FlexiCAVE: A Dynamically Configurable High-Resolution Display Facility

Zainab Aamir, Saeed Boorboor, Ahamed Shoaib, Chahat Kalsi, and Arie E. Kaufman, *Fellow, IEEE*

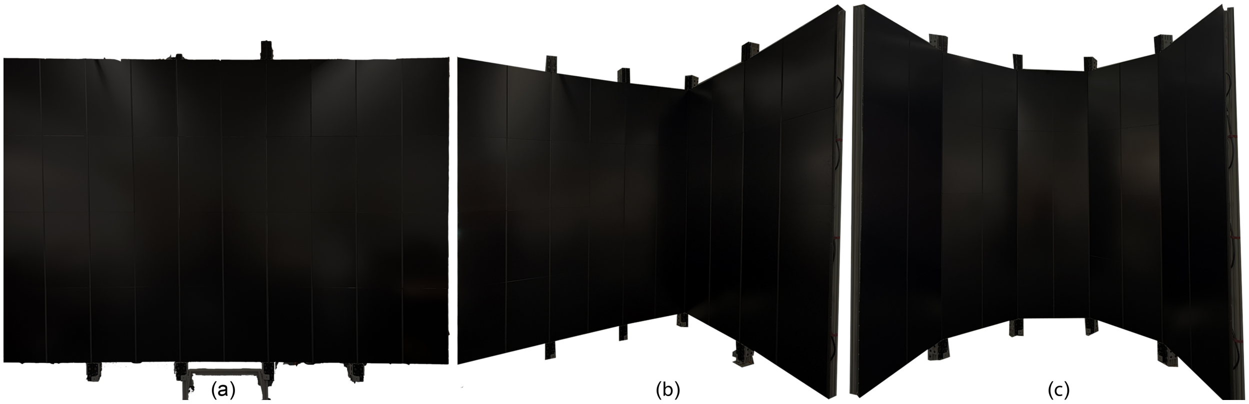

Fig. 1: The FlexiCAVE display facility can be dynamically transformed into various display configurations by adjusting the hinged display columns into many various configurations, such as: (a) flat, (b) flat two-sided, or (c) curved.

Abstract—Large high-resolution displays (LHRDs) have become essential tools in visualization and visual analytics, providing expansive and flexible visualization spaces and supporting simultaneous details and context and intuitive physical navigation for data exploration. Despite significant advancements in display technology, particularly regarding display resolution and form factors, a persistent challenge has been determining optimal display configurations for diverse analytical tasks. This has motivated us to design and construct the FlexiCAVE, a novel LHRD facility uniquely designed with rotatable display columns, enabling dynamic horizontal curvature adjustments. Comprising forty high-pixel-density displays arranged across ten columns, the FlexiCAVE offers an active stereo system with a total resolution of approximately 83 million pixels, and the columns can rotate inward up to 90° per hinge. Unlike existing static or single-configuration curved displays, the FlexiCAVE dynamically adapts its curvature in real-time, supported by our custom-developed rendering engine that synchronizes camera views with the changing display layouts. To demonstrate the utility of the FlexiCAVE, we present several application scenarios showcasing the facility significance and flexibility, including interactive radial slicing in volume rendering through physical column rotation and the dynamic switching between multivariate plots. With its adaptability and innovative design, we believe that FlexiCAVE represents the next generation in LHRD technology, setting new standards and design spaces for future large data visual systems.

Index Terms—Immersive Visualization, High-resolution Tiled Display, Curved Display, Foldable Display, Stereo, CAVE

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 [2, 4, 5, 54]. Compared to traditional desktop settings, where a display screen occupies only a limited field of view (FoV) of the user, LHRD facilities and devices 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 [4], simultaneously providing details and context.

During the three decades since the introduction of the first generation CAVE [15] and PowerWall [53], the proliferation of inexpensive GPU power and high-pixel-density displays has driven the evolution of LHRD facilities from mega- to giga-pixel resolutions [42], and with novel display configurations, spanning from planar and rectangular setups to more immersive, dome-like [24] and cylindrical [8, 22] arrangements. However, in parallel to these developments, there has been an ongoing question of display form factor and configuration [29]. Particularly, there have been studies comparing *flat* versus *curved* LHRDs, as well as different display curvatures, for various tasks, such as visual search [28,29,31,46], and data visualization [30] and interaction [7,48].

Naturally, since data come in all shapes and sizes, there is no one-size-fits-all display configuration. For example, studies have shown that curved displays provide higher levels of immersion [8, 22], and outperform flat displays in search and comparison tasks [29]. In contrast, flat holds an advantage over curved displays in overview tasks [29], in multi-user collaboration settings [42], and when the data presented have lower visual density [30]. Moreover, some studies have indicated that a significant difference in user performance cannot be conclusively identified between flat and varying curvature display configurations [31]. However, it is worth mentioning that many works have highlighted an overall preference for curved LHRDs, which offer enhanced immersive viewing experiences and improved task performance.

Motivated by this scenario, the lessons learned, and the vision for the future of LHRDs [17], we have designed and constructed the FlexiCAVE, an LHRD facility with rotatable hinged display columns that

• *All authors are with the Center for Visual Computing, Stony Brook University, Stony Brook, NY.*
Corresponding E-mail: ari@cs.stonybrook.edu

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

allow for dynamic layout configurations. It comprises forty high-pixel-density displays with an active stereo system in a 3.14m × 2.14m arrangement (when *flat*), offering a resolution of 7620 × 10800 pixels, totaling approximately 83 million pixels. The screens are distributed across 10 columns where every two columns can be rotated inwards up to 90°. To the best of our knowledge, the FlexiCAVE is the only dynamically configurable (in terms of horizontal curvature layout) stereo display system with the highest resolution. Fig. 1 demonstrates three different configurations the facility can be folded into, and Fig. 2 shows a diagram of the FlexiCAVE, along with some of its hardware components.

In this paper, we outline the hardware and engineering design decisions that went into constructing and setting up the facility. At the application level, while most parallel and distributed rendering middlewares assume a static screen layout, we have leveraged the idea of flexible distributed execution [20] to develop a rendering engine that updates camera views and gesture tracking, in real-time, to adapt to the FlexiCAVE dynamic configuration ability. To demonstrate the FlexiCAVE, we show several use-case scenarios that explore novel design spaces and data interactions based on varying spatial alignment of the displays. Specifically, in addition to the naïve use case of setting up the facility, such as flat, “L"-shaped, “U"-shaped, or curved, based on different application needs, we (1) exhibit a volume rendering visualization that performs interactive radial slicing by physically rotating the display columns to the desired radial slice position, (2) explore multivariate data analysis by switching through scatter plot and parallel plot visualizations based on screen configuration, and (3) replicate a user study [30] for studying the layout of small multiples based on curvature, in our case, using a physical LHRD facility. Lastly, we discuss perceptual aspects and give an outlook on how the next generation FlexiCAVE can be extended in the future.

2 Related Works

LHRDs can take the form of powerwalls or cave-like facilities and have evolved over the past few decades into various immersive and high-resolution display configurations. One of the earliest and influential designs for an LHRD facility is the CAVE [32] - a cube-shaped system originally conceptualized by Cruz et al. [15] The setup enables 3D visualizations by projecting images onto surrounding walls, ceilings, and floors. Multiple systems have since evolved from this original CAVE concept [16]. One such advancement is CAVE2 [22], which used LCD technology with 72 displays in a cylindrical configuration. A non-traditional configuration was introduced with StarCAVE [18], a five-sided pentagonal square structure. The Reality Deck at Stony Brook University further expanded the possibilities of LHRDs with a four-walled, room-sized configuration with 416 high-resolution displays [42]. They also introduced SILO, a fully immersive, high resolution cylindrical tiled-display facility that provides an almost 360 degree field-of-regards [8]. NASA's hyperwall [44] offers a unique approach with its flat configuration of 49 screens, each independently adjustable along the z-axis, enabling non-planar arrangements by modifying pitch, yaw, and translational adjustments between rows and columns. Similarly, Dataspace [12] is a reconfigurable mixed –reality environment consisting of 15 high resolution OLED displays mounted on 7-degree-of-freedom robotic arms attached to the ceiling. The setup allows the screens to dynamically move, rotate and reposition. Additionally, the setup includes two HD projectors, augmented reality (AR) headsets and a virtual reality (VR) extension, which includes a digital twin of the Dataspace environment and an interactive central table. More recently, Mayer et al. [34] introduced a novel LED-based five-sided CAVE, significantly reducing space requirements and maintenance complexities compared to the traditional projector-based setups.

Several earlier works have investigated how variations in display configurations, including form factor, size and curvature, influence user performance and perception [29]. Shupp et al. [45] examined user performance across varying display sizes and curvatures, finding that larger, curved displays enhance spatial performance and reduce physical navigation effort during multiscale geo-spatial search tasks. Similarly, Kyung et al. [28] explored the interactive effects of the display curvature radius and display size on visual search tasks, analyzing performance metrics such as accuracy, speed and fatigue. Sultana and Alam [46] investigated how curved display FoVs impact reading comprehension and visual search performance on large displays in VR environments. Liu et al. [30] explored adapting small multiples data visualization techniques to immersive 3D spaces, introducing a "shelves" metaphor and examining how layout curvature impacts user performance. Their findings suggest flat layouts are more efficient for smaller datasets, whereas semi-circular arrangements become preferable as the dataset size increases. In a subsequent study, Liu [31] further evaluated the influence of flat, semi-circular, and fully wraparound display geometries on spatial memory within immersive environments Other works have explored user interaction efficiency across varying display geometries. Amanat et al. [48] analyzed pointing performance on virtual displays with different curvatures, finding flatter configurations improved pointing speed, while curved displays offered a more consistent FoV alignment. Bihani et al. [7] extend this by exploring how user position influences pointing performance on large curved displays, emphasizing the need for adaptive interaction models. Zannoli and Banks [55] investigated the perceptual consequences of curved screens, noting that curvature marginally improved FoV but substantially reduced reflections and supported a broader range of user positions. However, accurately rendering visualizations on their 65-inch display with a curvature radius of 4.18 meters requires more advanced techniques due to its geometric complexity. Comparisons of flat and curved immersive environments further indicate that curved displays enhance depth perception and navigation, outperforming traditional powerwall setups in visual exploration tasks [8]. Wei et al. [52] examined the impact of concave, convex, and planar display geometries on reading performance in VR, revealing that curvature significantly influences legibility, reading speed, and comprehension.

With the growing complexity of LHRD facilities, there is a growing need for scalable and immersive rendering tools to support them. One of the earlier works addressing this is SAGE (Scalable Adaptive Graphics Environment) [43], a middleware solution designed to support collaborative visualization on high resolution tiled displays. Equalizer [20], a parallel rendering framework built on OpenGL, provides an API to develop scalable graphics applications for immersive environments and display walls. Additionally, CGLX [19], a cross-platform graphics library, provides a high-performance visualization framework specifically tailored for networked display environments. CGLX enables transparent porting of single-node OpenGL applications to visualization clusters. CGLX allows for adding more displays dynamically with their interface. SAGE2 [33] advances to address limitations in collaboration workflows and utilizes web and cloud technology to facilitate co-located and remote collaboration on LHRDs. Building on these foundations, UniCAVE [47] introduced an accessible Unity3D plugin specifically aimed at non-head-mounted immersive display systems. It provides an available and easy-to-use Unity3D extension package for distributed, immersive tiled, or projection-based VR display systems. Similarly, the development of the nDisplay plugin [27] for Unreal Engine 4 has expanded the software's capability as a visualization platform for LHRDs and powerwalls. It should be noted that most of these frameworks are currently not actively maintained, highlighting a gap in continued development and support for visualization tools.

3 The FlexiCAVE

The construction of the FlexiCAVE was guided by the following key design principles:

- **D1** A high-resolution display facility that provides an optimal resolution target, where the pixel density matches the visual acuity of the human visual system [22, 42]
- **D2** The ability to *modify* the display setup such that it can accomodate varying curvatures [29].
- **D3** Support for stereoscopic rendering.

Belkasim et al. [6] envisioned LHRDs as “a coherent 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

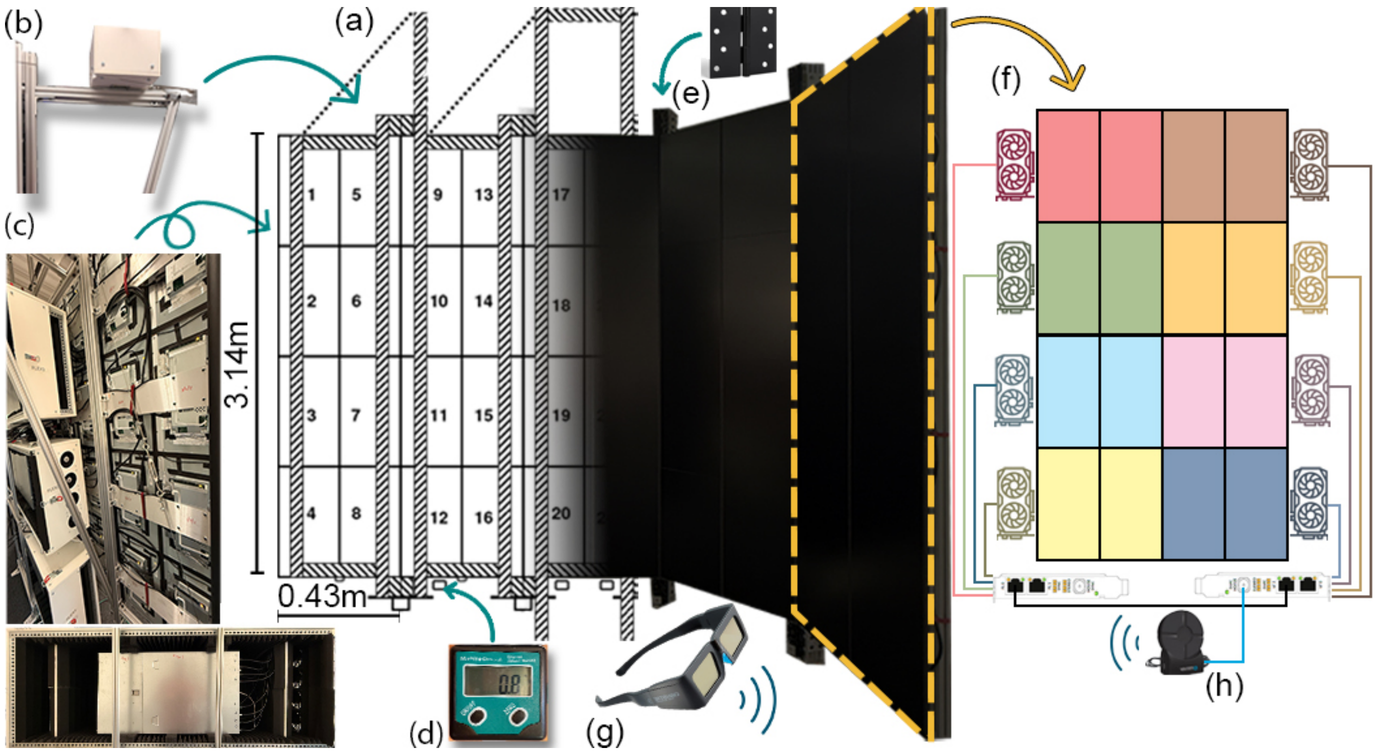

Fig. 2: (a) Partial 2D-diagram of the FlexiCAVE structure showing the aluminum support frames, hardware components, and display layout, followed (to the right) by a demonstrated configuration. (b) Counter-weight to maintain stability of the suspended display column pairs as they are rotated. (c) A rear view of the FlexiCAVE showing the mounted displays and our custom noise-canceling cabinets for the visualization server nodes. A view of the inside is shown at the *bottom* of (c). (d) Angle sensors mounted on all the hinges. (e) Hinges used to rotate the display column-pairs. (f) A diagram illustrating the GPU configuration for each node (2 display columns). (g) Active stereo glasses controlled by (h) an RF-emitter connected to the GPU server timing to synchronize frame swaps during active stereo.

scale, Shupp et al. [29] highlighted that the positioning of displayed information within the environment and its relationship to the user influences the manner in which 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. Specifically, advancements in display technologies have allowed for enhanced visual acuity (the quality of the visuals a display can deliver), enabling users to approach the display surfaces and naturally perform multiscale exploration rather than perceiving the resolution limits of the display technology. Furthermore, researchers have investigated configurations and curvatures that best adapt to user interaction and environmental context. However, most current LHRD facilities are restricted to *fixed* or *static* configurations. Therefore, there remains a need for a design that effectively offers the flexibility of varying curvatures [17, 29] to adapt to different user needs and viewing environments.

In an effort to address this gap in LHRD design, we have constructed a flexible powerwall called the FlexiCAVE. This facility comprises 10 columns of high-resolution display screens (**D1**), with every two columns pivoted to allow flexible inward bending (**D2**). Fig. 2 (a) shows a to-scale rendered diagram of the facility. Each column of the FlexiCAVE is 3.14m high and 0.214m wide, equipped with four 1920×1080 FHD monitors mounted in portrait orientation. Moreover, it has been shown that using stereoscopic technology in virtual environments improves depth perception accuracy and enhances interpretation of 3D models, thereby increasing productivity when carrying out visualization tasks [11, 40]. Therefore, to maintain a high pixel resolution and density for stereo rendering (**D3**), we chose displays with high refresh rates that can support an active stereo system.

3.1 Hardware Setup

When building an LHRD facility, each component, from the display system to the visualization cluster, influences its overall performance and user experience. This section outlines several hardware and engineering design decisions we made during the construction of the FlexiCAVE.

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 outlined in the literature [16, 22, 42] when choosing the display hardware:

- *Image quality*: The display panels should be high-quality with good contrast, backlight uniformity, and large viewing angles.
- *Resolution target*. To achieve a high-resolution target, the monitors should provide ∼100 pixels per inch (PPI).
- *Bezel size*. Ideally, the bezel should be as narrow as possible; however, it should not exceed 8mm for a 23in display and 15mm for a 30in display.
- *Stereo support*. Stereo is very desirable, but not at the cost of significantly reduced pixel density.

We evaluated several commercially available displays based on their bezel sizes and panel technologies. To best meet the design goals and display criteria, we eliminated projector technology because of its low resolution and the frequent need for calibration after screen movement. Furthermore, we did not proceed with micrometer LED or OLED panels because of their high initial and maintenance costs, given the necessary PPI and laser alignment required to maintain the proposed facility. Additionally, we could not find off-the-shelf O/LED monitors compatible with commercially available active stereo systems that synchronize with GPU output frames. Among LCD technology, the key challenge then was balancing bezel size with stereo support.

Balancing all factors, we opted for the ViewSonic XG2431 panels [49], a professional 24-inch FHD In-Plane Switching (IPS) panel with 1920×1080 resolution. IPS panels are a class of LCD technology known for their superior color accuracy, wide viewing angles, and consistent image quality compared to Twisted Nematic (TN) and Vertical Alignment (VA) panels, hence their prevalence in gaming monitors and laptops. Moreover, the particular ViewSonic IPS panel has a reasonably high response time (1ms) with a refresh rate of 240 Hz, which allowed us to employ an active stereo system without compromising pixel resolution. Finally, although the bottom bezel is 6mm, we modified the monitors with a custom mount that reduced the bezels to 4mm.

Display Layout Given the available physical space and design goal **D2**, we constructed the FlexiCAVE with 10 display columns. The displays are mounted on a custom light-weight aluminum frame. In its design, the center two columns are fixed to the ceiling and floor, and the remaining column pairs are suspended to the central frame, as shown in the drawing in Fig. 2 (a). The hinges between each column pair can freely rotate the screens at an inward angle ranging from 90° to 180°. Consequently, the spatial dimensions of the FlexiCAVE can vary from 2.14m wide when flat and approximately a radius of 1m when fully enclosed. Fig. 1 shows examples of three possible configurations the facility can be angled into. A high-precision inclinometer sensor (Fig. 2 (d)) is connected to each hinge that transmits the rotation angle of the column pair to the visualization system, which in turn is used to update the visualizations. Moreover, a counterweight is installed on each pair of rotatable columns to provide stability and balance while they are rotated, as shown in Fig. 2 (b).

Tracking Interactions in LHRDs typically adopt embodied interactions, such as walking and gestural interfaces [36], and by integrating hand-held controllers or pointing devices [3]. Moreover, off-axis stereoscopic rendering requires tracking the user's gaze, which can be estimated by tracking the head position and orientation. For effective outside-in tracking, we opted to install HTC SteamVR base stations [50] to utilize HTC controllers as wands, along with OptiTrack Flex 13 infrared (IR) cameras [38] for body and head tracking. While both utilize IR technology, an OptiHub synchronization hardware enables a seamless interface that avoids overlapping of the IR tracking lights [39].

One challenge in setting up the IR cameras was finding an optimal layout. Since there is significant manual and computational overhead to re-callibrate the camera extrinsic parameters every time a camera is displaced, they cannot be mounted on the rotatable FlexiCAVE frame. Therefore we mounted 8 IR cameras along the ceiling and floor of the FlexiCAVE room such that the IR markers would be visible in at least three cameras given any FlexiCAVE configuration. Similarly, we mounted four HTC base stations overhead along the ceiling, as suggested in the guidelines.

3.2 Visualization System

The visualization cluster driving the FlexiCAVE consists of four nodes: Two nodes power the suspended columns (end-nodes), each driving 16 display panels (2 columns), one node powers the central fixed column (center-node), which drives 8 displays, and one node is dedicated as the head node (head-node). Each node is equipped with the following:

- 2×Intel Xeon Ice Lake Silver Processor(16 Core, 2.40Ghz)
- 8×NVIDIA RTX A5000 (24 GB GDDR6)
- 2×NVIDIA Quadro Sync II Board
- 12×16GB DDR4 RAM

In the end-nodes, two monitors are connected to one GPU, whereas in the center-node, each monitor is connected to a single GPU. This difference in the monitor-to-GPU ratio is solely a configuration decision since we maintain a consistent hardware specification across all nodes. In all nodes, four GPUs are connected to one NVIDIA Quadro Sync board. A schematic layout of the configuration of the end-nodes is illustrated in Fig. 2 (f). The Sync board facilitates a highly synchronized and scalable system, aligning display refresh rates across multiple systems and mitigating imaging artifacts in multi-display configurations.

For setting up an active stereo system, the Sync boards are daisy-chained, and one node is designated as the timing server. A Volfoni ActiveHub RF emitter [51], shown in Fig. 2 (h), is connected to the timing server to synchronize the GPU left- and right-eye framebuffer swaps with its eye-wear, the Volfoni Edge RF glasses. Boorboor et al. [8] discovered when constructing the Silo immersive facility that

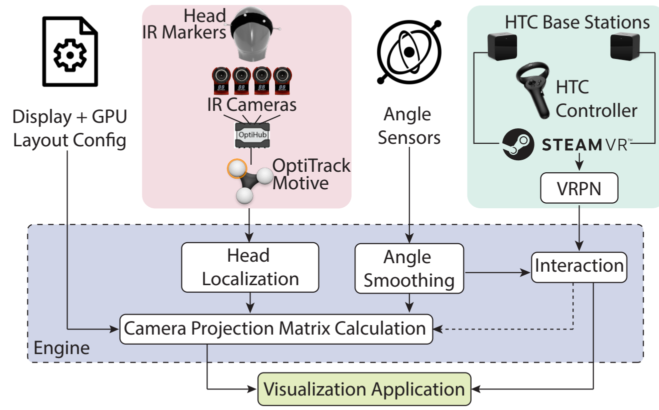

Fig. 3: System design of our modular Engine framework for the multi-node and multi-GPU FlexiCAVE setup. Using display and system configurations, along with user and controller tracking, Engine updates camera views and visualizations, abstracting them from the application design.

fiber-optic cables are inadequate for a high-framerate stereo system, as the graphics card driver may fail to receive coherent signals. Therefore, we followed suit and used high-quality copper cables to connect the monitors to the GPUs and placed the nodes next to the FlexiCAVE, enclosed in a noise-cancelling cabinet, to avoid signal loss or delay due to the extended cable length, as shown in Fig. 2 (c). Lastly, the monitors in the FlexiCAVE are positioned in portrait orientation since the direction of polarized light in landscape orientation is perpendicular to the RF eyewear lens' polarization filters.

Software Architecture We have implemented an OpenGL-based system that follows a replicated execution model [21], which we refer to as *Engine*. That is, for a multi-node and multi-GPU setup, instances of a target application are launched individually, and a master instance synchronizes components such as camera updates, interaction, tracking, and communications across all instances and nodes. Specifically, the Engine takes the system configuration, column angles from the inclinometer, and the head and controller 6 degrees of freedom (DoF) as input, and subsequently updates the application views and visualizations across all target instances. In our setup, the master instance is deployed on the dedicated head-node. Fig. 3 illustrates the Engine system flowchart.

The configuration (config) file describes the FlexiCAVE display dimensions, the physical positions of the rotating columns, the camera viewport for each target application instance, and the initial camera position in the virtual scene. Additionally, it includes system information such as the machine name and the assigned GPU. The GPU assignment ensures that the OpenGL context is initialized and uses the particular GPU for processing and rendering.

For head tracking, the IR cameras are synchronized using the OptiHub, which, paired with the OptiTrack Motive software [37], provides real-time 6 DoF pose estimation of the head. Likewise, button triggers and controller localization data are sent to the Engine over the virtual-reality peripheral network (VRPN) protocol.

At each frame, the Engine updates the camera view matrix for each target application OpenGL context based on the viewport defined in the config file and the *head* position, controller inputs, and display arrangement. The head position can either be static (using a constant 3D position defined in the config) or dynamic (from the IR head tracking system). The display arrangement is determined using the inclinometer readings. To adjust for rapid fluctuations in inclinometer readings due to noisy sensors or slight unintentional movement of the columns, the Engine performs a real-time smoothing technique based on rotational momentum. Rotational momentum is defined as an integer value that increments or decrements each frame based on the direction of angle change, bounded by a specified maximum momentum limit. A momentum threshold is established, and the output angle remains unchanged until the momentum surpasses this threshold. Once exceeded, the angle value is smoothly adjusted at a defined speed, measured in degrees per

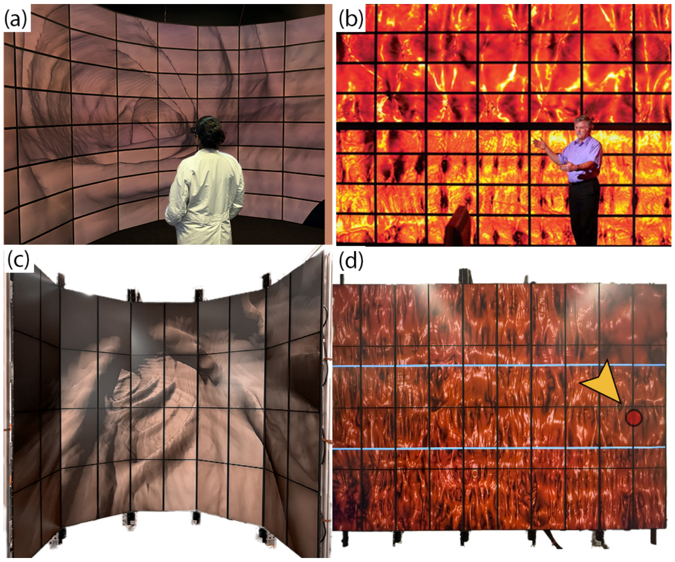

Fig. 4: (a) Immersive Virtual Colonoscopy (VC) visualized in a curved display facility. (b) 2D view of the flattened VC data visualized in a planar display facility. The reconfigurable display layout of the FlexiCAVE can facilitate both immersive and flat VC visualization techniques. (c) Immersive VC of the user-selected region (red cursor pointed by the yellow arrow) from the flattened colon view, shown in (d).

frame, towards the target angle, continuing this adjustment until either the target angle is reached or momentum falls back below the threshold. Additionally, in certain instances, the controller may be used to perform scene rotation and translation relative to the head.

Our design of the Engine abstracts application design from implementing camera view updates due to head tracking and display layout reconfiguration. Nevertheless, programmers have the ability to use angle thresholds as interaction triggers for updating visualizations, which we demonstrate in Sec. 4.

4 Use Case Applications

We present a few applications that demonstrate the utility of the FlexiCAVE and explore a new design space where the physical layout reconfiguration of the powerwall influences visual analysis. In this manuscript, all figures show a monoscopic view of the 3D applications for picture-quality purposes to avoid stereo ghosting. A video exhibiting our applications in stereoscopic view is included in the Supplementary Materials.

Use Case I: Layout Arrangements for Applications

One major constraint for any LHRD facility design is its fixed display layout. As discussed in Sec. 2, each unique layout enhances perception for a specific subset of applications. For instance, curved screens are best suited for immersive 3D applications where the displays approach the sphere of influence and perception of a human [1]. Conversely, flat screens facilitate a larger collaborative space and balance physical navigation and effective performance in visual comparison tasks for planar visualizations [30]. FlexiCAVE alleviates this limitation by supporting various layouts (flat, two-sided *L*-shape, three-sided *U*-shape) and display curvatures.

Consider the example of Virtual Colonoscopy (VC). It is a noninvasive screening method that allows an expert to examine the surface of the colon, similar to optical colonoscopy, by reconstructing computed tomographic (CT) images of a patient's abdomen. Given the inherent tubular shape of the colon, visualizing its 3D model in immersive [35] and curved display [8] settings has been shown to be an effective modality for screening for polyps (the precursor to colorectal cancer). Moreover, due to the length of the colon, alternative methods, such as colon flattening [25], have been introduced to improve inspection time and minimize misreads by presenting the entire inner surface of the colon as a single 2D image. Fig. 4 (a) shows immersive VC in the Silo, a curved display facility, and Fig. 4 (b) shows a flattened colon in the Reality Deck (RD), a planar display facility. Figs. 4 (c) and (d) demonstrate that it is possible to effectively deploy both visualization techniques in our single FlexiCAVE facility by rearranging the facility to optimal curved and flat arrangements, respectively.

To further improve visual analysis, we have implemented an interactive transitioning feature between the flattened and immersive VC visualizations. When inspecting the flattened colon in the flat layout, users can use the controller to point to a specific region. By bending the FlexiCAVE into a curved arrangement, the Engine transitions the application to immersive VC at the user-selected point. Users can then navigate through the 3D model and by bending out the FlexiCAVE back to the flat layout, Engine would transition to the flattened colon view, updating the new cursor position.

Use Case II: Dynamic 3D Scene View Update

One aspect of virtual navigation in LHRDs involves interactively rotating the virtual camera, either using a ubiquitous device or gestures, to bring hidden regions of the 3D scene into the user's view or the display space. Similar to the use case described above, the camera movement for each display viewport is uniform with respect to the head, constrained by the fixed LHRD shape. The flexibility to dynamically bend display columns along their hinges in the FlexiCAVE allows users to create unique visual perspectives.

We demonstrate this use case with an urban flooding visualization application, Submerse [10], utilized for studying flood progression and resilience and evacuation planning, shown in Fig. 5. To evaluate the utility of the FlexiCAVE, we invited domain science experts and emergency managers who previously participated in Submerse workshops and user studies in both flat (RD) and curved (Silo) LHRD facilities, for an informal study to explore the dynamically reconfigurable screen arrangement and provide feedback. To this end, the participants were first asked to familiarize themselves with the reconfigurable nature of the displays. They were also provided with a controller for 6 DoF scene navigation. After the warm-up, the participants were asked to study a flooding scenario for a new urban scene rather than prior studies to avoid memory recall, followed by an informal interview.

The feedback highlighted by most participants was that there was a learning curve involved in grasping the idea that the columns can be adjusted to change the views. Since in previous settings, the navigation was either automatic or controlled by a gamepad, they shared that they used to feel more like a *passenger* observing the flooding scenario from inside a defined *vessel*. However, once they became adept at reconfiguring the columns, they felt more in control of the viewing directions. One particular advantage brought up was that since displays were physically movable, the position of the display arrangement gave them a clearer sense of direction. This is illustrated in Fig. 5. When navigating looking ahead, as in Fig. 5 (a), they were able to bend one whole side by 90° from the center column to understand what landmarks are immediately perpendicular to the other side, as in Fig. 5 (b). Similarly, bending each column independently would render the scene in the respective direction, as shown in Fig. 5 (c). The insets from each figure highlight the view changes of the reference building and its surroundings. Moreover, by observing the physical layout of the FlexiCAVE, they shared that not only did it enhance the sense of viewing direction, but it also allowed them to return to a previous view, which, they reflected, is more difficult in flat or curved LHRDs, as they would lose their perception of "*north*" when rotating the virtual camera.

Lastly, all participants agreed that there is a physical demand in bending the heavy display columns, and it is only optimal in collaborative settings, where one can assist in reconfiguring while another is guiding. We appreciated this comment and address this issue in Sec. 5.

Motivated by the comment that the display arrangements improve the sense of direction, we conducted a brief study to analyze the gain in 3D view perception when users have the ability to change the layout, compared to a static flat layout. In our study design, we arrange cubes in a 3D scene, at various yaw angles from a fixed camera position. In each trial, the participants were asked to give an estimated angle between two cubes – reference and target – where the reference cube

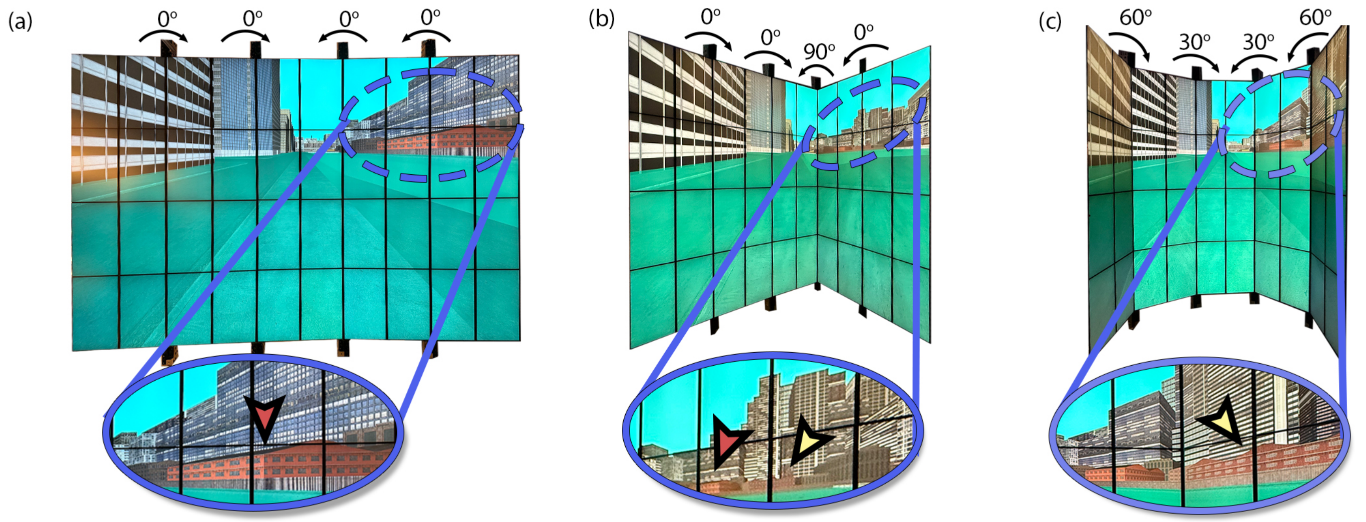

Fig. 5: Visualizing Submerse [10], an urban flood simulation visualization application, in the FlexiCAVE. The figure shows the virtual flooding scene in (a) flat, (b) L-shaped, and (c) curved layouts. The illustration on the top shows a top-down view of the FlexiCAVE layout with the angle differences between the column-pairs. The insets at the bottom highlight changes in the view, such as the building position and orientation pointed by the arrows, due to different display configurations. These views are updated in real-time as the FlexiCAVE layout is dynamically reconfigured.

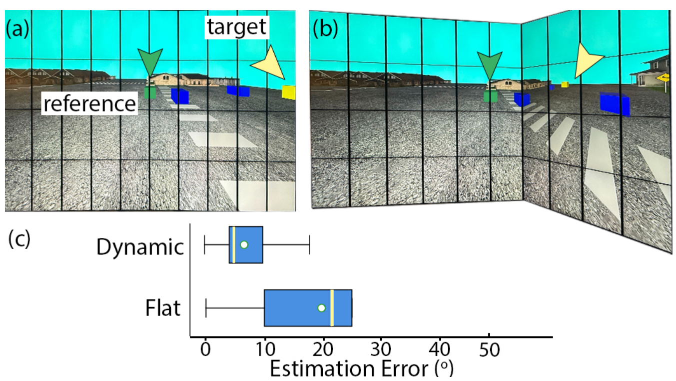

Fig. 6: Direction perception study for (a) the flat layout (target at 45°, 2m from reference), and (b) an example user-set dynamic layout (target at 30°, 6m from reference). (c) Angle estimation error for both layouts, with the white bar and dot representing median and mean errors, respectively.

was always placed at a fixed distance from the camera along the starting camera *lookat* direction. The target cube was positioned at 30°, 45°, and 90° from the reference, at distances 2m and 6m from the camera position. The study was conducted with 12 participants. Each participant estimated two angles for both distances for each layout type: static flat (flat) and dynamic. In the flat, the participants were shown the 3D scene in the flat FlexiCAVE configuration, and using a controller, they were allowed to rotate the camera yaw until they found the target cube in the scene and answered the estimated angle. In the dynamic tasks, participants were only allowed to rotate the FlexiCAVE displays to find the target cube (no controller). An example for each layout is shown in Fig. 6 (a) and (b) respectively. At the start of the study, users were shown angle differences using real-world objects to ensure that they were familiar with the angle separations.

For all trials, we calculated the difference between the angle estimated by the participant and the actual yaw angle between the reference and target cubes. Fig. 6 (c) shows the cumulative angle estimation error plot. With a median error of 7.5° (SD=5.36°), the participants were better at estimating the angle when they could interactively bend the displays. In contrast, the estimation error with the flat layout was 22° (SD=14.51°). Moreover, the mean angle estimation error for the target placed at 2m and 30°, 45°, and 90° in the flat layout was 25.6°, 10.3°, and 13.2°, respectively, and 8.6°, 5.2°, and 5.8° in the dynamic layout. For the target placed at 6m and 30°, 45°, and 90°, the estimation errors in the flat layout was 22.3°, 8.1°, and 10.8°, respectively, and 7.7°, 6.2°, and 5.1° in the dynamic layout. Therefore, the study supports the premise that physically configuring the display column layouts facilitates a better sense of viewing direction for 3D scenes.

Use Case III: Mapping Physical to Virtual Interaction

In their analysis of interactions, Jacob et al. [26] have identified *environmental awareness and skills* as an essential theme for enhancing human-computer interaction in emerging technologies. This refers to people having a sense of their surroundings and the skills necessary to negotiate, manipulate, and navigate within their environment. While static LHRD layouts facilitate navigation within the 3D virtual environment, they are often perceived as windows looking out to the virtual world, albeit in fixed directions. Similarly, in planar representations for information visualization, the layout of plots must conform to the constraints of the LHRD design. In contrast, the FlexiCAVE maintains a coherent spatial awareness between its physical layout and the corresponding virtual display, enabling users to effectively negotiate and manipulate visualizations. Utilizing this flexibility, we introduce PIVoT (Physical Interaction to Virtual Transformation), an interaction design space where rotating the FlexiCAVE display column-pairs facilitates tangible visualization updates. We illustrate PIVoT using two applications, one for scientific visualization using volume rendering and another for information visualization (InfoVis).

Volume Rendering. Scientific visualization techniques, such as volume rendering, have been well-demonstrated in immersive and LHRD environments [9, 23, 41, 42], where large and complex data can be efficiently processed and effectively visualized in high-resolution. A key interaction in volume rendering is using single or multiple cutting planes to slice through the volume, allowing users to examine its internal structures. However, configuring and manipulating these planes can be technical, especially in LHRDs where standard GUIs and keyboard-and-mouse interactions are limited. To this end, we design PIVoT such that the physical rotation of the display columns functions as interactive cutting planes. Specifically, each column of the FlexiCAVE acts as a tangible interface for a virtual cutting plane, which users can manipulate by rotating about their vertical pivot, translating physical actions into intuitive virtual plane movements.

In our implementation of this design, we initialize our volume renderer by placing the volume with origin, $\mathbf{O}_v$, at the center of the virtual FlexiCAVE center column-pair. Users may use the controller to translate and rotate the volume. Each display column C_i is associated with a cutting plane π_i in 3D space, where $i \in [-2, -1, 0, 1, 2]$, such that, $i = 0$ denotes the center column-pair, $i < 0$ are columns to the left of

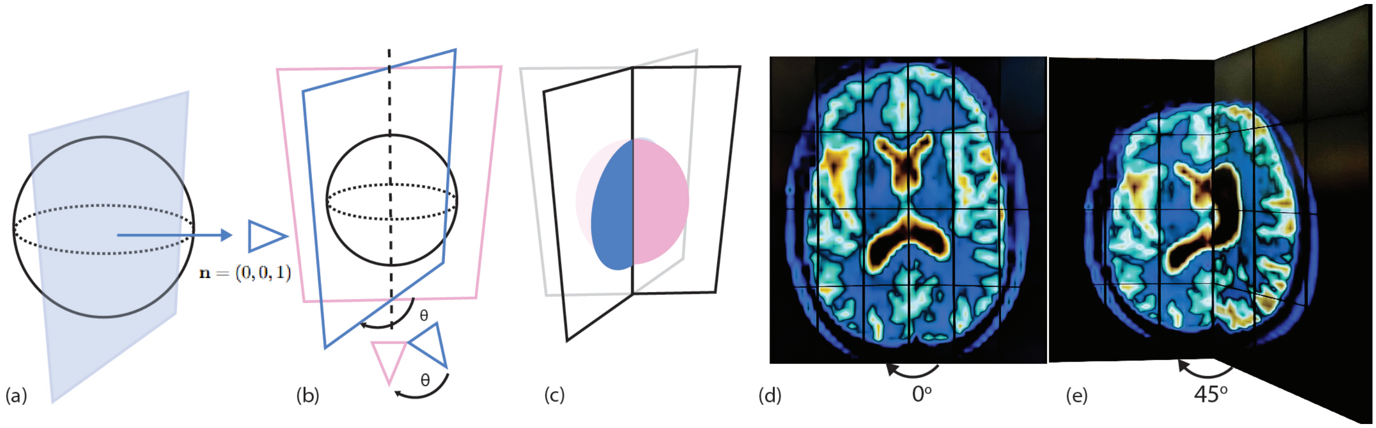

Fig. 7: PIVoT as a tangible cutting plane for volume rendering. (a) A volume is placed at the center of the FlexiCAVE, with the display acting as a cutting plane and an orthographic camera placed perpendicular to the display. (b) Rotating a column-pair virtually rotates the neighboring cutting plane. (c) Volumes from both cutting planes are aligned for seamless visualization. (d) Brain MRI volume rendering with the volume placed at the center, when the FlexiCAVE is flat. (c) Volume rendering result when one side of the FlexiCAVE is rotated by 45°.

the center and $i > 0$ are columns to the right of the center. For each plane, its origin, $\mathbf{O}_i$, is defined as:

$$\mathbf{O}_i = \mathbf{O}_x + w.\mathbf{d}_x \,, \qquad \text{where} \quad x = \begin{cases} i-1 & \text{for } i > 0 \\ i+1 & \text{for } i < 0 \end{cases} \tag{1}$$

where w is the width of the virtual column-pair displays, and d_i is the unit direction vector along the column-pair's hinge axis. An orthographic camera c_i for each column is placed at a fixed distance, d_c from the displays, with *lookat* v_i in the direction of the plane, parallel to $\mathbf{n}_i$. Therefore, in the initialized flat layout, all FlexiCAVE displays act as a cutting plane with normal, $\mathbf{n}_i = (0,0,1)$, perpendicular to the plane, as illustrated in Fig. 7 (a).

On receiving an angle update from the Engine, θ_{t+1}, we update the cutting plane such that it is rotated $R_{\mathbf{u}}(\Delta\theta)$around the camera up vector, where $\Delta\theta = \theta_{t+1} - \theta_t$. Thus, the new camera position and *lookat* vector for the rotated column is updated as

$$c_i' = \mathbf{O}_x + R_{\mathbf{u}}(\Delta\theta)\,(c_i - \mathbf{O}_x) \tag{2}$$

$$\mathbf{O}_i' = \mathbf{O}_x + R_{\mathbf{u}}(\Delta\theta)\,(\mathbf{O}_i - \mathbf{O}_x) \tag{3}$$

$$n_i = \frac{\mathbf{O}_i' - c_i'}{\|\mathbf{O}_i' - c_i'\|} \tag{4}$$

The translation vector ensures that only the visible portion of the dataset is modified by the interaction to maintain a seamless slicing experience. This is illustrated in Figs. 7 (b) and (c).

We invited a neuroscientist collaborator, who is knowledgable in using volume rendering tools and often utilizes cutting planes to observe brain MRI scans, to comment on the utility of this interaction mechanism. Fig. 7 (d) shows the brain MRI volume presented to the scientist, aligned in the transverse direction and positioned at the center. Fig. 7 (e) shows an example of a cutting plane rotated at 45°. After spending some time exploring the volume by translating it and using the FlexiCAVE cutting plane metaphor, they appreciated the physical-to-virtual mapping of the cutting planes. Unlike the Submerse participants in Use Case III, the scientist felt that they grasped the interaction early on. They commented that the radially rotating cutting plane has the potential to allow scientists to explore internal volumes in a new way. Rotating the cutting planes is largely uncommon due to the technical overhead of effectively manipulating the plane position and orientation. Consequently, scientists commonly rely on the three orthogonal cutting planes (sagittal, transverse, and coronal).

A scenario highlighted as most beneficial was to center the volume and bend both sides to compare the internal structures. This was particularly useful because not only was it a seamless experience to orient the cutting plane, compared to a desktop application, but they were also able to clearly visualize the volume by standing in between

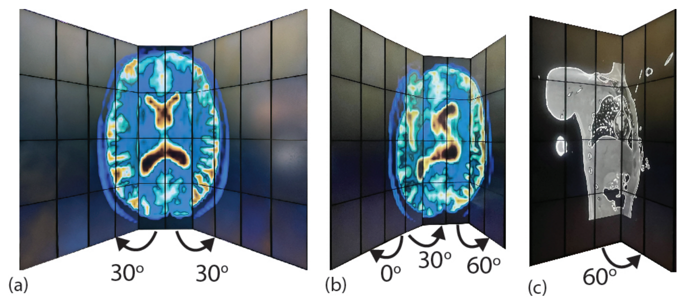

Fig. 8: (a) A "U"-shaped layout resulting in a center (flat) cutting plane and two planes rotated at 30°. (b) An example where the column-pairs from the center are rotated at 30° and 60°, respectively. (c) Volume rendering of a full-body CT data, showing a limitation where rotating the plane > 60° exceeds the volume bounds.

the *U*-shaped FlexiCAVE, which would otherwise be impossible on the desktop given the 2D projection on a single screen. This layout is shown in Fig. 8 (a).

Finally, the scientist expressed that they see the utility of rotating one entire side as a large cutting plane. However, utilizing both pivots on a single side, as shown in Fig. 8 (b), while it produces interesting visual angles, was challenging to comprehend. This was also partially attributed to the fact that such a visual exploration tool is not a common practice. We would also wish to mention one limitation here that for this interaction to work effectively, the volume must be large enough such that it covers the full 90° rotation of the FlexiCAVE side. For example, a full body CT volume rendering will only cover a radial slicing of up to 60°, as can be seen in Fig. 8 (c).

Multivariate Data Visualization. In their work, ImAxes [13], Cordeil et al. have introduced an immersive system for multivariate data visualization where the type of visualization depends on the proximity and relative orientation of the axes with respect to one another. Similar to designing interactions for LHRDs, ImAxes is designed as a modeless visualisation system where GUIs are not required and user actions influence direct manipulation. We take motivation from this work and apply one of the embodied spatial mapping rules for PIVoT.

An ImAxes workspace involves a set of axes, and the arrangement of the axes in 3D space results in constructing an InfoVis visualization. Briefly, relevant to our design space, a lone axis in ImAxes represents a Histogram of a data attribute, combining two axes creates a 2D scatterplot, and positioning a series of axes near each other produces a parallel coordinates plot (PCP). We translate these rules as follows. Since we cannot *detach* the display columns of the FlexiCAVE, our interaction starts directly with 2D scatterplots. Each vertical axis (y-axis) separating the FlexiCAVE display column-pairs is assigned an

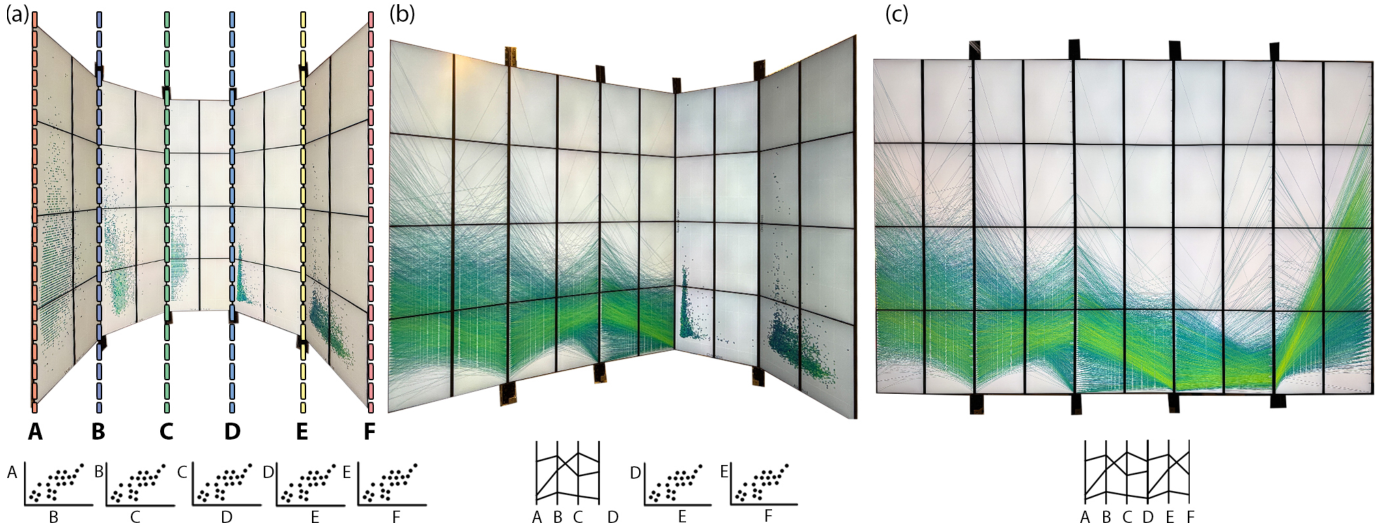

Fig. 9: PIVoT for dynamically updating InfoVis plots. (a) Each column-opair vertical axis is assigned an attribute (A to F). When column-pairs are $> 15°$, a 2D scatterplot of the neighboring attributes is shown. (b) By *flattening* neighboring column-pairs, a PCP of the respective attributes is generated. (c) A fully flat FlexiCAVE showing a PCP for all attributes. This figure shows plots for the Wine Quality dataset [14] with attributes fixed acidity, volatile acidity, citric acid, residual sugar, chlorides, and sulfur dioxide, from A to F, respectively.

attribute, and the attribute of the horizontal axis (x-axis) is assigned based on the neighboring y-axis attribute. Therefore, in the *rest* position, each column-par starts with a 2D scatterplot. We define the *rest* as the FlexiCAVE layout that initializes the visualization application. In our implementation, we specify this as the curved layout, where the angles between two column-pairs should be $> 15°$. This is shown in Fig. 9 (a), replicating the ImAxes use case for visualizing the Wine Quality dataset [14]. All column-pairs in the figure are curved, and scatterplots are visualized based on the neighboring attributes. For instance, the vertical column of the first column-pair is assigned the *fixed acidity* attribute, and the following column is assigned the *volatile acidity* attribute. Therefore, in the curved layout, the first column-pair shows a 2D scatterplot for (*fixed acidity*, *volatile acidity*). In our design, users can assign attributes of a dataset to each vertical axis using a web interface. Since the FlexiCAVE has 6 vertical axes, we are limited to visualizing 6 attributes at once. At any point, users can move attributes across columns using either the web interface or the controller.

Next, we translate the action of generating a PCP as *flattening* the desired column-pairs (angle between column-pairs to be $< 15°$). That is to say, with reference to Fig. 9 (b) as an example layout, with the third vertical axis assigned the *citric acidity* attribute, straightening the first two column-pairs will result in a PCP for (*fixed acidity*, *volatile acidity*, *citric acidity*), respectively. Moreover, since the remaining two column-pairs are still in *rest*, they continue displaying the scatterplots. Finally, as shown in Fig. 9 (c), if all the column-pairs are flattened, the FlexiCAVE will show a PCP for all the assigned attributes.

We invited three InfoVis researchers, experts in designing dimensionality reduction techniques, to give feedback on this interaction design. They commented that while utilizing LHRDs for InfoVis is not novel, the ability to fold the facility "around you" while analyzing data is conceptually interesting. However, they did discern that, other than the degree of wrapping the data around the user for efficient visual exploration, the utility of dynamic layout arrangement in InfoVis is not very clear. Specifically regarding the PIVoT-based interaction, they similarly expressed that switching between scatterplot and PCP can be triggered by a controller button, and therefore, it is uncertain if the physical load of rotating the screens to switch between the visualizations is an effective interaction mechanism. In conclusion, they expressed that LHRDs have an advantage over VR head-mounted displays (HMDs), especially since users have better spatial awareness, they support natural collaboration, and the large displays offer higher resolution. Therefore, while translating the interaction design from HMD to LHRD is conceptually intriguing, further studies are required to explore effective PIVoT designs.

Use Case IV: Replicating the Evaluation of Small Multiples Data Visualization

Liu et al., [30] have conducted a study to evaluate the effect of curvature for visualizing a "shelves" metaphor for the layout of small multiples, comparing three layout configurations: flat, half-circle, and full circle. While their user studies concluded no significant differences between the different layouts with regard to time and accuracy, they reported that participants preferred a half-circle layout in carrying out the tasks provided. This was further elucidated that *it was a good compromise between walking and rotation, and that it allows for an overview at a glance by taking a step back.* We have replicated this study further to examine user accuracy, preference, and patterns, given that the users have an additional option to configure any layout supported by the FlexiCAVE. In describing the study, we call this layout *flex*.

In this study, participants were asked to identify the maximum value and trends from 40 multiples (expressed as an extreme design case in the original study [30]). The layouts for our study were flat, half-circle, and flex. The plot multiples were 2D bar plots (the original study had 3D bar plots) of the world indicator dataset. Each multiple represented the value of 5 indicators for 5 countries for a specific year. We used a within-subjects design, which consisted of the 3 layouts for the task, and there were 3 repetitions for each combination, yielding 102 trials for 12 participants. As with the original study, we ensured that the multiples for performing the tasks were at 7 or 8 Manhattan distance in all layouts. In our replication of this study, we do not restrict the user to a completion time, as the temporal demand could influence the user's effort in reconfiguring the displays during visual exploration.

To get acquainted with the FlexiCAVE and its layout configuration process, we first gave the participants an overview of the facility and provided them with a 3D scene navigation task that encouraged them to explore various layouts. Before each flex task, the FlexiCAVE was reverted to a flat layout. We recruited 12 participants (8 male, 3 female; mean age=26.1 with SD=1.3). All participants were knowledgeable in InfoVis. To quantify our findings, we tracked the users' head to measure movement trajectory and recorded their responses, along with their layout order of preference.

Results for the layout preference, movement distribution, and variation in the display column-pairs' angles (compared to half-circle) are presented in Fig. 10. Participants were fairly accurate in all layouts, with only the flat layout having two incorrect answers. Overall, as shown in Fig. 10 (a), flex was most preferred, with no one ranking it as least preferred. While 2 participants preferred flat first, in contrast to 1, who preferred half-circle first, overall, half-circle was the second preferred layout.

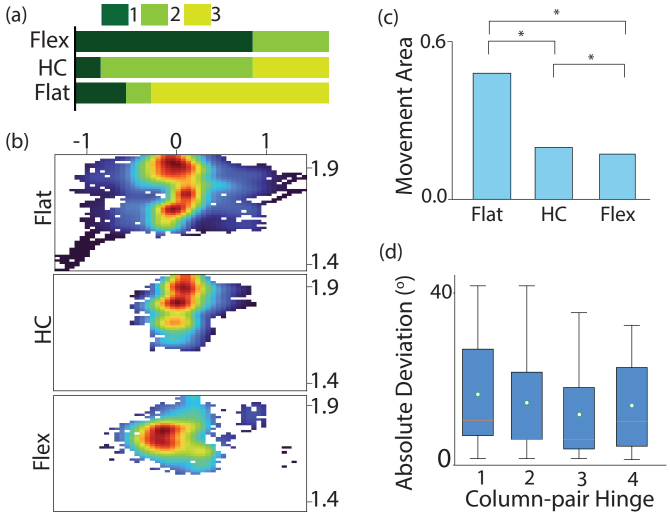

Fig. 10: Results from the replicated user-study [30], comparing the Flat, Half-Circle (HC), and dynamic FlexiCAVE (Flex) layouts. (a) Preference ranking for the three layouts. (b) Heatmap of the participants' movements. The horizontal (parallel to the flat FlexiCAVE) and vertical axes are distances from the FlexiCAVE in meters. (c) Normalized movement area. (d) Angle difference between [30]'s half-circle and the custom angles set by the participants in the Flex, per column-pair. The orange bars are the median and the white dots are the mean angle differences in degrees.

Fig. 10 (b) shows a heatmap of the participants' movements while performing the tasks, calculated by discretizing the total area covered by all participants across all tasks into bins of 0.01m and measuring the frequency of the participants' head-tracked data for each bin. Fig. 10 (c) is a plot showing the normalized movement area coverage for the three layouts. The participants covered significantly less area when performing the task in flex, compared to both half-circle and flat. They also walked significantly more with flat than half-circle, corroborating the original study.

Translating the half-circle arc length of 2m from the original study resulted in a 45° angle between each column-pair in the FlexiCAVE. Fig. 10 (d) shows the variation in angles for each FlexiCAVE column-pair, based on the participants' configuration in the flex layout, with respect to the half-circle. It can be observed that the participants preferred varying curvatures than the half-circle curvature examined in the original study. We noted that 8 participants configured curved layouts, with varying curvature, 3 participants configured a combination of planar and curved, and 1 participant configured the flex similar to the half-circle. It is important to note that users were only allowed to adjust the layout to their preference before starting their task.

While the FlexiCAVE did not significantly improve accuracy, nor did it improve performance efficiency for visualizing a large number of InfoVis multiples, it can be concluded from this study that users have their individual preferences in adjusting the curvature or the overall layout for visual exploration. Thus, as highlighted by previous studies, there is no definitive *one size fits all* layout , however, we believe that FlexiCAVE can prove to be the next generation LHRD design that can support diverse curvature design spaces for diverse applications.

5 Lessions Learned and Future Directions

During the design, construction, and application development processes for the FlexiCAVE, several valuable insights emerged, particularly in two main thematic areas: the future of LHRDs and the design space for flexible or dynamically reconfigurable LHRDs.

Future of LHRDs Our goal in building FlexiCAVE was to envision the next generation of LHRDs beyond existing designs primarily focusing on display form factor and resolution scale. This motivated our design and construction of a flexible facility, marking the first step towards dynamically reconfigurable LHRDs. An initial practical challenge we noticed in this endeavor was the physical manipulation of the FlexiCAVE columns. We realized that while users were receptive to dynamically configuring the displays as a method for a tangible form of interaction, the physical load required to rotate the displays negatively impacted their interaction over time. Moreover, this action of rotating and fine-adjusting the display columns momentarily detached the users from the visualization task. To this end, as a future update, we are planning to introduce a mechanism for the motorized rotation of the FlexiCAVE columns. Introducing this automation not only addresses ease-of-use but can also lead to new avenues for interaction and layout management. Particularly, methods can be explored for displays to automatically adjust their orientation based on application requirements, scene context, number of collaborators, and user ergonomics, which could significantly enhance usability and versatility.

We acknowledge that advancements in HMD technology is facilitating novel methods in visualization and embodied interaction. Nevertheless, they still fall short in their ability to efficiently process and render large, dense, and complex datasets where pixel-level precision is essential. Moreover, collaboration in HMDs remains an ongoing area of research. A next generation of dynamically reconfigurable LHRDs can emerge as promising visual modalities that balance high-fidelity visualization and co-located collaboration.

In terms of form factor, we envision future designs that extend to flexible vertical curvature and outward bending capabilities, further pushing the boundaries of current LHRD technology. This will be more achievable as LED display technologies become more affordable.

Design Space for Interactive Reconfigurability of LHRDs The feedback received regarding our PIVoT prototypes highlighted an interesting conceptual design, however, requiring , we feel that comprehensive studies and techniques are needed for space for flexible and dynamically reconfigurable displays. Specifically, we observed that while dynamically updating camera views for 3D visualization provided intuitive benefits, their significant potential remains unexplored for InfoVis layouts.

6 Conclusion

Acknowledgments

The authors wish to thank A, B, and C. This research was supported in part by NSF award IIS2107224 and ONR award N000142312124.

References

[1] C. Andrews, A. Endert, B. Yost, and C. North. Information visualization on large, high-resolution displays: Issues, challenges, and opportunities. *Information Visualization*, 10(4):341–355, 2011. 5

[2] C. Andrews and C. North. The impact of physical navigation on spatial organization for sensemaking. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2207–2216, 2013. 1

[3] T. Babic, H. Reiterer, and M. Haller. Understanding and creating spatial interactions with distant displays enabled by unmodified off-the-shelf smartphones. *Multimodal Technologies and Interaction*, 6(10):94, 2022. 4

[4] R. Ball and C. North. Analysis of user behavior on high-resolution tiled displays. *Human-Computer Interaction*, pp. 350–363, 2005. 1

[5] R. Ball, C. North, and D. A. Bowman. Move to improve: promoting physical navigation to increase user performance with large displays. *SIGCHI Human Factors in Computing Systems*, pp. 191–200, 2007. 1

[6] I. Belkacem, C. Tominski, N. Médoc, S. Knudsen, R. Dachselt, and M. Ghoniem. Interactive visualization on large high-resolution displays: A survey. *Computer Graphics Forum*, p. e15001, 2022. 2

[7] D. Bihani, A. A. Ullah, C.-O. Dufresne-Camaro, W. Delamare, P. Irani, and K. Hasan. Exploring pointer enhancement techniques for target selection on large curved display. *ACM Human-Computer Interaction*, 8(ISS):214–235, 2024. 1, 2

[8] S. Boorboor, D. Gutiérrez-Rosales, A. Shoaib, C. Kalsi, Y. Wang, Y. Cao, X. Gu, and A. E. Kaufman. Silo: Half-gigapixel cylindrical stereoscopic immersive display. *IEEE Conference on Virtual Reality*, 2025. 1, 2, 4, 5

[9] S. Boorboor, S. Jadhav, M. Ananth, D. Talmage, L. Role, and A. Kaufman. Visualization of neuronal structures in wide-field microscopy brain images. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):1018–1028, 2018. 6

[10] S. Boorboor, Y. Kim, P. Hu, J. M. Moses, B. A. Colle, and A. E. Kaufman. Submerse: Visualizing storm surge flooding simulations in immersive

display ecologies. *IEEE Transactions on Visualization and Computer Graphics*, 30(9):6365–6377, 2023. 5, 6

[11] X. Brioso, C. Calderon-Hernandez, J. Irizarry, and D. Paes. Using immersive virtual reality to improve choosing by advantages system for the selection of fall protection measures. *American Society of Civil Engineers*, pp. 146–153, 2019. 3

[12] M. Cavallo, M. Dholakia, M. Havlena, K. Ocheltree, and M. Podlaseck. Dataspace: A reconfigurable hybrid reality environment for collaborative information analysis. In *IEEE Conference on Virtual Reality and 3D User Interfaces*, pp. 145–153, 2019. 2

[13] M. Cordeil, A. Cunningham, T. Dwyer, B. H. Thomas, and K. Marriott. ImAxes: Immersive axes as embodied affordances for interactive multivariate data visualisation. *Proceedings of the ACM Symposium on User Interface Software and Technology*, pp. 71–83, 2017. 7

[14] P. Cortez, A. Cerdeira, F. Almeida, T. Matos, and J. Reis. Modeling wine preferences by data mining from physicochemical properties. *Decision Support Systems*, 47(4):547–553, 2009. 8

[15] C. Cruz-Neira, D. J. Sandin, T. A. DeFanti, R. V. Kenyon, and J. C. Hart. The CAVE: Audio visual experience automatic virtual environment. *Communications of the ACM*, 35(6):64–73, 1992. 1, 2

[16] T. DeFanti, D. Acevedo, R. Ainsworth, M. Brown, S. Cutchin, G. Dawe, K.-U. Doerr, A. Johnson, C. Knox, R. Kooima, et al. The future of the cave. *Open Engineering*, 1(1):16–37, 2011. 2, 3

[17] T. DeFanti, D. Acevedo, R. Ainsworth, M. Brown, S. Cutchin, G. Dawe, K.-U. Doerr, A. Johnson, C. Knox, R. Kooima, F. Kuester, J. Leigh, L. Long, P. Otto, V. Petrovic, K. Ponto, A. Prudhomme, R. Rao, L. Renambot, D. Sandin, J. Schulze, L. Smarr, M. Srinivasan, P. Weber, and G. Wickham. The future of the CAVE. *Open Engineering*, pp. 16–37, 2011. 1, 3

[18] T. A. DeFanti, G. Dawe, D. J. Sandin, J. P. Schulze, P. Otto, J. Girado, F. Kuester, L. Smarr, and R. Rao. The StarCAVE, a third-generation CAVE and virtual reality optiportal. *Future Generation Computer Systems*, 25(2):169–178, 2009. 2

[19] K.-U. Doerr and F. Kuester. CGLX: a scalable, high-performance visualization framework for networked display environments. *IEEE Transactions on Visualization and Computer Graphics*, 17(3):320–332, 2010. 2

[20] S. Eilemann, M. Makhinya, and R. Pajarola. Equalizer: A scalable parallel rendering framework. *IEEE Transactions on Visualization and Computer Graphics*, 15(3):436–452, 2009. 2

[21] S. Eilemann, M. Makhinya, and R. Pajarola. Equalizer: A scalable parallel rendering framework. *IEEE Transactions on Visualization and Computer Graphics*, 15(3):436–452, 2009. 4

[22] A. Febretti, A. Nishimoto, T. Thigpen, J. Talandis, L. Long, J. Pirtle, T. Peterka, A. Verlo, M. Brown, D. Plepys, et al. CAVE2: a hybrid reality environment for immersive simulation and information analysis. *The Engineering Reality of Virtual Reality, SPIE*, 8649:9–20, 2013. 1, 2, 3

[23] C. Hänel, B. Weyers, B. Hentschel, and T. W. Kuhlen. Visual quality adjustment for volume rendering in a head-tracked virtual environment. *IEEE Transactions on Visualization and Computer Graphics*, 22(4):1472–1481, 2016. 6

[24] T. Höllerer, J. Kuchera-Morin, and X. Amatriain. The Allosphere: a large-scale immersive surround-view instrument. *Emerging displays technologies: the future of displays and interacton*, pp. 3–es, 2007. 1

[25] W. Hong, X. Gu, F. Qiu, M. Jin, and A. Kaufman. Conformal virtual colon flattening. *Proceedings of the ACM Symposium on Solid and Physical Modeling*, pp. 85–93, 2006. 5

[26] R. J. Jacob, A. Girouard, L. M. Hirshfield, M. S. Horn, O. Shaer, E. T. Solovey, and J. Zigelbaum. Reality-based interaction: a framework for post-wimp interfaces. *Proceedings of the SIGCHI Conference on Human factors in Computing Systems*, pp. 201–210, 2008. 6

[27] M. Krüger, D. Gilbert, T. W. Kuhlen, and T. Gerrits. Game engines for immersive visualization: using Unreal Engine beyond entertainment. *PRESENCE: Virtual and Augmented Reality*, 33:31–55, 2024. 2

[28] G. Kyung and S. Park. Curved versus flat monitors: Interactive effects of display curvature radius and display size on visual search performance and visual fatigue. *Human Factors*, 63(7):1182–1195, 2021. 1, 2

[29] S. Lauren, A. Christopher, D.-K. Margaret, Y. Beth, and N. Chris. Shaping the display of the future: The effects of display size and curvature on user performance and insights. *Human–Computer Interaction*, 24(1-2):230–272, 2009. doi: 10.1080/07370020902739429 1, 2, 3

[30] J. Liu, A. Prouzeau, B. Ens, and T. Dwyer. Design and evaluation of interactive small multiples data visualisation in immersive spaces. *IEEE Virtual Reality and 3D User Interfaces*, pp. 588–597, 2020. 1, 2, 5, 8, 9

[31] J. Liu, A. Prouzeau, B. Ens, and T. Dwyer. Effects of display layout on spatial memory for immersive environments. *ACM Human-Computer Interaction*, 6(ISS):468–488, 2022. 1, 2

[32] S. Manjrekar, S. Sandilya, D. Bhosale, S. Kanchi, A. Pitkar, and M. Gondhalekar. Cave: an emerging immersive technology–a review. *International Conference on Computer Modelling and Simulation*, pp. 131–136, 2014. 2

[33] T. Marrinan, J. Aurisano, A. Nishimoto, K. Bharadwaj, V. Mateevitsi, L. Renambot, L. Long, A. Johnson, and J. Leigh. Sage2: A new approach for data intensive collaboration using scalable resolution shared displays. *International Conference on Collaborative Computing: Networking, Applications and Worksharing*, pp. 177–186, 2014. 2

[34] E. Mayer, T. Odaker, D. Kolb, S. Müller, and D. Kranzlmüller. LED CAVE-new dimensions for large-scale immersive installations. *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces Workshop*, pp. 515–519, 2024. 2

[35] S. Mirhosseini, I. Gutenko, S. Ojal, J. Marino, and A. Kaufman. Immersive virtual colonoscopy. *IEEE Transactions on Visualization and Computer Graphics*, 25(5):2011–2021, 2019. 5

[36] M. Nancel, J. Wagner, E. Pietriga, O. Chapuis, and W. Mackay. Mid-air pan-and-zoom on wall-sized displays. *ACM SIGCHI Conference on Human Factors in Computing Systems*, pp. 177–186, 2011. 4

[37] OptiTrack. OptiTrack Motive Software. `https://optitrack.com/software/motive/`. Accessed: Dec 2024. 4

[38] OptiTrack. OptiTrack Primex 13. `https://www.optitrack.com/cameras/flex-13/`. Accessed: 2025-03-15. 4

[39] OptiTrack. 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: 2025-03-15. 4

[40] D. Paes, J. Irizarry, and D. Pujoni. 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*, 130:103849, 2021. 3

[41] A. Panagiotidis, S. Frey, and T. Ertl. Exploratory performance analysis and tuning of parallel interactive volume visualization on large displays. *EuroVis Short Papers*, pp. 13–17, 2015. 6

[42] C. Papadopoulos, K. Petkov, A. E. Kaufman, and K. Mueller. The Reality Deck–an immersive gigapixel display. *IEEE Computer Graphics and Applications*, 35(1):33–45, 2015. 1, 2, 3, 6

[43] L. Renambot, A. Rao, R. Singh, B. Jeong, N. Krishnaprasad, V. Vishwanath, V. Chandrasekhar, N. Schwarz, A. Spale, C. Zhang, et al. Sage: the scalable adaptive graphics environment. *Proceedings of WACE*, 9(23):2004–09, 2004. 2

[44] T. A. Sandstrom, C. Henze, and C. Levit. The hyperwall. *Proceedings of the International Conference on Coordinated and Multiple Views in Exploratory Visualization*, pp. 124–133, 2003. 2

[45] L. Shupp, R. Ball, B. Yost, J. Booker, and C. North. Evaluation of viewport size and curvature of large, high-resolution displays. *Graphics Interface*, pp. 123–130, 2006. 2

[46] S. Sultana and M. A. Alam. 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. *Graphics Interface Conference*, pp. 1–11, 2024. 1, 2

[47] R. Tredinnick, B. Boettcher, S. Smith, S. Solovy, and K. Ponto. Uni-CAVE: A unity3d plugin for non-head mounted vr display systems. *IEEE Virtual Reality*, pp. 393–394, 2017. 2

[48] A. A. Ullah, W. Delamare, and K. Hasan. Exploring users' pointing performance on virtual and physical large curved displays. *ACM Symposium on Virtual Reality Software and Technology*, pp. 1–11, 2023. 1, 2

[49] ViewSonic. ViewSonic V3D245 LCD Display. `https://www.viewsonic.com/global/products/lcd/XG2431`. Accessed: 2025-02-20. 3

[50] H. Vive. SteamVR Base Station. `https://www.vive.com/us/accessory/base-station2/`. Accessed: 2025-03-15. 4

[51] Volfoni RF Active 3D System. `http://volfoni.com/en/activhub-rf50/`. Accessed: July 2024. 4

[52] C. Wei, D. Yu, and T. Dingler. Reading on 3d surfaces in virtual environments. *IEEE Conference on Virtual Reality and 3D User Interfaces*, pp. 721–728, 2020. 2

[53] P. R. Woodward, D. H. Porter, M. R. Knox, S. T. Andringa, and A. Stender. A system for interactive volume visualization on the PowerWall. `http:`

//www.lcse.umn.edu/research/powerwall/powerwall.html. Accessed: 2025-01-21. 1

[54] B. Yost and C. North. The perceptual scalability of visualization. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):837–844, 2006. 1

[55] M. Zannoli and M. S. Banks. The perceptual consequences of curved screens. *ACM Transactions on Applied Perception*, 15(1):1–16, 2017. 2