BridgeVis: Linking Quantitative Analysis and Qualitative Experience for Data Visualization in Virtual Reality

Yuxuan Guo[a], Haowei Xiong[b], Wai Tong[c,*]

[a] *Shanghai Jian Qiao University, Shanghai, China*
[b] *University of the Arts London, London, United Kingdom*
[c] *Texas A&M University, TX, United States of America*

Abstract

Virtual Reality (VR) can enhance qualitative understanding of data by making measures and scales experiential, but often sacrifices quantitative precision. In contrast, traditional visualizations such as charts excel at quantitative analysis but fail to convey natural experiences of data. We introduce BridgeVis, an approach that bridges this gap by integrating interactive 2D charts into 3D VR environments, combining the strengths of both representations. We define a design space comprising three dimensions to guide the creation of immersive data experiences: interaction modes, linking mechanisms, and mapping directions. Building on this design space, we developed four prototypes spanning diverse data types and narrative contexts. Findings from a user study (N=12) indicate that BridgeVis effectively supports both quantitative analysis and qualitative experience, while not imposing additional cognitive load. This work contributes a validated design framework, empirical validation, and design implications for creating hybrid immersive tools that bridge precision and immersion.

Keywords: Virtual Reality, Data Visualization

1. Introduction

A widely adopted approach to enhancing public understanding of data is visualization. By converting abstract data into graphical elements, visualization helps audiences grasp complex information, discern relationships, and uncover potential associations [1]. For example, data-driven storytelling [2], as a common approach to enhancing users' understanding of data, has been extensively applied in fields such as journalism [3], business [4], and education [5]. In recent years, the forms in which the public receives information have diversified alongside advances in technology. For instance, the growing accessibility of virtual reality(VR) has provided users with an immersive way of experiencing virtual content, thereby extending the boundaries of data visualization. Building on this characteristic, Lee et al. proposed the concept of *Data Visceralization* [6]. Compared with conventional data storytelling approaches, this method enables users to experience reconstructed information within a VR environment, thereby enhancing their comprehension of measurements and units in a more engaging and immersive manner. However, this approach may forfeit some advantages of traditional visualization, it can be less effective in conveying precise values and supporting comparative analysis. Moreover, for data that do not represent physical quantities, data visceralization remains limited, as it still requires users to engage in mental conversion to fully grasp the meaning.

Therefore, a natural question arises: can the advantages of traditional 2D visualization and virtual reality be combined to complement each other? A substantial body of research has explored the integration of 2D and 3D visualization through extended reality (XR) technologies. For example, some approaches allow seamless switching between 2D and 3D visualizations in mixed reality environments [7], while others create hybrid visualizations by compositing 2D charts into 3D contexts [8, 9]. These efforts largely focus on transforming or comparing data in the form of charts, using XR to strengthen users' quantitative analytical capabilities. However, their contribution to qualitative understanding and experiential engagement remains limited. In contrast, the work of Zhou et al. [10] offers a new direction, by leveraging the strengths of 2D screens (effective for presenting precise data) and 3D environments (capable of delivering concrete, experiential immersion) they propose a hybrid data storytelling method designed to enhance persuasiveness, immersion, and user experience. Nonetheless, this study primarily outlines potential design directions and challenges, without conducting detailed testing or analysis.

To further explore the potential of combining 2D visualization with immersive 3D scenes in supporting data comprehension, we propose BridgeVis, an experience that embeds interactive 2D charts within virtual reality experiences. In this approach, while data is reconstructed or narrative experiences are designed in a VR environment, an interactive 2D chart is simultaneously added to the scene, enabling users to connect quantitative analysis with qualitative engagement via interactions. For example, when exploring a reconstructed skyscraper model in VR, users can quickly fly to the top of the building to gain an intuitive sense of its height. However, precise height comparisons within the scene is less convenient due to perspective

*Corresponding author

distortion. Similarly, when the fluctuations of the NASDAQ index are represented as a roller coaster ride experience, users unfamiliar with the stock market may vividly experience the thrills of its volatility, yet their perception of macro level trends remains less clear than with traditional 2D charts. To address these issues, we designed an interactive panel that follows the user's left hand during the VR experience. The panel displays a 2D chart visualization of the scene's corresponding data, along with supplementary information. Compared to static annotations embedded in the environment, our approach allows users to engage through a series of interactions, providing a more intuitive experience while strengthening the connection between qualitative understanding of the data and quantitative analysis.

In this study, we introduce a design space to better characterize BridgeVis, consisting of three dimensions: Interaction Mode, Linking Mechanism, and Mapping Direction. Interaction Mode captures the ways in which users interact with 2D charts and 3D environments within VR. Linking Mechanism describes how qualitative experiences and quantitative analysis can be connected through different strategies in immersive settings. Mapping Direction reflects whether user behaviors are primarily driven by 2D quantitative information or 3D qualitative experience, thereby illustrating the direction of interaction flow. Building on this design space, we developed four VR prototypes using real-world datasets. Among them, S1 and S2 focus on reconstructing physical measurement data at scale, while S3 and S4 explore the integration of immersive data-driven storytelling with interactive charts for abstract data. Each prototype demonstrates the application of our design space from different categories and perspectives. Specifically, we instantiated the dimensions of Interaction Mode, Linking Mechanism, and Mapping Direction to tailor the experience for each data type, ensuring a structured connection between the 2D analytical charts and the 3D immersive environments.

To evaluate the effectiveness of BridgeVis, we conducted a mixed-method user study with 12 participants from diverse backgrounds. The study followed a within-subject comparative design, combining task-based experiments, free exploration, questionnaires, and post-study interviews. Our aim was to investigate three core questions: whether BridgeVis increases cognitive load, whether it can simultaneously support quantitative analysis and qualitative experience, and how interaction modes and linking mechanisms influence user experience. Results from task performance and NASA–TLX scores showed that introducing interactive 2D charts into immersive environments did not significantly increase cognitive workload. Questionnaire ratings further indicated that participants perceived BridgeVis as highly effective in supporting both quantitative accuracy and qualitative understanding, with many highlighting the complementary relationship between 2D charts (as analytical anchors) and 3D environments (as experiential mediums). Interviews reinforced these findings, offering rich insights into user preferences for direct selection, semantic associations, and quantitative-to-qualitative mappings. Together, these results provide empirical evidence that hybrid immersive visualization can balance analytical precision with experiential immersion, while suggesting design strategies for optimizing user interaction. From the collected data, we derived six design implications, which provide actionable guidance for future implementations of BridgeVis. Finally, we discuss broader aspects of the system, outline its limitations, and future work.

2. Related Work

2.1. Data-Driven Storytelling

Data visualization has long played the role of supporting storytelling, finding broad applications across diverse domains. Segel and Heer [2], through their study of cases such as reports, blogs, and academic papers, introduced the concept of narrative visualization. Previous research has introduced narrative design patterns to systematically construct data stories [11, 12] . These patterns encompass narrative intents such as engagement, empathy, and pacing, which can be applied individually or in combination within storytelling practices [13]. Scholars have extended this work by examining data-driven storytelling in various formats and domains, including data videos [14], data comics [15], and animated data-gifs [16] and so on.

With technological advances, the media through which data stories are delivered has also evolved, from early print, to desktop and mobile screens, and more recently to immersive devices such as AR [17] and VR [18, 19]. Méndez et al. [20] summarized developments in immersive data storytelling by reviewing recent research as well as applications in journalism and gaming. For instance, Casamayou et al. [21] transformed data into first-person sensory experiences within VR by designing a virtual roller coaster, replacing the traditional practice of visual inspection. In our work, data-driven storytelling provides essential support for designing qualitative experiences. We build on related design methods to enhance users' understanding of data, and through interaction, we map narrative-driven experiences onto users' qualitative engagement with data.

2.2. (2D + 3D) Hybrid Interface

Many applications have attempted to embed familiar 2D panels into 3D environments [22, 23, 24, 25]. This approach makes it possible to preserve spatial structure while at the same time conveying abstract attributes [26]. Foundational work on interface integration strategies distinguishes among several types of connections: interaction linkage [27], visual linkage [28], and state transitions between 2D and 3D elements [26, 29, 30]. Situated brushing-and-linking studies evaluate highlights (halos, outlines, arrows) that help users find 3D referents after brushing 2D marks [31, 32]. In geo-VR, coordinated multiple views arranged as stacks/grids can outperform single-view toggles for multilayer comparisons [33]. Animated transformations preserve mental maps when morphing between representations [7].

However, prior research also highlights that introducing 2D panels into VR is not cognitively neutral. When 2D interfaces are poorly integrated into the 3D context, such as floating overlays that occlude the scene, require frequent context switching, or provide visually dense information, they can increase extraneous cognitive load by forcing users to divide attention between incompatible representational formats. For example, traditional 2D text overlays in VR impose higher mental effort compared to more spatially integrated representations [34]. Likewise, Rzayev et al. found that head-fixed paragraph-style 2D text increased perceived task load and reduced usability in VR reading tasks, demonstrating that poor placement of 2D elements can hinder comprehension [35]. Reviews of immersive learning environments similarly warn that interface elements that are not aligned with the spatial logic of VR, such as dense charts or flat HUD-style displays—tend to elevate users' cognitive burden and disrupt immersion [36]. These findings collectively suggest that while hybrid 2D–3D interfaces are powerful, their benefits depend strongly on design decisions that minimize unnecessary mental effort. In this work, our approach seeks to retain the clarity and precision of traditional data visualization while enhancing immersion and experiential engagement, without imposing additional cognitive load.

2.3. Understanding Data in Immersive Environments

The rise of immersive technologies such as VR and AR has provided powerful tools for enhancing users' understanding of data. By reconstructing objects in virtual environments, these technologies can rapidly build immersive worlds that enable users to interact with and better grasp abstract concepts or scenarios. This potential has also been explored in data visualization. For example, Lee et al. introduced *Data Visceralization* [6], a VR-based approach designed to restore intuitive understanding of physical measurement units. Testing showed that participants found this method effective in improving data comprehension. However, challenges remain in presenting precise quantitative information. Distortions caused by perspective and other spatial factors in 3D environments can bias users' perception of data. Thus, traditional chart-based visualizations remain necessary to complement immersive representations and ensure quantitative accuracy. Moreover, the study primarily addressed data grounded in physical measurements. Whether similar approaches can extend to virtual or abstract data remains unexplored. In addition, techniques such as scale perception [37], data physicalization [38, 39, 40, 41], and embodied measures and perception [42, 43, 44, 45] can be rapidly instantiated and embedded within virtual environments, offering users vivid and concrete experiences of otherwise abstract information, thereby facilitating more intuitive and profound data understanding.

In this study, we propose a hybrid visualization approach designed to synergize qualitative immersive experiences with quantitative analytical precision. Rather than relying on scene-based annotations—which often suffer from clutter and ambiguity—we introduce an interactive 2D panel to structure user engagement. This design preserves the immersive qualities of VR while ensuring distinct accessibility to precise quantitative information, thereby effectively bridging the gap between experiential understanding and analytical interpretation. Furthermore, we expand the scope of immersive representation beyond physical measures (e.g., speed, height) to encompass abstract data. By integrating strategies such as *Concrete Scales* [46] and *Concrete Re-Expression* [47], we map abstract values onto familiar analogies (e.g., equating volume to everyday objects), grounding numerical concepts in relatable experiences.

2.4. Visualization Task and Interaction

Task-centered frameworks effectively scaffold visualization design, ranging from high-level user intents [48, 49] to atomic analytic operations [50]. Comprehensive typologies further map these abstract goals to concrete interaction techniques [51], categorizing intents [52] and dynamics [53] into actionable patterns. Recently, *Immersive Analytics* [54] has extended these goals to XR, utilizing novel multimodal or cross-device interactions to enhance analytical efficiency [55, 56].

However, designing hybrid interactions that balance quantitative analysis and qualitative immersion in 3D settings remains underexplored. To address challenges such as perspective distortion and navigation, we adapt the taxonomies of Yi et al. [52] and Heer and Shneiderman [53] for the general public. We prioritize familiarity by retaining standard 2D operations (e.g., filtering, sorting) for quantitative control [48, 53], while restricting 3D interactions to basic modalities like teleportation and ray-based selection to minimize learning costs.

3. Design Space

In developing our design space, we adopted an iterative semi-systematic process. We began by taking interaction as the primary entry point, brainstorming possible interaction patterns that could be applied in data explorations. From this process, we identified three initial dimensions: interaction modes, linking mechanisms, and mapping directions, which served as the foundational categories for our framework. These concepts were embodied in a series of prototypes that allowed us to examine their practical applicability (Section 4). Based on the evaluation of these prototypes, the design space was further refined to better capture effective strategies while addressing observed limitations. Our goal is that this design space functions as a conceptual starting point for future creators of BridgeVis, offering both inspiration and guidance by bridging theoretical insights with practical design strategies.

3.1. Dimension I: Interaction Modes

Interaction constitutes the core component of the immersive experience: it enables users to explore the environment and, beyond visual and auditory perception, serves as the initial step between qualitative engagement and quantitative understanding. In this dimension, we considered common interaction patterns in both 2D views and 3D environments, and from these we developed three representative interaction modes.

Direct Selection is among the most common and intuitive forms of interaction, broadly adopted in both 2D and 3D contexts. In VR, both 3D objects and 2D interfaces can be selected through ray or touch. This form of interaction is well-suited for simple tasks, providing immediate and easily interpretable feedback. For example, selecting a bar in a 2D bar chart may trigger corresponding changes in the 3D scene, thus linking quantitative data with qualitative experience.

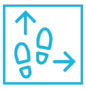

Spatial Navigation represents the core set of interactions that support movement within immersive environments. These include walking, teleportation, climbing, and flying. Through spatial navigation, users can explore the virtual environment while receiving dynamic feedback from linked 2D visualizations. For instance, users move within the 3D environment, the associated 2D chart can update in real time to present location-specific information, thereby strengthening the link between embodied exploration and data understanding.

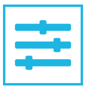

Parametric Control involves adjusting parameters, which is more complex than direct selection. Given users' greater familiarity with 2D operations, this interaction mode primarily relies on 2D interfaces, such as sliders, toggles, or sorting mechanisms. These interactions allow users to manipulate numerical values directly. For example, adjusting a timeline slider on a 2D chart can induce corresponding changes within the VR scene, thereby linking parameter-based manipulation with immersive qualitative experience.

3.2. Dimension II: Linking Mechanisms

This dimension reflects the methods by which qualitative experiences and quantitative information are mapped to one another. Unlike interaction, which serves as the entry point, linking mechanisms occur after user interaction and form the core bridge between qualitative understanding and quantitative analysis. Based on analysis of prior work and our own design practices, we propose four representative linking mechanisms.

Spatial Mapping links positions in 2D charts with corresponding locations in the virtual environment. It is particularly effective for values expressed in spatial or physical units (e.g., length or height), and can also support object rearrangement within the scene as well as rapid spatial navigation. For example, by selecting a specific height value on a 2D chart, the user can be directly transported to the corresponding height in the 3D environment.

Visual Encoding establishes correspondence between 2D data and 3D representations through visual changes. By altering features such as color and shape, this mechanism enables users to quickly identify relevant data and perform tasks such as observation or categorization. For instance, selecting a group of data in a 2D chart can change the associated 3D objects' color, drawing user's attention.

Temporal Synchronization connects data and experiential changes through event-based or time-dependent relationships. It is most suitable for datasets with temporal attributes. For example, as a user selects different time points in a 2D line chart, the state of a 3D object or environment may change accordingly, allowing the user to experience temporal dynamics in a more intuitive manner.

Semantic Association introduces additional qualitative information to facilitate the interpretation of abstract data. In 2D, this often appears as text or graphical annotations, while in 3D environments, it can be implemented by constructing related virtual objects or employing analogy models such as those in the Concrete Scales framework [46]. For example, when a user selects a volume value in a chart, the VR scene displays a familiar object of equivalent capacity, making the numerical information easier to comprehend.

3.3. Dimension III: Mapping Direction

In our work, 2D charts primarily support quantitative analysis, while 3D environments emphasize qualitative understanding. Mapping direction reflects the flow of interaction, from its starting point to its destination, which is determined once the user initiates an interaction. Within this dimension, we identify three mapping directions: Quantitative to Qualitative, Qualitative to Quantitative, and Bidirectional Synchronization.

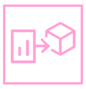

Quantitative to Qualitative: In this direction, interactions begin with the 2D panel and are then propagated to the 3D environment through a linking mechanism, providing users with corresponding qualitative experiences of the selected data. Since interactions originate from the 2D interface, suitable interaction modes primarily include Direct Selection and Parametric Control.

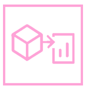

Qualitative to Quantitative: In this way, interactions are initiated in the 3D environment and reflected back in the 2D chart. Through linking mechanisms, users can access complementary quantitative information to supplement their immersive experiences. Interactions originating in the 3D environment are best supported by Direct Selection and Spatial Navigation.

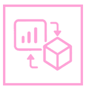

Bidirectional Synchronization allows 2D and 3D interfaces to update in parallel, maintaining consistency across both modes of representation. Achieving this effect typically requires a common medium, with time being the most common choice, especially for datasets with continuous temporal dimensions. For example, as the timeline progresses, both the 2D chart and 3D environment evolve simultaneously to reflect the changing data. While initial synchronization may be triggered by a simple interaction such as Direct Selection, subsequent updates occur automatically, ensuring a linkage between quantitative and qualitative representations.

4. Prototype

To further explore BridgeVis, we developed a series of VR prototypes. All prototypes were implemented in Unity 3D and tested on the Meta Quest 3 headset. The datasets used in these prototypes were adapted from real-world sources, including academic publications, news articles, and publicly available data. For clarity, we refer to the four prototypes as S1, S2, S3, and S4. Both S1 and S2 emphasize clear 1:1 mappings with physical units, and prove effective in conveying qualitative understanding. Each of these scenes contained multiple data objects, allowing for strong comparative relationships that supported quantitative analysis. Moreover, their datasets could be straightforwardly represented in both 2D interfaces and immersive 3D environments. Prototypes S3 and C4, by contrast, introduced additional objects and narrative experiences beyond strict physical unit mappings. Instead of directly reconstructing physical quantities, these prototypes employed immersive storytelling and analogy-based models to provide users with a qualitative sense of quantitative data. This design approach allowed us to examine how narrative immersion and metaphorical representations can complement traditional data-driven visualization. In addition to the 3D scenes, each prototype incorporated a 2D interface attached to the user's left-hand controller. This panel primarily featured interactive charts, designed to provide users with precise quantitative information that complemented the immersive experience of the virtual environment.

S1: Skyscraper Height Comparison

For decades, countries have competed to construct ever-taller skyscrapers, repeatedly breaking height records. While the exact numerical heights of these buildings are readily available through news reports or online encyclopedias, most people rarely gain a firsthand sense of scale unless they physically visit such structures. Even then, comparing one building with another in real world is often impractical, long distance intervals making it difficult to retain a clear sense of relative height.

To address the difficulty of perceiving vertical scale and perspective distortion in reality, we reconstructed four skyscrapers (e.g., Burj Khalifa, Merdeka 118, Shanghai Tower, and Tokyo Skytree) at a 1:1 scale in a virtual 2×2 grid. While this setup enhances visceral scale perception, precise comparison remains challenging. We addressed this by integrating an interactive bar chart on the user's controller. In the chart, the X-axis represents the skyscrapers, and the Y-axis their heights, with horizontal gridlines marking intervals of 100 meters. From the perspective of our design space, the system primarily utilizes Direct Selection and Parametric Control with Spatial Mapping. Clicking on height intervals in the chart teleports the user to the corresponding altitude in the 3D scene (Fig. 1 S1), while a vertical slider masks portions of the bars to hide corresponding building segments (Fig. 1 S1). This establishes a clear Quantitative → Qualitative mapping, turning precise numeric values into an embodied sense of altitude. Conversely, as users navigate via joystick, real-time height is reflected on the chart (Fig. 1 S1), supporting a loose Qualitative → Quantitative loop.

S2: Olympic Men's 100-Meter Sprint

The men's 100-meter sprint is one of the most iconic events in the Olympic Games. Athletes' performances in this discipline often far exceed the capacities of ordinary individuals, which makes it difficult for the users to develop a visceral sense of such extraordinary speed. Although audiences can access race videos and recorded results, their experience of the data remains numerical and detached from embodied perception.

To bridge the gap between abstract timing data and the visceral experience of speed, we reconstructed 15 Olympic medalists (2008–2024) as animated 3D avatars. High data density and rapid motion made visual tracking difficult, so we integrated an interactive 2D panel to support Parametric Control and Direct Selection. The x-axis represents the athletes' names and competition years, while the y-axis encodes their finishing times . Users can switch the sorting mode from chronological to performance-based, which simultaneously reorders the chart and physically rearranges the athletes on the 3D track (Fig. 1 S2). Selecting a bar or an aggregated medal count highlights the corresponding avatar by revealing their authentic race outfit (Fig. 1 S2). Additionally, filtering by name hides unselected models to reduce clutter (Fig. 1 S2). This hybrid design allows users to initiate analysis in the 2D panel (Quantitative → Qualitative) to locate specific athletes, then navigate to the corresponding area in the 3D scene to experience the race.

S3: NASDAQ Index Trend

The stock market is closely tied to economic performance, yet for audiences without direct investment experience, the fluctuations often reported in financial news are difficult to grasp intuitively. The Wall Street Journal, for instance, once translated 21 years of NASDAQ's price-to-earnings ratios into a VR roller coaster experience, enabling users to perceive cyclical rises and falls in a visceral manner [57]. Inspired by this idea, we reconstructed the weekly average NASDAQ index values for the first 28 weeks of 2025 as a line chart, where the x-axis represents calendar weeks and the y-axis denotes the index values. This line chart was reconstructed in VR as a roller coaster track.

To bridge this embodied experience with analytical precision, we implemented Temporal and Bidirectional Synchronization: a pointer on the 2D line chart moves in tandem with the user's first-person VR view, ensuring real-time correlation between the "ride" and the data (Fig. 1 S3). The design employs Visual Encoding by color-coding both the track and chart segments (e.g., yellow for volatility, red for decline) to maintain context. Users can navigate via Direct Selection on chart segments (Fig. 1 S3) or Parametric Control using a timeline slider (Fig. 1 S3) to replay specific phases. This configuration allows users to situate the physical sensation of fluctuations within a quantitative context supported by narrative news annotations.

S4: Volumes of Sperm Whale Organs

Many phenomena in the real world cannot be observed in their idealized state. For example, the internal organs of a whale are rarely accessible for examination, and even in dissection they cannot be viewed in a clean and fully structured form.

To visualize internal structures that are physically inaccessible or irregular in shape, we constructed a 1:1 anatomical model paired with an interactive pie chart. The design centers on Direct Selection and Semantic Association to bridge quantitative data and qualitative scale.Selecting an organ slice on the chart triggers Spatial Mapping, automatically teleporting the user to the optimal viewing angle (Fig. 1 S3), while Visual Encoding highlights the 3D organ in consistent colors (Fig. 1 S3). To ground abstract volumetric data, we applied Concrete Re-Expression strategies [47]: selecting the heart shows a 3D model of a large trash bin for comparison (Fig. 1 S3). This Quantitative → Qualitative mapping allows users to transition from proportional analysis to metaphor-enhanced perception.

4.1. Dimension Combinations

Based on our design space and prototyping practice, we identified four combinations of dimensions used in the prototypes. In this section, we elaborate on these combinations in detail,

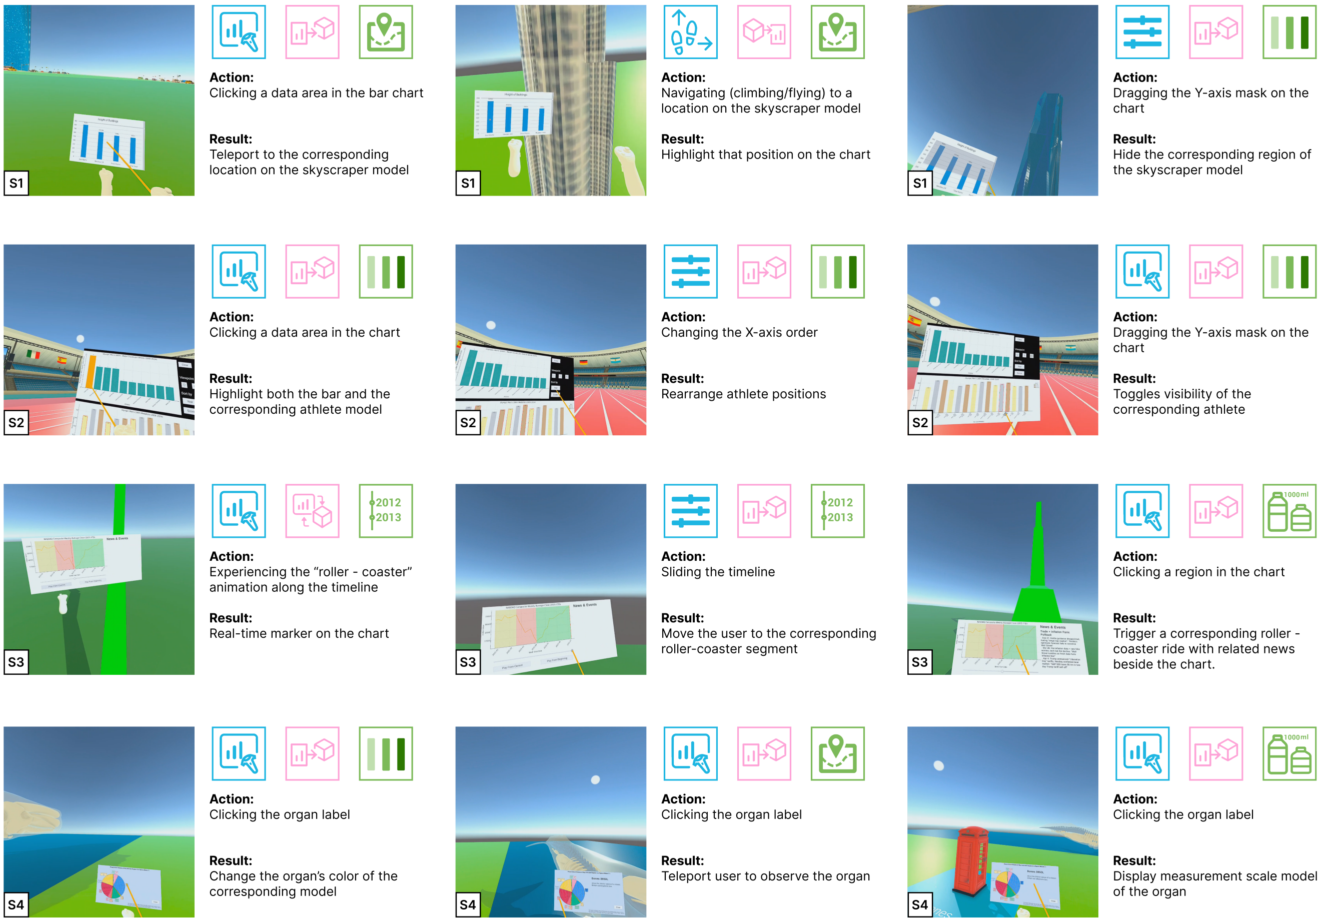

Figure 1: Explanations of different actions and results in the prototypes, including representative screenshots and icons used to indicate the corresponding design dimensions.

outlining their design rationale, suitable application scenarios, and expected outcomes. The correspondence between specific interactions and the dimensions of the design space is illustrated in Fig. 1. Our goal is to provide a set of references that can guide future creators in selecting appropriate strategies for designing immersive data experiences.

Direct Selection + Spatial Mapping / Visual Encoding / Semantic Association + Quantitative → Qualitative / Qualitative → Quantitative: This combination provides users with the most intuitive connection between detailed quantitative data and corresponding qualitative experiences. On the one hand, users can begin from a 2D chart and obtain immersive feedback in the 3D scene. On the other hand, by selecting objects in the 3D environment, users can access supplementary details and precise data in the 2D interface. This bidirectional link makes the combination especially effective for tasks that require both experiential immersion and access to granular information.

Parametric Control + Spatial Mapping / Visual Encoding + Quantitative → Qualitative: This combination excels in comparative analysis tasks involving larger datasets. Parametric operations, such as sliders or sorting allow users to efficiently filter and lock onto data objects of interest in 2D views, while spatial or visual linkage maps these selections into the 3D scene. As a result, the environment can highlight only the relevant objects or reorganize spatial configurations to suit the user's exploration. This setup supports focused, user-driven exploration while maintaining contextual clarity.

Spatial Navigation + Spatial Mapping + Qualitative → Quantitative: This combination is particularly well-suited for geospatial or spatially structured datasets. It allows users to navigate freely within a rich 3D environment through immersive movements such as walking or teleportation. At the same time, the 2D interface dynamically updates to provide precise numerical information corresponding to the user's current position. This design supports exploratory experiences while ensuring access to detailed quantitative data when needed.

Direct Selection + Temporal Synchronization + Bidirectional Sync: This combination is suited for time-continuous datasets. Users can immerse themselves in experiencing events over a given time span, while both the 2D chart and the 3D scene remain synchronized in real time. This also enables users to directly compare quantitative information with lived immersive experiences. Furthermore, temporal synchronization allows users to control, pause, or revisit specific time segments, making it suitable for applications where repetitive exploration of temporal dynamics is critical, such as historical event reconstruction or dynamic system simulations.

5. User Study

To investigate users' perceptions of BridgeVis and evaluate its effectiveness, we conducted a mixed-method user study. Participants were invited to directly experience our BridgeVis prototypes, allowing us to systematically assess the potential benefits and challenges of integrating 2D and 3D data representations in immersive environments.

Through this study, we aimed to address the following research questions:

- RQ1: Does BridgeVis increase users' cognitive load, potentially leading to cognitive overload?
- RQ2: Can BridgeVis effectively support both quantitative analysis and qualitative experience simultaneously?
- RQ3: How do different interaction modes and linking mechanisms influence users' experiences?

By addressing these research questions, our study seeks to clarify the cognitive, analytical, and experiential implications of hybrid data visualization in immersive environments, providing insights for the design of BridgeVis.

5.1. Participants

We recruited 12 participants (P1–P12) through online advertisements. The group included 7 females and 5 males, with ages ranging from 20 to 35 years. Among them, 9 participants had prior experience using virtual reality devices. In terms of data literacy, 7 participants reported regularly reading data journalism and stories, whereas 3 participants indicated that they rarely engaged with such content. Participants came from diverse backgrounds, including interaction design, computer science, management, and sports science, ensuring a range of perspectives. All participants signed an informed consent form prior to the study, and received a payment of $8–$12.

5.2. Study Procedure

The study consisted of three main phases: Introduction, a within-subject comparative experiment, a free-form exploration session, and a post-study interview. Each participant spent approximately 90 minutes completing the study.

Introduction. At the beginning of the study, participants were introduced to the research background and the experimental procedure. They were then equipped with a VR headset and given time to familiarize themselves with its basic operations as well as the interaction methods used in our prototypes.

Within-Subject Comparative Experiment. To examine whether BridgeVis may impose cognitive overload, we designed a within-subject comparative experiment using prototypes S1 and S2, along with their control versions (Fig. 2). In the control condition of S1, participants were not provided with a 2D chart; instead, building heights were annotated directly within the scene. In the control condition of S2, participants were not given a 2D chart either; instead, they distinguished athletes by their clothing colors and annotations.

We designed four tasks (T1–T4) simulating potential user behaviors in the scenes. While real-world analysis often involves complex workflows, such as multi-attribute filtering or deriving trends from large datasets, these complex activities are built upon low-level tasks [50]. Therefore, we selected representative atomic tasks (Retrieve Value, Filter, and Find Extremum) to establish a baseline for cognitive load assessment.

For tasks T1 and T2, participants were asked to compare the structural height of two specific buildings, an attribute that is difficult to gauge from 2D charts or a ground-level perspective and then navigate to the taller one to experience its height. Tasks T3 and T4 simulated a dynamic event, requiring participants to rapidly locate a specific athlete's spatial position at the exact moment the winner crossed the finish line. To evaluate interface efficacy, T1 and T3 utilized a chart-based condition, while T2 and T4 only relied on in-scene annotations. By isolating these fundamental interaction units, we aimed to assess the cognitive cost of the hybrid interface itself, minimizing confounding variables associated with complex reasoning.

To control for learning effects, the experiment order was counterbalanced using a Latin square design. Task completion times were recorded, and after each task, participants completed the NASA–TLX [58] questionnaire to measure cognitive load. Before recruiting participants, we conducted a pilot study with two experienced VR users to ensure comparable task difficulty across conditions and to refine the study procedure.

Free-Form Exploration. Following the comparative experiment, participants were introduced to all four prototypes (S1–S4) along with their background and interaction methods. They were given 20 minutes to freely explore the prototypes. After exploration, participants completed a short questionnaire consisting of four five-point Likert-scale items, evaluating their experiences in terms of immersion, quantitative and qualitative understanding, and integration between the two.

Post-Study Interview. Finally, a semi-structured post-study interview was conducted with each participant. The interview included 13 questions covering three major themes: (1) connections between quantitative and qualitative understanding, (2) interaction experience, and (3) perceptions of immersion and cognitive load. A complete list of the interview questions is provided in the supplementary materials.

5.3. Results

Cognitive Load. During the within-subject comparative experiment, we collected two types of data: task completion time (including time spent on failed attempts until success) and NASA–TLX scores. To examine whether BridgeVis increases cognitive load, and whether this effect differs across scenes, we conducted a 2×2 repeated-measures analysis of variance with Scene (S1 vs. S2) and Chart Presentation (with vs. without chart) as independent variables, and NASA–TLX scores and completion times as dependent variables.

As shown in Fig. 3, the main effect of chart presentation on perceived cognitive load was not significant, $F(1, 11) =$

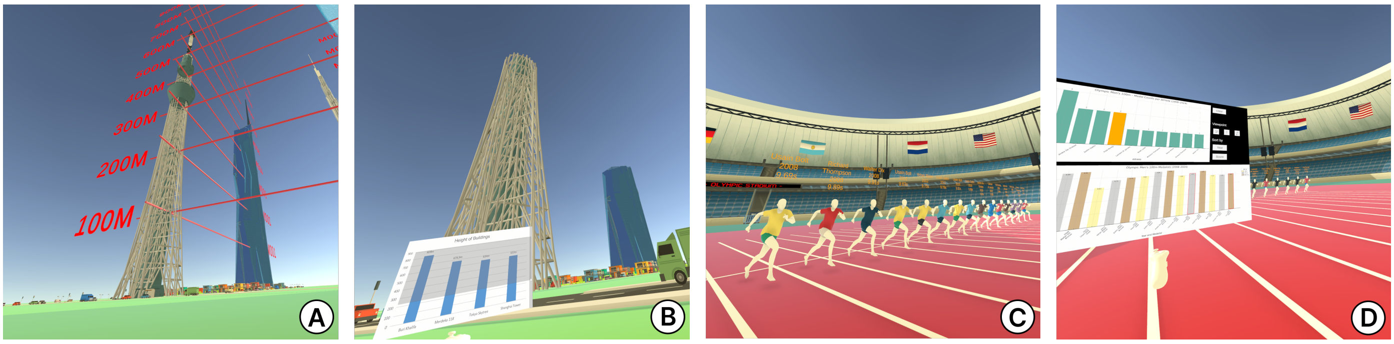

Figure 2: Comparative experiment scenarios: (A) S1 control condition, building height information is provided through in-scene annotations; (B) S1 experimental condition, building height is obtained via interactive chart with masking; (C) S2 control condition, athletes are distinguished by annotations above their heads and different clothing colors; (D) S2 experimental condition, athletes are distinguished through interactive chart with highlighting.

0.04, p = .85, with a negligible effect size (η_p^2 = 0.003). The mean TLX scores for the "With Chart" condition (Scene 1: M = 28.95; Scene 2: M = 24.10) were almost identical to the "Without Chart" condition (S1: M = 28.39; S2: M = 23.54). This lack of difference is statistically reinforced by the highly overlapping 95% confidence intervals (CIs). For instance, in S1, the CI for "With Chart" [19.11, 38.78] was nearly indistinguishable from "Without Chart" [19.60, 37.18]. This constitutes a core finding of our study: adding interactive charts in VR did not lead to a statistically significant increase in perceived workload. The main effect of scene was also not significant, $F(1, 11)$ = 2.86, p = .12, suggesting that the two scenarios did not differ in inherent perceived difficulty. Importantly, the interaction effect was negligible ($F \approx 0.00$, $p > .99$), indicating that the neutral influence of charts on cognitive load was consistent across both scenarios.

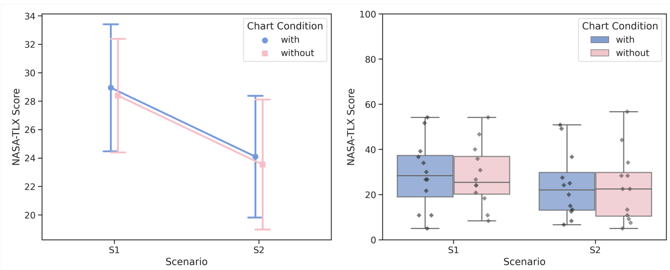

Figure 3: Subjective workload analysis using NASA-TLX scores. Left: Interaction plot displaying mean scores with 95% confidence intervals (CIs). Right: Boxplots showing the distribution of raw scores.

We also analyzed task completion time (Fig. 4) as a proxy indicator of cognitive load, under the assumption that higher workload tends to prolong task execution. Results showed no significant main effect of chart presentation, $F(1, 11)$ = 0.034, p = .86. Crucially, the 95% confidence intervals for time were notably wide and overlapped substantially across chart conditions. In S1, the CI for "With Chart" [93.49, 206.84] completely encompassed the range of "Without Chart" [124.91, 200.75]. This high variability (wide CIs) suggests that individual differences played a dominant role, masking the specific effects of the charts. Observations indicated that this variance was largely driven by users' varying VR proficiency. While VR-proficient users utilized charts efficiently to speed up analysis (contributing to the lower bound of the CIs), novice users struggled with basic interactions (contributing to the upper bound), which inflated completion times regardless of the chart condition.

Interestingly, although the interaction was not statistically significant, distinct trends appeared in the means. In the simpler S2, the presence of charts slightly increased average task time (from 96.17s to 115.08s), whereas in the complex S1, charts appeared to decrease time (from 162.83s to 150.17s). This suggests a nuanced trade-off: when tasks impose low intrinsic cognitive demand, additional chart information might act as a distraction; however, in complex tasks, the cognitive offloading benefits of charts may begin to outweigh the interaction costs.

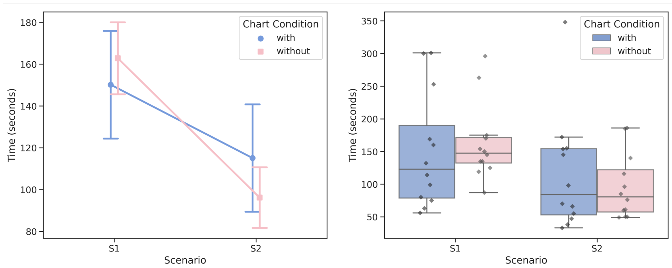

Figure 4: Analysis of task completion time. Left: Interaction plot of mean completion times with 95% confidence intervals. Right: Boxplots showing the distribution of data.

Taken together, these findings strongly support RQ1: in VR, carefully designed 2D charts can serve as an effective medium for information presentation. The overlapping CIs in both subjective scores and objective time confirm that BridgeVis does not impose additional cognitive burden on users' mental resources, provided the user has sufficient familiarity with VR.

Questionnaire Results. To evaluate the effectiveness of BridgeVis in VR environments, we collected participants' subjective feedback using a five-point Likert scale (1 = strongly disagree, 5 = strongly agree) on four key dimensions: Immersion, Quantitative Analysis, Qualitative Understanding, and Quant–Qual Integration. Results are summarized in Fig. 5.

In terms of quantitative analysis capability, 11 out of 12 participants (91.7%) rated the experience positively, with 7 par-

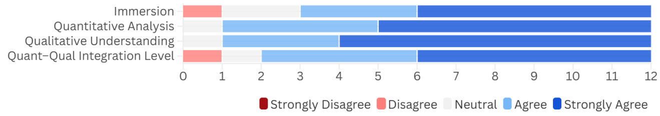

Figure 5: This figure shows participants' ratings on "Immersion", "Quantitative Analysis", "Qualitative Understanding", "Quant–Qual Integration Level" for BridgeVis experience

ticipants (58.3%) giving the highest rating. Only one participant remained neutral. Similarly, for qualitative understanding, 11 participants (91.7%) provided positive ratings, again with 7 participants (58.3%) selecting "strongly agree." These results strongly suggest that BridgeVis successfully supports both precise quantitative interpretation and immersive qualitative comprehension, providing empirical support for RQ2.

To determine whether participants perceived the 2D and 3D views as a cohesive whole rather than separate components, we evaluated integration between quantitative and qualitative information. Ten participants (83.3%) agreed that the system effectively combined the two perspectives (6 "strongly agree," 4 "agree"). This indicates that participants regarded BridgeVis as a seamlessly integrated analytic tool, rather than two disconnected modalities. Such integration was identified as a key factor enabling users to transition smoothly between quantitative and qualitative modes of reasoning.

As immersion is a critical factor for usability in VR-based systems, we also assessed participants' sense of immersion. Ten participants (83.3%) provided positive evaluations (6 "strongly agree," 4 "agree"), demonstrating that the design not only fulfilled analytical needs but also delivered an engaging, distraction-free experience. Participants reported being able to focus on the data itself without being hindered by interface complexity or environmental disruptions.

Overall, participants' feedback provides strong evidence that BridgeVis is an effective solution for hybrid visualization in VR. It not only enables immersive user experiences but also delivers dual support for quantitative and qualitative reasoning, with a high level of integration between the two. These findings will be further elaborated through qualitative insights from the post-study interviews presented in the next section.

Interview Analysis. In this section, we present an analysis of the post-study interviews conducted with 12 participants. The interviews were designed to complement the quantitative findings by providing deeper insight into participants' subjective experiences and perceptions of BridgeVis. The semi-structured format allowed us to capture nuanced feedback across three major themes: qualitative understanding, quantitative analysis, and the integration of the two, as well as reflections on interaction design, immersion, and cognitive load.

The 3D environment primarily supported qualitative understanding: Reconstructing physical quantities at a 1:1 scale in virtual reality can enhance users' grasp of the underlying meaning of data [6]. For instance, one participant reflected: "*In the skyscraper comparison scene (S1), it recreated measurements that are almost impossible to experience in real life. By placing several buildings side by side and allowing me to view them from different perspectives, it gave me a very direct sense of scale*" [P8]. Similarly, another remarked: "*In the sprinting scene (S2), I could clearly feel how much of a difference 0.1 seconds makes, which is something I could never experience from TV*" [P4]. Beyond reconstruction, participants also noted that alternative strategies could effectively enhance their sense of data. P6 explained: "*In the whale organ comparison, because you provided familiar everyday objects as analogies, this approach really helped me understand the scale better*".

2D Panel as Anchors for Quantitative Analysis and Navigational Exploration: Participants consistently regarded the interactive 2D chart as the foundational "anchor" for exploration. It served as a unified interface that prevented disorientation, with P3 noting it provided a "sense of orientation" and security during operations. The chart acted as a "shortcut" for rapid teleportation [P8] and efficient information retrieval via filtering [P10]. As P8 summarized, the chart offered a necessary overview, allowing users to "directly navigate to the corresponding object" for detailed inspection.

The perceived utility of the chart scaled with complexity. While participants noted that simple annotations sufficed for basic tasks like height comparison [P10], the chart became indispensable for complex analytical goals. P11 emphasized that multi-dimensional tasks, such as filtering athletes by country, would be "impossible... purely in the 3D scene." This suggests that while simple tasks may not strictly require BridgeVis, its value maximizes in sophisticated analytical workflows where 2D precision complements 3D immersion.

Complementary Relationship Between Charts and Environments: Participants emphasized the synergy between 2D quantitative analysis (e.g., exact values) and 3D qualitative experience (e.g., visceral scale). While the majority (9/12) focused primarily on the immersive environment, they relied on the chart to "*interpret the numbers in context*" [P11] and another noted that attention could shift dynamically: "*At the beginning I was mainly attracted by the environment, but after some time I still wanted to use the chart to check relationships and comparisons, especially in scenes that interested me*" [P8]. A minority (2/12) reported being so absorbed in the 3D scene that they initially overlooked the chart. However, even they acknowledged its indispensability for structured exploration: "*With the chart I was able to explore the scene more thoroughly. It helped me find perspectives and details that I would have missed on my own. I tried to imagine whether these scenes could work without the chart, but I don't think so, the chart is still indispensable*" [P9]. Participants described the 2D panel as a necessary "entry point" or outline: it allows users to "*jump straight*" to data of interest, preventing the inefficiency of manual navigation and making comparisons obvious [P8].

Cognitive Load Stemming from Inappropriate Information Presentation: Participants confirmed the appropriateness of the controller-attached interface, reporting minimal fatigue. Even with multiple charts, the experience remained manageable due to "clear" logical relationships [P12]. However, cognitive load increased with information density. Participants expressed reluctance to engage with dense text, with P5 noting they “*would not choose to read so much text in a VR,*” highlighting the inefficiency of textual content in immersion. Visual clutter also contributed to strain; P6 found the "red auxiliary height lines" in S1 initially "overwhelming." These findings indicate that cognitive load stems not from the charts themselves, but from their presentation, suggesting a need for lightweight, multimodal cues over dense textual or visual overlays.

5.4. Reflection on Design Space

Interaction Modes: Direct Selection was overwhelmingly preferred (8/12) for its efficiency. In contrast, Spatial Navigation, while essential for orientation, was deemed less efficient for specific retrieval. P12 highlighted a key tension between immersion and analysis: “*When I am fully immersed in the environment, I no longer want to look back at the numbers on the chart,*” suggesting that deep spatial engagement can reduce the motivation for quantitative checking. Parametric Control served as a valuable auxiliary layer, which P8 described as a "supporting method" to filter distractions once targets were identified. Taken together, these findings indicate that direct selection best satisfies users’ needs for quick and intuitive interaction, while spatial navigation supports immersive orientation and parametric control provides targeted refinement.

Linking Mechanisms: Semantic Association and Visual Encoding were praised (5/12 each) for grounding abstract data. In dense scenes like the sprint track, P11 noted that visual highlighting “*really reduced my burden,*” enabling immediate identification of targets from a distance. While Spatial Mapping was fundamental, participants pointed out potential disorientation after teleportation, suggesting a need for better post-travel visual guidance [P12]. Although participants expressed different preferences for linking mechanisms, they consistently affirmed that, across all four prototypes, the feedback after interaction aligned with their expectations and mirrored the conventions they were accustomed to when working with 2D platforms.

Mapping Directions: Quantitative → Qualitative was favored (7/12) as the most natural flow. P5 provided a compelling metaphor for this preference: “*If I think of this experience as a kind of game, the chart is like the instruction manual... clicking on the chart and seeing a response feels like a natural reaction.*” Bidirectional Synchronization received mixed feedback: while offering flexible perspective switching [P1], dynamic updates (e.g., S3’s roller coaster) risked dividing attention [P8].

Beyond these individual preferences, the study also revealed how participants gravitated toward specific combinations of dimensions in practice. The most frequently enacted pattern corresponded to the “Direct Selection + Spatial/Visual Linking + Quantitative → Qualitative” combination described in Section 4.1. Participants often began by clicking on items or ranges in the 2D panel to teleport to, highlight, or filter the corresponding objects in the 3D scene. This pattern allowed them to first establish a precise quantitative reference (e.g., an exact height or time) and then immediately experience its qualitative meaning in situ. A second common pattern combined Parametric Control with Spatial Mapping or Visual Encoding, especially in scenes with many similar objects. Here, sliders and filters on the panel were used to reduce clutter and reorganize the scene, after which participants switched to more embodied exploration—illustrating how parametric controls and spatial navigation worked together as a refinement layer on top of direct selection. In contrast, combinations centered on Spatial Navigation and Qualitative → Quantitative mappings were favored by participants who valued orientation and free exploration, while the “Direct Selection + Temporal Synchronization + Bidirectional Sync” combination in S3 was perceived as powerful yet demanding: it enabled repeated replay and cross-checking of a time series, but also risked dividing attention between the moving chart and the roller-coaster motion. Overall, these patterns show that participants did not use the dimensions independently; instead, they appropriated a small set of recurring combinations that balanced efficiency, immersion, and cognitive effort in different ways.

6. Implications for Design

Based on the findings of our study, we distill six key design implications for creating the BridgeVis experience.

I1: Use 3D for “Explorers” and 2D for “Analysts.”

A successful VR data experience must be flexible, adjusting the relative weight of 2D versus 3D interfaces, interaction modes, and information presentation depending on the user profile. Design must balance between “explorers” and “analysts.” For novices or the general public (explorers), curiosity and intuitive impressions are the primary motivations. As seen in our study, some participants became so absorbed in the 3D environment that they “*forgot the chart existed.*” For such users, the 3D environment should play the leading role, with storytelling-driven design, simplified interactions, and direct semantic cues, while the 2D chart should be reduced to a simple chart or a navigation menu. Conversely, for experts or analysts, precision and efficiency are paramount. These users require a powerful 2D dashboard with advanced filtering and multi-dimensional comparison capabilities, while the 3D environment serves as a contextual “field site” for validating and experiencing data.

I2: Use Annotations in 3D for Simple Datasets, Prioritize 2D Chart for Complex Datasets

Effective design must also take into account the type of the dataset, as this directly shapes both user experience and modes of presentation. From the perspective of data complexity, design must shift between environment-first and chart-first strategies. For simple datasets (e.g., skyscraper comparisons), the 3D environment is already intuitive, and annotations alone may suffice. Here, complex 2D charts may even add unnecessary cognitive load. In contrast, in scenarios with abundant or complex data (e.g., athlete comparisons or the NASDAQ roller coaster), 3D environments are full of distractions, making it hard to locate and compare. In these cases, the 2D chart becomes indispensable as a data manager, providing filtering, highlighting, and overview functions that enhance efficiency and perception.

I3: Unify 2D Precision and 3D Immersion into One Seamless System

Our results confirm that 2D charts and 3D environments play complementary roles rather than competing ones. Designers should treat them as an integrated system. The 3D environment excels at providing qualitative, embodied understanding, such as recreating physical scales or offering metaphoric analogies. Participants highlighted the value of such experiences. Meanwhile, the 2D chart ensures precise quantitative analysis, offering values, classifications, and filtering tools. Together, they answer the dual questions: "How much is it?" and "What does it feel like?" The design goal, therefore, should not be to balance attention equally, but to enable seamless switching between quantitative overview (2D) and qualitative immersion (3D), thereby forming a complete cycle of data comprehension.

I4: Make Simplicity the Default, Complexity Optional

Interaction choices directly shape the fluency and efficiency of the user experience. Our study shows that direct selection was overwhelmingly preferred (8/12 participants) for its efficiency and intuitiveness, enabling users to rapidly link charts and environments. Other modes, such as spatial navigation and parametric control, are best treated as secondary tools tailored for specific tasks (e.g., fine-grained exploration, filtering). A layered interaction design is recommended: direct selection as the default mode, supplemented by context-sensitive support for navigation and filtering. Experience should also anticipate user intent, for example, when a user lingers near a 3D object, the interface could automatically highlight corresponding chart items or suggest parameter controls.

I5: Be Faithful to Reality, Be Prudent with Metaphors

For concrete data (e.g., entities with physical measurements), the core design principle is to remain faithful to reality. The strength of such experiences lies in the high-fidelity, one-to-one reconstruction of real-world scales, enabling users to perceive phenomena that would otherwise be inaccessible due to physical limitations. In these cases, the main design challenges involve ensuring the precision of the models, the naturalness of interactions, and the smoothness of navigation. Because the experience is grounded in users' existing world knowledge, the cognitive cost of understanding is relatively low.

By contrast, for abstract data (e.g., financial trends or social statistics), the design challenge shifts toward creating metaphors. Designers must translate invisible numerical relationships into perceptible spatial or physical experiences. Such metaphors create strong emotional connections and memorable impressions. However, they must be applied with caution, as not all users can tolerate or enjoy intense physical simulations. For instance, in S3 while some participants reported excitement, others, particularly those prone to motion sickness, experienced discomfort. For these users, the metaphor not only failed to communicate information but actively disrupted data interpretation, diverting attention from understanding data to coping with physical unease. Therefore, when designing embodied metaphors for abstract data, it is essential to assess the preferences and physiological tolerance of the target audience. High-intensity or provocative metaphors should be used sparingly, or offered as optional modes rather than defaults.

I6: Replace Dense Text with Icons, Images, and Voice in VR

One of the most consistent findings is that VR environments are poorly suited to dense text-based content. Designers should avoid long textual descriptions and instead rely on icons, concise keywords, voice narration, or visual legends. Visual elements must also be carefully calibrated: for example, strongly salient features like the red auxiliary height lines in the skyscraper scene caused unnecessary stress for some participants [P6]. Overall, information presentation in VR should prioritize visual clarity and comfort, converting complex textual data into graphical or auditory formats. By doing so, designers can deliver rich information while minimizing cognitive load, thereby preserving the quality of immersive experiences.

7. Discussion

7.1. Use BridgeVis for Future Design

Our study demonstrates that BridgeVis effectively bridges the gap between quantitative analysis and qualitative experience in VR by integrating interactive 2D charts with immersive 3D scenes. The design space we propose is not merely descriptive but generative. It offers a systematic method for combining interaction modes, linking mechanisms, and mapping directions into effective hybrid designs. For example, Direct Selection × Visual Encoding represents the most intuitive combination, enabling rapid highlighting and comparison tasks. In contrast, Spatial Navigation × Spatial Mapping emphasizes exploration, making it particularly suitable for cultural heritage or geographic datasets. Other combinations extend the design possibilities: Temporal Synchronization × Bidirectional Sync is well-suited for time-series contexts such as climate change or financial data, where interactions in either charts or immersive environments update one another in real time. Meanwhile, Semantic Association × Parametric Control allows abstract data to be connected to familiar metaphors, giving users impressions of ratios and concrete scales. By articulating these combinations, our design space illustrates how BridgeVis can guide the creation of new prototypes, enabling designers to flexibly tailor configurations to dataset type, narrative goals, and user profiles. The design space and its dimension combinations provide not only a way to design future prototypes, but also a lens to interpret how different hybrid configurations channel users toward particular patterns of attention, movement, and reasoning.

7.2. Application Domains

domains. In education, BridgeVis can help learners connect abstract concepts with concrete experiences, for example, linking physics equations to embodied visualizations. In journalism, immersive narratives could combine data-driven reports with experiential metaphors, deepening public engagement with complex issues such as climate change or economic inequality. In scientific visualization, BridgeVis could help researchers explore multidimensional data by pairing quantitative dashboards with embodied spatial representations, fostering both accuracy and intuition. Similarly, in public exhibitions [59] or tourism [60], hybrid experiences could make data-rich content more engaging for general audiences by combining storytelling, exploration, and interactive analysis.

7.3. Limitations

Although our findings are promising, several limitations should be acknowledged. First, the study was conducted with a relatively small sample (N=12), which restricts the generalizability of the results. We observed notable individual differences in participants' responses, suggesting that a larger and more diverse participant pool is needed to validate the robustness of our conclusions. Second, our experimental tasks were intentionally designed as simple, atomic operations (e.g., identify, locate) to control for confounding variables in measuring cognitive load. However, this simplicity may underestimate the benefits of 2D charts, which typically excel in complex analytic scenarios involving multi-attribute filtering, sorting, and pattern recognition [51, 53]. Consequently, our results likely represent a conservative estimate of BridgeVis's efficiency gains. Future studies should incorporate open-ended, complex problem-solving tasks to fully evaluate the system's performance in high-level analytic workflows. Third, issues of VR usability and comfort remain. Some participants reported distraction or discomfort in certain embodied experiences, such as the roller coaster metaphor, underscoring the need to carefully balance immersion with user tolerance.

7.4. Future Work

Future research should address these limitations while also exploring new opportunities. Expanding the scope of the study through larger-scale experiments with more diverse participant groups would strengthen the external validity of our findings. Additional prototypes could be developed in different data domains to further examine the versatility and applicability of the design space. From a technical perspective, integrating multimodal input and output channels (e.g., haptic feedback [61], spatial audio) could enhance the complementarity between 2D charts and 3D environments by engaging additional sensory modalities. Another promising direction lies in AI-driven adaptive design [62], where systems dynamically adjust elements such as chart styles or metaphorical models according to user profiles, task complexity, or cognitive state. Such adaptivity could also extend to augmented reality (AR) contexts, where interactive quantitative panels are overlaid on the physical world to support hybrid analysis in real environments. In addition, while our study reveals clear user preferences and tendencies regarding hybrid visualization, there is still a lack of rapid, practical evaluation methods to help designers assess the effectiveness of data storytelling and data understanding in immersive contexts. Developing such evaluation frameworks could greatly accelerate design iteration and make immersive data storytelling more accessible to practitioners.

8. Conclusion

In this paper, we introduced BridgeVis, an approach that integrates interactive 2D charts into immersive virtual reality environments to bridge the gap between quantitative precision and qualitative experience. We proposed a design space comprising three dimensions and instantiated it through four prototypes covering diverse data types and narrative contexts. A user study with twelve participants demonstrated that BridgeVis effectively supports both quantitative analysis and qualitative understanding without causing cognitive overload. Participants described the 2D panel as an analytical "anchor" and the 3D environment as an experiential "medium," confirming the complementary relationship between the two modalities. Based on these findings, we derived five design implications. By emphasizing how 2D visualisation and 3D representations can be integrated into a unified experience, this work lays a foundation for future research and practice in immersive analytics.

References

[1] S. K. Card, J. D. Mackinlay, B. Shneiderman (Eds.), Readings in information visualization: using vision to think, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1999.

[2] E. Segel, J. Heer, Narrative visualization: Telling stories with data, IEEE Transactions on Visualization and Computer Graphics 16 (6) (2010) 1139–1148. doi:10.1109/TVCG.2010.179.
URL http://ieeexplore.ieee.org/document/5613452/

[3] D. Showkat, E. P. S. Baumer, Where do stories come from? examining the exploration process in investigative data journalism, Proceedings of the ACM on Human-Computer Interaction 5 (2021) 1–31. doi:10.1145/3479534.
URL https://dl.acm.org/doi/10.1145/3479534

[4] C. N. Knaflic, Storytelling with Data: A Data Visualization Guide for Business Professionals, John Wiley & Sons, 2015, google-Books-ID: retRCgAAQBAJ.

[5] A. Zentner, R. Covit, D. Guevarra, Exploring effective data visualization strategies in higher education, SSRN Electronic JournalPublisher: Elsevier BV (2019). doi:10.2139/ssrn.3322856.
URL https://www.ssrn.com/abstract=3322856

[6] B. Lee, D. Brown, B. Lee, C. Hurter, S. Drucker, T. Dwyer, Data visceralization: Enabling deeper understanding of data using virtual reality, IEEE Transactions on Visualization and Computer Graphics 27 (2) (2021) 1095–1105. doi:10.1109/TVCG.2020.3030435.
URL https://ieeexplore.ieee.org/document/9229242/

[7] B. Lee, M. Cordeil, A. Prouzeau, B. Jenny, T. Dwyer, A design space for data visualisation transformations between 2d and 3d in mixed-reality environments, in: CHI Conference on Human Factors in Computing Systems, ACM, 2022, pp. 1–14. doi:10.1145/3491102.3501859.
URL https://dl.acm.org/doi/10.1145/3491102.3501859

[8] Q. Zhu, T. Lu, S. Guo, X. Ma, Y. Yang, 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.
URL https://ieeexplore.ieee.org/document/10669767/

[9] Y. Yang, T. Dwyer, K. Marriott, B. Jenny, S. Goodwin, Tilt map: Interactive transitions between choropleth map, prism map and bar chart in immersive environments, IEEE Transactions on Visualization and Computer Graphics 27 (12) (2021) 4507–4519. doi:10.1109/TVCG.2020.3004137.
URL https://ieeexplore.ieee.org/document/9123548/

[10] X. Zhou, Y. Yang, F. Ortega, A. U. Batmaz, B. Lee, Data-driven storytelling in hybrid immersive display environments (2023). arXiv:2308.13015[cs], doi:10.48550/arXiv.2308.13015.
URL http://arxiv.org/abs/2308.13015

[11] R. Kosara, J. Mackinlay, Storytelling: The next step for visualization, Computer 46 (5) (2013) 44–50. doi:10.1109/MC.2013.36.
URL http://ieeexplore.ieee.org/document/6412677/

[12] J. Hullman, N. Diakopoulos, Visualization rhetoric: Framing effects in narrative visualization, IEEE Transactions on Visualization and Computer Graphics 17 (12) (2011) 2231–2240. doi:10.1109/TVCG.2011.255.
URL http://ieeexplore.ieee.org/document/6064988/

[13] B. Bach, M. Stefaner, J. Boy, S. Drucker, L. Bartram, J. Wood, P. Ciuccarelli, Y. Engelhardt, U. Köppen, B. Tversky, Narrative Design Patterns for Data-Driven Storytelling, in: N. H. Riche, C. Hurter, N. Diakopoulos, S. Carpendale (Eds.), Data-Driven Storytelling, 1st Edition, A K Peters/CRC Press, Boca Raton, Florida : Taylor & Francis/CRC Press, [2018], 2018, pp. 107–133. doi:10.1201/9781315281575-5.
URL https://www.taylorfrancis.com/books/9781315281568/chapters/10.1201/9781315281575-5

[14] F. Amini, N. Henry Riche, B. Lee, C. Hurter, P. Irani, Understanding data videos: Looking at narrative visualization through the cinematography lens, in: Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, ACM, 2015, pp. 1459–1468. doi:10.1145/2702123.2702431.
URL https://dl.acm.org/doi/10.1145/2702123.2702431

[15] B. Bach, Z. Wang, M. Farinella, D. Murray-Rust, N. Henry Riche, Design patterns for data comics, in: Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, ACM, 2018, pp. 1–12. doi:10.1145/3173574.3173612.
URL https://dl.acm.org/doi/10.1145/3173574.3173612

[16] X. Shu, A. Wu, J. Tang, B. Bach, Y. Wu, H. Qu, What makes a data-gif understandable?, IEEE Transactions on Visualization and Computer Graphics 27 (2) (2021) 1492–1502. doi:10.1109/TVCG.2020.3030396.

[17] X. Wang, L. Besançon, D. Rousseau, M. Sereno, M. Ammi, T. Isenberg, Towards an Understanding of Augmented Reality Extensions for Existing 3D Data Analysis Tools, in: Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, ACM, Honolulu HI USA, 2020, pp. 1–13. doi:10.1145/3313831.3376657.
URL https://dl.acm.org/doi/10.1145/3313831.3376657

[18] J. Bastiras, B. H. Thomas, Combining Virtual Reality and Narrative Visualisation to Persuade, in: 2017 International Symposium on Big Data Visual Analytics (BDVA), 2017, pp. 1–8. doi:10.1109/BDVA.2017.8114623.
URL https://ieeexplore.ieee.org/document/8114623/

[19] D. Hayatpur, H. Xia, D. Wigdor, DataHop: Spatial data exploration in virtual reality, in: Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology, ACM, 2020, pp. 818–828. doi:10.1145/3379337.3415878.
URL https://dl.acm.org/doi/10.1145/3379337.3415878

[20] J. Méndez, W. Luo, R. Rzayev, W. Büschel, R. Dachselt, Immersive data-driven storytelling: Scoping an emerging field through the lenses of research, journalism, and games, IEEE Transactions on Visualization and Computer Graphics (2025) 1–13doi:10.1109/TVCG.2025.3531138.
URL https://ieeexplore.ieee.org/document/10844576/

[21] V. Casamayou, Y. Jansen, P. Dragicevic, A. Prouzeau, Ride your data: Raise your arms, scream, and experience your data from a roller coaster cart, workshop paper (2022).
URL https://altvis.github.io/papers/2022/ride-your-data.pdf

[22] C. Zhang, W. Zeng, L. Liu, UrbanVR: An immersive analytics system for context-aware urban design, Computers & Graphics 99 (2021) 128–138. arXiv:2107.00227[cs], doi:10.1016/j.cag.2021.07.006.
URL http://arxiv.org/abs/2107.00227

[23] L. Zhao, T. Isenberg, F. Xie, H.-N. Liang, L. Yu, SpatialTouch: Exploring spatial data visualizations in cross-reality, IEEE Transactions on Visualization and Computer Graphics 31 (1) (2025) 897–907. doi:10.1109/TVCG.2024.3456368.
URL https://ieeexplore.ieee.org/document/10670539/

[24] M. R. Seraji, W. Stuerzlinger, HybridAxes: An immersive analytics tool with interoperability between 2d and immersive reality modes, in: 2022 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct), 2022, pp. 155–160, ISSN: 2771-1110. doi:10.1109/ISMAR-Adjunct57072.2022.00036.
URL https://ieeexplore.ieee.org/document/9974239/

[25] M. Cavallo, M. Dolakia, M. Havlena, K. Ocheltree, M. Podlaseck, Immersive insights: A hybrid analytics system forCollaborative exploratory data analysis, in: 25th ACM Symposium on Virtual Reality Software and Technology, ACM, 2019, pp. 1–12. doi:10.1145/3359996.3364242.
URL https://dl.acm.org/doi/10.1145/3359996.3364242

[26] J. Hong, R. Hnatyshyn, E. A. D. Santos, R. Maciejewski, T. Isenberg, A survey of designs for combined 2d+3d visual representations, IEEE Transactions on Visualization and Computer Graphics 30 (6) (2024) 2888–2902. arXiv:2401.05438[cs], doi:10.1109/TVCG.2024.3388516.
URL http://arxiv.org/abs/2401.05438

[27] A. Buja, J. McDonald, J. Michalak, W. Stuetzle, Interactive data visualization using focusing and linking, in: Proceeding Visualization '91, 1991, pp. 156–163. doi:10.1109/VISUAL.1991.175794.

[28] M. Sereno, L. Besançon, T. Isenberg, Supporting Volumetric Data Visualization and Analysis by Combining Augmented Reality Visuals with Multi-Touch Input, EuroVis 2019 - Posters (2019) 3 pagesArtwork Size: 3 pages ISBN: 9783038680888 Publisher: The Eurographics Association Version Number: 021-023. doi:10.2312/EURP.20191136.
URL https://diglib.eg.org/handle/10.2312/eurp20191136

[29] E. A. D. Santos, J. Hong, T. Isenberg, A design space for linked 2d and 3d visual representations (2022).
URL https://inria.hal.science/hal-03763560

[30] A. Assor, A. Prouzeau, P. Dragicevic, M. Hachet, Exploring Augmented Reality Waste Data Representations for Eco Feedback., in: Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems, ACM, Hamburg Germany, 2023, pp. 1–4. doi:10.1145/3544549.3583905.
URL https://dl.acm.org/doi/10.1145/3544549.3583905

[31] C. Quijano-Chavez, N. Doerr, B. Lee, D. Schmalstieg, M. Sedlmair, Brushing and linking for situated analytics, in: 2024 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW), IEEE, 2024, pp. 597–603. doi:10.1109/VRW62533.2024.00116.
URL https://ieeexplore.ieee.org/document/10536364/

[32] N. Doerr, B. Lee, K. Baricova, D. Schmalstieg, M. Sedlmair, Visual highlighting for situated brushing and linking, Computer Graphics Forum 43 (3) (2024) e15105, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.15105. doi:10.1111/cgf.15105.
URL https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.15105

[33] M. Spur, V. Tourre, E. David, G. Moreau, P. Le Callet, Exploring multiple and coordinated views for multilayered geospatial data in virtual reality, Information 11 (9) (2020) 425. doi:10.3390/info11090425.
URL https://www.mdpi.com/2078-2489/11/9/425

[34] S. Baceviciute, A. Mottelson, T. Terkildsen, G. Makransky, Investigating representation of text and audio in educational vr using learning outcomes and eeg, in: Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI '20, Association for Computing Machinery, New York, NY, USA, 2020, p. 1–13. doi:10.1145/3313831.3376872.
URL https://doi.org/10.1145/3313831.3376872

[35] R. Rzayev, P. Ugnivenko, S. Graf, V. Schwind, N. Henze, Reading in vr: The effect of text presentation type and location, in: Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems, CHI '21, Association for Computing Machinery, New York, NY, USA, 2021, p. 1–10. doi:10.1145/3411764.3445606.
URL https://doi.org/10.1145/3411764.3445606

[36] A. Skulmowski, K. M. Xu, Understanding cognitive load in digital and online learning: a new perspective on extraneous cognitive load, Educational Psychology Review 34 (1) (2022) 171–196. doi:10.1007/s10648-021-09624-7.
URL https://doi.org/10.1007/s10648-021-09624-7

[37] Y. Hou, H. Cui, R. Chen, W. Zeng, Understanding the impact of referent design on scale perception in immersive data visualization, in: Extended Abstracts of the

CHI Conference on Human Factors in Computing Systems, 2024, pp. 1–7. arXiv:2403.16018[cs], doi:10.1145/3613905.3650783.
URL http://arxiv.org/abs/2403.16018

[38] Y. Jansen, P. Dragicevic, P. Isenberg, J. Alexander, A. Karnik, J. Kildal, S. Subramanian, K. Hornbæk, Opportunities and challenges for data physicalization, in: Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, ACM, 2015, pp. 3227–3236. doi:10.1145/2702123.2702180.
URL https://dl.acm.org/doi/10.1145/2702123.2702180

[39] Y. Jansen, P. Dragicevic, J.-D. Fekete, Evaluating the efficiency of physical visualizations, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, ACM, Paris France, 2013, pp. 2593–2602. doi:10.1145/2470654.2481359.
URL https://dl.acm.org/doi/10.1145/2470654.2481359

[40] F. Taher, Y. Jansen, J. Woodruff, J. Hardy, K. Hornbæk, J. Alexander, Investigating the Use of a Dynamic Physical Bar Chart for Data Exploration and Presentation, IEEE Transactions on Visualization and Computer Graphics 23 (1) (2017) 451–460. doi:10.1109/TVCG.2016.2598498.
URL https://ieeexplore.ieee.org/document/7542185/

[41] Y. Wang, A. Segal, R. Klatzky, D. F. Keefe, P. Isenberg, J. Hurtienne, E. Hornecker, T. Dwyer, S. Barrass, An Emotional Response to the Value of Visualization, IEEE Computer Graphics and Applications 39 (5) (2019) 8–17. doi:10.1109/MCG.2019.2923483.
URL https://ieeexplore.ieee.org/document/8809393/

[42] R. S. Renner, B. M. Velichkovsky, J. R. Helmert, The perception of egocentric distances in virtual environments - A review, ACM Computing Surveys 46 (2) (2013) 1–40. doi:10.1145/2543581.2543590.
URL https://dl.acm.org/doi/10.1145/2543581.2543590

[43] E. Ebrahimi, L. S. Hartman, A. Robb, C. C. Pagano, S. V. Babu, Investigating the Effects of Anthropomorphic Fidelity of Self-Avatars on Near Field Depth Perception in Immersive Virtual Environments, in: 2018 IEEE Conference on Virtual Reality and 3D User Interfaces (VR), 2018, pp. 1–8. doi:10.1109/VR.2018.8446539.
URL https://ieeexplore.ieee.org/document/8446539/

[44] S. Jung, G. Bruder, P. J. Wisniewski, C. Sandor, C. E. Hughes, Over My Hand: Using a Personalized Hand in VR to Improve Object Size Estimation, Body Ownership, and Presence, in: Proceedings of the Symposium on Spatial User Interaction, ACM, Berlin Germany, 2018, pp. 60–68. doi:10.1145/3267782.3267920.
URL https://dl.acm.org/doi/10.1145/3267782.3267920

[45] S. A. Linkenauger, M. Leyrer, H. H. Bülthoff, B. J. Mohler, Welcome to Wonderland: The Influence of the Size and Shape of a Virtual Hand On the Perceived Size and Shape of Virtual Objects, PLoS ONE 8 (7) (2013) e68594. doi:10.1371/journal.pone.0068594.
URL https://dx.plos.org/10.1371/journal.pone.0068594

[46] F. Chevalier, R. Vuillemot, G. Gali, Using concrete scales: A practical framework for effective visual depiction of complex measures, IEEE Transactions on Visualization and Computer Graphics 19 (12) (2013) 2426–2435. doi:10.1109/TVCG.2013.210.
URL https://ieeexplore.ieee.org/document/6634143/

[47] J. Hullman, Y.-S. Kim, F. Nguyen, L. Speers, M. Agrawala, Improving comprehension of measurements using concrete re-expression strategies, in: Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, ACM, 2018, pp. 1–12. doi:10.1145/3173574.3173608.
URL https://dl.acm.org/doi/10.1145/3173574.3173608

[48] The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations, VL ’96, IEEE Computer Society, USA, 1996.

[49] S. Wehrend, C. Lewis, A problem-oriented classification of visualization techniques, in: Proceedings of the First IEEE Conference on Visualization: Visualization ‘90, 1990, pp. 139–143. doi:10.1109/VISUAL.1990.146375.
URL https://ieeexplore.ieee.org/document/146375/

[50] Low-Level Components of Analytic Activity in Information Visualization, INFOVIS ’05, IEEE Computer Society, USA, 2005. doi:10.1109/INFOVIS.2005.24.
URL https://doi.org/10.1109/INFOVIS.2005.24

[51] M. Brehmer, T. Munzner, A multi-level typology of abstract visualization tasks, IEEE Transactions on Visualization and Computer Graphics 19 (12) (2013) 2376–2385. doi:10.1109/TVCG.2013.124.
URL http://ieeexplore.ieee.org/document/6634168/

[52] J. S. Yi, Y. A. Kang, J. Stasko, J. Jacko, Toward a deeper understanding of the role of interaction in information visualization, IEEE Transactions on Visualization and Computer Graphics 13 (6) (2007) 1224–1231. doi:10.1109/TVCG.2007.70515.
URL https://ieeexplore.ieee.org/document/4376144/

[53] J. Heer, B. Shneiderman, Interactive dynamics for visual analysis, Communications of the ACM 55 (4) (2012) 45–54. doi:10.1145/2133806.2133821.
URL https://dl.acm.org/doi/10.1145/2133806.2133821

[54] K. Marriott, F. Schreiber, T. Dwyer, K. Klein, N. H. Riche, T. Itoh, W. Stuerzlinger, B. H. Thomas, Immersive analytics, Vol. 11190, Springer, 2018.

[55] S. Hubenschmid, J. Zagermann, S. Butscher, H. Reiterer, STREAM: Exploring the Combination of Spatially-Aware Tablets with Augmented Reality Head-Mounted Displays for Immersive Analytics, in: Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems, ACM, Yokohama Japan, 2021, pp. 1–14. doi:10.1145/3411764.3445298.
URL https://dl.acm.org/doi/10.1145/3411764.3445298

[56] F. Zhu, T. Grossman, BISHARE: Exploring bidirectional interactions between smartphones and head-mounted augmented reality, in: Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, ACM, 2020, pp. 1–14. doi:10.1145/3313831.3376233.
URL https://dl.acm.org/doi/10.1145/3313831.3376233

[57] R. Kenny, A. A. Becker, Is the nasdaq in another bubble? a virtual reality tour of the nasdaq, http://graphics.wsj.com/3d-nasdaq, accessed: 2025-07-23 (2015).

[58] S. G. Hart, L. E. Staveland, Development of nasa-tlx (task load index): Results of empirical and theoretical research, in: P. A. Hancock, N. Meshkati (Eds.), Human Mental Workload, Vol. 52 of Advances in Psychology, North-Holland, 1988, pp. 139–183. doi:https://doi.org/10.1016/S0166-4115(08)62386-9.
URL https://www.sciencedirect.com/science/article/pii/S0166411508623869

[59] R. A. Khot, J. Lee, D. Aggarwal, L. Hjorth, F. F. Mueller, TastyBeats: Designing Palatable Representations of Physical Activity, in: Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, ACM, Seoul Republic of Korea, 2015, pp. 2933–2942. doi:10.1145/2702123.2702197.
URL https://dl.acm.org/doi/10.1145/2702123.2702197

[60] D. A. Guttentag, Virtual reality: Applications and implications for tourism, Tourism Management 31 (5) (2010) 637–651. doi:10.1016/j.tourman.2009.07.003.
URL https://linkinghub.elsevier.com/retrieve/pii/S0261517709001332

[61] F. P. Brooks, M. Ouh-Young, J. J. Batter, P. Jerome Kilpatrick, Project GROPEHaptic displays for scientific visualization, ACM SIGGRAPH Computer Graphics 24 (4) (1990) 177–185. doi:10.1145/97880.97899.
URL https://dl.acm.org/doi/10.1145/97880.97899

[62] H. Romat, N. Henry Riche, C. Hurter, S. Drucker, F. Amini, K. Hinckley, Dear Pictograph: Investigating the Role of Personalization and Immersion for Consuming and Enjoying Visualizations, in: Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, ACM, Honolulu HI USA, 2020, pp. 1–13. doi:10.1145/3313831.3376348.
URL https://dl.acm.org/doi/10.1145/3313831.3376348