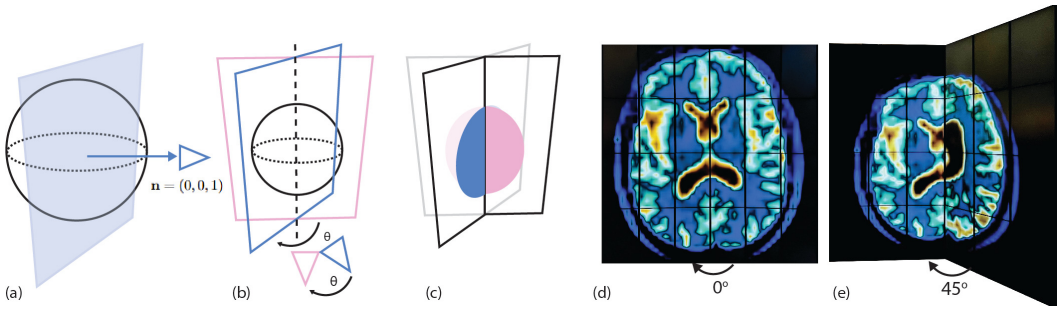


Fig. 10. PIVoT as a tangible cutting plane for volume rendering. (a) A centered volume with each display acting as a cutting plane, viewed with an orthographic camera. (b) Rotating a column-pair rotates the corresponding plane. (c) Combined slices align for seamless visualization. (d) Brain MRI volume rendered in flat layout. (e) Result when one side is rotated by 45°.

Participants and Apparatus. To gather formative feedback on this interaction, we invited five domain experts (3 male, 2 female, age range 24-35) including neuroscientists and researchers experienced in medical and volume visualization, to compare the PIVoT-based interaction with the widely used 3D Slicer application [1]. 3D Slicer represents the primary tool used by these experts in their everyday workflows and served as a familiar baseline for comparison. Our goal was not to compete with or replace such desktop tools, but to explore how physical display reconfiguration could be mapped to slicing operations in a large, immersive display environment.

Task and Procedure. Participants were evaluated individually. Each session lasted approximately 30-45 minutes and began with a 10-minute orientation to the FlexiCAVE and the PIVoT cutting plane interaction, including a demonstration of column rotation and its effect on slice orientation. Participants then explored the MRI brain volume at their own pace, using physical column reconfiguration to inspect internal structures, before transitioning to the same task in 3D Slicer on a desktop workstation. Participants were asked to explore a 3D MRI brain volume containing a spherical tumor, locate the tumor, and point out its spatial extent while inspecting surrounding anatomical structures in both environments. Following both conditions, participants provided structured comparisons and open-ended qualitative feedback through a brief semi-structured interview. This feedback was intended to assess the interpretability, intuitiveness, and perceived utility of mapping physical display reconfiguration to volume slicing, rather than to serve as a summative performance evaluation. Sessions were audio recorded and researcher notes were taken throughout.

Findings. Our inductive analysis yielded three themes, which we report below.

Spatial Comprehension Through Embodied Slicing. Participants reported that the curved layout of FlexiCAVE, combined with the radial slicing interaction, supported depth perception and spatial awareness of the volume. Two experts highlighted the value of physically manipulating the display, noting that the tangible coupling between column rotation and slice orientation made radial slicing: an operation uncommon in their current workflows, feel intuitive and easy to reason about. Participants further remarked that this embodied interaction improved their understanding of internal volumetric structures by aligning physical motion with spatial exploration.

Complementary Strengths of Physical and Desktop Interaction. In comparison, participants valued 3D Slicer as simpler and more immediately usable, valuing its precise boundary rendering, familiar interaction patterns, and conventional orthogonal slicing views. However, several experts noted that FlexiCAVE enabled a complementary mode of exploration: by centering the volume and bending both sides of the display into a *U*-shaped configuration, they could examine internal structures

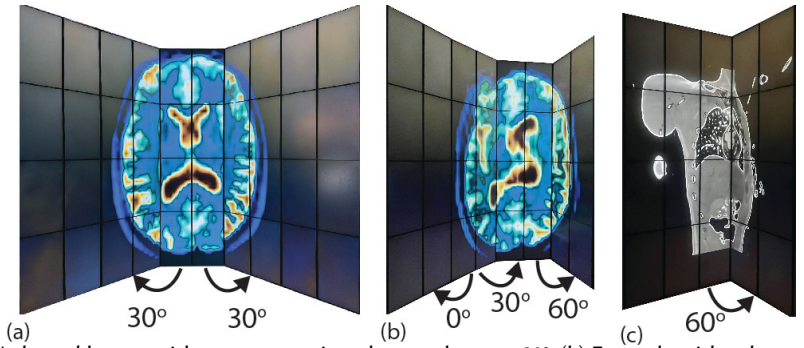


Fig. 11. (a) *U*-shaped layout with a center cutting plane and two at 30°. (b) Example with column-pairs rotated 30° and 60°. (c) Full-body CT rendering showing a limitation when rotations >60° exceed volume bounds.

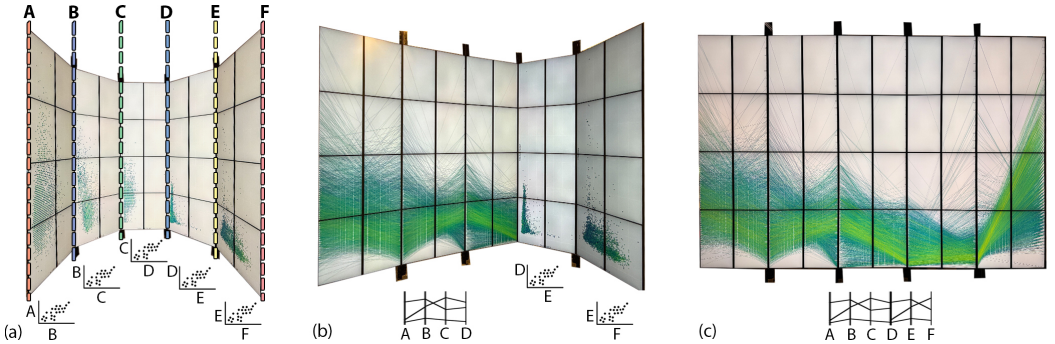


Fig. 12. PIVoT for dynamically updating InfoVis plots. (a) Each column-pair vertical axis is assigned an attribute (A to F). When column-pairs are > 15°, a 2D scatterplot of the neighboring attributes is shown. (b) By *flattening* neighboring column-pairs, a PCP of the respective attributes is generated. (c) A fully flat FlexiCAVE showing a PCP for all attributes. This figure shows plots for the Wine Quality dataset [20] with attributes fixed acidity, volatile acidity, citric acid, residual sugar, chlorides, and sulfur dioxide, from A to F, respectively.

across multiple slice orientations while standing within the display space (Fig. 11a). This spatial arrangement reduced the need for frequent viewpoint and plane adjustments, allowing participants to reason about the volume's internal organization through bodily position and physical display configuration rather than repeated interface manipulation.

Interaction Limitations of Radial Slicing. Participants also identified interaction limitations. Utilizing both pivots on one side of the display (Fig. 11b) produced novel viewing angles that were initially difficult to interpret, reflecting the unfamiliarity of such radial slicing perspectives in current analysis workflows. In addition, effective use of the full range of display rotation assumes approximately isotropic volume dimensions. For elongated datasets, such as full-body CT scans, radial slicing is practically limited to smaller angular ranges (e.g., ~60°), as rotations beyond this range extend beyond the volume bounds (Fig. 11c).

5.2 Exploring Scope: Information Visualization

To characterize the conditions under which PIVoT is effective, we explore a case where spatial correspondence between display geometry and data structure is absent. We study this to sharpen our insights from previous studies and understand : when does physical reconfiguration provide meaningful interaction benefit, and when does it not? We explored a design where column rotation toggles between scatterplots and parallel coordinates plots (Fig. 12), inspired by ImAxes [19]. In this design, each vertical axis between column pairs is assigned a dataset attribute. When column pairs

are separated by more than 15°, neighboring axes define a 2D scatterplot. Straightening columns below 15° merges attributes into a parallel coordinates plot (PCP), and fully flattening the display produces a global PCP across all six attributes. Attribute assignments can be reassigned through a simple web interface provided on a handheld tablet.

Participants, Task, and Procedure. We invited three InfoVis experts (1 male, 2 female, age range 25-30), specializing in dimensionality reduction, to evaluate this interaction and explore the data. Participants were evaluated individually in sessions lasting approximately 15-20 minutes, following a brief introduction to the FlexiCAVE and the PIVoT interaction. Each participant was asked to freely explore the Wine Quality dataset [20], using physical column reconfiguration to navigate between scatterplot and parallel coordinates views, with the goal of identifying correlations or trends across attributes. Sessions were audio recorded and researcher notes were taken.

Findings. Our inductive analysis revealed a fundamental distinction from volume visualization. In volume visualization, column rotation establishes a direct spatial mapping:the angle of rotation corresponds to the cutting plane angle, and display geometry serves as a physical proxy for the slice orientation. In contrast, the InfoVis design used column rotation as an arbitrary threshold trigger: curved layouts (>15°) show scatterplots while flat layouts show parallel coordinates. This mapping lacks spatial correspondence between physical geometry and data structure. Experts noted that the design rationale was unclear: the physical reconfiguration provided no perceptual benefit for understanding data relationships.

This finding sharpens the spatial correspondence principle suggested by the navigation and volume rendering scenarios above, it suggests that physical reconfiguration is most effective when display geometry provides a meaningful spatial proxy for the operation, not merely an alternative trigger mechanism. Future work should identify visualization transformations where physical structure can represent data structure in spatially meaningful ways. We examine this principle and its implications for interaction design in detail in Section 6.

6 Design Implications, Discussion and Future Work

Building and deploying FlexiCAVE has provided valuable insights into the design, implementation, and use of dynamically reconfigurable LHRDs. Through our controlled studies in spatial navigation, domain expert evaluations, and exploratory design studies, we have identified key opportunities and challenges that inform the future development of adaptive display environments. We organize our findings into four areas: when physical reconfiguration helps, implications for interaction design, technical infrastructure, and the broader evolution of reconfigurable display spaces.

6.1 When Physical Reconfiguration Helps: Spatial Correspondence

Our findings suggest that the effectiveness of physical reconfiguration as an interaction technique is not uniform: it depends critically on the relationship between display geometry and the underlying operation being performed. Across our three evaluation contexts, a consistent pattern emerged: reconfiguration was most beneficial when the physical arrangement of the display provided a meaningful spatial proxy for the task or data operation at hand.

In the angle estimation study, physically rotating display columns to face a target direction established a direct correspondence between display geometry and spatial orientation. Participants could use the visible column arrangement as a persistent external reference frame, reducing the cognitive demand of maintaining orientation through purely mental imagery. In the volume rendering scenario, column rotation mapped one-to-one onto cutting plane orientation, making an otherwise abstract slicing operation tangible and spatially interpretable. In both cases, the physical structure of the display did not merely serve as an input device, it externalized a spatial relationship that was otherwise internal to the user's mental model. In contrast, the information visualization

scenario, where column rotation served as a threshold trigger to toggle between scatterplot and parallel coordinates views, lacked this correspondence. The display geometry carried no inherent meaning relative to the data structure, and experts found the mapping unclear and unconvincing. This boundary condition sharpens our understanding of when physical reconfiguration adds value: not as a general-purpose input modality, but specifically when display geometry can represent a spatial aspect of the task or data.

These findings align with theories of embodied cognition [29], which posit that physical action can scaffold abstract reasoning. By making display geometry manipulable, FlexiCAVE externalizes spatial reference into the physical environment, potentially reducing the cognitive load of maintaining orientation purely through mental imagery. Critically, this benefit appears contingent on the degree to which the physical manipulation carries spatial meaning relative to the task, a constraint that has direct implications for how reconfigurable displays should be designed and evaluated going forward.

This pattern suggests that spatial correspondence between display geometry and data structure is a promising design principle for reconfigurable displays. It is important to note, however, that our findings are based on a limited set of interaction metaphors. We explored one case where spatial correspondence was present and one where it was absent, but we did not systematically explore the full range of possible mappings. It remains an open question whether alternative metaphors that lack direct spatial correspondence could be made effective through careful design, training, or habituation. Spatial correspondence should therefore be understood as a useful design lens, not a necessary condition, and future work should explore whether other forms of meaningful mapping — such as temporal navigation or abstract data transformation — can similarly support embodied interaction on reconfigurable displays.

6.2 Designing Interactions for Reconfigurable Displays

Our findings across the angle perception study, domain expert evaluation and PIVoT explorations suggest three design principles for reconfigurable display environments that extend beyond FlexiCAVE to inform the broader design of adaptive LHRDs.

Managing Physical Effort. Across all studies, manual column manipulation emerged as a powerful but effortful form of interaction. Participants valued the tangible, embodied nature of physical reconfiguration, particularly for coarse layout changes and spatial exploration. At the same time, sustained manual adjustment was less suitable for prolonged or fine-grained control, with some participants noting brief shifts of attention away from the analytical task. In multi-user scenarios, this effort was often naturally distributed, with collaborators informally taking on complementary roles, for example, one user adjusting columns while another navigated the virtual scene.

These observations motivate the use of *hybrid interaction models* that combine embodied engagement with ergonomic sustainability. FlexiCAVE's motorized actuation supports this balance by enabling precise, repeatable adjustments and quick return to predefined configurations while preserving opportunities for direct physical interaction. Future work should explore how users transition between manual and motorized control, how preferences vary by task and duration, and whether adaptive assistance can reduce physical effort while maintaining the cognitive and perceptual benefits of embodied reconfiguration. This suggests that reconfigurable display systems should provide a spectrum of reconfiguration modalities: from direct manual manipulation for embodied exploration to motorized or programmatic control for sustained or precise adjustments, rather than committing to a single interaction style.

Supporting Collaborative Reconfiguration. Emergent collaborative behaviors in our domain expert evaluations suggest that reconfigurable displays can function as shared spatial interfaces.

Participants naturally divided labor, used physical reconfiguration to establish shared viewpoints, and reported that column positions helped coordinate spatial reference between team members. These observations motivate more systematic investigation of collaborative reconfiguration. Future work should examine how coordination strategies vary with group size, task structure, and display layout, and whether physical reconfiguration can support more explicit collaborative workflows such as defining shared workspace regions or enabling role-specific views. Designing interface mechanisms that help teams coordinate layout changes and communicate spatial intent will be important for realizing the full collaborative potential of reconfigurable display environments. This points to a broader design opportunity: reconfigurable displays can function as shared spatial interfaces in multi-user settings, and future systems should explicitly support role distribution and coordination mechanisms that leverage physical reconfiguration as a collaborative resource.

Physical Geometry Should Be Visually Salient. Participants in our domain expert evaluation relied on seeing column positions peripherally to maintain orientation, suggesting that asymmetry in display structure provides persistent spatial cues that support navigation even when users are focused on central content. For reconfigurable display designers, this implies that layout changes should be perceptually legible — users should be able to read the current configuration at a glance without diverting attention from their primary task. This may favor designs where reconfiguration produces visible, asymmetric changes to display structure over those where layout shifts are subtle or uniform.

6.3 Technical Infrastructure and System Design

Dynamic Environments Tracking. A key layout decision was mounting tracking cameras on the room rather than the reconfigurable display frame. While this avoided the need for recalibration after each layout change, it introduced coverage challenges, ensuring that head and controller markers remained visible to sufficient cameras across all configurations required careful camera placement and validation. For future systems, we recommend early investment in tracking infrastructure design. Simulation-based coverage analysis can identify optimal camera placements before installation. For displays with more complex reconfiguration capabilities (e.g., vertical curvature, multiple degrees of freedom per column), on-frame tracking with automatic recalibration may be necessary despite its complexity.

Scalability and Form Factor Extensions. FlexiCAVE supports horizontal curvature adjustment through column rotation, but vertical curvature and outward bending remain constrained by the current mechanical design. Extending to two-dimensional curvature adjustment would expand the configuration space significantly, enabling a wider range of spatial layouts beyond planar arrangements. Such extensions introduce substantial engineering challenges. Supporting vertical curvature requires more complex hinge mechanisms and structural reinforcement to maintain stability under load. Outward bending demands counterweight reconfiguration or alternative stabilization strategies. The rendering pipeline must accommodate more degrees of freedom in camera projection calculations, and tracking coverage becomes more difficult to ensure. However, advances in display technology may provide alternative pathways. As large-format LED panels and flexible display technologies become more affordable and higher in resolution, mechanically actuated tiled displays could give way to truly deformable displays. Such systems would eliminate bezel gaps, support continuous curvature profiles, and potentially enable simpler actuation mechanisms. Investigating the interaction design implications of continuously deformable displays represents an exciting direction for future work.

6.4 Reconfigurable Displays in the Broader Ecosystem

Complementing Immersive Technologies. Head-mounted displays (HMDs) have advanced substantially in recent years, offering portability, wide fields of view, and increasingly high visual fidelity. Prior work has demonstrated their effectiveness for immersive exploration and individual analysis, particularly when mobility or personal viewpoint control is important. At the same time, HMD-based workflows continue to face challenges when working with large, dense datasets that benefit from pixel-level precision, as well as in co-located collaboration where shared spatial reference and mutual awareness are critical. Rather than positioning reconfigurable LHRDs as alternatives to HMDs, we view them as complementary components within a broader ecosystem of immersive technologies. FlexiCAVE supports high-fidelity, wall-scale visualization and co-located collaboration through a shared physical reference, while enabling fluid transitions between individual and group work without requiring users to don or doff headsets. Building on existing hybrid and cross-device approaches, future work should further explore coordinated use of HMDs and reconfigurable displays.

Toward Context-Aware Adaptation. Our current system treats reconfiguration as user-initiated: participants manually adjust columns or invoke motorized control through the web interface. However, the availability of real-time sensing opens opportunities for context-aware adaptation, where the display automatically adjusts its geometry based on task requirements, content characteristics, user position, or group dynamics. Yet automatic reconfiguration raises a fundamental tension. A central argument for physical reconfiguration is that manually moving the display engages embodied cognition [29]: the physical act of reshaping the environment may itself scaffold spatial awareness, memory, and orientation. Our findings support this, with participants reporting stronger directional awareness precisely because they had physically manipulated the columns. Automatic adaptation could deprive users of these proprioceptive cues, reducing reconfiguration to a passive environmental change rather than an intentional interaction act.

Whether this tradeoff is consequential likely depends on context. For repetitive or ergonomically demanding transitions, motorized assistance may reduce physical burden without sacrificing spatial understanding. For exploratory or spatially demanding tasks, however, automatic reconfiguration may be counterproductive if it bypasses the embodied process through which users build their spatial model of the environment. We therefore frame this as an open question: *under what conditions does automatic reconfiguration support versus undermine the embodied benefits of physical display manipulation?* Realizing such capabilities requires not only sensing infrastructure but also design knowledge—specifically, understanding when and how the display should adapt, which configurations best support particular tasks and contexts, and how transitions can be coordinated with ongoing interaction without disrupting users' cognitive flow. Investigating these challenges through participatory design, longitudinal deployments, and machine learning–based adaptation strategies will be essential for developing reconfigurable displays that respond intelligently to changing analytic and collaborative needs.

Establishing Design Guidelines. Our studies provide initial evidence that dynamic physical reconfiguration can enhance spatial navigation and support specific visualization tasks, but they also highlight the breadth of the design space for reconfigurable displays. In particular, our findings point to key dimensions that warrant further investigation, including which visualization types benefit most from dynamic geometry, how task characteristics (e.g., search vs. comparison, overview vs. detail, individual vs. collaborative) interact with display configuration, and which interaction techniques most effectively leverage physical reconfiguration beyond the PIVoT examples presented in this work. Addressing these dimensions requires systematic investigation across diverse tasks, datasets, and user populations. We envision a research agenda that combines controlled experiments,

longitudinal deployments capturing authentic use, and design exploration of new interaction techniques. Together, such efforts can establish evidence-based guidelines that inform the design of future reconfigurable display environments and help practitioners determine when and how these systems are most appropriate for their application domains.

6.5 Limitations

While the domain evaluation with Submerse addresses ecological validity, the small expert sample (N=5) and qualitative analysis limit generalizability. We did not measure task completion time, cognitive load, or performance outcomes in the Submerse evaluation, focusing instead on subjective experience and navigation strategies. The evaluation was conducted as part of a research workshop where participants worked together collaboratively, which allowed us to observe natural collaborative behaviors but precluded systematic comparison between individual and group performance conditions. Systematic evaluation of collaboration dynamics across different group sizes and controlled comparison of individual versus collaborative use remains future work.

The controlled angle estimation study compared three specific display configurations: flat, a single static curved layout, and dynamic - leaving open how performance might vary across a broader range of fixed curvatures. Future work should investigate whether the benefits of dynamic reconfiguration hold across different baseline curvatures. Finally, our PIVoT explorations examined a limited set of interaction metaphors, and the generalizability of our spatial correspondence findings to other visualization types and data domains remains to be established.

7 Conclusion

We presented FlexiCAVE, a dynamically reconfigurable large high-resolution display (LHRD) facility whose layout can be adjusted in real time by horizontally rotating hinged display columns. FlexiCAVE combines wall-scale visual fidelity (approximately 83 million pixels with active stereo) with a sensing-and-rendering pipeline that synchronizes camera perspectives with physical layout changes, enabling display geometry to be manipulated during active analysis rather than fixed at construction time. Through a controlled perceptual study and domain expert evaluation, we demonstrated that treating physical reconfiguration as an interaction technique yields measurable and practical benefits. In particular, dynamic layouts substantially improved spatial judgment and supported a stronger sense of direction during 3D navigation compared to static flat or curved configurations. These results suggest that reconfigurable geometry can externalize orientation cues that are difficult to maintain through purely virtual navigation alone. We further introduced PIVoT, a tangible interaction approach that maps physical display reconfiguration to spatially meaningful visualization updates. Through applied scenarios in 3D navigation, scientific visualization, and information visualization, we illustrated how dynamically coupling display geometry with visual transformations can support embodied interaction and spatial reasoning. In conclusion, we hope this work inspires the next generation of LHRDs, that for long have been constrained to a static shape with varying resolution scales.

Acknowledgments

This work was supported in part by NSF grants IIS-2107224 and IIS-2529207 and ONR grant N000142312124.

References

[1] 3D Slicer Community. 2026. 3D Slicer. https://www.slicer.org/. Accessed: 2026-02-05.

[2] Christopher Andrews, Alex Endert, Beth Yost, and Chris North. 2011. Information visualization on large, high-resolution displays: Issues, challenges, and opportunities. *Information Visualization* 10, 4 (2011), 341–355. doi:10.1177/1473871611415997

[3] Christopher Andrews and Chris North. 2013. The Impact of Physical Navigation on Spatial Organization for Sensemaking. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2207–2216. doi:10.1109/TVCG.2013.205

[4] Maliha Ashraf, Alexandre Chapiro, and Rafał K Mantiuk. 2025. Resolution limit of the eye—how many pixels can we see? *Nature Communications* 16, 1 (2025), 9086. doi:10.1038/s41467-025-64679-2

[5] Teo Babic, Harald Reiterer, and Michael Haller. 2022. Understanding and Creating Spatial Interactions with Distant Displays Enabled by Unmodified Off-The-Shelf Smartphones. *Multimodal Technologies and Interaction* 6, 10 (2022), 94. doi:10.3390/mti6100094

[6] Robert Ball and Chris North. 2005. Analysis of User Behavior on High-Resolution Tiled Displays. In *Proceedings of INTERACT 2005: IFIP TC13 International Conference on Human-Computer Interaction*. Springer, 350–363. doi:10.1007/11555261_30

[7] Robert Ball and Chris North. 2005. Effects of tiled high-resolution display on basic visualization and navigation tasks. In *CHI Extended Abstracts on Human Factors in Computing Systems*. 1196–1199. doi:10.1145/1056808.1056875

[8] Robert Ball, Chris North, and Doug A Bowman. 2007. Move to improve: promoting physical navigation to increase user performance with large displays. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2007), 191–200. doi:10.1145/1240624.1240656

[9] Ilyasse Belkacem, Christian Tominski, Nicolas Médoc, Søren Knudsen, Raimund Dachselt, and Mohammad Ghoniem. 2024. Interactive Visualization on Large High-Resolution Displays: A Survey. *Computer Graphics Forum* (2024), e15001. doi:10.1111/cgf.15001

[10] Anastasia Bezerianos and Petra Isenberg. 2012. Perception of Visual Variables on Tiled Wall-Sized Displays for Information Visualization Applications. *IEEE Transactions on Visualization and Computer Graphics* 18, 12 (2012), 2516–2525. doi:10.1109/TVCG.2012.251

[11] Dhruv Bihani, AKM Amanat Ullah, Charles-Olivier Dufresne-Camaro, William Delamare, Pourang Irani, and Khalad Hasan. 2024. Exploring Pointer Enhancement Techniques for Target Selection on Large Curved Display. *Proceedings of the ACM on Human-Computer Interaction* 8, ISS (2024), 214–235. doi:10.1145/3698135

[12] Saeed Boorboor, Doris Gutiérrez-Rosales, Ahamed Shoaib, Chahat Kalsi, Yue Wang, Yuyang Cao, Xianfeng Gu, and Arie E. Kaufman. 2025. Silo: Half-Gigapixel Cylindrical Stereoscopic Immersive Display. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*. doi:10.1109/VR59515.2025.00072

[13] Saeed Boorboor, Shreeraj Jadhav, Mala Ananth, David Talmage, Lorna Role, and Arie Kaufman. 2018. Visualization of neuronal structures in wide-field microscopy brain images. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2018), 1018–1028. doi:10.1109/TVCG.2018.2864852

[14] Saeed Boorboor, Yoonsang Kim, Ping Hu, Josef M Moses, Brian A Colle, and Arie E Kaufman. 2023. Submerse: Visualizing storm surge flooding simulations in immersive display ecologies. *IEEE Transactions on Visualization and Computer Graphics* 30, 9 (2023), 6365–6377. doi:10.1109/TVCG.2023.3332511

[15] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology* 3, 2 (2006), 77–101. doi:10.1191/1478088706qp063oa

[16] Xavier Brioso, Claudia Calderon-Hernandez, Javier Irizarry, and Daniel Paes. 2019. Using immersive virtual reality to improve choosing by advantages system for the selection of fall protection measures. In *Proceedings of the International Conference on Computing in Civil Engineering*. American Society of Civil Engineers, 146–153. doi:10.1061/9780784482421.019

[17] Marco Cavallo, Mishal Dholakia, Matous Havlena, Kenneth Ocheltree, and Mark Podlaseck. 2019. Dataspace: A Reconfigurable Hybrid Reality Environment for Collaborative Information Analysis. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*. 145–153. doi:10.1109/VR.2019.8797733

[18] Jian Chen, Haipeng Cai, Alexander P Auchus, and David H Laidlaw. 2012. Effects of stereo and screen size on the legibility of three-dimensional streamtube visualization. *IEEE transactions on Visualization and Computer Graphics* 18, 12 (2012), 2130–2139. doi:10.1109/TVCG.2012.216

[19] Maxime Cordeil, Andrew Cunningham, Tim Dwyer, Bruce H Thomas, and Kim Marriott. 2017. ImAxes: Immersive axes as embodied affordances for interactive multivariate data visualisation. In *Proceedings of the ACM Symposium on User Interface Software and Technology*. 71–83. doi:10.1145/3126594.3126613

[20] Paulo Cortez, António Cerdeira, Fernando Almeida, Telmo Matos, and José Reis. 2009. Modeling wine preferences by data mining from physicochemical properties. *Decision Support Systems* 47, 4 (2009), 547–553. doi:10.1016/j.dss.2009.05.016

[21] Carolina Cruz-Neira, Daniel J Sandin, Thomas A DeFanti, Robert V Kenyon, and John C Hart. 1992. The CAVE: Audio visual experience automatic virtual environment. *Commun. ACM* 35, 6 (1992), 64–73. doi:10.1145/129888.129892

[22] Thomas DeFanti, Daniel Acevedo, Richard Ainsworth, Maxine Brown, Steven Cutchin, Gregory Dawe, Kai-Uwe Doerr, Andrew Johnson, Chris Knox, Robert Kooima, et al. 2011. The future of the CAVE. *Open Engineering* 1, 1 (2011), 16–37. doi:10.2478/s13531-010-0002-5

[23] Thomas A DeFanti, Gregory Dawe, Daniel J Sandin, Jurgen P Schulze, Peter Otto, Javier Girado, Falko Kuester, Larry Smarr, and Ramesh Rao. 2009. The StarCAVE, a third-generation CAVE and virtual reality OptIPortal. *Future Generation Computer Systems* 25, 2 (2009), 169–178. doi:10.1016/j.future.2008.07.015

[24] Stefan Eilemann, Maxim Makhinya, and Renato Pajarola. 2009. Equalizer: A Scalable Parallel Rendering Framework. *IEEE Transactions on Visualization and Computer Graphics* 15, 3 (2009), 436–452. doi:10.1109/TVCG.2008.104

[25] Dylan Fafard, Ian Stavness, Martin Dechant, Regan Mandryk, Qian Zhou, and Sidney Fels. 2019. Ftvr in vr: Evaluation of 3d perception with a simulated volumetric fish-tank virtual reality display. In *Proceedings CHI Conference on Human Factors in Computing Systems*. 1–12. doi:10.1145/3290605.3300763

[26] Alessandro Febreti, Arthur Nishimoto, Terrance Thigpen, Jonas Talandis, Lance Long, JD Pirtle, Tom Peterka, Alan Verlo, Maxine Brown, Dana Plepys, et al. 2013. CAVE2: a hybrid reality environment for immersive simulation and information analysis. *Proceedings of SPIE - The Engineering Reality of Virtual Reality* 8649 (2013), 9–20. doi:10.1117/12.2005484

[27] Tobias Höllerer, JoAnn Kuchera-Morin, and Xavier Amatriain. 2007. The Allosphere: a large-scale immersive surround-view instrument. *Emerging Display Technologies: The Future of Displays and Interaction* (2007), 3–es. doi:10.1145/1278240.1278243

[28] Claudia Hänel, Benjamin Weyers, Bernd Hentschel, and Torsten W. Kuhlen. 2016. Visual Quality Adjustment for Volume Rendering in a Head-Tracked Virtual Environment. *IEEE Transactions on Visualization and Computer Graphics* 22, 4 (2016), 1472–1481. doi:10.1109/TVCG.2016.2518338

[29] Robert JK Jacob, Audrey Girouard, Leanne M Hirshfield, Michael S Horn, Orit Shaer, Erin Treacy Solovey, and Jamie Zigelbaum. 2008. Reality-based interaction: a framework for post-WIMP interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 201–210. doi:10.1145/1357054.1357089

[30] Mikkel R Jakobsen, Yonas Sahlemariam Haile, Søren Knudsen, and Kasper Hornbæk. 2013. Information visualization and proxemics: Design opportunities and empirical findings. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2386–2395. doi:10.1109/TVCG.2013.166

[31] Mikkel R Jakobsen and Kasper Hornbæk. 2015. Is moving improving? some effects of locomotion in wall-display interaction. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. 4169–4178. doi:10.1145/2702123.2702312

[32] Yvonne Jansen, Jonas Schjerlund, and Kasper Hornbæk. 2019. Effects of locomotion and visual overview on spatial memory when interacting with wall displays. In *Proceedings CHI Conference on Human Factors in Computing Systems*. 1–12. doi:10.1145/3290605.3300521

[33] Torsten Wolfgang Kuhlen and Bernd Hentschel. 2014. Quo vadis CAVE: Does Immersive Visualization Still Matter? *IEEE Computer Graphics and Applications* 34, 5 (2014), 14–21. doi:10.1109/MCG.2014.97

[34] Gyouhyung Kyung and Sungryul Park. 2021. Curved versus flat monitors: Interactive effects of display curvature radius and display size on visual search performance and visual fatigue. *Human Factors* 63, 7 (2021), 1182–1195. doi:10.1177/0018720820922717

[35] Jiazhou Liu, Arnaud Prouzeau, Barrett Ens, and Tim Dwyer. 2020. Design and evaluation of interactive small multiples data visualisation in immersive spaces. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*. 588–597. doi:10.1109/VR46266.2020.00081

[36] Jiazhou Liu, Arnaud Prouzeau, Barrett Ens, and Tim Dwyer. 2022. Effects of display layout on spatial memory for immersive environments. *Proceedings of the ACM on Human-Computer Interaction* 6, ISS (2022), 468–488. doi:10.1145/3567729

[37] Siddhesh Manjrekar, Shubhrika Sandilya, Deesha Bhosale, Sravanthi Kanchi, Adwait Pitkar, and Mayur Gondhalekar. 2014. CAVE: an emerging immersive technology–a review. In *Proceedings of the International Conference on Computer Modelling and Simulation*. IEEE, 131–136. doi:10.1109/UKSim.2014.20

[38] Elisabeth Mayer, Thomas Odaker, Daniel Kolb, Simone Müller, and Dieter Kranzlmüller. 2024. LED CAVE-New Dimensions for Large-Scale Immersive Installations. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces Workshops*. 515–519. doi:10.1109/vrw62533.2024.00099

[39] John P McIntire, Paul R Havig, and Eric E Geiselman. 2014. Stereoscopic 3D displays and human performance: A comprehensive review. *Displays* 35, 1 (2014), 18–26. doi:10.1016/j.displa.2013.10.004

[40] Mathieu Nancel, Julie Wagner, Emmanuel Pietriga, Olivier Chapuis, and Wendy Mackay. 2011. Mid-air pan-and-zoom on wall-sized displays. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*. 177–186. doi:10.1145/1978942.1978969

[41] Ava Nederlander, Zainab Aamir, and Arie E Kaufman. 2026. Scale-Aware Navigation of Astronomical Survey Imagery Data on High Resolution Immersive Displays. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (IEEE VRW)*.

[42] OptiTrack. 2026. OptiTrack Flex 13. https://www.optitrack.com/cameras/flex-13/. Accessed: 2026-01-24.
[43] OptiTrack. 2026. OptiTrack Sync Configuration with an HTC Vive System. https://docs.optitrack.com/virtual-reality/vr-hmd-setup/sync-configuration-with-an-htc-vive-system. Accessed: 2026-01-26.
[44] Daniel Paes, Javier Irizarry, and Diego Pujoni. 2021. An evidence of cognitive benefits from immersive design review: Comparing three-dimensional perception and presence between immersive and non-immersive virtual environments. *Automation in Construction* 130 (2021), 103849. doi:10.1016/j.autcon.2021.103849
[45] Alexandros Panagiotidis, Steffen Frey, and Thomas Ertl. 2015. Exploratory Performance Analysis and Tuning of Parallel Interactive Volume Visualization on Large Displays.. In *EuroVis (Short Papers)*. 13–17. doi:10.2312/eurovisshort.20151118
[46] Charilaos Papadopoulos, Kaloian Petkov, Arie E Kaufman, and Klaus Mueller. 2015. The Reality Deck–an Immersive Gigapixel Display. *IEEE Computer Graphics and Applications* 35, 1 (2015), 33–45. doi:10.1109/MCG.2014.80
[47] Feng Qiu, Bin Zhang, Kaloian Petkov, Lance Chong, Arie Kaufman, Klaus Mueller, and Xianfeng David Gu. 2008. Enclosed Five-Wall Immersive Cabin. In *Advances in Visual Computing*, George Bebis et al. (Eds.). Springer Berlin Heidelberg, 891–900.
[48] Timothy A Sandstrom, Chris Henze, and Creon Levit. 2003. The hyperwall. In *Proceedings of the International Conference on Coordinated and Multiple Views in Exploratory Visualization*. IEEE, 124–133. doi:10.1109/CMV.2003.1215010
[49] Lauren Shupp, Christopher Andrews, Margaret Dickey-Kurdziolek, Beth Yost, and Chris North. 2009. Shaping the display of the future: The effects of display size and curvature on user performance and insights. *Human–Computer Interaction* 24, 1-2 (2009), 230–272. doi:10.1080/07370020902739429
[50] Lauren Shupp, Robert Ball, Beth Yost, John Booker, and Chris North. 2006. Evaluation of viewport size and curvature of large, high-resolution displays. In *Proceedings of Graphics Interface*. 123–130. doi:10.5555/1143079.1143100
[51] Mahsa Sinaei Hamed, Pak Kwan, Matthew Klich, Jillian Aurisano, and Fateme Rajabiyazdi. 2024. The Elephant in the Room: Expert Experiences Designing, Developing and Evaluating Data Visualizations on Large Displays. *Proceedings of the ACM on Human-Computer Interaction* 8, ISS (2024), 301–329. doi:10.1145/3698139
[52] Samiha Sultana and Md Ashraful Alam. 2024. Exploring the Impact of Display Field-of-view and Display Type on Text Readability and Visual Text Search on Large Displays in Virtual Reality: Impact of Display Field-of-view and Type on Text Readability and Search in Large VR Displays. *Proceedings of Graphics Interface* (2024), 1–11. doi:10.1145/3670947.3670984
[53] AKM Amanat Ullah, William Delamare, and Khalad Hasan. 2023. Exploring users' pointing performance on virtual and physical large curved displays. *Proceedings of the ACM Symposium on Virtual Reality Software and Technology* (2023), 1–11. doi:10.1145/3611659.3615710
[54] HTC Vive. 2026. SteamVR Base Station. https://www.vive.com/us/accessory/base-station2/. Accessed: 2026-01-22.
[55] Volfoni RF Active 3D System. 2026. http://volfoni.com/en/activhub-rf50/. Accessed: 2026-01-30.
[56] Colin Ware and Glenn Franck. 1996. Evaluating stereo and motion cues for visualizing information nets in three dimensions. *ACM Transactions on Graphics (TOG)* 15, 2 (1996), 121–140. doi:10.1145/234972.234975
[57] Chunxue Wei, Difeng Yu, and Tilman Dingler. 2020. Reading on 3D surfaces in virtual environments. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*. 721–728. doi:10.1109/vr46266.2020.00095
[58] Matt Whitlock, Stephen Smart, and Danielle Albers Szafir. 2020. Graphical perception for immersive analytics. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE, 616–625. doi:10.1109/VR46266.2020.00084
[59] Paul R Woodward, David H Porter, Michael R Knox, Steven T Andringa, and Aaron Stender. 2026. A System for Interactive Volume Visualization on the PowerWall. http://www.lcse.umn.edu/research/powerwall/powerwall.html. Accessed: 2026-01-18.
[60] Bradley Yost, Yigit Haciahmetoglu, and Chris North. 2007. Beyond visual acuity: The perceptual scalability of information visualizations on large displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 101–110. doi:10.1145/1240624.1240639
[61] Beth Yost and Chris North. 2006. The Perceptual Scalability of Visualization. *IEEE Transactions on Visualization and Computer Graphics* 12, 5 (2006), 837–844. doi:10.1109/TVCG.2006.184

Received 12 February 2026; revised 30 April 2026; accepted 11 May 2026