Beyond Links: Exploring Visual Representations of Multi-View Relations in Mixed Reality

Weizhou Luo
weizhou.luo@tu-dresden.de
Interactive Media Lab Dresden, TUD Dresden University of Technology
Dresden, Germany

Rufat Rzayev*
rufat.rzayev@acm.org
Interactive Media Lab Dresden, TUD Dresden University of Technology
Dresden, Germany

Benjamin Russig
benjamin.russig@tu-dresden.de
Computer Graphics and Visualization, TUD Dresden University of Technology
Dresden, Germany

Sivanon Visutarporn
gam.sivanon@gmail.com
Interactive Media Lab Dresden, TUD Dresden University of Technology
Dresden, Germany

Marc Satkowski
msatkowski@acm.org
Fraunhofer Institute for Process Engineering and Packaging IVV
Dresden, Germany

Stefan Gumhold*†
stefan.gumhold@tu-dresden.de
Computer Graphics and Visualization, TUD Dresden University of Technology
Dresden, Germany

Raimund Dachselt*†
dachselt@acm.org
Interactive Media Lab Dresden, TUD Dresden University of Technology
Dresden, Germany

Abstract

This paper investigates associations, explicit representations of relations between multiple views in Mixed Reality (MR). While research on 2D desktop environments offers extensive recommendations for communicating relations between multiple views, MR environments lack such systematic guidance, necessitating adapted solutions that consider their spatial affordances. To address this gap, we systematically explored association techniques in existing research. Building on established 2D multi-view literature and refining insights from prior design principles, we developed a codebook to describe view relations and their representations. Applying it to a corpus of 44 immersive multi-view approaches, we identified recurring design strategies and synthesized them into a design space of visual association techniques adapted for immersive contexts. Based on a lightweight prototyping framework, we validate the utility of the design space through three envisioning scenarios, demonstrating how associations can support exploration, coordination, and sensemaking in MR applications. Our results inform the design of MR multi-view environments.

*Also with Centre for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI) Dresden/Leipzig, Germany
†Also with Centre for Tactile Internet with Human-in-the-Loop (CeTI), TUD Dresden University of Technology

CHI '26, Barcelona, Spain

ACM ISBN 979-8-4007-2278-3/2026/04
https://doi.org/10.1145/3772318.3791398

CCS Concepts

• **Human-centered computing → Mixed / augmented reality**.

Keywords

Spatiality, content organization, Multiple Views, Multi-view relations, Mixed Reality

ACM Reference Format:
Weizhou Luo, Rufat Rzayev, Benjamin Russig, Sivanon Visutarporn, Marc Satkowski, Stefan Gumhold, and Raimund Dachselt. 2026. Beyond Links: Exploring Visual Representations of Multi-View Relations in Mixed Reality. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26), April 13–17, 2026, Barcelona, Spain.* ACM, New York, NY, USA, 19 pages. https://doi.org/10.1145/3772318.3791398

1 Introduction

Multi-view systems are often used in 2D display setups for data visualization and exploration, enabling users to examine data objects through diverse and multiple representations [96, 99]. Approaches such as coordinated multiple views (CMV) [78], small multiples [64], as well as linking and brushing [42] can improve sensemaking processes by explicitly encoding relations between views [87, 93]. These approaches typically rely on different visual encodings, like links, highlights, and enclosures, to convey relations such as overview and detail or focus and context between views [78, 92]. Visually representing these relations can help users manage, organize, and understand multiple views effectively [39].

Mixed Reality (MR) enables virtual content to be embedded within three-dimensional environments, affording new forms of spatial layout, direct interaction with data representations, and embodied navigation [38]. In MR settings, coordinated views can be distributed volumetrically around the user or placed and overlaid within the environment, rather than being confined to a single screen. This opens up a range of challenges related to the number of views, occlusion, and the field of view of current MR devices, as well as opportunities such as using spatial grouping and volumetric anchors to represent relationships between views that are not possible in traditional 2D desktop systems.

Recent research highlights the potential of immersive multi-view environments. For example, Lee et al. [48] demonstrated the use of links and spatial groupings to show relationships between sticky notes as views to address problems in Virtual Reality (VR). Similarly, Zhu et al. [107] showcased use cases for composite visualizations in Augmented Reality (AR) and VR, focusing on the relationships between visualization views. Wen et al. [99] emphasized layouts for representing multi-view relations. Despite these promising examples, little work has systematically explored how the relations between multiple virtual views are actually represented in MR settings. A body of research on multi-view relations for 2D displays [15, 36, 39, 87, 92] highlights how structural choices, such as view layout, visual encodings, and proximity, can support the user's ability to perceive, interpret, and integrate information across multiple views. Such a systematic treatment for multi-view relations in MR is missing. Moreover, the relation among views might depend on factors like user needs, preferences, applicable tasks, and scenarios [49, 107]. Thus, providing a fixed set of techniques might not be suited to all users, necessitating the need for systematic guidance concerning the visual representation of multi-view relations in MR.

In this paper, we systematically investigate the visual representation techniques of relations between multiple views in MR environments. We refer to such a representation as an **association** that conveys a relationship between two or more views in MR. Building on the rich body of work on multiple views and multi-view relations in 2D desktop settings, we first constructed an initial codebook of association techniques (e.g., spatial proximity, linking). The codebook was iteratively refined through author discussions, considering design principles and the authors' prior experience in designing immersive systems. Using the final codebook as an analytical lens, we extracted concrete instances of association techniques in publications about multi-view immersive environments and identified recurring strategies, gaps, and opportunities to better leverage MR for representing multi-view relations. From this evidence base, we distilled a design space for visual association techniques that is specific to immersive multi-view environments. To demonstrate the utility of the design space, we proposed a design workflow for multi-view immersive environments, built a lightweight prototyping tool that allows researchers to configure and experiment with different association techniques, and illustrated three representative scenarios that apply both the workflow and the tool. Altogether, our work makes three contributions:

- We identify and describe seven **association strategies and patterns** that illustrate how view relations are commonly represented in immersive multi-view applications.
- We introduce a **design space** consisting of three components and 27 dimensions that enumerate and describe specific visual association techniques.
- We present a lightweight **prototyping framework** along with three illustrative **envisioning scenarios** that demonstrate how to apply the design space effectively.

We hope our contribution will serve as a practical reference for the Mixed Reality community, facilitating the systematic creation of multi-view experiences and inspiring future research in this direction.

2 Related Work

In this section, we review previous research on multiple views in 2D displays and immersive environments. We also consider the relevant literature on the layout of views, their relationships, and the representation of these relationships.

2.1 Multiple Views

Although the term "view" is widely used, there is no consensus on its definition, especially in MR settings. Based on the literature on 2D desktop settings, Sun et al. [92] stated that the view is usually described using four perspectives: a process-centric perspective (a visual mapping from data), a model-centric perspective (a set of data and specifications for displaying them), a perception-centric perspective (an independent, bounded area where data are displayed), and a task-centric perspective (an analytic scaffold that supports user tasks). They noted a lack of semantic coherence across these perspectives and defined a view as a collection of visual elements organized spatially in a semantically coherent manner to support specific analytical tasks. Furthermore, Wen et al. [99] defined a view in an AR environment as the mapping from data to visual representation with distinct boundaries that serve to separate one view from another and from the physical world.

Multiple views enable viewing one dataset from various perspectives or analyzing multiple datasets through a common viewpoint [77, 99]. The concept of multiple views has been extensively explored in recent decades [68, 79, 87]. Early studies highlighted the benefits of presenting complementary perspectives and advocated for a wider adoption of multi-view interfaces [77]. For instance, Roberts et al. [79] studied the use of the terminology "multiple views" and highlighted its significance in the visualization literature. In another publication, Roberts [78] discussed the state of the art of coordinated and multiple views. As a result, there is a growing body of research that focuses on design guidelines and recommendations for implementing multiple views [10, 75, 96]. For example, Wang Baldonado et al. [96] presented eight guidelines for the design of multiple view systems based on aspects, like diversity, complementarity, and consistency, that can positively impact the utility of multiple views. In addition, Chen et al. [10] analyzed 360 instances of multiple views to provide a recommendation system for their design.

Recently, researchers have explored the extension of multi-view applications into immersive settings. Roberts et al. [80] discussed challenges and opportunities of using multiple views in immersive visualizations and offered ten practical lessons. Among these lessons, they emphasized the importance of displaying alternative views simultaneously, linking views through interaction or visual effects to facilitate easy navigation through the data space, and addressing occlusion challenges in immersive environments. Knudsen and Carpendale [40] highlighted additional challenges related to collaboration and coordination techniques in immersive environments. In a specific example, Spur et al. [89] presented a system featuring multiple coordinated views that displayed tilted layers of geospatial data on a vertical stack in VR. Similarly, Satriadi et al. [86] proposed using multi-view visualizations in immersive environments to explore maps at various scales and depths, demonstrating the potential of multiple views in the immersive context.

2.2 Multiple View Layouts

Research on multiple views has also devoted considerable attention to layout design. In one of such research, Al-maneea and Roberts [1] examined 491 multi-view instances from the visualization literature and identified five layout options adopted by designers: (1) a predetermined scaffold fixed by the developer, suitable for a limited number of views; (2) a data-driven arrangement that positions views based on intrinsic data attributes; (3) a coordination-based grouping that clusters linked views; (4) a screen-size-driven scheme that scales and places panels according to available real estate; and (5) a user-controlled layout that allows users to arrange views interactively. Shaikh et al. [87] introduced a novel approach to the layout design of multi-view systems based on user perception and view content, arguing that the content of views and the relations between views should be key considerations in layout design. The authors proposed six perception-driven layouts, such as the stacked layout, which arranges views on top of each other, and the focused layout, which provides a central view while positioning other views in the periphery. Additionally, they identified three content-aware layouts, including the hierarchical layout, which creates a tree-like structure among views, and the composite layout, which maximizes space utilization by presenting content in vacant areas or integrating views, for example, by juxtaposition, superimposition, overloading, or nesting [36]. Beyond simple layouts, a few complex multi-view systems have been proposed that utilize flexible views. For example, Jigsaw [91] and SightBi [93] can accommodate many flexible views within their display space, allowing views to be repositioned or resized, potentially leading to overlaps. In addition, Vistribute [30] allows for the automatic distribution of views across various heterogeneous devices based on heuristics.

Layout design in 3D space can enhance memory and help reveal relations between different views [87]. For instance, a juxtaposed layout that keeps multiple visual representations close to each other is often used to facilitate comparison or, in general, when they are needed simultaneously for decision-making [29, 60]. Moreover, Liu et al. [53] explored a shelves-inspired layout for small multiples in immersive environments. They found that as the number of multiples increased, users preferred a semi-circular arrangement over flat or fully surrounding layouts. In comparison, Wen et al. [99] developed a prototype that automatically adapts the layout of multiple views in a 3D immersive environment using a cylindrical egocentric reference frame and a force-directed method to optimize spatial relationships and visibility. A user study demonstrated that this layout significantly improved task performance and received positive feedback regarding usability, workload, confidence, and satisfaction. Additionally, Li et al. [49] introduced a method for flexible multi-view visualization layouts that adjust to the user's position, viewing angle, and focus. By optimizing spatial arrangements, this approach ensures a balanced presentation and promotes intuitive, user-friendly comparisons between views. In summary, not only views and their content, but also the relations between them can influence a user's interpretation, insights, and performance, ultimately affecting the decision-making process when using multiple views [87]. Next, we will discuss the relations between views and their visual representations.

2.3 Relations between Views and Their Representations

Visually and interactively representing relationships between views can help individuals manage, organize, and understand connections between multiple views. Building on the notion of a visual element, the graphical unit that represents a data attribute, both Shaikh et al. [88] and Sun et al. [92] distinguish three families of cross-view relations: relations that connect visual elements across separate views, relations that bind a visual element in one view to another view that supplies its context or detail, and higher-level ties that relate whole views to one another. Moreover, Roberts [78] described five canonical view-to-view relations: overview+detail shows an abstracted, aggregated overview with a zoomed-in detail view; focus+context allocates a high resolution focus region while preserving surrounding context in another view; difference views merge two or more visualizations to make differences explicit; master-slave configurations link one view's interaction to drive another; and immersive "world-in-miniature" techniques provide a miniature control of the main view. In addition, Knudsen and Carpendale [39] stated that relational encodings can help users construct a mental overview of the views, compare data between different views, and convey provenance and narrative structure.

Researchers have developed various visual encodings to clarify the relations between different views. These methods include techniques that intentionally intersect view boundaries, align visual encodings across views, or overlay visual marks within views to highlight relationships [39]. These visual encodings not only indicate the existence of a relation but also can convey additional information. For example, Knudsen and Carpendale [39] noted that the thickness of a line representing the relationship between two views can signify the strength of that relation. In addition, relations can also have cardinality, defining the number of views or visual elements in a relation, including one-to-one, one-to-many, and many-to-many relations [67, 92].

Previous work on displaying relations between views in 2D desktop environments has identified three core design choices: linkage, coordination, and proximity [88, 93]. Linkage, also known as connection, makes relations explicit by drawing visible lines between related views or their visual elements (e.g., [13, 42, 93]). However, this design technique can quickly become cluttered and does not scale well for many-to-many relationships, leading to the need for aggregation techniques like edge bundling (e.g., [74, 92]). Coordination, or visual highlighting, relies on interaction-driven updates. For instance, brushing or filtering in one view highlights corresponding visual elements in others, as demonstrated in tools like Jigsaw [91] or Snap-Together [67]. This technique requires continuous user attention, so additional highlighting methods, such as color or outline emphasis along with links or arrows, can be beneficial to guide users' focus on the view relations [16]. Proximity encodes similarity through spatial arrangement. For example, views that are spatially close to one another may be considered related. This technique is particularly common in sensemaking scenarios, such as those that provide a "space to think" [2, 51]. It minimizes visual clutter but requires users to interpret spatial layouts effectively. Additionally, Sun et al. [92] proposed a systematic approach to design considerations for cross-view data relations in 2D visualizations, focusing on three key concepts: adding relationship marks, updating channels of existing visual elements, and enriching the appearance of existing visual elements. Furthermore, they provided design alternatives to visualize these concepts. For example, alternatives for modifying existing visual elements included changing the size, shape, or texture to make them more visually salient.

2.4 Summary

Despite the existence of extensive taxonomies and design guidelines for linking, coordinating, and arranging multiple views in 2D desktop environments, there is a lack of explicit design recommendations or taxonomies specifically tailored for MR. Current research often attempts to adapt 2D paradigms to 3D spaces or focuses on isolated techniques without systematically considering how spatial, perceptual, and interaction constraints in MR impact the design of visual representations of relations between views. This gap highlights the need for structured guidance in representing view relations within MR contexts.

3 Terminology

To establish the terminology that we use in this paper, we refined the definition of view by Sun et al. [92] for immersive contexts. This refinement incorporates insights from the literature on view relations [39], contextualizes them with the semantics of user actions and tasks [92], and integrates findings from research on immersive environments.

- A **view** is a self-contained, spatially bounded region that organizes visual elements to represent data. A view must be a semantic unit and perceptually distinct from others (e.g., via spatial separation or visual framing) and the environment.
- A **view relation** is the semantic or contextual connection between two or more views. Relations can be, for example, content-driven (based on shared content or data subsets, synchronized encodings, or computational dependencies), task-driven (defined by the analytical goals of the task, e.g., comparison, causality), or user-driven (shaped by individual preferences, cognitive styles, accessibility needs, or situational factors).
- An **association** is an explicit representation of view relations, which helps users comprehend relations among views, including their spatial, temporal, and further attributes (e.g., type and strength).

The definition of a view aligns with perception-centric and task-centric perspectives [92] as well as the definition by Wen et al. [99] while emphasizing spatial distinctiveness, a critical factor in immersive environments where views may be distributed across a 3D space. Besides, there might be an intersection between different relation types (e.g., a user's preference for a specific layout might also improve task performance).

4 Methodology

We systematically explore the design of associations in multi-view MR environments by grounding our approach in prior visualization research and immersive system designs (similar to prior approaches [3, 14, 46]). After collecting and screening relevant literature (Sec. 4.1), we constructed a codebook from 2D visualization research and adapted it to immersive contexts. Through iterative literature coding, result validation, and synthesis (Sec. 4.2), we derive actionable strategies (see Sec. 5) and define a design space (see Sec. 6).

4.1 Collecting and Screening

To identify relevant papers for our review, we utilized a snowball sampling technique similar to that employed in previous studies on design spaces (e.g., [14, 52]). Given the extensive research on multi-view visualization in traditional 2D display settings, we first collected papers addressing design considerations, design spaces, and frameworks for multi-view visualization [78, 87], between-view and cross-view relations [13, 39, 42, 63, 88, 92], and composite visualization [36]. From this initial set, we performed forward citation searches and reviewed immersive studies. However, this approach yielded limited discussions explicitly related to multiple-view relations in immersive environments. Consequently, we expanded our search through both backward and forward citation analyses, beginning with relevant immersive works [21, 22, 45, 52, 53, 57, 86] that emphasize spatial organization and layout, closely related to Gestalt principles [41], particularly proximity, to convey semantic relatedness. Additionally, similar to Lee et al. [46], we reviewed authors' long-standing collection of immersive research over the past decade, adding relevant items identified based on their titles and abstracts. We stopped collecting at thematic saturation, when further searches no longer yielded new papers with new insights, similar to Danyluk et al. [14]. This process resulted in an initial set of 98 papers, consisting of 17 in two-dimensional and 81 in immersive environments. Although not exhaustive, we believe this corpus offers a broad and representative sample of the design space.

While interactivity is critical for cross-view coordination, this work focuses on visual and spatial representations (e.g., layout, linking lines) to establish a foundational design space. We acknowledge interactivity as a complementary area for future study. To ensure that our review focused on informing immersive association design, we developed and applied specific inclusion criteria to the immersive corpus iteratively. We reviewed the collected papers and excluded those that did not meet all of the following criteria:

- **Use of augmented, mixed, or virtual reality technology.** The paper needed to incorporate immersive technologies. While work that solely focused on 3D visualization in desktop settings contributed to the development of our codebook, it was excluded from the reviewed corpus.

- **Presentation of at least two virtual views simultaneously.** The discussed system must render two or more virtual views within a head-mounted display or be spatially co-located in an immersive environment. Views displayed on external physical devices (e.g., tablets, smartphones) did not satisfy this criterion.
- **Concrete content with definable inter-view relations.** Relations among views were either explicitly described or could be reliably interpreted from accompanying materials (e.g., task descriptions in follow-up user studies). Papers that were object-agnostic and merely proposed broader spatial layouts of views informed the codebook but were excluded from the reviewed corpus.

Two authors collaboratively screened each collected paper based on its content, resolving any disagreements through discussion until they reached a consensus. We ended with a total of 44 core papers, the majority of which are from venues such as CHI, ISMAR, ISS, and TVCG, containing immersive design instances for subsequent coding.

4.2 Coding, Synthesis, and Validation

Before coding, we discussed and iterated on the preliminary codebook derived from our 2D visualization research among six coauthors, focusing on the affordances and requirements of immersive technologies. Two authors then divided the corpus and coded it individually using a web-based spreadsheet tool, while documenting their insights and interpretations in free-text fields. Once all coding was completed, we conducted a validation step by cross-checking each other's codes to ensure consistency and accuracy in classification. Any disagreements were resolved through discussion. Through this coding process, we identified common strategies and patterns for representing relationships between views in immersive settings.

We subsequently refined and consolidated the final set of codes, decomposing overly broad dimensions and merging redundant or overlapping categories. For instance, we adjusted the category *Temporal* to distinguish between sequential and simultaneous representations, rather than relying on absolute chronology. Additionally, we introduced *Grid* as a specific type of spatial partitioning that we frequently observed in multi-view immersive environments.

Through this bottom-up thematic analysis [7] process, we derived a design space consisting of three components and 27 dimensions. To validate its robustness, we re-applied the finalized coding scheme across the entire corpus of 44 papers to confirm coverage. We report our observed strategies and patterns in Sec. 5 and elaborate on the resulting design space in Sec. 6. The complete list of collected papers and our analyses is included as supplementary material.

5 Observed Association Strategies and Patterns

We present a catalog of seven strategies and patterns distilled from 44 instances of MR multi-view applications in the literature. These strategies serve as a classification scheme for common association designs. Individual cases can be categorized into multiple strategies, highlighting how the challenge of visualizing relations between different views in immersive environments has been typically addressed. Each strategy is accompanied by an illustration adapted from the literature, and yellow annotations indicate how the associations can be conveyed. We provide the complete classification and detailed description in the supplementary material.

5.1 Link

A *Link* (Fig. 1, left; 16 instances) is an association design strategy in which relations between views are represented explicitly through visual connections, such as lines, curves, bands, or arrows. These links can connect related views at different granularities, such as between views or between visual elements of distinct views. In immersive environments, these links may stretch between views or float in space, augmented with cues such as color and textual labels to encode the type or direction of relation. This strategy is particularly useful in contexts **where views are freely positioned** within the environment, needing visual aids for orientation or sequential guidance [22, 34, 105]. It is also beneficial when additional relational layers cannot be encoded by spatial layout alone, especially if that channel is already used for other purposes, such as representing temporal relationships [43, 104] or is intentionally reserved for other reasons, such as visibility [4, 85]. Furthermore, this approach is advantageous when one view has multiple relations to another [48, 102] or when **the number of relations** is too large to effectively convey using other strategies [4, 9, 24, 74, 86, 100, 107]. It provides a scalable and flexible means to represent large quantities of overlapping, multi-layered relations across views.

5.2 Planar and Cylindrical Uniform Grid

A *Planar Grid* (8 instances) organizes views on a regular lattice of intersecting horizontal and vertical lines, partitioning the virtual space into uniform cells. This simple scaffold ensures that each view maintains the same footprint, orientation, and consistent styling. This approach embodies the classic techniques of **small multiples or faceting**, where each view presents data (such as a category, time interval, or scenario) while maintaining consistent scales and design. The resulting uniformity and tight alignment enable immediate comparison of patterns and differences across multiple views, as demonstrated in examples comparing species distributions [61], pedestrian volumes over time in bar charts [52], and the sugar content of various cereal brands in scatterplots [107]. A variation of this strategy forms a *cylindrical grid* (Fig. 1, right; 4 instances) centered on the user to enhance visibility [53, 107]. However, it has been found that the flat grid layout is more efficient than the curved one for a small number of views [53].

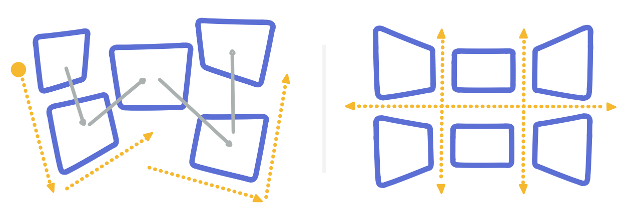

Figure 1: Link; Cylindrical Uniform Grid

5.3 Slideshow with Thumbnails

A *Slideshow with Thumbnails* (Fig. 2, left; 7 instances) strategy expresses a **heterogeneous relationship** between a dominant primary view and a set of smaller thumbnail views or tiles. The primary view serves either as an overview that displays the overall structure of the data or as a focused view that highlights **a particular element of interest**. The thumbnails, which can be arranged in a row, grid, radial fan, or free-form, provide contextual or detailed snapshots that the user can access on demand. For example, a flat map [86] or a globe [95] may feature multiple small maps alongside it. In another example, the primary view may present the conductor's perspective while thumbnails display individual instruments [37].

This strategy can allow efficient task switching between views, where the primary view is active and in focus, while the secondary views are only referenced occasionally [11, 71]. Placing thumbnails close to the user's body can enable quick switching of content in the primary view, similar to DataDancing [52]. Furthermore, a *proxy* approach [52, 70] can enhance this strategy, where thumbnails are exact duplicates of active views in the scene. This method can also be refined into a *Worlds-in-Miniatures* approach with multiples [14], which involves multiple replicas as views at the same scale showing different states of the environment, or it may feature one larger main replica with several smaller replicas displaying subsets of the environment.

5.4 Cluster and Containment

The *Cluster and Containment* (Fig. 2, right; 4 instances) refers to a strategy in which views that **share categorical relations**, such as a semantic meaning, are grouped and represented either via spatial clustering (e.g., proximity [57, 94, 102]) or via nested containment (i.e., areas or shapes that enclose view members [48, 94, 102]). This strategy provides a flexible approach to building associations, commonly used in ideation and sensemaking, where processes are dynamic, iterative, and require constant rebuilding of categorical relations. Containment, such as a planar container [48, 94], can provide strong visual cues as a perceptually salient grouping mechanism to represent categorical structure while organizing member views and facilitating batch operations.

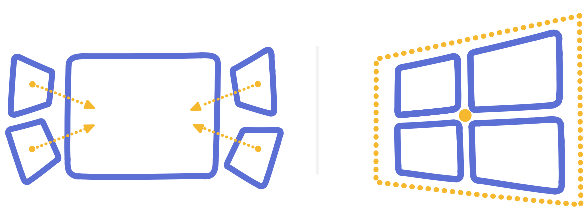

Figure 2: Slideshow; Cluster & Containment

5.5 Stack and Ripple

The *Stack and Ripple* (9 instances) illustrates how the third spatial dimension of MR can be used to encode **hierarchical**, **granular**, or contextual relations among views. Mapping relations along depth provides clear cues that enhance spatial reasoning.

A *Stack* (Fig. 3, left; 6/9 instances) is a pile of views arranged on top of another, where the arrangement is ordered or sequential, representing different layers or facets at the same location or subject. In immersive environments, actual overlap of the views is not always required to imply the relations. Instead, views are often uniform in size, aligned to a common axis, and parallel to each other. Examples include various layers on geographic maps (e.g., traffic, satellite, terrain) [5, 61, 89], zooming or resolution levels [103], different depths in brain imaging [5], or slide elements organized by layer [5], all of which are associated with literal location. Moreover, a stack can represent different facets of the same subject, for instance, a 3D parallel coordinates visualization showing multidimensional user data [9]. This strategy can convey progressive levels of detail and support rapid mental mapping between adjacent layers.

A *Ripple* (Fig. 3, right; 3/9 instances) adopts an egocentric concentric layout, positioning views as rings that radiate outward from the user's current location. The innermost ring typically hosts the primary focus view, while successive outer rings present weaker or less-frequently accessed panels, such as context views [55]. This arrangement can be inverted: a compact thumbnail ring near the user supplies quick navigation, while a larger, distant focus view occupies the outermost ring (e.g., [70]). The radial geometry of the ripple can make the distance between a view and its associated ring an explicit indicator of relational strength, such as views close to the ring implying tight coupling and distant ones suggesting looser ties. Researchers have also discussed the trade-off between prioritizing task accessibility (e.g., views close to the user are easier to interact with) and maintaining clarity of view relations (e.g., views placed near their related views preserve semantic proximity) [99].

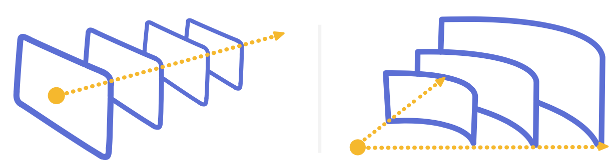

Figure 3: Stack and Ripple

5.6 Temporal Reading Order

Temporal Reading Order (Fig. 4, left; 8 instances) exploits the established reading order bias, people's tendency to scan visual information in the direction dictated by their native language (e.g., left-to-right, top-to-bottom), to make view relations explicit. By arranging views along this axis, designers provide an immediately recognizable "storyline" that guides the user from earlier to later states, from input to output, or from source to derived artifacts. This strategy is beneficial for indicating temporal relations, including **chronology** (events in their order of occurrence) [62, 95, 100], **workflow** (the sequence of steps or processes through user actions) [35, 43], and **provenance** (the documentation of the origin, transformations, and processes a view has undergone) [19, 104].

This strategy is often combined with branching to trace paths from the same origin to different destinations, with temporal order encoded either left-to-right [19, 35, 43, 104] or top-to-bottom [95].

To favor the viewing angle, views can be laid out with the user in the center of the reference frame to form a circular layout [19, 35]. However, reading order can become confusing if spatial views are not properly aligned, potentially leading to misinterpretation of view relations. For example, the ranking of search results can be misinterpreted in a pure grid layout [97].

5.7 Reference and Anchor

Reference and Anchor (Fig. 4, right; 11 instances) is a strategy in which virtual views are **positioned in relation to physical entities** such as the user's body, physical screens, objects, or the surrounding environment. In this strategy, the reference frame becomes a center of gravity, almost like a magnet, attracting views to be arranged around it. One motivation is to improve visibility and interaction accessibility [31, 107], especially when the user is mobile in the immersive environment and the views must remain referenced to them. Another is to align views with physical objects that act as landmarks or structuring elements [4, 57, 85], often when those objects are directly related to the content (e.g., situating explanations or information about the object itself). These two motivations may sometimes compete, leading to a design trade-off between user-centered and object-centered anchoring. Finally, when a physical display is part of the workflow, multiple virtual views are commonly arranged near it [11, 33, 44, 71, 73, 76, 100], leveraging the tangibility of the display to invite interaction and benefiting from the familiarity of users with established screen-based practices.

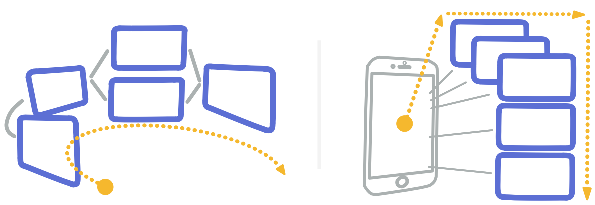

Figure 4: Temporal Reading Order; Reference & Anchor

6 A Design Space For View Association

We aim for a comprehensive yet actionable design space that focuses on key design considerations. We consolidate insights from 2D and immersive multi-view research, Gestalt principles [41], and our experience in immersive environments. Using a bottom-up approach, we scrutinized all codes in the codebook and clustered them by similarity into subgroups and higher-level groups. This resulted in a design space of view association in Mixed Reality (see Fig. 5) with three components: the *view* (Sec. 6.1), the *relation* that connects or distinguishes views (Sec. 6.2), and the *association* strategies making these relations visible in MR (Sec. 6.3). This organization captures the solution space in a logical progression from individual entities to relations between them and to methods for operationalizing these relations, aligning with prior work on between-view visualizations in 2D settings [39] and with visualization design frameworks that move from what (data, datasets, attributes) to how to visualize [66]. For a comprehensive understanding, we provide a structured listing of the design space in Appendix A, detailing each dimension and its values.

6.1 View: "What are the views?"

A view can be described along three fundamental dimensions: content modality, dimensionality, and state. Together, these attributes capture what a view is composed of, how it occupies space, and how it behaves over time.

Content modality refers to the type of information that makes up the view, such as *text, images, video, visualizations, or objects*. Views may be unimodal, relying on a single type, or mixed-modal, incorporating multiple types of media. However, one modality usually acts as the primary carrier of meaning. **Dimensionality** concerns the spatial extent of the view. Some views are purely *2D*, confined to a single plane, while others are in the form of *2.5D*, where flat surfaces are given slight extrusion to suggest depth. *3D* views are volumetric, possessing depth and orientation in immersive space. Finally, the **state** of a view captures its temporal and behavioral dynamics. A *static* view is time-invariant during use. An *interactable* view can be manipulated by the user but does not update automatically. In contrast, a *dynamic* view changes independently, often driven by streaming data, simulation, or narrative progression.

6.2 Relation: "How are the views semantically linked or distinguished from other views?"

The relation component characterizes the semantic connections among views (or their sub-regions), describing how they belong together or differ without specifying any visual treatment. Such relations can be understood through four main attributes: structural, qualitative, positional, and temporal.

Structural attributes. It defines the composition of a relation. Relations may differ in **type**, such as *binary* relations that merely indicate the presence or absence of a connection, *categorical* distinctions that assign discrete but unordered labels, or *ordered* structures that express either ranked categories or quantitative magnitudes. Structural attributes are also shaped by **cardinality**, specifying how many views participate on each side: a simple *one-to-one* coupling, a branching *one-to-many* relation, or a *many-to-many* configuration forming a network. Finally, structural scope varies in **granularity**, as relations may operate *within-view*, extend *between-views*, or span both levels in a *mixed-scope*.

Qualitative attributes. It concerns how relations are valued or weighted. A common distinction is **priority**, where views are arranged along a spectrum of dominance, with some being *primary* or *secondary*, others standing on *equal* footing, and still others serving *supplementary* or *peripheral* roles. Beyond priority, **relevance** captures the extent to which one view depends on another. Some relations are merely *related* by context, while others are *dependent* (enabling or causing one another), or express *belonging* as part–whole connections. Additionally, **strength** reflects the intensity of the connection between views, ranging from *weak* or *loose* to *firm* or *tight*.

Positional & Temporal attributes. It specifies where and how the relations are oriented. **Direction** describes the flow of influence,

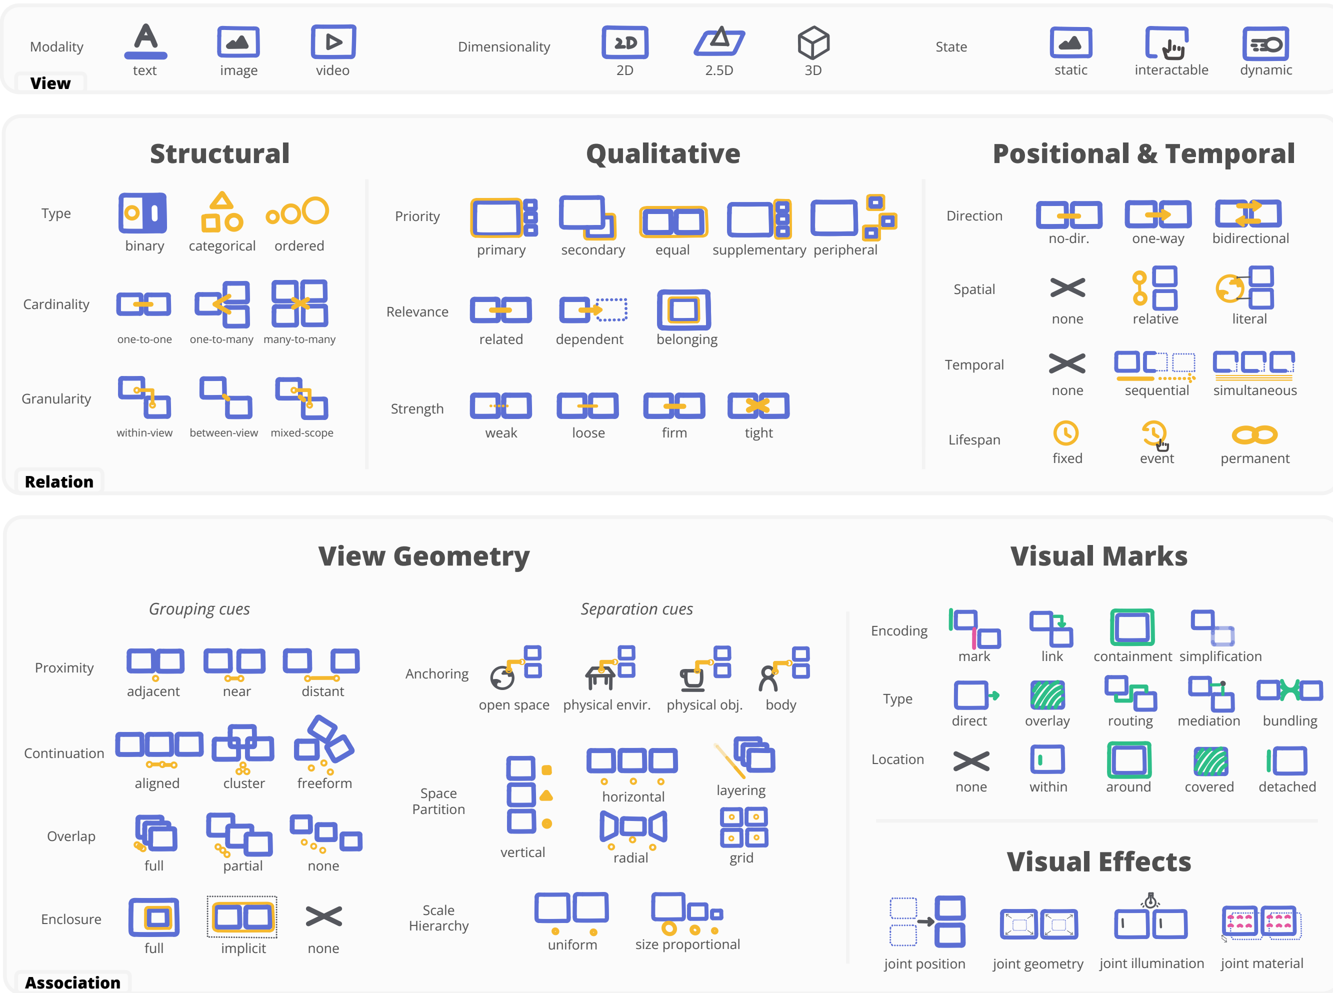

Figure 5: A design space for association techniques between multiple views in Mixed Reality. View component (Top) characterizes what the views are. Relation component (Middle) describes how views are semantically connected or distinguished. Association component (Bottom) specifies how these relations are made visible to users.

indicating whether it is absent in *undirected* connections, asymmetrical in *one-way* couplings, or reciprocal in *bidirectional* relations. **Spatiality** adds another dimension, where relations can be spatially *irrelevant*, *relative* in an abstract topological layout, or *literal*, reflecting concrete geographic or geometric positions. Finally, **temporal** dimension indicates when relations occur. A relation can be temporally *irrelevant*, unfold in *sequential* order, or take place *simultaneously* with others. In addition, relations vary in **lifespan**. Some are *fixed* for a predefined duration, others are *event-based* and appear or disappear in response to triggers, while still others are *permanent*, persisting throughout the entire experience.

6.3 Association: "How are the relations made visible in Mixed Reality?"

The association component specifies how abstract relations are translated into perceivable visual structures in MR. This translation is achieved through three design approaches: the geometric layout of views, the graphical marks that encode connections, and the dynamic effects that guide attention over time. Together, these approaches form a layered vocabulary through which relational semantics become visible, legible, and actionable.

View Geometry. Associations are formed through layout strategies that either foster cohesion within groups (*grouping cues*) or enforce boundaries between them (*segregation cues*). We adapt prior design spaces to this context (e.g., [21, 52, 53, 56, 87]).

Grouping cues reduce ambiguity by applying Gestalt principles [41] such as proximity, similarity, and continuation.

Using **proximity** as a cue, the perceived relations between views shift with their spacing. When views are *adjacent* (touching or separated by only a thin gap), they seem to share the same bounding hull, giving the impression of a single, unified entity. Views that are merely *near* each other, distinct but within an accessible cluster, are perceived to belong to the same local context. By contrast, *distant* views lie outside this perceptual unit, requiring eye or head shifts.

The **continuation** principle can also be applied in MR. *Aligning* views along a common axis or curve suggests a trajectory or conceptual chain, while *clustered* arrangements form compact local units, and *freeform* distributions signal a lack of explicit ordering.

Overlap acts as a visual cue that instantly signals how tightly two views are coupled. The *full overlays* obscure one view almost entirely, often implying substitution or synonymy. The *partial overlays* allow glimpses of the underlying content, hinting at temporal succession or layered alternatives. When views are *fully disjoint*, they occupy separate regions of space, preserving their independence.

Enclosure adds yet another layer of cohesion. A view can act as a *nested* container for other views. Multiple views can *implicitly share the same surface or spatial inclusion* without explicit boundaries. Finally, views can remain completely *separate* with no shared region. Each of these strategies accelerates pre-attentive recognition of clusters and reduces cognitive cost during navigation, as items are perceived as a cohesive unit rather than scattered fragments.

Separation cues complement grouping by highlighting boundaries, keeping clusters distinct, and preventing unintended perceptual merging.

Anchoring is one such strategy that arranges views to be grouped *freely in world coordinates*, attached to *physical environments*, including large static surfaces (e.g., walls, floors, or tables) that serve as prominent landmarks, bound to portable *physical objects*, traveling with those objects as they move, or locked to the user's *body* and arm-reach zone for high reachability and immediate interaction. Each anchoring reference helps to distinguish the views from others.

Space partitions add further granularity by dividing the environment into organizational zones. Views placed along the *vertical axis* (up-down) can suggest categorical contrast, *horizontal* rows (left-right) can stress equality, *layering* (foreground–background) along depth can create hierarchical planes, *radial or cylindrical wraps* place views in arcs or arrays around the user, and *grids* subdivide the space as a plane with regular cells.

Scale hierarchy differentiates groups by size: *uniformly scaled* views signal equal importance, while *size-proportional* views increase salience of primaries and downweight secondaries.

Visual Marks. Once a spatial foundation is established, visual marks provide static encodings that map relational semantics to graphical cues. Building on the visual grammar of 2D elements [92, 98], we define the following dimensions.

Marks can take many forms of **visual encoding**: *colors or glyphs* that encode attributes directly, explicit *links* in the form of lines, ribbons, or tubes that tie views together, and *containment* boundaries, including boxes, hulls, or brackets, that visually demarcate related sets. *Simplification* strategies such as transparency, blur, or reduced saturation lower the visual weight of secondary items, letting more important ones stand out.

Different **encoding types** further shape perception. *Direct* encodings attach to views directly, while *overlays* integrate encodings onto the view's surface. *Routing* detours encodings around to minimize occlusion, whereas *mediating* adds intermediate anchors to help organize links. *Bundling* aggregates multiple links into one, clarifying the number of connections while reducing clutter.

Encoding placement also modulates meaning: encoding may originate *inside* a view, attach *around* its perimeter as halos or outlines, *fully cover* the view as frames, or sit *detached* at an offset. The choice of mark, type, and placement together determines how relations can be further highlighted, traced, and compared.

Visual Effects. Finally, *visual effects* introduce time-varying dynamics that give associations a temporal rhythm. By exploiting Gestalt principle of common fate, coordinated changes across related views or marks allow groups to stand out as coherent perceptual units. Despite its potential [82], we observed only one instance (i.e., [102]) in our corpus.

Joint position changes, such as synchronous translations, oscillations, or rotations, signal that items belong together and often imply shared trajectories. **Joint geometry changes** like coordinated growth, shrinkage, or pulsation, provide salient attention cues while avoiding occlusion. They can also encode semantic dimensions such as relational strength (larger, faster pulses for tighter links) or priority (primary views pulse while secondary ones remain steady). **Joint illumination changes** are an option unique to 3D and MR. For example, specific views can be illuminated by spotlights while ambient light is dimmed down, providing a particularly immersive focus+context dynamic suitable for e.g. storytelling. **Joint material property changes** interact with scene illumination to add nuance; for instance, brief blinks for event cues, momentary brightness lifts to highlight clusters, and gradual opacity fades for focus+context that keeps relations visible in the background. In MR, color/brightness generalizes to albedo and emission and extends to other parameters (e.g., roughness, metalness). Subtle or dramatic, transient or sustained, these effects add temporal meaning beyond static encodings.

7 Design Space Demonstration

In this section, we demonstrate how our design space can be utilized by describing three existing examples of associations identified in the reviewed papers (Sec. 7.1), proposing a design workflow (Sec. 7.2), developing an enabling prototype (Sec. 7.3), and applying the design space to three use cases (Sec. 7.4)[1].

7.1 Associations in Related Work

One usage of our design space is descriptive, providing essential vocabulary and a practical scaffold for researchers to characterize the design of view associations, as illustrated in our literature coding. In the following we use *italics* to denote instantiated values of our design dimensions and highlight three cases from the reviewed approaches.

In AR linked panels [43], code snippets (*text*) are represented as *2D* and *interactive* views. Views are related by their code sources *(categorical)* and levels of dependence *(ordered)*, and can be coupled with one or multiple other views *(one-to-one/to-many, between-view)*. All views are of *equal weight*, without additional directional, positional, or temporal relations. Associations are realized through *aligned horizontal* rows and *vertical* columns that encode provenance and depth, respectively. Views are in *uniform size* and *directly* connected with *links* to specify one-to-many relations.

[1]Video recordings of each use case are available in the supplementary material.

In Multiple Coordinated Spaces [61], *scatter plots* of species and habitat *heatmaps* of environmental variables are shown as *2D*, relatively *static* views. Views show different *categories* of data about the same locations, related through *literal position*. Accordingly, their associations *align* and *layer* these views as stacks of *the same size*, *anchoring* them to a *large physical screen* that presents an overview of the data, while *color marks* mutually coordinate the views.

In VR concert [37], live-streaming videos of music group performances are shown as *2D* and *dynamic* views. Views are *categorical* and *prioritized*, with one camera capturing the whole group and several cameras focusing on individuals, while all views share the same *temporal* span. Associations are designed as one *larger* view for the group and several *equal* smaller views for individuals, arranged either around the primary view or around the user in a *circular layout*.

7.2 Design Workflow

Another value of the design space is its generative power, which guides designers in recognizing choices, revealing trade-offs, exploring alternative configurations, and iteratively refining solutions. To support navigation across the 27 dimensions, we draw from Tamara Munzner's nested model for visualization design [65] and propose a four-step workflow (see Fig. 6) guided by four questions.

"What are the views?" The first step is to characterize each view in terms of content modality, dimensionality, and state. For example, text-heavy views can raise legibility concerns that differ from those of image- or video-based views. Higher dimensionality can introduce viewpoint management and parallax considerations. Finally, highly dynamic views (e.g., live streams) require more careful attention management than static or interactive ones. Explicitly defining these attributes helps designers anticipate legibility needs, placement strategies, and attention constraints early.

"What relations do the views have?" The next step is to abstract relations across the structural, qualitative, positional, and temporal dimensions. This abstraction yields an inventory that can be sorted and ranked to determine which relations are primary and which are supporting and can remain implicit, thereby informing subsequent association choices.

"Which techniques make relations perceivable?" Once relations are identified, designers can visualize them via association techniques using three components. Visual geometry structures the overall topology of how views are placed spatially. Marks inscribe static cues that make relations explicit and legible at a glance. Effects animate these ties in time, guiding attention and highlighting transitions.

"How do we reduce clutter while preserving meaning?" The final step involves deploying and iteratively refining the design in the target environment by revisiting design choices to minimize occlusion and clutter. The goal is to balance salience and legibility with the constraints and affordances of the situated physical context.

Although presented linearly for clarity, these stages are typically iterative and interdependent. Outputs at one level become inputs to the next, and early mis-specifications (e.g., in relational abstraction) can cascade, limiting the effectiveness of otherwise careful association design and refinement.

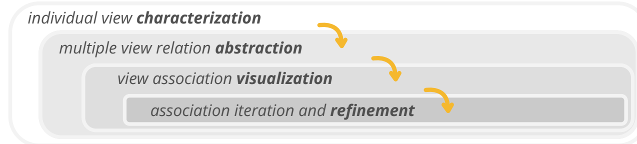

Figure 6: Inspired by [65], we propose a four-stage workflow to operationalize the design space: (1) *characterize* views, (2) *abstract* and define relevant relations, (3) *visualize* them with suitable association techniques, and (4) *refine* the association design in situ.

7.3 Enabling Prototype for Exploring Associations

We implemented a lightweight prototype in Unity for Meta Quest 3 that supports authoring multi-view scenarios with natural hand interaction via the Meta XR Interaction SDK and allows for customizing associations in the Unity Editor.

Authoring starts with templates common in prior work (e.g., *Grid*, *Slideshow*, *Circular* 180°/360°). Each view can be assigned content (images, videos, 3D models) and transformed in space (position, scale, orientation), and views can be added or removed at any time. Users can quickly orient selected views to themselves ("Look At Me"), toggle content previews, lock transforms, and bookmark views for later access. Selection supports both a viewpoint-based lasso and a snap-to-hand technique for fast repositioning. When several views are selected, they can be grouped for joint manipulation, duplicated, saved as a template, or reset to a predefined arrangement, supporting quick reconfiguration and iterative design.

Associations can be further customized in the Unity editor, offering users a selection of five visual marks (see Fig. 7) sampled from our design space: *Line Connectors*: tube-shaped 3D polylines with configurable anchors, curviness (whether the link follows a curved or straight/shortest path), style (solid/dashed, uni-/bidirectional) and radius; *Cube Covers*: outlined bounding boxes that mark a set of views, with margin and outline thickness; *Capsule Covers*: piecewise capsules for sequences or continuity, with configurable margin; *Metaballs*: iso-surfaces of a potential field [6] generated around associated views creating smooth, organically shaped covers, with configurable iso-value and field strength; and *Border Highlights*: thin outlines for lightweight emphasis with minimal occlusion. All encodings expose global color and opacity to tune salience and manage clutter.

7.4 Use Case Illustration

To demonstrate the practical value of our design space, we applied the design workflow to develop three representative MR scenarios: traffic monitoring, storytelling, and sensemaking. For each scenario, we first characterized the views, then abstracted the relations between them, selected appropriate association techniques from the design space, and finally refined them in situ. We use **boldface** and *italics* to denote **association** and *relation* components, respectively.

7.4.1 Traffic Control and Intersection Monitoring. Urban traffic is typically managed from control centers that integrate multiple

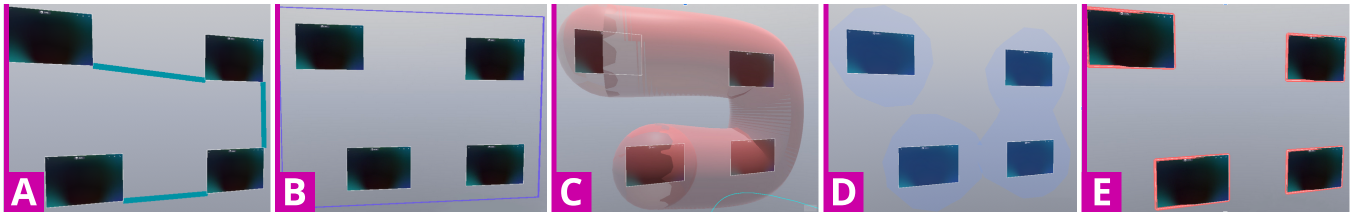

Figure 7: The visual marks for the associations that we prototyped, applied to a set of views. (A) *Line Connectors.* (B) *Cube cover.* As the views are co-planar, the outline reduces to a flat rectangle. (C) *Capsule cover.* (D) *Metaballs.* Three nearby views are covered by an organic shape, while the distant view is decoupled from them. (E) *Border highlights.*

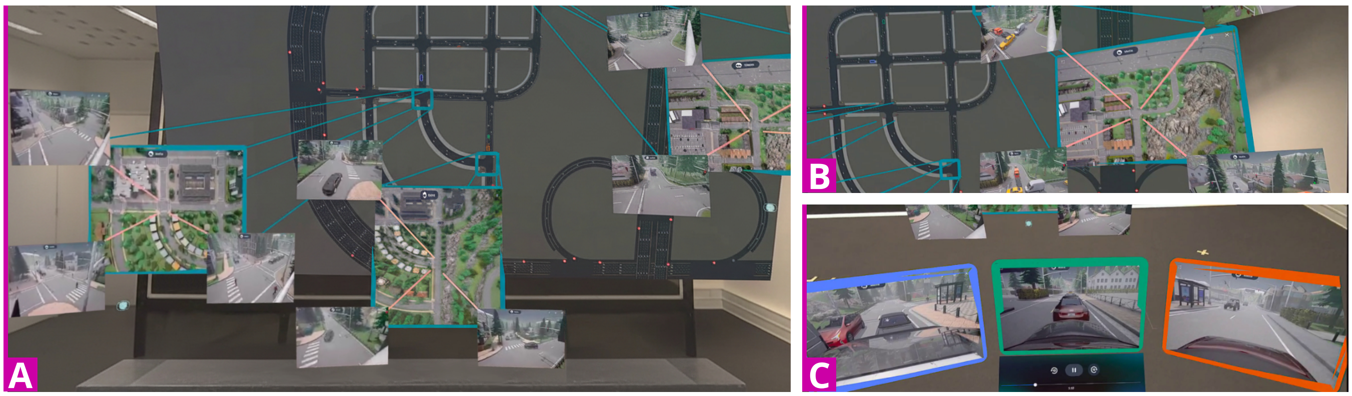

Figure 8: *Traffic Control and Intersection Monitoring*: MR control center as realized in our prototype. (A) Overview of the setup with the big city map in the back, moved up to reduce occlusion by other views. The spatial distribution of the view clusters for each monitored crossing roughly matches their map locations. (B) Solid links between the CCTV views and the intersection orthophoto allow operators to quickly associate specific roads with monitored events. (C) Autonomous car video feeds are anchored to the user's body and are only weakly linked onto the map through their color coding and west-east ordering.

semantically related views [81], including closed-circuit television (*CCTV*) feeds. In the future, autonomous vehicles are expected to operate under human oversight[2]. We envision an MR multi-view control center to support this cognitively demanding and time-sensitive task. To illustrate this idea, we created a 5-minute traffic simulation using CARLA [17], featuring CCTV feeds from major intersections and telemetry video from vehicles (see Fig. 8).

Scenario: Carlos, an MR control center operator, monitors autonomous cars in a suburb of about 1 km^2. The control center combines image, video, and dynamic visualization. A large map provides spatial context (e.g., car positions). In front of it, video feeds include front-facing cameras from autonomous cars and CCTV streams from key crossings, plus close-up orthophotos of those crossings to contextualize the CCTV views. CCTV views are **clustered** around the orthophoto of the crossing they *belong* to, and **partially overlap** when viewed head-on, similar to a slideshow with thumbnails. The orthophotos are **anchored in open space** and their *relative spatial locations* mirror the crossings' distribution on the map. View placements are further refined to form a slight **radial continuation** to improve visibility. Together, these views form a hierarchy with the town map at the root, arranged in **layered space partitions** and connected by **non-directional links** that indicate the actual locations inside the parent view.

Carlos also **anchors** *prioritized* vehicle camera feeds to **his body** near the lower edge of his field of view. They are arranged in an **aligned row**, outlined with **colored frames** that match **colored icons** moving on the map, and ordered dynamically from left to right according to their *geographic position* from west to east.

Reflection: The main design challenge in this use case is that monitoring multiple dynamic, temporally synchronized views is cognitively demanding. We therefore design associations that help users follow and switch between views more smoothly by decomposing positional relations and clustering them into layered spatial partitions. Prioritized vehicle camera feeds are anchored close to the body to maintain constant awareness, and visual marks are only introduced when necessary to avoid clutter.

7.4.2 Criminal Investigation Storytelling. An evidence board is a visual tool used to organize critical evidence, key facts, and relationships, helping piece together complex situations [12]. We adapted The Kelly Murder [90][3], a fictional problem-solving exercise, and reframed it into a storytelling scenario that mirrors how derived insights are communicated to stakeholders in reality. In addition to narrative text fragments, we incorporated 3D models and images to represent entities in the story (see Fig. 9).

Scenario: Bob, a museum enthusiast, visits a seasonal MR exhibition on the Kelly Murder. After entering the reconstructed investigator's office, holographic photographs and documents emerge

[2]As mandated by the EU AI Act: https://artificialintelligenceact.eu/article/14/

[3]Text extracts and connections are inspired by Novakova et al. [69].

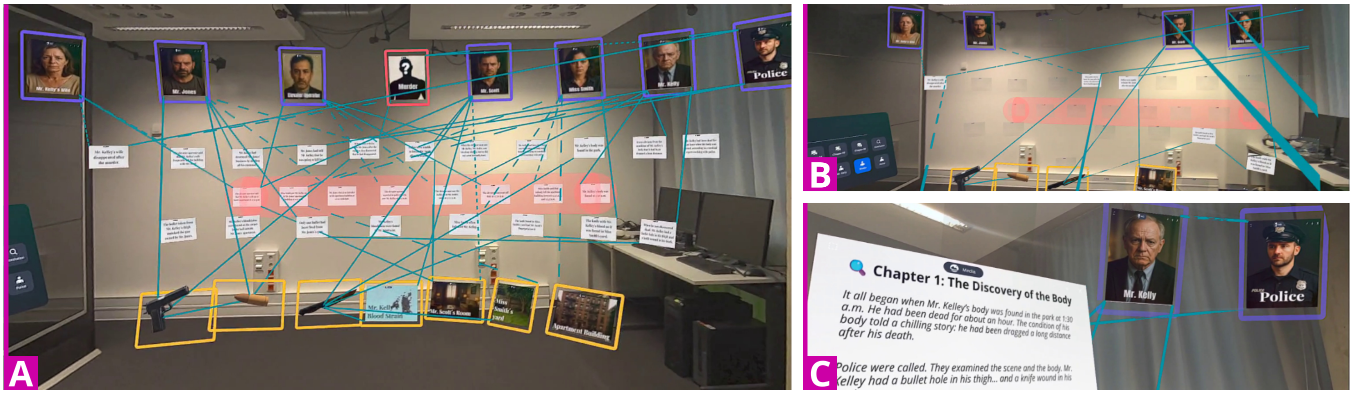

Figure 9: *Storytelling about a Criminal Investigation*: A virtual evidence board with interactive storytelling features realized in our prototype. (A) Overview of the layered structure of *primary* views showing the clues of "Who", "What", "When", and "with What" from top to bottom. Physical clues use 3D models for immersion. (B) Depending on the narration state, some clues and their links are faded out jointly as visual effect. (C) The *secondary* view is attached to the user's arm to guide the narrative.

from the physical evidence board and reorganize into a spatial structure. In the *primary views*, three *categories* of the clues are arranged on **three vertical layers**: top "Who" (people involved), middle "What" (key events and evidence), and bottom "with What" (objects such as weapons rendered as 3D models). The clues are grouped by **proximity**, **aligned along gentle arcs**, and outlined in **color** to support viewing from any angle. A **capsule** covers temporal clues ("When") showing the *sequence of events*. View placements are further refined to form a slight **convex layering** (forward and backward) to improve visibility. *Binary* relations between clues are indicated by **solid uni-directional** links for *causal/functional relations* (e.g., person A fires gun B), and **dotted non-directional links** for *implicit connections* (e.g., shared context).

Bob raises his left arm to reveal a *secondary view* containing overarching narration anchored to **his forearm**. The narration comprises six chapters navigable via "Next" and "Back." As the story advances, associated entities (both clues and links) undergo **material changes** to become fully opaque while others fade to semi-transparent, highlighting *related* clues and guiding attention.

Reflection: The main design challenge is to structure numerous view relations into a coherent story while preserving access to raw sources for alternative interpretations, in line with the goals of evidence boards. We therefore organize associations by categorizing clues into vertical partitions. Visual marks, links, and containments are used extensively, and sometimes redundantly, to keep relations perceptible. Joint visual effects align the temporal progression of views with the narration, mitigating visual clutter and supporting comprehension.

7.4.3 Scientific Literature Sensemaking. Thematic analysis [7] is a widely used qualitative method for identifying, analyzing, and reporting relations and patterns within data. We adapted a published literature survey on designing conversational AI for aging [32] in MR, leveraging the large space for collaboration and sensemaking (see Fig. 10).

Scenario: Lizzy and Martin, UI designers, are developing a smart-TV system with AI capabilities for older adults. They begin by reviewing the conversational-AI literature and, after screening major digital libraries, select 14 papers for thematic analysis. These papers are instantiated as text views (title and abstract) in an MR workspace, arranged in a **2D grid** anchored to a **wall**, with the horizontal axis encoding *publication year (ascending)* and the vertical axis encoding *the category of the underlying AI technology*.

After coding, nine *primary themes* emerge. To further separate them *categorically*, themes are anchored to **room landmarks** in the physical environment and objects such as tables, chairs, and whiteboards, leveraging spatial memory for orientation and collaboration. Each theme has two *secondary sub-themes*; for example, Lack of Trust includes Questioning AI Competence and Credibility Concerns. These *belonging relations* are conveyed by **adjacent proximity** and further refined and reinforced by **clustered continuation**. *Cross-theme relatedness* is expressed through **directed links** (*one-way* or *bi-directional*).

For provenance, the thematic map **links** back to the wall of papers. To reduce visual clutter, links are **bundled** and **mediated**: nine hub knots aggregate connections from papers contributing to each theme. Paths run along the floor before rising to their themes, evoking climbing vines. The resulting spatial map reveals how research is distributed across time, technology, and design themes, informing their system design.

Reflection: The design goal in this scenario is collaborative sensemaking across categorical structures and chains of related views. We therefore design associations so that theme categories are anchored to distributed furniture as landmarks within an immersive thematic map, enabling collaborators to orient themselves and jointly refer to views. Additional relatedness is expressed through links that are bundled and mediated to show overall trends at a glance while preserving provenance to the original papers, with ordered relations still encoded in the wall grid for cross-checking.

8 Discussion and Future Work

Our findings highlight the promise and richness of mixed reality for designing view associations. We reflect on our results, distill implications, and outline directions for future research.

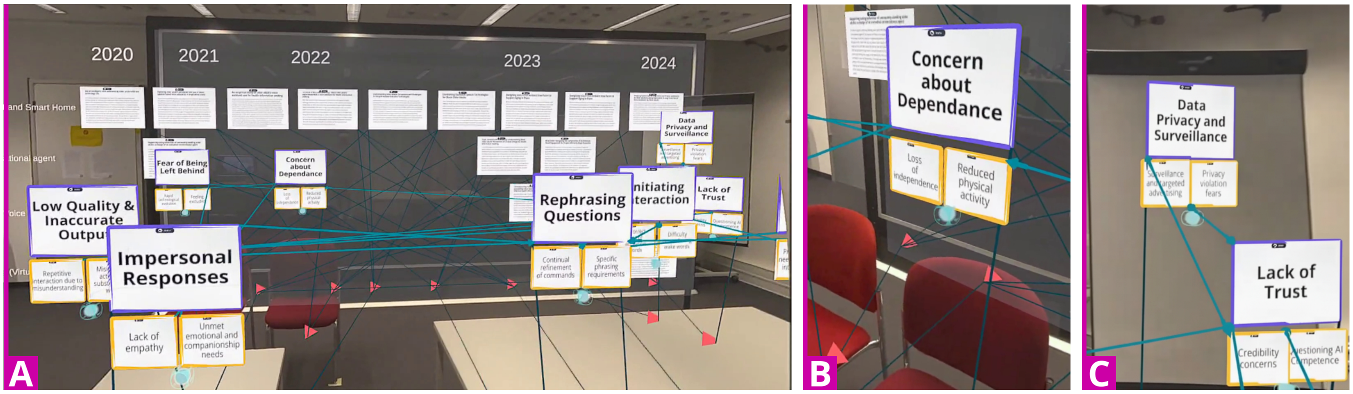

Figure 10: *Sensemaking about Scientific Literature*: A MR thematic map as realized in our prototype. (A) Overview – the relevant papers are presented in a matrix on the wall according to AI technology and publication date, with the identified themes clustered around the room. Links back to the papers are routed through hubs that bundle links according to theme. (B, C) Anchoring topics to various real objects in the room helps navigating the thematic map by leveraging spatial memory.

8.1 Towards Unique Strategies in MR

Compared to traditional 2D displays, where spatial layouts take into account the limited screen real estate (e.g., virtual displays in desktops), **spatiality** in immersive environments becomes a design resource for communicating relations between views. Depth can encode rich hierarchical structure (e.g., by stacking views) and support visibility (e.g., by orienting views in slideshows). Association strategies can be treated as flexible units that can be combined and formalized into nested strategies (e.g., stack plus cylindrical uniform grid [53]). With user mobility and embodied interaction, view groups can be scattered, distributed, or transferred between layouts (e.g., [56]). This also allows multiple strategies to coexist, simultaneously or sequentially, as users navigate with their bodies and interactions, and supporting content-, task-, and user-driven associations on demand. VR can additionally offer immersive "view diving" (e.g., [105]), while rendered environments and objects can also serve as view anchors (e.g., [99]).

In MR, **physicality** extends this space uniquely: views can be anchored to physical environments [4, 22, 52, 57, 85], real objects [4, 22, 57, 61, 85], users' bodies [21, 22, 52, 99], or displays [9, 11, 25, 44, 71, 73, 76, 100]. Aside from forming multi-device environments combining different device affordances (also see Sec. 8.3), this physical anchoring creates rich opportunities for **situatedness**, that is, organizing and associating multiple views to enhance existing workflows that users are familiar with (e.g., [11, 71, 100]), to augment tasks where the real context is integral to the experience (e.g., situated analytics [20, 46, 59, 101]), or to enrich the collaboration between collocated users (e.g., [9, 44, 57]). Our work provides an initial synthesis and set of guidelines, and we encourage researchers to contextualize this knowledge in concrete domains and explore new association patterns and strategies that account for real-world context.

8.2 Design Opportunities of Association in MR

MR provides new ways to express associations. First, beyond space partitioning, the user's body offers new loci for association. For instance, cockpit belts [21], foot-reach regions [52], and forearm/hand-mounted panels [54] can cluster related views. The physical environment also affords opportunistic structure like implicit enclosure via shared surfaces or 3D spatial inclusion [57, 58]. Moreover, advanced graphics channels, such as consistent material or shader families and shared illumination (rim/spotlight), can encode subtle associations, for example, rusted text for importance and unshaped matter for explicitness [108]. The moving viewpoint itself can also become a design resource. For instance, stacking does not require a perfect overlay, while slight offsets remain legible when users "peek" with head motion. View-dependent visibility can also make designated views appear while others gently fade, so conditional opacity marks relatedness. Finally, multimodal cues such as spatial audio (directional beacons, auditory icons/earcons) and, where appropriate, lightweight haptics provide cross-modal guidance. We call for research to explore meaningful affordances (cf. [84]) and establish a broader, MR-native repertoire for expressing associations.

8.3 Associations between Physical and Virtual Displays

In addition to associating fully virtual views in immersive environments, MR enables virtual views to be anchored to digital views on physical displays, as also discussed in cross-device environments [8] and hybrid user interfaces [23, 33]. Virtual views offer high flexibility in positioning, scaling, and presenting high-dimensional content, whereas physical displays often remain preferable for interaction due to their tangible, tactile nature. Leveraging these complementary strengths also requires that view association account for device heterogeneity.

Key considerations include user mobility and the form factors of physical displays. When users are stationary, virtual views can extend static physical displays (e.g., conventional monitors) with additional context or detail [11, 71, 73]. On the move, MR provides an expansive canvas, while small mobile devices act as tangible controllers or situated reference points [76], suggesting dichotomous workflows that are either digital-view-centric or virtual-view-centric, similar to BISHARE [106]. In such hybrid setups, the distributed roles should be clearly communicated. For instance, *priority* relation between views can be indicated through *scale hierarchies* of virtual views relative to physical views, as well as their *relevance* and *strength* can be supported by using *proximity*, *continuation*, *overlap*, and *enclosure*. Another design goal is to guide user attention across heterogeneous views and coordinate transitions between them, where visual marks and visual effects can be used to emphasize currently relevant views.

8.4 Gestalt Principles and Perception in MR

Visual perception in MR differs significantly from 2D due to changing viewpoints, world lighting, and cluttered backgrounds. In this work, we revisit Gestalt principles [41] and operationalize them within our design space: *similarity* can be conveyed via color, shape, or size of views; *common fate* is underused yet promising for ephemeral or on-demand associations, but must be carefully balanced against potential loss of content legibility and increased visual distraction; and *closure* is deliberately avoided as a primary cue, given its viewpoint-dependent nature. Based on the reviewed literature and our prototypes, we also observe provisional heuristics: *figure–ground* (via layered depth) conditions the perception of all other cues (e.g., [5, 55]); *common fate* can temporarily override most static cues (e.g., [102]); explicit *connectedness* often dominates *proximity* and *similarity* (e.g., [85, 105]); and *common region* typically outperforms *proximity* (e.g., [48, 94]). Nevertheless, despite their influence in 2D UI design, these conflict or combination rules and thresholds require confirmation through controlled experiments under realistic MR conditions.

Likewise, the relative strengths and compositional effects of visual marks (e.g., using thickness and color) in MR remain underexplored [18, 28], particularly under illumination shifts, occlusion, and attentional competition from the physical environment [26, 83]. We advocate systematic user studies to compare distinct association techniques and quantify cue precedence and reliability across environmental conditions.

8.5 Interaction For View Associations

Although this paper focuses on associations as visual representations, interaction is integral to making associations usable. Classic brushing and linking remain relevant but must account for visual clutter as well as content outside the field of view. Consequently, out-of-view designs [49, 50] (e.g., halos, trails, directional cues) and other attention-guiding techniques are needed to keep remote views perceptually connected. Moreover, because view relations can be driven by content, task, or user, associations will evolve across workflow phases and perspectives. Designers and end-users need in-situ tools to compose and revise them, including compress/spread layout operations [56], snap-to-depth-band, anchor switching [47, 72], bundle/unbundle [74], and promote/demote priority. Lastly, unlike 2D desktops, immersive settings also allow users to move, enabling new ways of interaction with associations. For instance, proximity-based interaction [27] can disclose relations by revealing related views when users cross distance thresholds. Future research should leverage user mobility as an interaction resource for exploring associations.

9 Conclusion

In this paper, we explored the design of multi-view environments in Mixed Reality by making relations between views explicit through visual associations. Through a systematic review and synthesis of literature, we developed a codebook that characterizes essential association attributes, identified seven recurring strategies and patterns in current practice, and distilled them into a design space comprising three components and 27 dimensions. The design space clarifies how spatial layout, visual encodings, and motion can be combined to convey hierarchy, comparison, provenance, or workflow in immersive contexts. It serves as both an analytical lens for evaluating existing systems and a creative toolkit for developing future MR applications in complex information spaces. We further implemented a lightweight prototyping framework and three envisioning scenarios to illustrate how this design space can guide practical MR applications and highlight the benefits of designs centered on associations. Our contributions bridge established 2D research with emerging immersive approaches, revealing both continuities and gaps in how associations are represented. We provide a design vocabulary and structured guidance that enable researchers and practitioners to analyze relations systematically and design for them intentionally. Ultimately, by highlighting the role of associations in structuring information across multiple views, we lay the foundation for creating MR environments that are more legible, coordinated, and expressive.

Acknowledgments

This work was funded by the German Research Foundation (DFG, Deutsche Forschungsgemeinschaft) as part of Germany's Excellence Strategy – EXC 2050/2 – Project ID 390696704 – Cluster of Excellence "Centre for Tactile Internet with Human-in-the-Loop" (CeTI) of Technische Universität Dresden, DFG grant 389792660 as part of TRR 248 (see https://perspicuous-computing.science), the German Federal Ministry of Research, Technology and Space (BMFTR, SCADS22B) and the Saxon State Ministry for Science, Culture and Tourism (SMWK) by funding the competence center for Big Data and AI "ScaDS.AI Dresden/Leipzig", and by the BMFTR in the program of "Souverän. Digital. Vernetzt.". Joint project 6G-life, project identification numbers: 16KISK001K & 16KIS2413K.

References

[1] Hayder M. Al-maneea and Jonathan C. Roberts. 2019. Towards Quantifying Multiple View Layouts in Visualisation as Seen from Research Publications. In *2019 IEEE Visualization Conference (VIS)*. 121–121. doi:10.1109/VISUAL.2019.8933655
[2] Christopher Andrews, Alex Endert, and Chris North. 2010. Space to think: large high-resolution displays for sensemaking. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Atlanta, Georgia, USA) *(CHI '10)*. Association for Computing Machinery, New York, NY, USA, 55–64. doi:10.1145/1753326.1753336
[3] Ambre Assor, Arnaud Prouzeau, Martin Hachet, and Pierre Dragicevic. 2024. Handling Non-Visible Referents in Situated Visualizations. *IEEE Transactions on Visualization and Computer Graphics* 30, 1 (2024), 1336–1346. doi:10.1109/TVCG.2023.3327361
[4] Mohamed Reda Benkhelifa, Emmanuel Dubois, Louis-Pierre Bergé, and Marcos Serrano. 2025. Augmented urban models: Investigating how data layout and link to the physical referent impact visual search in augmented reality. *International*

Journal of Human-Computer Studies 203 (2025), 103556. doi:10.1016/j.ijhcs.2025.103556

[5] Verena Biener, Daniel Schneider, Travis Gesslein, Alexander Otte, Bastian Kuth, Per Ola Kristensson, Eyal Ofek, Michel Pahud, and Jens Grubert. 2020. Breaking the Screen: Interaction Across Touchscreen Boundaries in Virtual Reality for Mobile Knowledge Workers. *IEEE Transactions on Visualization and Computer Graphics* 26, 12 (2020), 3490–3502. doi:10.1109/TVCG.2020.3023567

[6] James F. Blinn. 1982. A Generalization of Algebraic Surface Drawing. *ACM Transactions on Graphics* 1, 3 (July 1982), 235–256. doi:10.1145/357306.357310

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

[8] Frederik Brudy, Christian Holz, Roman Rädle, Chi-Jui Wu, Steven Houben, Clemens Nylandsted Klokmose, and Nicolai Marquardt. 2019. Cross-Device Taxonomy: Survey, Opportunities and Challenges of Interactions Spanning Across Multiple Devices. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) *(CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–28. doi:10.1145/3290605.3300792

[9] Simon Butscher, Sebastian Hubenschmid, Jens Müller, Johannes Fuchs, and Harald Reiterer. 2018. Clusters, Trends, and Outliers: How Immersive Technologies Can Facilitate the Collaborative Analysis of Multidimensional Data. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) *(CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3173574.3173664

[10] Xi Chen, Wei Zeng, Yanna Lin, Hayder Mahdi AI-maneea, Jonathan Roberts, and Remco Chang. 2021. Composition and Configuration Patterns in Multiple-View Visualizations. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2021), 1514–1524. doi:10.1109/TVCG.2020.3030338

[11] Yi Fei Cheng, Ari Carden, Hyunsung Cho, Catarina G. Fidalgo, Jonathan Wieland, and David Lindlbauer. 2025. Augmented Reality In-the-Wild: Usage Patterns and Experiences of Working with AR Laptops in Real-World Settings. *IEEE Transactions on Visualization and Computer Graphics* 31, 10, 9195–9212. doi:10.1109/TVCG.2025.3592962

[12] Rob Coley. 2017. The case of the speculative detective: Aesthetic truths and the television 'crime board'. *NECSUS – European Journal of Media Studies* Spring 2017 (2017).

[13] Christopher Collins and Sheelagh Carpendale. 2007. VisLink: Revealing Relationships Amongst Visualizations. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1192–1199. doi:10.1109/TVCG.2007.70521

[14] Kurtis Danyluk, Barrett Ens, Bernhard Jenny, and Wesley Willett. 2021. A Design Space Exploration of Worlds in Miniature. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 122, 15 pages. doi:10.1145/3411764.3445098

[15] Dazhen Deng, Weiwei Cui, Xiyu Meng, Mengye Xu, Yu Liao, Haidong Zhang, and Yingcai Wu. 2023. Revisiting the Design Patterns of Composite Visualizations. *IEEE Transactions on Visualization and Computer Graphics* 29, 12 (2023), 5406–5421. doi:10.1109/TVCG.2022.3213565

[16] Nina Doerr, Benjamin Lee, Katarina Baricova, Dieter Schmalstieg, and Michael Sedlmair. 2024. Visual Highlighting for Situated Brushing and Linking. *Computer Graphics Forum* 43, 3 (2024), e15105. doi:10.1111/cgf.15105

[17] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. 2017. CARLA: An Open Urban Driving Simulator. In *Proceedings of the 1st Annual Conference on Robot Learning*. 1–16.

[18] Charles-Olivier Dufresne-Camaro, Yumiko Sakamoto, and Pourang Irani. 2024. Evaluating the effects of colour blending on optical-see-through displays for ubiquitous visualizations. In *Proceedings of the 50th Graphics Interface Conference* (Halifax, NS, Canada) *(GI '24)*. Association for Computing Machinery, New York, NY, USA, Article 9, 13 pages. doi:10.1145/3670947.3670982

[19] Raku Egawa and Takashi Ijiri. 2021. Multi-Window Web Browser with History Tree Visualization for Virtual Reality Environment. In *Adjunct Proceedings of the 34th Annual ACM Symposium on User Interface Software and Technology* (Virtual Event, USA) *(UIST '21 Adjunct)*. Association for Computing Machinery, New York, NY, USA, 32–34. doi:10.1145/3474349.3480221

[20] Neven ElSayed, Bruce Thomas, Kim Marriott, Julia Piantadosi, and Ross Smith. 2015. Situated Analytics. In *2015 Big Data Visual Analytics (BDVA)*. 1–8. doi:10.1109/BDVA.2015.7314302

[21] Barrett Ens, Juan David Hincapié-Ramos, and Pourang Irani. 2014. Ethereal planes: a design framework for 2D information space in 3D mixed reality environments. In *Proceedings of the 2nd ACM Symposium on Spatial User Interaction* (Honolulu, Hawaii, USA) *(SUI '14)*. Association for Computing Machinery, New York, NY, USA, 2–12. doi:10.1145/2659766.2659769

[22] Barrett Ens and Pourang Irani. 2017. Spatial Analytic Interfaces: Spatial User Interfaces for In Situ Visual Analytics. *IEEE Computer Graphics and Applications* 37, 2 (2017), 66–79. doi:10.1109/MCG.2016.38

[23] Steven Feiner and Ari Shamash. 1991. Hybrid user interfaces: breeding virtually bigger interfaces for physically smaller computers. In *Proceedings of the 4th Annual ACM Symposium on User Interface Software and Technology* (Hilton Head, South Carolina, USA) *(UIST '91)*. Association for Computing Machinery, New York, NY, USA, 9–17. doi:10.1145/120782.120783

[24] Stefan P. Feyer, Bruno Pinaud, Stephen Kobourov, Nicolas Brich, Michael Krone, Andreas Kerren, Michael Behrisch, Falk Schreiber, and Karsten Klein. 2024. 2D, 2.5D, or 3D? An Exploratory Study on Multilayer Network Visualisations in Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics* 30, 1 (2024), 469–479. doi:10.1109/TVCG.2023.3327402

[25] Travis Gesslein, Verena Biener, Philipp Gagel, Daniel Schneider, Per Ola Kristensson, Eyal Ofek, Michel Pahud, and Jens Grubert. 2020. Pen-based Interaction with Spreadsheets in Mobile Virtual Reality. In *2020 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. 361–373. doi:10.1109/ISMAR50242.2020.00063

[26] Zeinab Ghaemi, Barrett Ens, Ulrich Engelke, and Bernhard Jenny. 2023. Drawing Connections: Designing Situated Links for Immersive Maps. *Proc. ACM Hum.-Comput. Interact.* 7, MHCI, Article 205 (Sept. 2023), 26 pages. doi:10.1145/3604252

[27] Saul Greenberg, Nicolai Marquardt, Till Ballendat, Rob Diaz-Marino, and Miaosen Wang. 2011. Proxemic interactions: the new ubicomp? *Interactions* 18, 1 (Jan. 2011), 42–50. doi:10.1145/1897239.1897250

[28] Dongyun Han, Anastasia Bezerianos, Petra Isenberg, and Isaac Cho. 2025. Perception of Visual Variables on Virtual Wall-Sized Tiled Displays in Immersive Environments. *IEEE Transactions on Visualization and Computer Graphics* 31, 5 (2025), 3045–3055. doi:10.1109/TVCG.2025.3549190

[29] Jiayi Hong, Rostyslav Hnatyshyn, Ebrar A. D. Santos, Ross Maciejewski, and Tobias Isenberg. 2024. A Survey of Designs for Combined 2D+3D Visual Representations. *IEEE Transactions on Visualization and Computer Graphics* 30, 6 (June 2024), 2888–2902. doi:10.1109/TVCG.2024.3388516

[30] Tom Horak, Andreas Mathisen, Clemens N. Klokmose, Raimund Dachselt, and Niklas Elmqvist. 2019. Vistribute: Distributing Interactive Visualizations in Dynamic Multi-Device Setups. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) *(CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3290605.3300846

[31] Robin Horst, Anika Degreif, Marvin Mathy, and Ralf Dörner. 2019. Virtual Reality Callouts - Demonstrating Knowledge With Spatial-Related Textual Information. In *Computer Graphics and Visual Computing (CGVC)*, Franck P. Vidal, Gary K. L. Tam, and Jonathan C. Roberts (Eds.). The Eurographics Association. doi:10.2312/cgvc.20191252

[32] Yuanhui Huang, Quan Zhou, and Anne Marie Piper. 2025. Designing Conversational AI for Aging: A Systematic Review of Older Adults' Perceptions and Needs. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 181, 20 pages. doi:10.1145/3706598.3713578

[33] Sebastian Hubenschmid, Marc Satkowski, Johannes Zagermann, Julián Méndez, Niklas Elmqvist, Steven Feiner, Tiara Feuchtner, Jens Emil Grønbæk, Benjamin Lee, Dieter Schmalstieg, Raimund Dachselt, and Harald Reiterer. 2025. Hybrid User Interfaces: Past, Present, and Future of Complementary Cross-Device Interaction in Mixed Reality. arXiv:2509.05491 [cs.HC] https://arxiv.org/abs/2509.05491

[34] Brandon Huynh, Abby Wysopal, Vivian Ross, Jason Orlosky, and Tobias Höllerer. 2022. Layerable Apps: Comparing Concurrent and Exclusive Display of Augmented Reality Applications. In *2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. 857–863. doi:10.1109/ISMAR55827.2022.00104

[35] Sungwon In, Eric Krokos, Kirsten Whitley, Chris North, and Yalong Yang. 2024. Evaluating Navigation and Comparison Performance of Computational Notebooks on Desktop and in Virtual Reality. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 606, 15 pages. doi:10.1145/3613904.3642932

[36] Waqas Javed and Niklas Elmqvist. 2012. Exploring the design space of composite visualization. In *2012 IEEE Pacific Visualization Symposium*. 1–8. doi:10.1109/PacificVis.2012.6183556

[37] Minju Kim, Yuhyun Lee, and Jungjin Lee. 2022. Multi-view Layout Design for VR Concert Experience. In *Proceedings of the 30th ACM International Conference on Multimedia* (Lisboa, Portugal) *(MM '22)*. Association for Computing Machinery, New York, NY, USA, 818–826. doi:10.1145/3503161.3548347

[38] Karsten Klein, Michael Sedlmair, and Falk Schreiber. 2022. Immersive analytics: An overview. *IT - Information Technology* 64, 4-5 (2022), 155–168. doi:10.1515/itit-2022-0037

[39] Søren Knudsen and Sheelagh Carpendale. 2016. View Relations: An Exploratory Study on Between-View Meta-Visualizations. In *Proceedings of the 9th Nordic Conference on Human-Computer Interaction* (Gothenburg, Sweden) *(NordiCHI '16)*. Association for Computing Machinery, New York, NY, USA, Article 15, 10 pages. doi:10.1145/2971485.2971566

[40] Søren Knudsen and Sheelagh Carpendale. 2017. Multiple views in immersive analytics. In *IEEE VIS 2017 Workshop on Immersive Analytics*.

[41] Kurt Koffka. 1935. *Principles of Gestalt Psychology* (1 ed.). Routledge. doi:10.4324/9781315009292

[42] Philipp Koytek, Charles Perin, Jo Vermeulen, Elisabeth André, and Sheelagh Carpendale. 2018. MyBrush: Brushing and Linking with Personal Agency.

IEEE Transactions on Visualization and Computer Graphics 24, 1 (2018), 605–615. doi:10.1109/TVCG.2017.2743859

[43] Lucas Kreber and Stephan Diehl. 2023. A Comparative Evaluation of Tabs and Linked Panels for Program Understanding in Augmented Reality. In *2023 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. 29–38. doi:10.1109/ISMAR59233.2023.00017

[44] Ricardo Langner, Marc Satkowski, Wolfgang Büschel, and Raimund Dachselt. 2021. MARVIS: Combining Mobile Devices and Augmented Reality for Visual Data Analysis. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 468, 17 pages. doi:10.1145/3411764.3445593

[45] Benjamin Lee, Xiaoyun Hu, Maxime Cordeil, Arnaud Prouzeau, Bernhard Jenny, and Tim Dwyer. 2021. Shared Surfaces and Spaces: Collaborative Data Visualisation in a Co-located Immersive Environment. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2021), 1171–1181. doi:10.1109/TVCG.2020.3030450

[46] Benjamin Lee, Michael Sedlmair, and Dieter Schmalstieg. 2024. Design Patterns for Situated Visualization in Augmented Reality. *IEEE Transactions on Visualization and Computer Graphics* 30, 1 (2024), 1324–1335. doi:10.1109/TVCG.2023.3327398

[47] Joon Hyub Lee, Sang-Gyun An, Yongkwan Kim, and Seok-Hyung Bae. 2018. Projective Windows: Bringing Windows in Space to the Fingertip. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) *(CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–8. doi:10.1145/3173574.3173792

[48] Joon Hyub Lee, Donghyeok Ma, Haena Cho, and Seok-Hyung Bae. 2021. Post-Post-it: A Spatial Ideation System in VR for Overcoming Limitations of Physical Post-it Notes. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI EA '21)*. Association for Computing Machinery, New York, NY, USA, Article 300, 7 pages. doi:10.1145/3411763.3451786

[49] Jiayin Li, Lixiang Zhao, Hai-Ning Liang, and Lingyun Yu. 2023. ImmerView: Adaptive Multi-View Layout for Immersive Situated Visualizations. In *2023 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*. 108–112. doi:10.1109/ISMAR-Adjunct60411.2023.00030

[50] Tica Lin, Yalong Yang, Johanna Beyer, and Hanspeter Pfister. 2023. Labeling Out-of-View Objects in Immersive Analytics to Support Situated Visual Searching. *IEEE Transactions on Visualization and Computer Graphics* 29, 3 (2023), 1831–1844. doi:10.1109/TVCG.2021.3133511

[51] Lee Lisle, Kylie Davidson, Edward J.K. Gitre, Chris North, and Doug A. Bowman. 2021. Sensemaking Strategies with Immersive Space to Think. In *2021 IEEE Virtual Reality and 3D User Interfaces (VR)*. 529–537. doi:10.1109/VR50410.2021.00077

[52] Jiazhou Liu, Barrett Ens, Arnaud Prouzeau, Jim Smiley, Isobel Kara Nixon, Sarah Goodwin, and Tim Dwyer. 2023. DataDancing: An Exploration of the Design Space For Visualisation View Management for 3D Surfaces and Spaces. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 379, 17 pages. doi:10.1145/3544548.3580827

[53] Jiazhou Liu, Arnaud Prouzeau, Barrett Ens, and Tim Dwyer. 2020. Design and Evaluation of Interactive Small Multiples Data Visualisation in Immersive Spaces. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. 588–597. doi:10.1109/VR46266.2020.00081

[54] Feiyu Lu and Yan Xu. 2022. Exploring Spatial UI Transition Mechanisms with Head-Worn Augmented Reality. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 550, 16 pages. doi:10.1145/3491102.3517723

[55] Yu Lu, Yong Hu, Xukun Shen, and Zhaojia Chen. 2023. An Immersive Layout Framework for Web Design in Virtual Reality. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI EA '23)*. Association for Computing Machinery, New York, NY, USA, Article 36, 7 pages. doi:10.1145/3544549.3585889

[56] Weizhou Luo, Mats Ole Ellenberg, Marc Satkowski, and Raimund Dachselt. 2025. Documents in Your Hands: Exploring Interaction Techniques for Spatial Arrangement of Augmented Reality Documents. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 1218, 22 pages. doi:10.1145/3706598.3713518

[57] Weizhou Luo, Anke Lehmann, Hjalmar Widengren, and Raimund Dachselt. 2022. Where Should We Put It? Layout and Placement Strategies of Documents in Augmented Reality for Collaborative Sensemaking. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 627, 16 pages. doi:10.1145/3491102.3501946

[58] Weizhou Luo, Anke Lehmann, Yushan Yang, and Raimund Dachselt. 2021. Investigating Document Layout and Placement Strategies for Collaborative Sensemaking in Augmented Reality. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI EA '21)*. Association for Computing Machinery, New York, NY, USA, Article 456, 7 pages. doi:10.1145/3411763.3451588

[59] Weizhou Luo, Zhongyuan Yu, Rufat Rzayev, Marc Satkowski, Stefan Gumhold, Matthew McGinity, and Raimund Dachselt. 2023. Pearl: Physical Environment based Augmented Reality Lenses for In-Situ Human Movement Analysis. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 381, 15 pages. doi:10.1145/3544548.3580715

[60] Sehi LYi, Jaemin Jo, and Jinwook Seo. 2021. Comparative Layouts Revisited: Design Space, Guidelines, and Future Directions. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2021), 1525–1535. doi:10.1109/TVCG.2020.3030419

[61] Tahir Mahmood, Erik Butler, Nicholas Davis, Jian Huang, and Aidong Lu. 2018. Building Multiple Coordinated Spaces for Effective Immersive Analytics through Distributed Cognition. In *2018 International Symposium on Big Data Visual and Immersive Analytics (BDVA)*. 1–11. doi:10.1109/BDVA.2018.8533893

[62] Tiago Martins and Teresa Romão. 2025. Insights Into the Usability of a Digital Art Archive in Mixed Reality. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '25)*. Association for Computing Machinery, New York, NY, USA, Article 343, 7 pages. doi:10.1145/3706599.3720092

[63] F. McGee, M. Ghoniem, G. Melançon, B. Otjacques, and B. Pinaud. 2019. The State of the Art in Multilayer Network Visualization. *Computer Graphics Forum* 38, 6 (2019), 125–149. doi:10.1111/cgf.13610

[64] Wouter Meulemans, Jason Dykes, Aidan Slingsby, Cagatay Turkay, and Jo Wood. 2017. Small Multiples with Gaps. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (2017), 381–390. doi:10.1109/TVCG.2016.2598542

[65] Tamara Munzner. 2009. A Nested Model for Visualization Design and Validation. *IEEE Transactions on Visualization and Computer Graphics* 15, 6 (2009), 921–928. doi:10.1109/TVCG.2009.111

[66] Tamara Munzner. 2014. *Visualization Analysis and Design*. A K Peters/CRC Press, Boca Raton, FL, USA.

[67] Chris North and Ben Shneiderman. 2000. Snap-together visualization: a user interface for coordinating visualizations via relational schemata. In *Proceedings of the Working Conference on Advanced Visual Interfaces* (Palermo, Italy) *(AVI '00)*. Association for Computing Machinery, New York, NY, USA, 128–135. doi:10.1145/345513.345282

[68] Chris North and Ben Shneiderman. 2001. Component-based, user-constructed, multiple-view visualization. In *CHI '01 Extended Abstracts on Human Factors in Computing Systems* (Seattle, Washington) *(CHI EA '01)*. Association for Computing Machinery, New York, NY, USA, 201–202. doi:10.1145/634067.634188

[69] Neda Novakova, Veselin Todorov, Steven van den Broek, Tim Dwyer, and Bettina Speckmann. 2025. BlockSets: A Structured Visualization for Sets with Large Elements. arXiv:2512.18055 [cs.CG] https://arxiv.org/abs/2512.18055

[70] Leonardo Pavanatto, Jens Grubert, and Doug A. Bowman. 2025. Spatial Bar: Exploring Window Switching Techniques for Large Virtual Displays. In *2025 IEEE Conference Virtual Reality and 3D User Interfaces (VR)*. 186–194. doi:10.1109/VR59515.2025.00043

[71] Leonardo Pavanatto, Feiyu Lu, Chris North, and Doug A. Bowman. 2025. Multiple Monitors or Single Canvas? Evaluating Window Management and Layout Strategies on Virtual Displays. *IEEE Transactions on Visualization and Computer Graphics* 31, 3 (2025), 1713–1730. doi:10.1109/TVCG.2024.3368930

[72] Siyou Pei, David Kim, Alex Olwal, Yang Zhang, and Ruofei Du. 2024. UI Mobility Control in XR: Switching UI Positionings between Static, Dynamic, and Self Entities. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 611, 12 pages. doi:10.1145/3613904.3642220

[73] Gary Perelman, Emmanuel Dubois, Alice Probst, and Marcos Serrano. 2022. Visual Transitions around Tabletops in Mixed Reality: Study on a Visual Acquisition Task between Vertical Virtual Displays and Horizontal Tabletops. *Proc. ACM Hum.-Comput. Interact.* 6, ISS, Article 585 (Nov. 2022), 20 pages. doi:10.1145/3567738

[74] Arnaud Prouzeau, Antoine Lhuillier, Barrett Ens, Daniel Weiskopf, and Tim Dwyer. 2019. Visual Link Routing in Immersive Visualisations. In *Proceedings of the 2019 ACM International Conference on Interactive Surfaces and Spaces* (Daejeon, Republic of Korea) *(ISS '19)*. Association for Computing Machinery, New York, NY, USA, 241–253. doi:10.1145/3343055.3359709

[75] Zening Qu and Jessica Hullman. 2018. Keeping Multiple Views Consistent: Constraints, Validations, and Exceptions in Visualization Authoring. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (2018), 468–477. doi:10.1109/TVCG.2017.2744198

[76] Carolin Reichherzer, Jack Fraser, Damien Constantine Rompapas, and Mark Billinghurst. 2021. SecondSight: A Framework for Cross-Device Augmented Reality Interfaces. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI EA '21)*. Association for Computing Machinery, New York, NY, USA, Article 234, 6 pages. doi:10.1145/3411763.3451839

[77] J.C. Roberts. 1998. On encouraging multiple views for visualization. In *Proceedings. 1998 IEEE Conference on Information Visualization. An International Conference on Computer Visualization and Graphics*. 8–14. doi:10.1109/IV.1998.694193

[78] Jonathan C. Roberts. 2007. State of the Art: Coordinated & Multiple Views in Exploratory Visualization. In *Fifth International Conference on Coordinated and Multiple Views in Exploratory Visualization (CMV 2007)*. 61–71. doi:10.1109/CMV.2007.20

[79] J. C. Roberts, H. Al-maneea, P. W. S. Butcher, R. Lew, G. Rees, N. Sharma, and A. Frankenberg-Garcia. 2019. Multiple Views: different meanings and collocated words. *Computer Graphics Forum* 38, 3 (2019), 79–93. doi:10.1111/cgf.13673

[80] Jonathan C. Roberts, Peter W. S. Butcher, and Panagiotis D. Ritsos. 2022. One View Is Not Enough: Review of and Encouragement for Multiple and Alternative Representations in 3D and Immersive Visualisation. *Computers* 11, 2 (2022). doi:10.3390/computers11020020

[81] Emanuel Robinson, Daniela Barragan, Doug Dembowski, Todd Szymkowski, Sheryl Miller, Gary Golembiewski, and Sonia Ferezan. 2018. *Human Factors Guidelines for Transportation Management Centers*. Technical Report FHWA-HRT-16-060; HRDS-30/11-18(WEB)E. United States. Federal Highway Administration, McLean, VA. https://rosap.ntl.bts.gov/view/dot/37813 Prepared by Leidos; Westat, Inc.; Lakeside Engineers; Excalibur Associates, Inc. Contributors: Jimmy Chu; Brian H. Philips.

[82] Abdulhaq Adetunji Salako, Hannes Hagen, and Christian Tominski. 2025. Animated Transition between Node-Link and Parallel Coordinates Visualizations. In *Japan Visualization Symposium (JapanVis)* (Tokyo, Japan). doi:10.48550/arXiv.2507.16563

[83] Marc Satkowski and Raimund Dachselt. 2021. Investigating the Impact of Real-World Environments on the Perception of 2D Visualizations in Augmented Reality. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 522, 15 pages. doi:10.1145/3411764.3445330

[84] Marc Satkowski, Weizhou Luo, and Rufat Rzayev. 2025. Of Affordance Opportunism in AR: Its Fallacies and Discussing Ways Forward. doi:10.48550/arXiv.2503.20406 arXiv:2503.20406 [cs]

[85] Kadek Ananta Satriadi, Andrew Cunningham, Bruce H. Thomas, Adam Drogemuller, Antoine Odi, Niki Patel, Cathlyn Aston, and Ross T. Smith. 2022. Augmented Scale Models: Presenting Multivariate Data Around Physical Scale Models in Augmented Reality. In *2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. 54–63. doi:10.1109/ISMAR55827.2022.00019

[86] Kadek Ananta Satriadi, Barrett Ens, Maxime Cordeil, Tobias Czauderna, and Bernhard Jenny. 2020. Maps Around Me: 3D Multiview Layouts in Immersive Spaces. *Proc. ACM Hum.-Comput. Interact.* 4, ISS, Article 201 (Nov. 2020), 20 pages. doi:10.1145/3427329

[87] Abdul Rahman Shaikh, David Koop, Hamed Alhoori, and Maoyuan Sun. 2022. Toward Systematic Design Considerations of Organizing Multiple Views. In *2022 IEEE Visualization and Visual Analytics (VIS)*. IEEE Computer Society, Los Alamitos, CA, USA, 105–109. doi:10.1109/VIS54862.2022.00030

[88] Abdul Rahman Shaikh, Maoyuan Sun, Xingchen Liu, Hamed Alhoori, Jian Zhao, and David Koop. 2025. iTrace: Interactive Tracing of Cross-View Data Relationships. arXiv:2505.23079 [cs.HC] https://arxiv.org/abs/2505.23079

[89] Maxim Spur, Vincent Tourre, Erwan David, Guillaume Moreau, and Patrick Le Callet. 2020. Exploring Multiple and Coordinated Views for Multilayered Geospatial Data in Virtual Reality. *Information* 11, 9 (2020). doi:10.3390/info11090425

[90] Gene Stanford and Barbara Dodds Stanford. 1969. *Learning discussion skills through games*. Citation Press, New York. OCLC: 71972.

[91] John Stasko, Carsten Gorg, Zhicheng Liu, and Kanupriya Singhal. 2007. Jigsaw: Supporting Investigative Analysis through Interactive Visualization. In *2007 IEEE Symposium on Visual Analytics Science and Technology*. 131–138. doi:10.1109/VAST.2007.4389006

[92] Maoyuan Sun, Akhil Namburi, David Koop, Jian Zhao, Tianyi Li, and Haeyong Chung. 2022. Towards Systematic Design Considerations for Visualizing Cross-View Data Relationships. *IEEE Transactions on Visualization and Computer Graphics* 28, 12 (2022), 4741–4756. doi:10.1109/TVCG.2021.3102966

[93] Maoyuan Sun, Abdul Rahman Shaikh, Hamed Alhoori, and Jian Zhao. 2022. SightBi: Exploring Cross-View Data Relationships with Biclusters. *IEEE Transactions on Visualization and Computer Graphics* 28, 1 (2022), 54–64. doi:10.1109/TVCG.2021.3114801

[94] Ibrahim A. Tahmid, Lee Lisle, Kylie Davidson, Chris North, and Doug A. Bowman. 2022. Evaluating the Benefits of Explicit and Semi-Automated Clusters for Immersive Sensemaking. In *2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. 479–488. doi:10.1109/ISMAR55827.2022.00064

[95] Hansheng Wang, Xiaojian Chen, Zeyu Xia, Hailong Wang, Xiaohan Wang, and Richen Liu. 2022. iTDW: Immersive Tiled Display Wall with Clustering-Driven Layout. In *2022 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*. 512–517. doi:10.1109/ISMAR-Adjunct57072.2022.00109

[96] Michelle Q. Wang Baldonado, Allison Woodruff, and Allan Kuchinsky. 2000. Guidelines for using multiple views in information visualization. In *Proceedings of the Working Conference on Advanced Visual Interfaces* (Palermo, Italy) *(AVI '00)*. Association for Computing Machinery, New York, NY, USA, 110–119. doi:10.1145/345513.345271

[97] Austin Ward, Bogeum Choi, and Robert Capra. 2023. Eyes on Immersive Search: Eye-Tracking Study of Search Engine Result Pages in Immersive Virtual Environments. In *Proceedings of the 2023 Conference on Human Information Interaction and Retrieval* (Austin, TX, USA) *(CHIIR '23)*. Association for Computing Machinery, New York, NY, USA, 43–54. doi:10.1145/3576840.3578306

[98] Colin Ware. 2004. *Information Visualization: Perception for Design* (2nd ed.). Morgan Kaufmann Publishers Inc., San Francisco. doi:10.1016/B978-155860819-1/50001-7

[99] Zhen Wen, Wei Zeng, Luoxuan Weng, Yihan Liu, Mingliang Xu, and Wei Chen. 2023. Effects of View Layout on Situated Analytics for Multiple-View Representations in Immersive Visualization. *IEEE Transactions on Visualization and Computer Graphics* 29, 1 (2023), 440–450. doi:10.1109/TVCG.2022.3209475

[100] Jonathan Wieland, Hyunsung Cho, Sebastian Hubenschmid, Akihiro Kiuchi, Harald Reiterer, and David Lindlbauer. 2024. Push2AR: Enhancing Mobile List Interactions Using Augmented Reality. In *2024 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. 671–680. doi:10.1109/ISMAR62088.2024.00082

[101] Wesley Willett, Yvonne Jansen, and Pierre Dragicevic. 2017. Embedded Data Representations. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (2017), 461–470. doi:10.1109/TVCG.2016.2598608

[102] Ying Yang, Tim Dwyer, Zachari Swiecki, Benjamin Lee, Michael Wybrow, Maxime Cordeil, Teresa Wulandari, Bruce H Thomas, and Mark Billinghurst. 2024. Putting Our Minds Together: Iterative Exploration for Collaborative Mind Mapping. In *Proceedings of the Augmented Humans International Conference 2024* (Melbourne, VIC, Australia) *(AHs '24)*. Association for Computing Machinery, New York, NY, USA, 255–258. doi:10.1145/3652920.3653043

[103] Lei Zhang, Mason Prather, Devon Haynes, Feiyu Lu, Ying Chen, and Joy Li. 2025. Saving the Polar Bears: A Collaborative VR Storytelling Experience Enhanced by Multi-scale and Multi-view WiM Designs. In *2025 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*. 1546–1547. doi:10.1109/VRW66409.2025.00418

[104] Yidan Zhang, Barrett Ens, Kadek Ananta Satriadi, Ying Yang, and Sarah Goodwin. 2023. Embodied Provenance for Immersive Sensemaking. *Proc. ACM Hum.-Comput. Interact.* 7, ISS, Article 435 (Nov. 2023), 19 pages. doi:10.1145/3626471

[105] Yifei Zhang, Jingyan Zhang, Ziyue Lin, Lin-Ping Yuan, Jun Han, Yalong Yang, and Siming Chen. 2025. ST2VR: An Interactive Authoring System for SpatioTemporal STorytelling in Virtual Reality with Hierarchical Narrative Structure. In *2025 IEEE 18th Pacific Visualization Conference (PacificVis)*. 224–234. doi:10.1109/PacificVis64226.2025.00028

[106] Fengyuan Zhu and Tovi Grossman. 2020. BISHARE: Exploring Bidirectional Interactions Between Smartphones and Head-Mounted Augmented Reality. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–14. doi:10.1145/3313831.3376233

[107] Qian Zhu, Tao Lu, Shunan Guo, Xiaojuan Ma, and Yalong Yang. 2025. CompositingVis: Exploring Interactions for Creating Composite Visualizations in Immersive Environments. *IEEE Transactions on Visualization and Computer Graphics* 31, 1 (2025), 591–601. doi:10.1109/TVCG.2024.3456210

[108] Siyi Zhu and Joel Chan. 2025. Sense and Sensability: Exploring Future Immersive Environments for Scholarly Sensemaking. In *Proceedings of the 2025 Conference on Creativity and Cognition (C&C '25)*. Association for Computing Machinery, New York, NY, USA, 886–901. doi:10.1145/3698061.3726952

A Appendix

We provide detailed descriptions of our design space (see Fig. 5), including 27 dimensions and their values. The design space is organized as three components: the *view* (see Sec. A.1), the *relation* that connects or distinguishes views (see Sec. A.2), and the *association* strategies making these relations visible in MR (see Sec. A.3).

A.1 View: "What are the views?"

This component comprises three basic dimensions that describe a view: *content modality*, *dimensionality*, and *state*.

Content Modality. The information from which a view is composed, for instance, *text*, *image*, *video*, *visualization*, or *object*. Views can be unimodal or mixed-modal, but one modality normally acts as the primary carrier of meaning.

Dimensionality. The intrinsic spatial extent of the view, defining whether it occupies a plane or a volume.

- *2D*: confined to a single depth plane.
- *2.5D*: fundamentally planar with slight extrusion that suggests depth.
- *3D*: volumetric, with depth and orientation in immersive space.

State. The temporal and behavioral status of the view, describing what can change and who initiates the change.

- *Static*: time invariant during use.
- *Interactable*: can be altered by the user but does not self-update.
- *Dynamic*: changes without user action due to streaming data, simulation, or narrative progression.

A.2 Relation: "How are the views semantically linked or distinguished from other views?"

The relation component characterizes the semantic connections among views (or subregions of views) without prescribing any visual treatment. It covers structural, qualitative, positional, and temporal attributes.

The structural attribute defines "what the relation is made of".

Type. The logical nature of the relation among entities.

- *Binary*: either a connection exists or it does not.
- *Categorical*: a discrete label with no inherent order.
- *Ordered*: either ordinal rank (ordered categories without equal intervals) or quantitative magnitude (numeric values with consistent intervals or ratios).

Cardinality. The number of views that participate on each side of a relation.

- *One-to-one*: couples one view with exactly one other view.
- *One-to-many*: couples one view with multiple views, fanning out from a source.
- *Many-to-many*: involves multiple views on both sides, forming a networked coupling.

Granularity. The scope at which the relation semantics operate.

- *Within-view*: applies to a region or segment inside a single view.
- *Between-view*: exists between whole views treated as single units.
- *Mixed-scope*: spans both subview regions and whole views.

The qualitative attribute defines "how the relation is valued".

Priority. The positional ordering that ranks views along a dominance spectrum.

- *Primary*: central or most important.
- *Secondary*: relevant but lower in importance or presented after the primary.
- *Equal*: views share the same rank.
- *Supplementary*: adds optional details.
- *Peripheral*: of marginal relevance.

Relevance. The degree to which one view relies on another for meaning or function, distinct from causal, functional, and part–whole connections.

- *Related*: a general thematic or contextual connection.
- *Dependent (enables/causes)*: one view requires the other to function or make sense.
- *Belonging (part-of)*: a constituent part of a larger whole.

Strength. A scalar magnitude reflecting the binding intensity of a relation. We define it as *weak, loose, firm, tight*, with increasing degrees of relational intensity.

The positional attribute defines "where and in what orientation the relation sits".

Direction. The vector property of a relation.

- *No-direction*: symmetrical or without a preferred direction.
- *One-way*: flows in a single direction.
- *Bidirectional*: influence travels both ways.

Spatial. The spatial attribute of multiple relations.

- *None*: spatially irrelevant.
- *Relative (abstract)*: typological positioning that can be represented as relative geometry.
- *Literal*: concrete spatial relations that can be represented geographically.

The temporal attribute defines "when and for how long the relation occurs".

Temporal. The chronological arrangement of relations.

- *None*: temporally irrelevant.
- *Sequential*: follows a temporal order where one occurs before or after another.
- *Simultaneous*: occurs at the same time or substantially overlaps.

Lifespan. The persistence status of the duration of a relation.

- *Fixed*: exists for a predetermined duration.
- *Event-based*: appears or disappears in response to an event.
- *Permanent*: persists for the entire experience.

A.3 Association: "How are the relations made visible in Mixed Reality?"

The association component specifies the visual strategies that turn abstract relations into perceivable structure. It comprises three orthogonal approaches: *view geometry* for macro layout, *visual marks* for static graphical encodings, and *visual effects* for time-varying, attention-guiding cues.

A.3.1 View Geometry. View geometry is the layout-level attribute that sets where and how groups of related entities occupy 3D space. It governs where and how entire groups of related entities are positioned in 3D space, determining the gross topology of the scene. It answers the question "where and how groups of related entities occupy 3D space?", providing the spatial layouts onto which finer encodings can later be applied.

Grouping cues are spatial strategies that foster perceptual cohesion, causing entities to be experienced as a single, higher-order unit.

Proximity. The spacing between views; smaller gaps promote grouping.

- *Adjacent*: views touch or are separated by a thin gap and appear to share a bounding hull.
- *Near*: distinct objects within easy reach that read as a local cluster.
- *Distant*: separation large enough to require eye or head movement.

Continuation. The alignment of axes or paths across views.

- *Aligned*: views follow a common axis or curve.
- *Clustered*: views are clustered into a locally cohesive unit.
- *Freeform*: views are placed without an explicit ordering path.

Overlap. The degree of occlusion from the user's viewpoint.

- *Full overlay*: one view strongly occludes another, suggesting synonymy or temporal substitution.
- *Partial*: a view partially occludes another, yet the background view remains partly legible.
- *None*: views are disjoint and fully visible.

Enclosure. Whether views share a containing region.

- *Full closure (nested)*: a view acts as a container that encloses other views.
- *Implicit region (same surface)*: views are on the same surface or background without an explicit boundary.
- *None*: views do not share a perceived place.

Segregation cues are the complementary set of spatial strategies that accentuate between-group boundaries so that multiple clusters remain visually and cognitively distinct even when rendered in a shared 3D volume.

Anchoring. The spatial frame of reference to which a view is attached.

- *Open space*: placing freely in world coordinates.
- *Physical environment*: attached to real surfaces (e.g., walls, tables).
- *Physical object*: bound to portable real objects.
- *Body (on-body, arm-reach)*: follows the user's body or stays within arm reach.

Space Partition. The primary organizational axes that divide space into meaningful zones.

- *Vertical (up-down)*: stacked along the y axis.
- *Horizontal (left-right)*: arranged along the x axis.
- *Layering (foreground-background)*: placed on discrete depths along the z axis.
- *Radial/Cylindrical*: wrapped around the user in an arc or full circle.
- *Grid*: arranged in a regular matrix of spatial cells.

Scale hierarchy. How relative size expresses importance among related views.

- *Uniform*: equal size signals parity and neutral salience.
- *Size proportional*: larger for primaries and smaller for peripherals.

A.3.2 Visual Marks. Visual marks are static mark-level techniques applied to views or subviews to encode, clarify, or simplify perceived relations. They provide additional cues that can be read at a glance or on inspection.

Visual Encoding. The graphical element used to convey association.

- *Mark (color, glyph)*: color or glyph to encode attributes.
- *Link*: an explicit line, ribbon, or tube that joins views.
- *Containment*: brackets, boxes, or hulls that enclose related items.
- *Simplification (transparency, blurriness)*: transparency or blur to reduce visual weight.

Encoding Type. The geometric strategy by which encodings are applied.

- *Direct*: attached straight to the target view.
- *Overlay*: co-located on the view's surface.
- *Routing*: paths detour to avoid occlusion.
- *Mediating*: explicit waypoints structure links.
- *Bundling*: multiple links are aggregated into a shared route.

Encoding Location. Where a visual encoding attaches to a view.

- *None*: no explicit attachment; relation is implied elsewhere.
- *Within*: originates inside the view's bounds.
- *Around*: attaches to a perimeter halo or outline.
- *Fully covered*: overlays the view as a full outline or surface.
- *Detached (offset)*: starts at an offset node or hub.

A.3.3 View Effects. The time-varying, group-synchronous channels that guide attention and express the transient state without changing layout. They can be applied to views and to visual marks.

- **Joint position change (velocity and direction)**: movement with shared velocity or direction across related elements, implementing common fate.
- **Joint size change (growth and compression)**: synchronous growth or compression as an attention cue; can encode strength or priority.
- **Joint illumination change (blink, dim, and brighten)**: brightening, dimming, or selective lighting to focus attention while maintaining context.
- **Joint material property change**: brief modulation of albedo, emission, roughness, metalness, or opacity to highlight or background content.