

Evaluating View Management for Situated Visualization in Web-based Handheld AR

A. Batch,¹ S. Shin,¹ J. Liu,¹ P. W. S. Butcher,² P. D. Ritsos,² and N. Elmqvist¹

¹University of Maryland, USA

²Bangor University, UK

Abstract

As visualization makes the leap to mobile and situated settings, where data is increasingly integrated with the physical world using mixed reality, there is a corresponding need for effectively managing the immersed user's view of situated visualizations. In this paper we present an analysis of view management techniques for situated 3D visualizations in handheld augmented reality: a shadowbox, a world-in-miniature metaphor, and an interactive tour. We validate these view management solutions through a concrete implementation of all techniques within a situated visualization framework built using a web-based augmented reality visualization toolkit, and present results from a user study in augmented reality accessed using handheld mobile devices.

CCS Concepts

• **Human-centered computing** → **User interface management systems**; **Visualization systems and tools**; **Visualization theory, concepts and paradigms**;

1. Introduction

Advances in mobile and wearable display interfaces, positional sensing, and computer graphics have fueled recent efforts of displaying data *in situ* [BKT*21]. Current research themes such as ubiquitous [EI13], immersive [CCC*15], and situated analytics [TWD*18] (IA/UA/SA) explore a world where contextual data is readily available at the fingertips of the user, anywhere and anytime [Elm23]. Particularly exciting is the topic of *situated visualization*, where data relevant to a place is visualized directly in its physical space [WF09,ETM*16,SLC*19]. However, there are several challenges in making such situated visualization practically useful, such as the intrinsic augmented/mixed reality (AR/MR) challenges of registration mechanisms, power consumption, device ergonomics, etc.; recent work surveyed the grand challenges of immersive analytics [EBC*21]. One such challenge is *view and session management* [BFH01] of situated visualizations: effectively interacting with situated visualizations that co-inhabit the user's physical space in AR. While we focus on handheld AR in this work, these ideas also apply to head-mounted display (HMD) systems.

We (1) present an analysis of the properties and challenges of view management for situated visualizations in Augmented Reality, enumerating concerns such as physical distance and reach, orientation and legibility, and depth and occlusion. These challenges apply to both the components of a single situated visualization, as well as to multiple visualizations that exist in the same physical environment. We then (2) revisit existing techniques from the domains of computer graphics, Virtual Reality (VR), and visualization to eval-

uate implementations of interaction, layout, and presentation techniques accessible in handheld AR. Our goal is to assess how these can be designed and implemented to mitigate our identified challenges. We investigate the following techniques: (1) a *shadowbox* that eliminates effects of perspective foreshortening and occlusion in 3D visualizations, including *unfolding* to transform a 3D visualization into orthographic 2D views; (2) a *world in miniature (WIM)* technique for overviewing and accessing multiple visualizations in a 3D environment; and (3) a *data tour* for guiding the user through a 3D situated analytics environment to visit all views of interest.

While each of our techniques are derived from existing work, we claim that their application and implementation as web-based situated visualizations in handheld AR is novel. We also report on findings from a user study where 12 participants performed situated analytics tasks using our proposed techniques (<https://osf.io/ma72f/>). While these results highlight many of the typical challenges associated with AR and sensemaking, we also found convincing evidence supporting our view management techniques.

2. Related Work

2.1. Immersive and Ubiquitous Analytics

The domain of immersive analytics (IA) investigates how immersive interaction and display technologies can be used to support analytical reasoning and decision making [CCC*15]. IA builds on advances in a variety of technologies, including multimodal interaction, modern fabrication techniques, VR, and MR/AR (often collectively denoted XR). As discussed by Ens et al. [EBC*21] in their

survey on grand IA challenges, the degree to which data visualizations are spatially integrated with physical referents and support spatial interaction has resulted in concepts such as ubiquitous analytics [EI13] in 2013 followed by immersive analytics [CCC*15] and situated analytics [TWD*18] in 2015. Here we adopt the term *situated* due to the relatively relaxed registration requirements as per this taxonomy. In particular, we do not concern ourselves with precise registration, but rather the depiction of the visualizations in space when investigating our view management techniques.

Our approach adheres to the definitions of AR and MR from Azuma [Azu97] and Milgram and Kishino [MK94]. However, we acknowledge the multi-faceted nature and varying interpretations of MR [SHN19], in particular its expanded view by Barba et al. [BMM12]. The latter fits well with the handheld Augmented Reality we use in our work and is influenced by a less popular definition of AR from Mackay [Mac98], which describes an environment augmented by interactive, networked objects. This is also consistent with Weiser's vision for ubiquitous computing [Wei91] and thus ubiquitous analytics [EI13, Elm23]. We believe that this definition supports the notion of situating data in the user's own environment, as it implies an interconnected world where data is generated and consumed by a variety of interlinked and georeferenced sources in addition to being accessed immersively. We believe this vision of visualization beyond the desktop [RRB*14] provides new opportunities for situated, ubiquitous, and immersive data displays.

2.2. Situated Analytics

Situated analytics (SA) [ETM*16, TWD*18] deals with MR views of information that visually link virtual and physical objects of interest, registering abstract information to spatial locations and supporting analytical interactions. While IA and SA are relatively new areas of research, there is a multitude of existing implementations that could be called "immersive analytics" systems, many of them also "situated." Some of these address universally challenging issues, such as labeling [GLK*12, MTM*16], highlighting [EST16], impact of real-world background on visualization perception [SD21], and synergy of HMDs and handhelds [LSBD21]. However, most SA systems tend to be specialized for a use case.

The advent of light, handheld, multi-functional devices (e.g., smartphones, tablets etc.) and relatively affordable HMDs has made AR/MR accessible to a larger audience. This, consequently, has aided the emergence of situated visualization systems for the general public, such as for tourism [CRF16], entertainment [BCD*12], shopping [ETM*16], and sports [LSY*21]. In many analytical disciplines, the setting is "in the field" rather than an office or laboratory. A popular application is decision-making processes and situational awareness for manufacturing, construction [BK05], agriculture [ZC19], and infrastructure and utilities [SMK*09]. In this space, Whitlock, Wu, and Szafir [WWS20] conduct a design probe involving expert users from five such disciplines to evaluate the needs and challenges of existing SA systems for data analysis and collection, and demonstrate their resulting design recommendations in their FieldView implementation.

Empirical grounding for IA and its expanding design space has begun to emerge from both VR/MR [WSS20] and visualization

communities [BSB*18, BCC*20]. In particular, the use of spatial visualizations in immersive settings has been surveyed in various recent efforts [BYK*21, BKT*21]. Zollman et al. [ZGL*20] present a taxonomy for visualization in AR based on many such examples, and use it to extend the traditional information visualization pipeline to situated implementations. They identify six recurring design dimensions in their AR visualization taxonomy: purpose (for using AR), visibility (vs. occluded or out of view), depth cues, abstraction (for reducing data complexity), filtering (to a subset of observations to reduce clutter), and compositing (the method for modifying the user's non-augmented view of reality). While our work may be comparable to Zollman et al. [ZGL*20] in that our conclusions are based on a compendium of existing implementations, we focus more narrowly on the view management end of the visualization pipeline, and target handheld devices in particular.

2.3. View Management in 3D

As discussed above, because of their common lineage and technological platform, situated visualizations share many of the same challenges as that of XR. This means that in our work we can draw inspiration from established techniques for view management in immersive environments. World-in-Miniature (WIM) [SCP95] presents the user with an interactive overview of the entire virtual scene. This enables the user to see objects that may be occluded by other features of the view, offers a sense of where the user is located within the scene, and allows them to manipulate objects that fall well outside their natural reach. Go-Go Interaction [PBW196] offers another solution to the issue of reach by—as the *Inspector Gadget* reference in the name implies—enabling the user to virtually extend the length of their arm at will to interact with remote objects. Elastic-Arm [AGT*15] is a multimodal version that uses haptic feedback to guide the navigation of an elastic arm; it curves instead of appearing in a straight line based on angle of pose, and can reach around corners to unseen parts of the scene.

Attempts to resolve the problem of 3D view depth perception distorting the user's judgment of object size and relative scale commonly lead to a reduction in the dimensionality of the view itself. Real-world metaphors often appear in design decisions related to immersive view management [LJKM*17]; this is demonstrated by one early technique where users interact with a 3D object via its 2D shadow [HZR*92]. A visualization-specific adaptation of this metaphor is seen in ExoVis [Tor03, TAK*05], which involves the use of a 2D projection of a 3D object on the exterior sides of an opaque 3D cube. Yet another approach to managing depth and distance perception as well as occlusion, is seen in the *cutting plane* [HPGK94], a technique for interactively creating a slice of a 3D render of spatial data. A different tack for managing both distance and depth perception within an immersive scene is the use of visual indicators of depth and distance in the form of annotations, shadow planes, or pictorial cues [WH05].

3. View Management for Handheld SA

A *situated visualization* is a data visualization that has been embedded into an AR or MR environment [Azu97] to support situated [SH16] analytics. Inevitably, making efficient use of such sit-

uated visualizations requires significantly more overhead, navigation, and layout considerations than for traditional visualizations drawn on a 2D screen. Bell et al. [BFH01] define *view management* for AR and VR as “maintaining visual constraints on the projections of objects on the view plane, such as locating related objects near each other, or preventing objects from occluding each other.” Drawing on this definition, we refer to *view management for situated visualization* as optimizing the user’s view of (and access to) the visualizations in an IA/SA environment.

While there are significant and valid concerns with using 3D visualizations in the first place [Mun14], these are mostly moot when discussing SA on MR devices. In such settings, the user is by definition mobile in the real world, usually using hand-held (or head-mounted) displays, and thus representing visualizations spatially is inevitable even if the visualizations themselves are not 3D.

We here derive design properties based on the different challenges of view management for situated visualization. These properties can act as a classifying framework for existing techniques, but can also be used as a generative lens for designing future situated views. Here we enumerate the basic properties of a visualization inhabiting a situated analytics environment [BFH01]:

- **Position:** The visualization’s location in the environment;
- **Size:** Its geometric size in relation to the rest of the world;
- **Transparency:** The opacity of the visualization, which also incorporates its general geometry (i.e., some visualizations such as a 3D scatterplot are more sparse than a volume rendering);
- **Priority:** A relative priority for each individual visualization (potentially whether a visualization is selected or not);
- **Orientation:** The visualization’s 3D rotation;
- **Distance:** Its distance from the viewer and other visualizations;
- **Area of interest:** The area (often a 3D volume) from which to optimally view the visualization (i.e., the location of the user within which a visualization becomes relevant);
- **Spatial relation to the surrounding world:** The visualization’s relation to real objects in the physical world (e.g., the proximity of a visualization of accident data to dangerous stairs); and
- **Relative motion:** Dynamic motion of the visualization with respect to the viewer, e.g., when one or both are moving [YBVI22].

The above list is by no means an exhaustive or even minimal one, as some properties are derivatives of others (distance vs. position, for example). However, it is useful to distinguish each of these properties individually, as they all give rise to specific challenges. Furthermore, some of these challenges apply within a single visualization (e.g., occlusion within the points in a 3D bar chart), whereas others apply for multiple visualizations (e.g., an overview for all of the visualizations in an environment), and some apply to both (e.g., the priority of one object within the view relative to others, or occlusion between marks, as well as occlusion between two visualizations). An additional set of challenges arise from the limitations and affordances offered by handhelds. Screen size, interaction capabilities, and ergonomics are very different for handheld compared to, say, HMDs, resulting in a unique mixture of limitations. Yet, the omnipresence of smartphones makes handheld MR an interesting and pragmatic approach. Setting aside limitations related to form factor (which we discuss in Section 5), we outline below the main challenges that arise from SV. Again, we make no effort to mak-

ing these challenges orthogonal, but instead list them individually because they add reasoning power to our argument.

Visibility: Maintaining visibility of a situated visualization in the user’s field of view is a fundamental challenge [BFH01]. Many situated visualizations are just that, situated in a specific location in the physical world, which means that they easily fall outside the user’s vision, either by being too far away or above, below, or behind the user. In such situations, the visualizations cannot be moved to always be visible, and other mechanisms must be employed to make the user aware of their existence and location. This challenge is compounded when multiple visualizations are jostling for space on the limited field of vision afforded by handheld devices.

Occlusion: The 3D nature of SA environments means that a geometric object can be hidden by other objects even if they do not intersect in 3D space [ET08]. The problem is further exacerbated when they do intersect. This fundamental challenge affects both marks within a single visualization, such as a cluster of marks in a 3D scatterplot occluding an outlier, as well as between multiple visualizations, such as a 3D volume occluding a distant barchart.

Overview: Overview is a central aspect of data visualizations [Shn96], but gaining an overview in an SA environment is particularly challenging because of its 3D nature. This is not merely about the ability to access and read an individual visualization, but being aware of its existence in the first place; a visualization that is fully occluded by other visualizations, outside the user’s field of view, or too far away to see, will inevitably not be included in the overview. This means that many of the below challenges contribute to the overall *Overview* challenge. Moreover, limited screen sizes and rendering capabilities of handhelds exacerbate the problem.

Perspective Foreshortening: A more subtle aspect of 3D is the impact of perspective foreshortening due to visualizations being at different distances from the viewer. Perspective foreshortening arises from the non-linear 3D perspective, essentially making nearer items disproportionately larger than more distant ones. Besides having an impact on the *Occlusion* challenge, it also makes it difficult to compare between two visualization marks at different distances, such as two different bars in a 3D bar chart.

Legibility: A particular concern for situated visualizations that are not rotation invariant or not always facing the user, such as a billboard, is legibility, especially for text. In such situations, the slanted or rotated view of the visualization makes reading more difficult or even impossible. Similar legibility concerns arise when a visualization is far away from the viewer, making graphical features in general—and text in particular—too small to distinguish.

Physical Navigation: When a situated visualization is too far away to be legible or manipulated, either the visualization or the user will need to move. When this task falls on the user, such as when the visualization cannot be moved from a specific geographic position or real-world object, this translates to the user physically having to navigate to the object of interest. Unlike in dedicated VR spaces, such as open labs or even CAVEs, such navigation can be particularly tricky in a physical environment filled with slippery or uneven surfaces, physical barriers, and other people.

Temporal and Spatial Continuity: Finally, as observed by Bell et al. [BFH01], it is important to maintain continuity over time

and in space so that visualizations do not “jump around” due to discontinuous layouts that are calculated independently from frame to frame. Thus, situated visualizations should be rendered smoothly over time and space, without positional or rotational drift.

→ **Physical Reach:** For handheld AR, users at best have a stylus or touch cursor to interact with the visualization directly. This means that a situated visualization can be located in a position that is not physically accessible, even to a person otherwise able to move.

4. Prototyping Situated View Management for Handheld MR

Here we present our selection of techniques and approaches (World-In-Miniature, Shadowboxes, guided tours) through the lens of our design exploration in Section 3. Because we conducted our evaluation during the pandemic, our techniques were chosen so that they can be evaluated remotely (or outdoors). For each technique, we present how it modifies the properties of the situated visualizations, and discuss the challenges addressed (Table 1).

Note that we do not claim that our proposed techniques are novel, only that their implementation and application to handheld MR is. Furthermore, while we do not claim that our list is exhaustive for all view management situations, our choice of techniques was driven by the need to support all of the challenges we have identified in Section 3 (as evidenced by Table 1).

Table 1: Challenges. Challenges addressed by each technique.

Challenge \ Technique	🏠	📦	🗺️
Visibility	✓	✓	✓
Occlusion	✓	✓	✓
Overview	✓	✓	✓
Perspective Foreshortening		✓	
Legibility		✓	✓
Physical Navigation	ⓘ		✓
Physical Reach	ⓘ		✓
Temporal & Spatial Continuity	ⓘ		ⓘ

🏠 WIM 📦 ShadowBox 🗺️ Data Tour
 ⓘ partially or indirectly addressed

4.1. World-in-Miniature 🏠

Accessing all components of a 3D scene can be challenging when viewing the scene *in situ* from a first-person perspective. We adapt the World-In-Miniature [SCP95] technique—a miniaturized 3D view of the world that is controlled by the user, allowing them to see their surroundings from a third-person view—to give the user a third-person view of their 3D environment.

We apply the basic WIM approach to a situated visualization environment where the WIM is instantiated by the user and is represented by a box containing a miniature representation of all virtual features in the scene (including the user). The WIM may include contextual information such as a mesh of the landscape as detected by the device, or a map tile layer based on the user’s GPS coordinates. Finally, the WIM, as an object in the user’s view itself, has its own properties, in addition to affecting the properties of the situated visualizations. The visualization properties affected are:




Figure 1: World-in-miniature. Sketch of an example world-in-miniature (WIM) scenario.

- **Position:** Situated visualizations can be moved by dragging them in the WIM. Their positions within the WIM also reflect their relative positions in the world.
- **Size:** The WIM duplicates all situated visualizations at a significantly smaller scale within a space.
- **Transparency:** Making WIM elements semi-transparent allows the user to see the real world as well as spot virtual elements that are occluded by other virtual elements.
- **Orientation:** The orientation a situated visualization in the WIM should reflect that of its true orientation in the world.
- **Distance:** The relative distance between the user and situated objects is represented accurately by the WIM.
- **Spatial relation to the surrounding world:** The WIM allows for decoupling situated visualizations from the real world.

The WIM technique directly addresses the challenges of *Overview*, *Visibility*, and *Occlusion* by creating miniature copies of all objects in the scene, including those not visible to the user, and by giving the user the freedom to rotate the scene to discover and access hidden content. Furthermore, this also eliminates the need for *Physical Navigation* and *Physical Reach*, as the miniature allows for easy navigation and access. However, the technique does not address *Perspective Foreshortening* and *Legibility*, but can be designed to respect *Temporal and Spatial Continuity* by synchronizing the positions of the miniature visualizations with those of their true situated counterparts.

4.2. Shadowbox 📦

The 3D perspective and nature of a situated visualization can make it hard to efficiently and correctly access its data. A Shadowbox puts a given 3D object (situated visualization) inside a virtual “display case” represented as a 3D box, with 2D orthogonal projections of the object on each of its faces, and support for unfolding the box.

Our Shadowbox encapsulates the visualization inside a display case that makes viewing the data easier and more accurate. In this way, the Shadowbox is similar to ExoVis [TAK*05] but presents the user with either an exterior or interior view of the box and enables hiding the 3D object to mitigate occlusion. While variants of ExoVis feature a slicing mechanism, it does not present a flattened silhouette view of all features of the figure from six angles fixed with respect to the object. Therefore, we claim that the Shadowbox is a novel technique. We also propose “unfolding” the sides of the box to align multiple 2D projections in one plane, showing all projections at once (Figure 2c). The properties affected are:

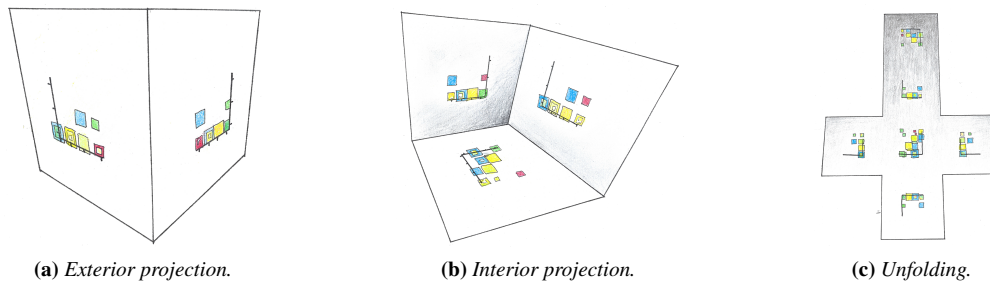


Figure 2: Shadowbox. (a) exterior and (b) interior projection modes, as well as (c) unfolding interaction.

- *Area of interest:* The Shadowbox provides optimal views of a situated visualization along each of the primary axes.
- *Spatial relation to the surrounding world:* The Shadowbox has the side-effect that it isolates and separates the situated visualization being examined from the rest of the world.

The Shadowbox was primarily designed to manage the *Perspective Foreshortening* typical in 3D environments by using an orthographic projection for each of the planes. It enables exact visual comparison for, e.g., the bars in a 3D bar chart. However, it can also help aid *Visibility* and *Occlusion* as well as support *Overview* of the 3D object. The axis alignment can also facilitate *Legibility*.

4.3. Data Tour 🗺️

Sometimes view management limits the awareness of the available points of interest. A Data Tour, i.e., a guided walk through all points of interest in an environment, can address this issue. The algorithm for placing navigation cues may vary. A simplified set of steps for implementing a data tour are as follows:

1. Select all nodes not already visited from the user.
2. Render a visual guide in the user's field of view to the currently nearest point of interest.
3. Once the user comes in close proximity to the point of interest, add it to the list of visited points and repeat from Step 1.

A more sophisticated implementation would implement actual wayfinding into the system, thus taking advantage of street-level maps and blueprints to guide the user on the most optimal path. In this case, the properties affected are:

- *Position:* Visualizations are not moved; instead, the user visits the visualizations through physical navigation.
- *Distance:* The Data Tour tries to bring each point of interest within optimal distance to the user over the entire tour.
- *Spatial relation to the surrounding world:* Significantly, the Data Tour preserves the spatial location and mapping of each situated visualization to the physical world.

The Data Tour facilitates accessing situated visualizations to avoid *Occlusion* and support *Visibility* by guiding the user's *Physical Navigation*. This also means that the *Legibility* and *Physical Reach* challenges can be addressed by guiding the user to the objects. In particular, the Data Tour provides *Temporal and Spatial Continuity* since it does not alter the environment at all.

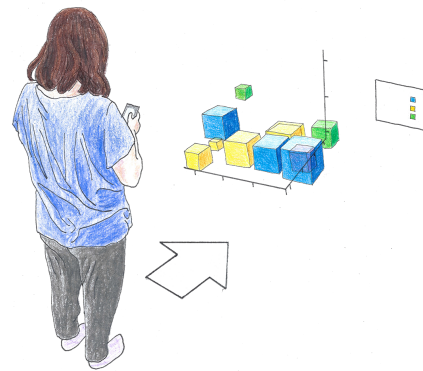


Figure 3: Guided tour. A user is guided to a location of interest by an arrow overlay positioned on the ground in front of her.

5. User Study

The goal of our study was to gauge how people might use a selection of our view management techniques described in Section 4 for navigation and analysis tasks in handheld Augmented Reality. We conducted an exploratory study on participant time and space usage, correctness for basic analysis tasks, and self-reported measures of task-related user experience factors. Since our goal was primarily to assess the general utility of view management techniques, we opted for a relatively small sample of participants as well as generic tasks with no specific real-world application. We provide all of the study materials here: <https://osf.io/ma72f/>

5.1. Prototype Implementation

We implemented all techniques described in this paper as a unified system demonstrating their synthesis using `three.js`, `VRIA` [BJR21], and `AFrame-React`. `VRIA` was used to create “staged” visualizations—the initial collection of 3D visualizations instantiated upon loading the page. We deployed this testing platform on Android phones using Mozilla’s WebXR viewer. We were unable to deploy on iPhones because at the time of our writing this paper, the only WebXR viewer available on Apple iPhones—Mozilla’s WebXR Viewer—was no longer being maintained and its final version suffered from major performance issues with one of the libraries upon which our implementation depends.

We use a MERN (MongoDB, Express, Angular, Node) stack as a backend. We set up an Express NodeJS server API endpoint that

posts to a MongoDB database to log user session events (navigation, multiple choice question responses), position and orientation during their experiment sessions. When the scene is initialized, an ID is assigned to the session and posted to the server from the client via the API. Each one-second interval after the session is posted, the client interface posts the user's scene camera position and rotation along with their session ID, and the server updates the database collection with these features, the server timestamp, and the user's IP. Whenever the user answers a question, the client interface posts the user's answer, the possible answers (with the correct answer indicated), the POI and task it corresponds to.

5.2. Participants

We recruited 12 participants (hereafter referred to as "users") with professional backgrounds in user experience, design, and interface or systems development, maintenance, and engineering, as this profile fits the intended users of this work. Moreover, all users were screened to possess AR-compatible Android mobile devices.

5.3. Apparatus and Data

Our study was conducted exclusively using handheld AR running on Android smartphones. All participants used their own devices, which we ensured were sufficiently new to fully support Mozilla's WebXR viewer that our implementation depended upon.

The questions in our experiment all correspond to features of a synthetic public safety dataset. This dataset comprised spatially-tagged falling accident events on a university campus with features of the time and setting in which the accident took place (weather, season, year), of the victim (gender, role within the institution), and of the accident itself (injury severity).

User experience survey data was collected via Google Forms following the session task completion experiment. The survey structure was based heavily on the NASA Task Load Index (TLX) [HS88]. All questions from the NASA TLX were asked for each technique individually, and the survey ended with two ranked voting questions—one for the techniques used for navigation and one for the techniques used for analysis—and three open-ended text entry questions about the user's experience. NASA TLX questions and scales were used verbatim in our survey, with one exception: We flipped the scale of the question "How successful were you in accomplishing what you were asked to do?" prior to conducting our formal study after several pilot users reported misinterpreting the scale's order for that question specifically. Qualitative data was also collected in the form of researcher notes, video and audio recording, and transcripts from the session.

5.4. Experimental Design and Procedures

We conducted user sessions using video conferencing on Zoom lasting approximately 30 minutes to 1 hour during which the user was given a summary description of what to expect during the session. The researchers verbally confirmed that the user was aware that video and audio recording would be collected throughout the session; the users were also informed once the recording begins. Users were asked to stand in the middle of the space they would be

using during the session, and to visit a randomly-generated URL for the experiment implementation using the Chrome browser on their mobile device. The experiment implementation prompted the user to find their way to a point-of-interest (POI) using one of three techniques represented in AR in their living space through their mobile device, and once they had arrived at their destination, to answer a multiple choice question using either a shadowbox technique or the VRIA figure alone. These tasks and the order in which they were encountered by users are described in Section 5.5.

Once the users completed this series of sequences, they were redirected to a page informing them that the study was complete and providing them with a link to the survey. The users were asked to complete their survey while they were on the video call with the researcher present and to discuss any final thoughts about the session that they did not feel were captured by the survey. After the user completed the survey and discussed any additional feedback about the techniques they wished to provide, the video and audio recording were halted and the session was ended. All user sessions were conducted within the span of four days.

5.5. Tasks

We asked users to perform a series of navigation tasks (see Figures 4e, 4d, and 4f), each of which initialized a sequence of analysis tasks in which the user was required to answer multiple choice questions (see Figures 4e, 4b, 4a, and 4c) about an observation in the synthetic dataset described in Section 5.3. The sequence and permutations of tasks requested of the user are detailed in Table 2.

The techniques used for navigation tasks include a guided data tour using an arrow on the floor (Figure 3), the WIM (Figure 1), and a control condition in which a blue ring marker was rendered directly beneath the target POI. When the user's viewport reached a position within 0.5 meters of the POI, the user was prompted by the implementation with the first of a series of three multiple choice questions about variables represented via the position of square marks relative to three axes, each question referencing a different variable from the synthetic dataset. This sequence—navigate, then answer three multiple choice questions—was repeated six times.

The multiple choice questions during the first four repetitions referred to only one analysis technique per repetition. The final two repetitions cycled through the analytical techniques, with a different technique being applied for each question. For the first two repetitions, the guided tour arrow (Figures 3, 4f) was used for the navigation task; the following two repetitions used the WIM (Figures 1, 4d); the final two questions used the control condition of a blue ring beneath the target point of interest (Figure 4e). The first of the four analysis conditions used was the control condition of a 3D VRIA chart (Figure 4e). This sequence was followed by another sequence using the exterior wall projection view of the Shadowbox (Figures 2a, 4a), followed by the interior wall projection (Figures 2b, 4b), followed by the unfolded view (Figures 2c and 4c).

The WIM and guided data tour both address or partly address the challenges of physical navigation, physical reach, and spatial/temporal continuity (Table 1), and so they meet the criteria of being appropriate for navigation tasks. Our navigation control condition—the use of a hovering ring beneath the navigation

Table 2: Sequence of tasks.

Step	Task Type	Variable	Technique	POI	Correct	Incorrect
1	Navigation	lat/longitude	📍	10	POI 10	
2	Analysis	injury severity (0-5)	📖	10	3	1, 2, 4
3	Analysis	weather	☁️	10	Clear	Snow, Sleet, Rain
4	Analysis	year	📅	10	2016	2017, 2018, 2019
5	Analysis	role	👤	10	Undergrad	Faculty, Grad, Staff
6	Navigation	lat/longitude	📍	11	POI 11	
7	Analysis	injury severity (0-5)	📖	11	4	1, 2, 3
8	Analysis	weather	☁️	11	Sleet	Snow, Clear, Rain
9	Analysis	year	📅	11	2018	2016, 2017, 2019
10	Analysis	role	👤	11	Undergrad	Faculty, Grad, Staff
11	Navigation	lat/longitude	📍	1	POI 1	
12	Analysis	injury severity (0-5)	📖	1	0	1, 2, 3
13	Analysis	weather	☁️	1	Sleet	Snow, Clear, Rain
14	Analysis	year	📅	1	2019	2016, 2017, 2018
15	Analysis	role	👤	1	Undergrad	Faculty, Grad, Staff
16	Navigation	lat/longitude	📍	9	POI 9	
17	Analysis	injury severity (0-5)	📖	9	1	0, 2, 3
18	Analysis	weather	☁️	9	Clear	Snow, Sleet, Rain
19	Analysis	year	📅	9	2019	2016, 2017, 2018
20	Analysis	role	👤	9	Grad	Faculty, Undergrad, Staff
21	Navigation	lat/longitude	📍	3	POI 3	
22	Analysis	injury severity (0-5)	📖	3	2	1, 3, 4
23	Analysis	weather	☁️	3	Snow	Clear, Sleet, Rain
24	Analysis	year	📅	3	2019	2016, 2017, 2018
25	Analysis	role	👤	3	Undergrad	Faculty, Grad, Staff
26	Navigation	lat/longitude	📍	4	POI 4	
27	Analysis	injury severity (0-5)	📖	4	2	1, 3, 4
28	Analysis	weather	☁️	4	Snow	Clear, Sleet, Rain
29	Analysis	year	📅	4	2018	2016, 2017, 2019
30	Analysis	role	👤	4	Grad	Faculty, Undergrad, Staff

📍 Hovering Marker (Control) 👤 WIM 📍 Data Tour
 📖 VRIA Chart (Control) 📖 ShadowBox (Folded, Interior Projection)
 📖 ShadowBox (Folded, Inverted/Exterior Projection) 📖 ShadowBox (Unfolded)

target—met the criteria for indicating that a point in space was a target of interest, but we did not feel that it was a good candidate for inclusion in our list of techniques by itself (Section 4), because it is more a standalone mark than a full-fledged technique. Finally, we chose the three states of the Shadowbox as our sole test condition as the Shadowbox is arguably a novel technique contributed by this work, whereas additional test conditions would push the task load for our users from reasonable to onerous.

In practice, we found that the time required by users to complete the tasks did indeed tend to hit near the maximum amount of time they were willing to commit for several users. The control condition for analytical tasks of having the user respond to questions using a VRIA figure, rather than a shadowbox, represents what we believe to be a reasonable default of using the targets of the Shadowboxes’ projections and removing the Shadowboxes.

5.6. Results

Each user was exposed to each technique multiple times, with a reasonable and clear improvement between task performance during the tutorial attempt and following task completion attempts. For this reason, we opted to evaluate observations after the tutorial for each technique. Despite the guided tour arrow technique being users’ very first method for navigation in the MR scene, users were able to locate the target POI significantly faster than when using the

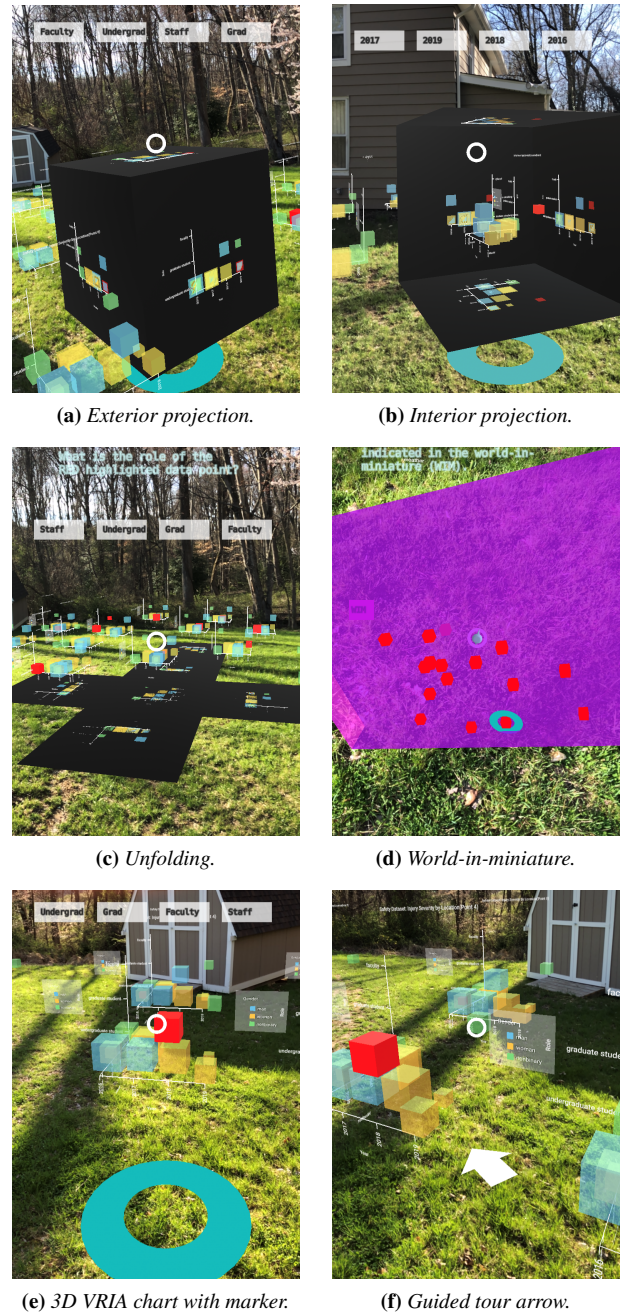


Figure 4: View management implementations. Based on 3D VRIA figures (e, f) and a Shadowbox with exterior and interior projection modes (a and b), as well as the unfolding technique (c), for the analysis stage. The WIM (d), target marker (e), and guided tour (f) were all used for navigation.







other two techniques (Figure 5a). Users also reported feeling most confident in their success at the navigation task when they were using the guided tour, and generally did not feel stressed, under time pressure, or as if they had to physically or mentally overexert themselves while using the technique. Users took significantly longer

finding their way to the target POI using the WIM; like the guided tour, the task completion time matches the users' responses to the NASA TLX, where they reported feeling least successful and generally negatively about the WIM's application to navigation tasks.





Correctness was slightly higher when using the Shadowbox's folded view with 2D charts projected onto the interior walls of the box (Figure 2b) relative to 3D charts, although the inverted, exterior projection had a long right tail, with several users correctly answering all questions using the exterior projection. The unfolded Shadowbox performed poorly relative to other techniques, with users answering fewer questions correctly using this view than other views. Nevertheless, despite its poor performance in correctness, the unfolded view did see a faster response time than all other analysis techniques (Figure 5c), followed by the interior projection and the exterior-wall projection views (roughly tied, with the interior projection having slightly faster mean times but a long tail of slower responses), while 3D view responses took significantly longer.

A Friedman test for duration in analysis tasks returned a value of 5.57 over $k = 30$ trials per technique with a p-value of 0.135, and a Friedman test for duration in navigation tasks with $k=18$ trials per technique returned a value of 2.33 with a p-value of 0.311, and so we cannot reject the null hypotheses that there is no difference in duration between at least two of the techniques at a 5% significance level. The Friedman test for correctness yielded a more encouraging value of 8.02 over $k = 30$ trials with a p-value of 0.0458, indicating that we can reject the null. For a more detailed investigation of the differences in correctness by technique, we performed a series of Wilcoxon signed-rank tests between the correctness of each of the techniques (Table 3). Significant differences were present between the unfolded ShadowBox and the folded ShadowBox and between the unfolded ShadowBox and the control (3D VRIA charts).

Table 3: Wilcoxon signed-rank tests of correctness by technique.

			
	21.0 ($p=0.052$)		
	34.0 ($p=0.046$)*	56.0 ($p=0.80$)	
	25.5 ($p=0.012$)*	48.0 ($p=0.44$)	22.0 ($p=0.53$)

* Significant at 5% significance level.

 VRIA Chart (Control)  ShadowBox (Folded, Interior Projection)
 ShadowBox (Folded, Inverted/Exterior Projection)  ShadowBox (Unfolded)

Participants' use of space is summarized in Figure 5d. Broadly, users preferred viewing the POIs with their viewport at a height between 1.2 and 1.5 meters. The exception to this is most notably the unfolded Shadowbox, which saw users' point of view angle diverge dramatically from that of all other techniques; they raised their devices higher while answering questions using the unfolded Shadowbox than they did with other devices. They also tended to view the unfolded Shadowbox at a greater distance; this was particularly true during their first sequence of questions using the technique, and then users tended to move closer to the center of the POI for later attempts at interpreting dataset values using this technique. In conjunction with the generally poor correctness (Figure 5b) associated with the unfolded Shadowbox, along with user feedback, we speculate that the reason for this was that users suffered difficulty in seeing the panels of the unfolded Shadowbox; they also encour-

tered minor occlusion issues as the remaining 3D VRIA objects were left in the view during this stage of the experiment.

Conversely, users tended to prefer viewing interior wall projection Shadowboxes at a closer distance than the other techniques, and moved closer to the POI during the first set of questions using the technique, but then spent more time farther away from the POI during the final question using this technique. In fact, this pattern of dramatically changing the distance and then reverting to a distance more similar to the one observed during the first sequence appears in all techniques except for the unfolded Shadowbox. When asked to rank their preferred techniques for navigation, users responded resoundingly against the WIM, with the target marker coming in slightly behind the guided tour as most preferred. When asked to rank their preferred techniques for analysis, users preferred the 3D chart, and—somewhat surprisingly—gave the interior projection view of the Shadowbox the fewest votes.

In the NASA TLX (see Figures 6a and 6b), users reported that the guided tour left them feeling less stressed and irritated, less pressed for time, and required less physical or mental exertion than the other navigation techniques, although it did see competition from the ring marker control condition. One user volunteered that they found the guided tour superior for finding their way to the target POI, but the ring marker did a better job of helping them pick the right POI when multiple POIs were situated near each other. On the other hand, other users did find the ring markers easy to pick out or fun to hunt down. The general consensus among users was that the WIM was not very effective for navigation, but that a large part of the problem was the mechanism by which the users controlled its rotation (namely, via the orientation of the phone). Users reported feeling the most successful in their task completion while using the guided tour, and the accuracy of this sentiment is strongly reflected in the task completion times shown in Figure 5a.

In general, users seemed to feel most comfortable with the 3D VRIA chart visualizations acting as our control condition, which they reported as requiring less mental exertion, and inflicting less stress and irritation upon them relative to other techniques. The 3D VRIA charts performed generally well in the users' survey responses for most categories, and they reported feeling most successful using this technique—despite answering more questions correctly using the interior wall projection view of the Shadowbox, and taking longer to complete their tasks using the 3D chart than using any other technique. This may be a result of the view itself being somewhat more common relative to the Shadowbox views. Despite this, users did report that the 3D charts required more physical exertion than any other technique, while the interior wall view of the Shadowbox required the least. They also reported that the interior wall projection made them feel least pressed for time. The unfolded Shadowbox was received negatively.

The Kruskal-Wallis H test yielded no significant values for either the navigation nor the analysis tasks for the TLX measures presented in Figure 6, with the highest test statistic at 3.8 ($p = 0.15$). When asked to directly compare techniques, however, the Kruskal-Wallis H test did yield significant values for rankings of both navigation ($H = 64.03$, $p < 0.001$) and analysis ($H = 64.75$, $p < 0.001$) techniques. There were no significant differences in the mean rank using the Games-Howell Test for analysis task rankings by respon-

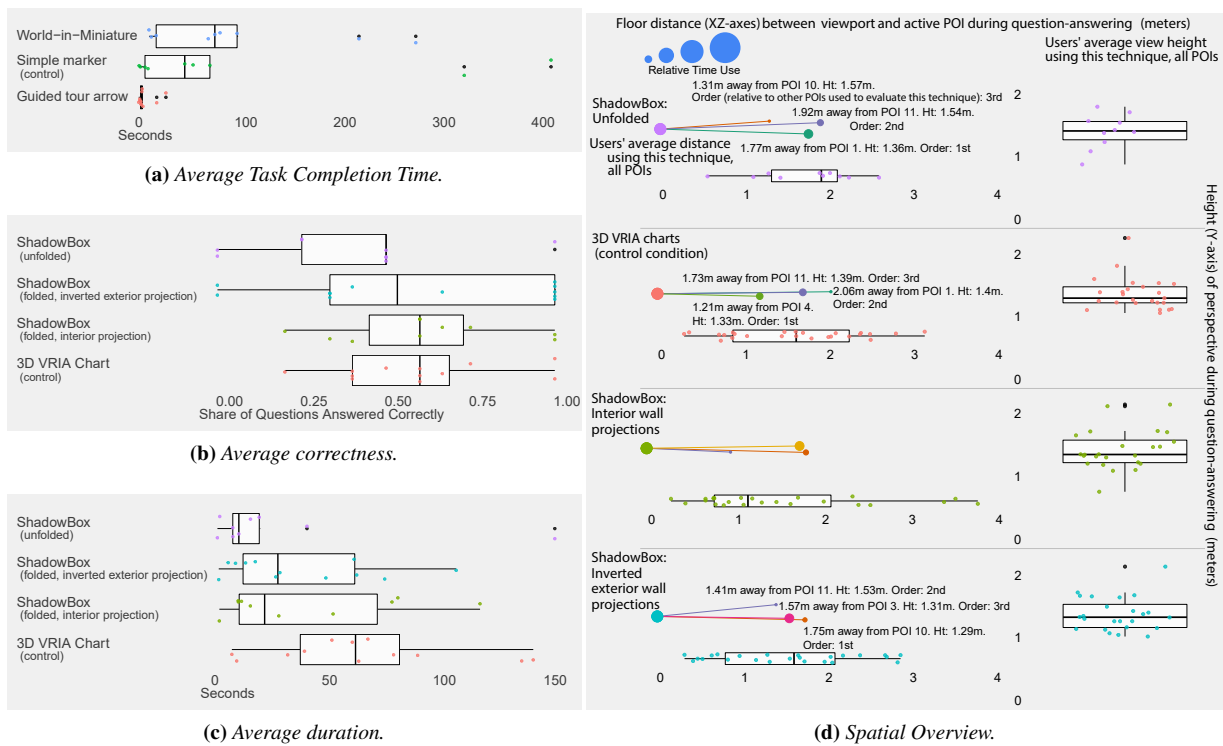


Figure 5: Distribution of user’s average task completion time for navigation tasks by technique (a), excluding the first navigation task using each technique. The white area of the box plots begin (left) at the 25th percentile, are split at the 50th percentile, and end at the 75th percentile. The whiskers extend from the 25th and 75th percentile hinges to the farthest observation within 1.5 times the inter-quartile range of either hinge Distribution of average correctness (b) and task completion time (c) for each user by technique, excluding the first attempted question response using each technique (steps 2 through 5 of Table 2). Distributions of user means in viewport heights, distance from the active point of interest (POI) during the question-answering period for each analysis technique, with mean distances by POI (d). Mean user-POI distances, heights, and the order in which the user interacted with each POI annotated by mark indicating distance.

dents, while the results for navigation tasks (Table 4) can be interpreted as meaning that the WIM was significantly ranked worse relative to the other techniques.

both navigation ($H = 64.03, p < 0.001$) and analysis ($H = 64.75, p < 0.001$) techniques.

Table 4: Games-Howell Test of Difference in Mean Ranks.

Groups	Mean Diff.	Std. Err.	t-value	p-value
	-0.357	0.184	1.376	0.369
	0.786	0.167	3.328	0.007*
	1.143	0.180	4.487	0.001*

* Significant at 5% significance level.

Hovering Marker (Control) WIM Data Tour

Finally, during our sessions, users commented on overall fun and enjoyment felt. In response to open-ended questions about their experience, four users described the ability to move around the scene as “fun” without being prompted, saying that “[it] was fun to navigate”, “I like [being able to move and look at charts] from all around—such fun!”, “hunting around for [the target ring marker] actually added some fun to the session.”, and “AR mode was fun!”.

When asked to directly compare techniques, however, the Kruskal-Wallis H test did yield significant values for rankings of

6. Discussion

There were a number of details in our results that bear discussion. One such detail was in the participant use of space. We suspect that there may be several factors at play in explaining the pattern of the differences in users’ view distance during the second task discussed in Section 5.6. The difference from first to second sequence may be the result of users becoming more competent with the system by the second sequence, correcting view distance issues encountered during the first sequence. The difference from second to third sequence may be a result of the user, now a seasoned expert, feeling free to take a leisurely stroll around the scene. Given the potential for confounds, we examined the data for outliers, but after reproducing Figure 5d several times with one potential outlier omitted each time, we found that the results did not significantly diverge from Figure 5d, which includes all users.

A different issue related to the remote nature of the study was user interpretation of subjective survey measures that could possibly have been avoided if the researchers and users had been sharing a room, as it may have been easier to notice the discrepancy be-

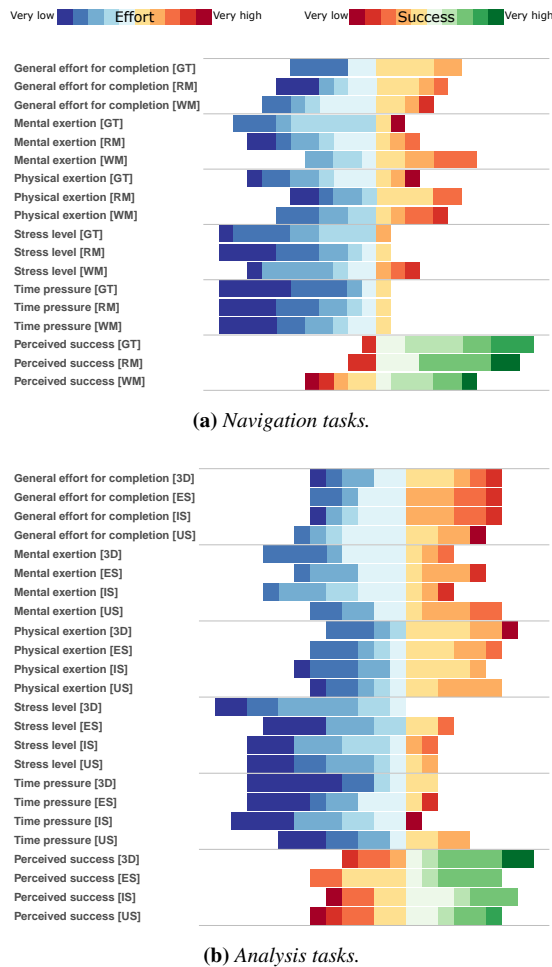


Figure 6: Survey responses for navigation and analysis tasks based on NASA TLX. A color legend is provided at the top. Color intensity represents response intensity, bar length indicates the share of users in each category. Abbreviations: [GT]: Guided Tour; [RM]: Ring Marker (control condition); [WM]: World-in-Miniature.

tween user responses relative to user remarks and researcher observations. The brief window of time during which user studies were conducted, and lack of opportunity for a chance observation, precluded a thorough review of survey responses as they were completed. Future remote research in WebXR should mind such pitfalls in data collection, but we are enthusiastic about its potential.

6.1. Generalizing the Results

We originally designed the techniques in this paper for handheld AR. However, while this study was restricted to handheld smartphones, many of these techniques are also applicable to HMD users, such as those using a Microsoft HoloLens 2. In fact, we suspect that the increased immersion and situation awareness arising from an HMD will enable better performance for many of these techniques. For one thing, when using an HMD, the user can simply look around rather than move their hand to change their viewpoint.

For another, they will not have to dedicate their hand to holding their device, leaving both hands free for rich 3D interaction.

Our participant population for the user study was a convenience sample. The study was conducted at the height of the COVID-19 pandemic and we faced significant challenges in recruiting participants during lockdown, even for a remote study such as this. The uniform and tech-savvy nature of our participants was almost a necessity in order to find people with Android phones and the ability to get the Mozilla WebXR browser running to enable them to run our software. We feel that the participant pool we used is adequate for a preliminary user evaluation of this nature. However, if our findings are to generalize, we will need to invite a broader cross-section of AR users to participate in future studies.





6.2. Limitations

Our study has several limitations that may impact our findings. First, we chose to study handheld AR devices because of the relative unavailability of consumer-facing AR HMD devices. Our focus on professionals willing to spend an hour of their time completing a user session restricts our population further. Second, the state of pandemic lockdown at the time that this study was conducted made in-person sessions impossible due to safety and ethical concerns, as well as organizational policy. Consequently, we were forced to limit the study to mobile users and simulated a situated view, and did not link the virtual objects with real-world ones.

There is always the possibility of measurement error causing some of the phenomena observed in this paper, and this may likely be the result of the limitations of the WebXR tools used in our implementation. However, our implementation represents the current state of the art in AR-in-the-browser applications, and it has several limitations. For example, when the user initializes an AR session, the WebXR scene camera no longer shares user position.

We also note that the generalizability of our results is somewhat curtailed by the small scale of our sample as well as the lack of real-world tasks and applications. Future experiments should involve more participants performing view management during actual immersive sensemaking tasks involving real data.

7. Conclusions

Our results indicate that the interior projection view of the **Shadowbox**  outperforms other techniques when the three factors of completion time, correctness of interpreting data, and user experience are all taken into consideration. It is closely followed by the 3D VRIA charts in correctness, slightly outperformed by 3D charts in the survey, and, like every other technique we have evaluated, beaten outright in task completion time by the unfolded **Shadowbox** . Our results also indicate that the **Guided Tour**  arrow outperforms the other techniques when completion time and user experience are taken into consideration, given that it received the most favorable ratings in most user navigation time was significantly faster using this technique than the other two evaluated in this study. The **WIM**  did not perform well for user navigation, but part of the blame for this result lays with the rotation mechanism in our implementation. Our results also indicated, anecdotally, that users find data exploration enjoyable in AR.

Acknowledgments

Special thanks to Chloe Batch for the illustrations. This work was supported partly by grant IIS-1901485 from the U.S. National Science Foundation, and the DSP Centre funded by the ERDF through the Welsh Government. Any opinions, findings, and conclusions or recommendations expressed here are those of the authors and do not necessarily reflect the views of the funding agencies.

References

- [AGT*15] ACHIBET M., GIRARD A., TALVAS A., MARCHAL M., LÉCUYER A.: Elastic-Arm: Human-scale passive haptic feedback for augmenting interaction and perception in virtual environments. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces* (Piscataway, NJ, USA, 2015), IEEE, pp. 63–68. doi:10.1109/VR.2015.7223325. 2
- [Azu97] AZUMA R. T.: A survey of augmented reality. *Presence: Teleoperators and Virtual Environments* 6, 4 (1997), 355–385. doi:10.1162/pres.1997.6.4.355. 2
- [BCC*20] BATCH A., CUNNINGHAM A., CORDEIL M., ELMQVIST N., DWYER T., THOMAS B. H., MARRIOTT K.: There is no spoon: Evaluating performance, space use, and presence with expert domain users in immersive analytics. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (2020), 536–546. doi:10.1109/TVCG.2019.2934803. 2
- [BCD*12] BALDUINI M., CELINO I., DELL’AGLIO D., DELLA VALLE E., HUANG Y., LEE T., KIM S.-H., TRESP V.: Bottari: An augmented reality mobile application to deliver personalized and location-based recommendations by continuous analysis of social media streams. *Journal of Web Semantics* 16 (2012), 33–41. doi:https://doi.org/10.1016/j.websem.2012.06.004. 2
- [BFH01] BELL B., FEINER S., HÖLLERER T.: View management for virtual and augmented reality. In *Proceedings of the ACM Symposium on User Interface Software and Technology* (New York, NY, USA, 2001), ACM, pp. 101–110. doi:10.1145/502348.502363. 1, 3
- [BJR21] BUTCHER P. W. S., JOHN N. W., RITSOS P. D.: VRIA: A web-based framework for creating immersive analytics experiences. *IEEE Transactions on Visualization and Computer Graphics* 27, 7 (2021), 3213–3225. doi:10.1109/TVCG.2020.2965109. 5
- [BK05] BEHZADAN A. H., KAMAT V. R.: Visualization of construction graphics in outdoor augmented reality. In *Proceedings of the Winter Simulation Conference* (Piscataway, NJ, USA, 2005), IEEE, p. 7. doi:10.1109/WSC.2005.1574469. 2
- [BKT*21] BRESSA N., KORSGAARD H., TABARD A., HOUBEN S., VERMEULEN J.: What’s the situation with situated visualization? a survey and perspectives on situatedness. *IEEE Transactions on Visualization and Computer Graphics* 28, 1 (2021), 107–117. doi:10.1109/TVCG.2021.3114835. 1, 2
- [BMM12] BARBA E., MACINTYRE B., MYNATT E. D.: Here We Are! Where Are We? Locating Mixed Reality in The Age of the Smartphone. *Proceedings of the IEEE* 100, 4 (2012), 929–936. doi:10.1109/JPROC.2011.2182070. 2
- [BSB*18] BACH B., SICAT R., BEYER J., CORDEIL M., PFISTER H.: The hologram in my hand: How effective is interactive exploration of 3D visualizations in immersive tangible augmented reality? *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (2018), 457–467. doi:10.1109/TVCG.2017.2745941. 2
- [BYK*21] BESANÇON L., YNNERMAN A., KEEFE D. F., YU L., ISENBURG T.: The state of the art of spatial interfaces for 3d visualization. *Computer Graphics Forum* 40, 1 (2021), 293–326. doi:https://doi.org/10.1111/cgfm.14189. 2
- [CCC*15] CHANDLER T., CORDEIL M., CZAUDERNA T., DWYER T., GLOWACKI J., GONCU C., KLAPPERSTUECK M., KLEIN K., SCHREIBER F., WILSON E.: Immersive analytics. In *Proceedings of the International Symposium on Big Data Visual Analytics* (Piscataway, NJ, USA, 2015), IEEE, pp. 1–8. doi:10.1109/BDVA.2015.7314296. 1, 2
- [CRF16] CAVALLO M., RHODES G. A., FORBES A. G.: Riverwalk: Incorporating historical photographs in public outdoor augmented reality experiences. In *Adjunct Proceedings of the IEEE International Symposium on Mixed and Augmented Reality* (Piscataway, NJ, USA, 2016), IEEE, pp. 160–165. doi:10.1109/ISMAR-Adjunct.2016.0068. 2
- [EBC*21] ENS B., BACH B., CORDEIL M., ENGELKE U., SER-RANO M., WILLETT W., PROUZEAU A., ANTHES C., BÜSCHEL W., DUNNE C., DWYER T., GRUBERT J., HAGA J. H., KIRSHENBAUM N., KOBAYASHI D., LIN T., OLAOSEBIKAN M., POINTECKER F., SAFFO D., SAQUIB N., SCHMALSTIEG D., SZAFIR D. A., WHITLOCK M., YANG Y.: Grand challenges in immersive analytics. In *Proceedings of the ACM Conference on Human Factors in Computing Systems* (New York, NY, USA, 2021), ACM, pp. 459:1–459:17. doi:10.1145/3411764.3446866. 1
- [EI13] ELMQVIST N., IRANI P.: Ubiquitous Analytics: Interacting with Big Data Anywhere, Anytime. *Computer* 46, 4 (2013), 86–89. 1, 2
- [Elm23] ELMQVIST N.: Anywhere & everywhere: A mobile, immersive, and ubiquitous vision for data analytics. *Communications of the ACM* (2023). To appear. 1, 2
- [EST16] ELSAYED N. A. M., SMITH R. T., THOMAS B. H.: HORUS EYE: See the invisible bird and snake vision for augmented reality information visualization. In *Adjunct Proceedings of the IEEE International Symposium on Mixed and Augmented Reality* (Piscataway, NJ, USA, 2016), IEEE, pp. 203–208. doi:10.1109/ISMAR-Adjunct.2016.0077. 2
- [ET08] ELMQVIST N., TSIGAS P.: A taxonomy of 3D occlusion management for visualization. *IEEE Transactions on Visualization and Computer Graphics* 14, 5 (2008), 1095–1109. doi:10.1109/TVCG.2008.59. 3
- [ETM*16] ELSAYED N. A. M., THOMAS B. H., MARRIOTT K., PI-ANTADOSI J., SMITH R. T.: Situated analytics: Demonstrating immersive analytical tools with augmented reality. *Journal of Visual Languages & Computing* 36 (2016), 13–23. doi:https://doi.org/10.1016/j.jvlc.2016.07.006. 1, 2
- [GLK*12] GRASSET R., LANGLOTZ T., KALKOFEN D., TATZGERN M., SCHMALSTIEG D.: Image-driven view management for augmented reality browsers. In *Proceedings of the IEEE International Symposium on Mixed and Augmented Reality* (Piscataway, NJ, USA, 2012), IEEE, pp. 177–186. doi:10.1109/ISMAR.2012.6402555. 2
- [HPGK94] HINCKLEY K., PAUSCH R., GOBLE J. C., KASSELL N. F.: Passive real-world interface props for neurosurgical visualization. In *Proceedings of the ACM Conference on Human Factors in Computing Systems* (New York, NY, USA, 1994), ACM, pp. 452–458. doi:10.1145/191666.191821. 2
- [HS88] HART S. G., STAVELAND L. E.: Development of nasa-tlx (task load index): Results of empirical and theoretical research. *Advances in Psychology* 52 (1988), 139–183. doi:https://doi.org/10.1016/S0166-4115(08)62386-9. 6
- [HZR*92] HERNDON K. P., ZELEZNIK R. C., ROBBINS D. C., CONNER D. B., SNIBBE S. S., VAN DAM A.: Interactive shadows. In *Proceedings of the ACM Symposium on User Interface Software and Technology* (New York, NY, USA, 1992), ACM, pp. 1–6. doi:10.1145/142621.142622. 2
- [LJKM*17] LAVIOLA JR J. J., KRUIJFF E., MCMAHAN R. P., BOWMAN D. A., POUPYREV I.: *3D User Interfaces: Theory and Practice*, second ed. Pearson Education, Inc., London, United Kingdom, 2017. 2
- [LSBD21] LANGNER R., SATKOWSKI M., BÜSCHEL W., DACHSELT R.: MARVIS: Combining mobile devices and augmented reality for visual data analysis. In *Proceedings of the ACM Conference on Human Factors in Computing Systems* (New York, NY, USA, 2021), ACM. doi:10.1145/3411764.3445593. 2

- [LSY*21] LIN T., SINGH R., YANG Y., NOBRE C., BEYER J., SMITH M. A., PFISTER H.: Towards an understanding of situated ar visualization for basketball free-throw training. In *Proceedings of the ACM Conference on Human Factors in Computing Systems* (New York, NY, USA, 2021), ACM, pp. 461:1–461:13. URL: <https://doi.org/10.1145/3411764.3445649>, doi:10.1145/3411764.3445649. 2
- [Mac98] MACKAY W. E.: Augmented reality: Linking real and virtual worlds: A new paradigm for interacting with computers. In *Proceedings of the ACM Conference on Advanced Visual Interfaces* (New York, NY, USA, 1998), ACM, pp. 13–21. doi:10.1145/948496.948498. 2
- [MK94] MILGRAM P., KISHINO F.: A Taxonomy of Mixed Reality Visual Displays. *IEICE Transactions on Information Systems E77-D*, 12 (1994), 1321–1329. pp. 1321–1329. 2
- [MTM*16] MADSEN J. B., TATZQERN M., MADSEN C. B., SCHMALSTIEG D., KALKOFEN D.: Temporal coherence strategies for augmented reality labeling. *IEEE Transactions on Visualization and Computer Graphics* 22, 4 (2016), 1415–1423. doi:10.1109/TVCG.2016.2518318. 2
- [Mun14] MUNZNER T.: *Visualization Analysis and Design*. A.K. Peters Visualization Series. A. K Peters, Natick, MA, USA, 2014. 3
- [PBW96] POUPYREV I., BILLINGHURST M., WEGHORST S., ICHIKAWA T.: The Go-Go Interaction technique: Non-linear mapping for direct manipulation in VR. In *Proceedings of the ACM Symposium on User Interface Software and Technology* (New York, NY, USA, 1996), ACM, pp. 79–80. doi:10.1145/237091.237102. 2
- [RRB*14] ROBERTS J. C., RITSOS P. D., BADAM S. K., BRODBECK D., KENNEDY J., ELMQVIST N.: Visualization beyond the desktop – the next big thing. *IEEE Computer Graphics and Applications* 34, 6 (2014), 26–34. doi:10.1109/MCG.2014.82. 2
- [SCP95] STOAKLEY R., CONWAY M. J., PAUSCH R.: Virtual reality on a WIM: Interactive worlds in miniature. In *Proceedings of the ACM Conference on Human Factors in Computing Systems* (New York, NY, USA, 1995), ACM, pp. 265–272. doi:10.1145/223904.223938. 2, 4
- [SD21] SATKOWSKI M., DACHSELT R.: Investigating the impact of real-world environments on the perception of 2D visualizations in augmented reality. In *Proceedings of the ACM Conference on Human Factors in Computing Systems* (New York, NY, USA, 2021), ACM. doi:10.1145/3411764.3445330. 2
- [SH16] SCHMALSTIEG D., HÖLLERER T.: *Augmented Reality: Principles and Practice*. Addison-Wesley, Boston, MA, USA, 2016. 2
- [Shn96] SHNEIDERMAN B.: The eyes have it: A task by data type taxonomy for information visualizations. In *Proceedings of the IEEE Symposium on Visual Languages* (Piscataway, NJ, USA, 1996), IEEE, pp. 336–343. doi:10.1016/b978-155860915-0/50046-9. 3
- [SHN19] SPEICHER M., HALL B. D., NEBELING M.: *What is Mixed Reality?* ACM, New York, NY, USA, 2019, p. 1–15. URL: <https://doi.org/10.1145/3290605.3300767>. 2
- [SLC*19] SICAT R., LI J., CHOI J., CORDEIL M., JEONG W.-K., BACH B., PFISTER H.: DXR: A toolkit for building immersive data visualizations. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2019), 715–725. doi:10.1109/TVCG.2018.2865152. 1
- [SMK*09] SCHALL G., MENDEZ E., KRUIJFF E., VEAS E., JUNG-HANNS S., REITINGER B., SCHMALSTIEG D.: Handheld augmented reality for underground infrastructure visualization. *Personal Ubiquitous Computing* 13, 4 (2009), 281–291. doi:10.1007/s00779-008-0204-5. 2
- [TAK*05] TORY M., ATKINS M., KIRKPATRICK A., NICOLAOU M., YANG G.-Z.: Eyegaze analysis of displays with combined 2D and 3D views. In *Proceedings of the IEEE Conference on Information Visualization* (Piscataway, NJ, USA, 2005), IEEE, pp. 519–526. doi:10.1109/VISUAL.2005.1532837. 2, 4
- [Tor03] TORY M.: Mental registration of 2D and 3D visualizations (an empirical study). In *Proceedings of the IEEE Conference on Information Visualization* (Piscataway, NJ, USA, 2003), IEEE, pp. 371–378. doi:10.1109/VISUAL.2003.1250396. 2
- [TWD*18] THOMAS B. H., WELCH G. F., DRAGICEVIC P., ELMQVIST N., IRANI P., JANSEN Y., SCHMALSTIEG D., TABARD A., ELSAYED N. A. M., SMITH R. T., WILLETT W.: Situated analytics. In *Immersive Analytics* (New York, NY, USA, 2018), vol. 11190 of *Lecture Notes in Computer Science*, Springer, pp. 185–220. doi:10.1007/978-3-030-01388-2_7. 1, 2
- [Wei91] WEISER M.: The computer for the 21st Century. *Scientific American* 265, 3 (1991), 94–104. doi:10.1145/329124.329126. 2
- [WF09] WHITE S., FEINER S.: *SiteLens: Situated Visualization Techniques for Urban Site Visits*. ACM, New York, NY, USA, 2009, p. 1117–1120. URL: <https://doi.org/10.1145/1518701.1518871>. 1
- [WH05] WITHER J., HÖLLERER T.: Pictorial depth cues for outdoor augmented reality. In *Proceedings of the IEEE International Symposium on Wearable Computers* (Piscataway, NJ, USA, 2005), IEEE, pp. 92–99. doi:10.1109/ISWC.2005.41. 2
- [WSS20] WHITLOCK M., SMART S., SZAFIR D. A.: Graphical perception for immersive analytics. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces* (Piscataway, NJ, USA, 2020), IEEE, pp. 616–625. doi:10.1109/VR46266.2020.00084. 2
- [WWS20] WHITLOCK M., WU K., SZAFIR D. A.: Designing for mobile and immersive visual analytics in the field. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (2020), 503–513. doi:10.1109/TVCG.2019.2934282. 2
- [YBVI22] YAO L., BEZERIANOS A., VUILLEMOT R., ISENBERG P.: Visualization in motion: A research agenda and two evaluations. *IEEE Transactions on Visualization and Computer Graphics* 28, 10 (2022), 3546–3562. doi:10.1109/TVCG.2022.3184993. 3
- [ZC19] ZHENG M., CAMPBELL A. G.: Location-based augmented reality in-situ visualization applied for agricultural fieldwork navigation. In *Adjunct Proceedings of the IEEE International Symposium on Mixed and Augmented Reality* (Piscataway, NJ, USA, 2019), IEEE, pp. 93–97. doi:10.1109/ISMAR-Adjunct.2019.00039. 2
- [ZGL*20] ZOLLMANN S., GRASSET R., LANGLOTZ T., LO W. H., MORI S., REGENBRECHT H.: Visualization techniques in augmented reality: A taxonomy, methods and patterns. *IEEE Transactions on Visualization and Computer Graphics* 29, 9 (2020), 3808–3825. doi:10.1109/TVCG.2020.2986247. 2