Maps Around Me: 3D Multiview Layouts in Immersive Spaces

KADEK ANANTA SATRIADI, Monash University, Australia
BARRETT ENS, Monash University, Australia
MAXIME CORDEIL, Monash University, Australia
TOBIAS CZAUDERNA, Monash University, Australia
BERNHARD JENNY, Monash University, Australia

Visual exploration of maps often requires a contextual understanding at multiple scales and locations. Multiview map layouts, which present a hierarchy of multiple views to reveal detail at various scales and locations, have been shown to support better performance than traditional single-view exploration on desktop displays. This paper investigates the extension of such layouts of 2D maps into 3D immersive spaces, which are not limited by the real-estate barrier of physical screens and support sensemaking through spatial interaction. Based on our initial implementation of immersive multiview maps, we conduct an exploratory study with 16 participants aimed at understanding how people place and view such maps in immersive space. We observe the layouts produced by users performing map exploration search, comparison and route-planning tasks. Our qualitative analysis identifies patterns in layout *geometry* (spherical, spherical cap, planar), *overview-detail relationship* (central window, occluding, coordinated) and *interaction strategy*. Based on these observations, along with qualitative feedback from a user walkthrough session, we identify implications and recommend features for immersive multiview map systems. Our main findings are that participants tend to prefer and arrange multiview maps in a spherical cap layout around them and that they often rearrange the views during tasks.

CCS Concepts: • **Human-centered computing** → *Human computer interaction (HCI).*

Additional Key Words and Phrases: Multiscale exploration, multiview maps, immersive space, virtual reality

ACM Reference Format:
Kadek Ananta Satriadi, Barrett Ens, Maxime Cordeil, Tobias Czauderna, and Bernhard Jenny. 2020. Maps Around Me: 3D Multiview Layouts in Immersive Spaces. 1, 1 (October 2020), 20 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

Exploring geospatial information often requires inspection of multiple locations across large areas, viewed at different scales [43]. The ability to view maps at different scales is important because of differences in abstraction [4]. For instance, visualisations of transportation networks show flights between continents at a global level and local terrestrial transport systems at a national level. Other applications such as the analysis of global pandemics and earthquake epicentres often

Authors' addresses: Kadek Ananta Satriadi, Monash University, Melbourne, Australia, kadek.satriadi@monash.edu; Barrett Ens, Monash University, Melbourne, Australia, barrett.ens@monash.edu; Maxime Cordeil, Monash University, Melbourne, Australia, max.cordeil@monash.edu; Tobias Czauderna, Monash University, Melbourne, Australia, tobias.czauderna@monash.edu; Bernhard Jenny, Monash University, Melbourne, Australia, bernie.jenny@monash.edu.

XXXX-XXXX/2020/10-ART $15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

Fig. 1. Our immersive interface allows users to create large hierarchies of multiple maps at different scales and arrange them in 3D space. In the exploratory implementation shown here, a user views a spherical cap layout surrounding a large central overview. Visual links and colour hues indicate hierarchical groupings.

require investigation of both the global distribution and individual geographical locations. While simple exploration can be done with a single visualisation, more complex tasks such as visual comparison benefit from multiple views [54]. The *overview + detail* technique and similar more complex techniques have been developed for supporting such task [69, 71]. These techniques use two or more distinctive levels of scales to show the global context and local details.

This paper explores 3D layouts of hierarchical map views for geospatial analysis with immersive technology (Fig. 1). Immersive virtual reality (VR) and augmented reality (AR) displays are now widely available at a low cost, and allow virtual content, such as maps, to be manipulated in three-dimensional space. Furthermore, application of VR and AR for data visualisation and analysis, known as *Immersive Analytics* [15], is gaining traction. This growing body of research has investigated other potential benefits of immersive interfaces, such as spatial memory, proprioception, and embodied interaction supported by spatial body movement and 3D direct manipulation. For instance, studies have demonstrated benefits of immersive environments in the ability to move one's head [26] and body [35] in 3D space or arranging document layouts in 3D [46]. Immersive 3D display space in VR/AR also allows intuitive visualisation of 3D geographical features such as terrain and building, as well as 3D geovisualisations, for example, 3D flow maps [75], bar graphics in virtual landscapes and on maps [55], choropleth and prism maps [76], and space-time cubes [70].

This work focuses on the question how to create effective map layouts in 3D space. Given the near-limitless possibilities of layout configurations, we take a qualitative approach of observing how users choose to arrange a set of hierarchical views. The aim of this work is to improve our understanding of how users arrange multiview maps in 3D space for different tasks. We conducted a study with 16 participants to understand how people layout map views in immersive space for a variety of common multiscale exploration tasks (i.e. search, comparison, and route planning) and summarised a list of design implications.

Our findings indicate that users have a strong preference towards occlusion avoidance and locality preservation (the closeness of detail-maps with the geographic space on their parent overview map). These two principles yielded central window layouts, where detail maps are placed around the overview. Users adopted an egocentric strategy by orienting maps towards their position, resulting in a spherical cap layout where maps are placed on a part of an invisible spherical surface. These results can inform the future design of manual and automated layout management in immersive multiview analysis tools. Based on these findings we developed an exploratory implementation with both automated and manual layout features. A user walkthrough revealed that users maintained similar interaction strategies but preferred a wider variety of layouts.

The contributions of this paper are: (*i*) the first exploration of the layout of immersive multiview maps by extending the existing multiview technique into 3D layouts; (*ii*) results of an in-depth study of multiview map layouts in VR that describe and categorise user patterns and behaviours; and (*iii*) a discussion on the design implications for future immersive multiview map systems.

2 RELATED WORK

2.1 Multiscale Exploration Techniques

The most prominent classification of multiscale[1] exploration techniques was proposed by Cockburn et al. [16]. They group multiscale techniques into three main categories: *zooming*, *overview + detail*, *focus + context*. *Zooming* alters the scale of a single view in time, and relies on memorisation of spatial information for the comparison of multiple locations [10]. *Overview + detail* techniques use multiple views for simultaneous visualisation at multiple scales. *Focus + context* techniques magnify a focus area on a single view in a surrounding context region by using non-linear magnification such as fisheye lenses [28] or the perspective wall [59].

While no technique is optimal for all multiscale exploration tasks (see [16, 42]), research supports the argument that multiple *overview + detail* views are useful for tasks that require attention on multiple locations. The multiple views (multiview) technique is faster than *zooming* for multiscale comparison when the task load is higher than the capacity of visual working memory [37, 53]. Multiview is also faster than *zooming* for multiscale search [37], as well as better for reading comprehension than *focus + context* [33]. Our work extends the multiview technique to immersive environments.

2.2 Studies on Multiview Techniques on 2D displays

Research has explored multiview techniques for 2D displays such as desktop displays [36, 37, 44] , tabletops [11, 14, 64], and large display walls [34, 41]. The PolyZoom technique [37] uses a hierarchical layout of top-down views to show multiple locations simultaneously. This technique supports multifocus interaction by providing a means to adjust individual focus views. Lekschas et al. [44] studied various layouts of a high number of detail views referred to as scalable insets. Butscher et al. [14] designed multiview lenses with physics collisions to avoid information overlaps. The Canyon technique [34] combines multiview layouts with a space-folding metaphor [14, 22].

Other studies investigated the use of multiple views to improve route visualisation on maps [40, 71]. For collaborative multiscale exploration, Rusnak et al. [64] conducted a guessability study to formulate touch gesture interaction guidelines for collaborative multiview exploration on a tabletop display. Bortolaso et al. [11] explored multiview techniques for collaborative map exploration on a tabletop display.

The major disadvantage of *overview + detail* techniques for 2D displays is the requirement for a large display space and the overhead time needed for the view management [72]. There have been attempts at automating layouts to optimise the use of display space and automate the management of views. Approaches include top-down hierarchy views [37], lenses with collisions [14], and automatic insets placement [30, 44]. Our work builds on these multiview techniques towards multiscale exploration in three-dimensional space.

2.3 Effect of View Size on Multiscale Exploration

Past studies on high-resolution wall displays (e.g. [45, 60, 61]) can also give insights on multiscale exploration beyond desktop interfaces. In this section, we specifically discuss studies that investigate the effect of a large view size on multiscale navigation on 2D displays. An early

[1]We use the term "multiscale" [29, 39, 53] instead of multi-level [16, 42].

study [31] suggested that a large view reduces the need for multiscale navigation and increases the usability of the interface. The superiority of a large view size for navigation was confirmed by latter studies [6, 8, 52, 66]. The performance benefit of a large view size can be attributed to a better spatial coverage of the information space, also referred to as physical field of view [21, 52, 68]. If more information can be seen in a single view, the need for virtual navigation (i.e. zooming and panning) and memorisation is decreased.

The large spatial coverage of a wall display allows users to physically navigate by walking, crouching or orienting their head, rather than only navigating virtually by panning and zooming [6, 7, 63, 66]. For large displays, physical navigation has been found to outperform virtual navigation in search, pattern finding [7], and visual exploration [56] tasks. Lastly, large displays also improve user engagement [56].

For immersive spaces, only a recent study investigated the effect of size and navigation of 3D scatterplots [74]; however we are not aware of any study evaluating the effect of display size on multiscale exploration. Nonetheless, inspired by studies with physical displays, we use a large overview map as a starting point of the exploration and in the design of our user study below. A large view is also suitable for an accurate pointing and selection using input modalities that rely on large muscle groups such as hand-held controller [77]. We expect that providing a large overview, coupled with the flexibility of arranging detail views in 3D space, will imbue users with the benefits even beyond those known of 2D wall displays (i.e. improved information visibility, increased engagement, and benefits of physical navigation). While this expectation remains to be shown by future research, our study aims at providing initial information to guide the design and management of 3D layouts for hierarchical multiview maps.

2.4 Studies in Views Arrangements

Data mountain [57] and The Task Gallery [58] are seminal papers on view management in 3D desktop virtual environments. Studies on desktop interfaces have mainly focused on the effect of spatial memory of 2D vs 3D documents or view arrangements with mixed results on performance [17, 18, 57, 58, 67] but showed good user preference over 3D arrangements [17]. A study by Cockburn and McKenzie [18] has a close resemblance to our study where they allow users to arrange physical printed views in a very limited physical 3D space. Their study was nonetheless different from ours in terms of type of information (standalone views vs multiscale views) and users interaction freedom (limited space vs large space).

An important design factor when placing 2D views in an immersive space is the point of reference of virtual views and whether reference points are viewed with an egocentric or an exocentric perspective [24]. With an egocentric perspective, views are arranged in regard to the user's position. With an exocentric perspective, views are arranged regardless of the user's position. For instance, the Personal Cockpit [26] imagined that a curved, egocentric view layout would be useful for a mobile user, but the same views could be placed in a flat exocentric arrangement on a nearby wall when arriving at a destination. Egocentric views can also be arranged concentrically relative to a user's shoulder, elbow or wrist joint to support ergonomic direct input [26, 50]. Exocentric layouts typically require additional constraints to be considered, such as the geometry and visual appearance of a surrounding room, the presence of physical obstacles, or the available visual sightlines [1, 25, 27].

Although, to the best of our knowledge, there is no previous study exploring how users arrange multiscale hierarchical views in immersive 3D space, we can gain insights from existing studies in the immersive analytics domain. In a study by Batch et al. [9] using the ImAxes immersive visualisation system [20], participants placed visualisations egocentrically and in close range for exploration tasks. Conversely, participants used more space to arrange visualisations in an exocentric way when presenting insights to others in a collaborative setting. Lisle et al. [46] demonstrated how the immersive version of the "space to think" [2] scenario can benefit users in a text comprehension task. Lisle et al. found that the participant created a dome-like structure when preparing the documents for writing. Liu et al. conducted a quantitative study on immersive small multiples visualisation layouts and found that users prefer a semi-circular layout over flat or full circle layouts [47]. Although these studies give insights on users' preferences, the layouts were predefined [47] or were based on open-ended data analysis tasks [2, 46]. We build on these works with a controlled study including several distinct analysis tasks, with a focus on multiscale maps with a hierarchical structure.

2.5 Summary

Our work is motivated by the need for design guidance for creating an immersive visualisation prototype for multiscale map exploration. In regard to designing the size of map views, the literature indicates that a large view displaying more information increases user performance, provides extra navigation strategies by walking or head rotation, and promotes an engaging user experience. Studies also suggest that multiview visualisation on 2D displays is beneficial for sensemaking and exploration tasks. Thus, it is sensible to use a large overview map as a starting point of the exploration process, and to use multiples views with the *overview + detail* technique. The view management cost of multiview techniques could be reduced by using automatic layouts, however, users may at times prefer to manipulate views directly for tasks such as side-by-side comparison. Related studies in immersive analytics showed that users place information in an egocentric perspective during the exploration process [9, 46]. This indicates that maps should be oriented towards the user, but the best way to structure 3D spatial layouts and organise multiscale views is not fully understood.

3 OBSERVATION STUDY

To guide the future development of immersive multiview maps systems, we designed a study to observe how users arrange multiscale maps in 3D space for use in different task scenarios. Participants were given a set of hierarchical maps and asked to place these relative to a large overview map in an immersive VR setting.

3.1 Design

Our map hierarchy included a single *overview* world map and six large maps of *metropolitan* areas. Each of the metropolitan maps had three child maps showing *city-level* information. To avoid making the map hierarchy overly complex, we chose to limit the maps to three levels of scale [42]. We set the size of the overview, metro-level, and city-level maps to 2×1 m, 0.6×0.6 m, and 0.5×0.5 m, respectively. In total, there were 24 maps to be arranged by participants. The world map could not be moved and had its centre fixed at 1.5 m above the ground.

We replaced links with labels to indicate hierarchical relationships between parent and child maps. This step was taken to reduce bias from participants assuming that they had to minimise the length of the links. We further avoided bias in the placement of maps by providing a nondescript VR environment, without any walls, furniture or other structures that might influence map placement. Thus, the only reference initially available is the fixed overview map.

3.2 Tasks

To encourage the emergence of a variety of layout strategies, we provided participants with different tasks that simulate various multiscale map interaction scenarios. These combine general multiview tasks (*naïve layout*, *general layout*) with adaptations of interactive geovisualisation objective primitives by Roth [62] (*search*, *comparison*, *route planning*). Detailed instructions of each of the tasks can be found in the Supplementary Material SP1.

Naïve layout – The first task asked participants to arrange the metropolitan and city-level maps freely, without any given task context. This task aimed at identifying how users arrange multiscale views in immersive space and is not specific to maps. To deter participants from creating a nonsensical layout, we specifically informed them that they had to explain the motivation behind the layout after the task. In the naïve layout task, we attached small "minimap" versions (5×5 cm) of the metropolitan-level and city-level maps to a virtual "maps stick" that participants held in their non-dominant hand (Fig. 2, right). The participants had to grab and drag the minimap away from the virtual stick to enlarge it to its full size. The motivation behind this setup was to avoid an arrangement bias due to the initial map positions.

In the following three tasks, participants were asked to repeat several task instances of a specific task. We included repetitions both to ensure they understood the task, and to allow us to observe a wider variety of approaches. To prevent bias from the previous task, the first instance of each task was reset to the initial naïve layout. Participants were offered frequent breaks between tasks. A single session of the study lasted on average 90 minutes.

Search – The search task was inspired by Roth's *identify* objective primitive [62]. We asked participants to perform map feature's searching and identification, which required them to examine each map view at least once. The task was to count the number of maps containing grey circles. At the start of each task instance, we defined five target maps in any level, showing one circle and four squares (Fig. 2, left, *Search* – left) and used other maps as distractors showing five squares (Fig. 2, left, *Search*). We randomly chose target maps but put a higher chance on maps that were partially or fully occluded from the participant's current view. We expected participants to minimise occlusion of views in their layout in this task.

Comparison – The third task was inspired by Roth's *compare* objective primitive [62]. This task required close examination to detect similarities and differences between pairs of maps. On a subset of ten city-level maps, we drew a set of shapes, and participants were asked to find map pairs that contained identical shapes. The shapes were a star, a circle, and a square. They varied in colour and were placed in a triangular arrangement. For example, the first two maps in Fig. 2, left, *Comparison*, are target maps with identical shapes whereas the other one is a distractor. This pattern is similar to the visual patterns used in a previous multiscale study [54]. Each task instance showed two pairs and eight distractors on random maps.

Route planning – The next task was a route-planning task where participants needed to plan a route between two locations. This task resembled Roth's *associate* objective primitive [62] task. It required participants to consider relationships between different maps, including correspondences between different zoom levels. The task design was inspired by a multiscale network visualisation on a 2D display [13]. We used a hypothetical transportation scenario and created a simple weighted geospatial network. On the overview map, we showed flight paths between metropolitan locations. On metropolitan-level maps, we showed connections between cities, representing train routes. Finally, we created a network of taxi locations, train stations, and airports on the city-level maps. We asked participants to find the lowest cost path between two nodes on two different city-level maps. This task required participants to inspect metropolitan-level maps and the overview map to connect the two target nodes. Fig. 2 (left, *Route*), shows the overview map and an example of a network on metropolitan-level maps and city-level maps (the two smallest maps). The numbers inside red circles indicated the cost of the paths.

General layout – The last task was to refine the naïve layout to a perceived "optimal" layout – i.e. a single layout that could be used to perform all of the above search, comparison, and route-planning tasks. Given the participant's experience gained in these tasks, we wanted to see how this informed layout would differ from the initial naïve layout. We expected that participants would come up with different variations of layouts that they perceived as optimal.

3.3 Software and Hardware Setup

We built the software for the study using modules of our exploratory implementation with the Unity Engine, the Mapbox plugin for maps, and the Virtual Reality Toolkit (VRTK) library for interactivity.

We aimed to make it as effortless as possible to control maps to limit any potential bias of our implementation on the constructed layouts. We provided interactions only for selection, translation, and rotation of the map views. We chose not to include any navigation functionality such as panning or zooming within views, as these could affect the strategies chosen during the tasks. We provided a ray-casting pointer for map selection. Once a map was selected participants could move it directly by moving the controller or adjust its relative depth and yaw-orientation using the controller pad. The up-down axis moved the selected view further or nearer. The left-right axis controlled the maps yaw rotation, which allowed participants to rotate a map on its local vertical axis, as one might swivel a desktop monitor. The interactions could be fully performed with a single controller, but we provided two controllers to support ambidextrous interaction.

The immersive VR display was an HTC Vive Pro running on a PC system with a GTX 1080 GPU, a Core i7 processor, and 16 GB of RAM. The dimension of the tracking area was 4×3 m. The experimenter sat in front of the screen with an interface to control the study.

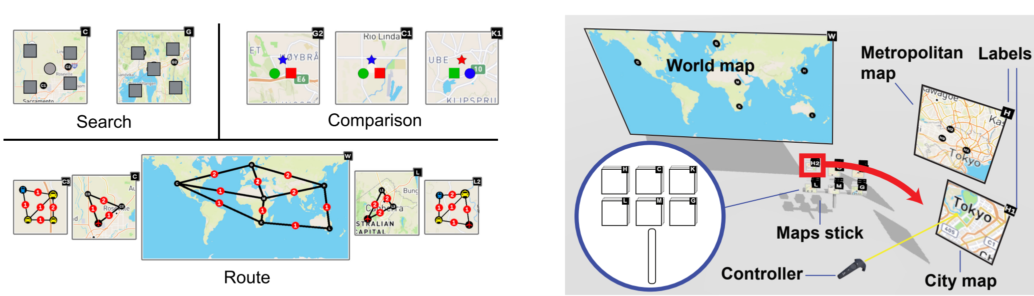

Fig. 2. Examples of target maps in search, comparison, and route-planning tasks (left). Study setup (right). The red arrow illustrates how the minimap is grabbed from the map stick and then grows to the full size. The H map is at metropolitan level and H1 (bottom right) is one of its three child maps.

3.4 Data Collection

We collected the following data:

- *Demographic and expertise data.* Standard demographic and expertise data (familiarity with VR and map reading).
- *Map layout data.* We stored all map positions and orientations as they appeared in the layout at the end of each task. The total number of layouts from 16 participants and 5 tasks is 80 layouts, each consisting of 25 maps.
- *Interviews.* We conducted interviews to gain insight in user strategies. We asked questions after all tasks were completed and the participants were still in the VR environment.
- *Interaction log.* Activity of the participants and all layout operations were captured for later playback for analysis.

As the focus of our study was observing the users' layouts and strategies in five different tasks, we did not use accuracy and performance measures of the user's interaction during the tasks.

3.5 Participants

We recruited 16 participants (8 male, 8 female) with an average age of 27 ($SD = 4$) from our university. Seven participants had less than 5 hours experience with VR while four had more than 20 hours experience. The average rating on a 5-point Likert-scale for maps understanding was 3.7 ($SD = 1.2$), with 1 being a novice and 5 being an expert.

3.6 Study Procedure

Training – In the first training we asked participants to visit four points around the tracked space sequentially to familiarise them with the surrounding space and encourage locomotion during the study. The second training was designed to make participants familiar with maps positioning and orienting techniques.

Briefing – We showed an instruction dialogue in front of the participants which disappeared once the participants confirmed that they understood the task. The participants were allowed to ask questions to clarify the tasks. Before the tasks started, the experimenter confirmed participants' understanding of the task by asking them to explain the task briefly. For the search, comparison and route-planning tasks, participants were also informed that they could rearrange views as they needed, but that the arrangement should be general enough to help anyone to answer a similar question in the future.

Layout authoring – For search, comparison and route-planning tasks, participants performed three to five task instances until the experimenter observed consistency in the strategy and layout. In the first instance of each task, participants started from their naïve layout. We deliberately chose not to use the thinking aloud protocol to avoid interfering in participants' sensemaking process [2].

Semi-structured interview – Following the completion of all tasks, we asked participants to briefly explain their motivation behind the layouts they created and how these helped to perform the given tasks.

4 RESULTS

For analysis we used an iterative single coder method, following an approach used in similar studies [48, 56]. At the initial stage, all authors explored structured and unstructured data to get an overall insight of possible classifications. The main author analysed the study results. During the coding process, results were iteratively discussed by all authors to resolve ambiguous cases and reduce subjective bias. Since the classifications were straightforward to define, we decided to use a single coder.

4.1 Layout Geometry

Participants created a total of 80 final layouts (16 participants × 5 tasks). Figures of all layouts can be seen in the Supplementary Material SP2. Of these, we discarded four layouts from two participants in the layout geometry analysis. Participant 6 decided not to perform the route task, and three layouts created by participant 13 did not fit any discernible layout pattern (e.g. P13 on Fig. 3, A). We analysed the 76 remaining layouts.

Three distinctive layout geometries emerged from visual inspection of static layouts: **spherical**, **spherical cap**, and **planar**. The spherical layout wraps around users in 360 degrees. Maps with a spherical cap layout are arranged on a portion of a sphere surface. The planar layout is a flat layout where maps are mostly aligned on or near the plane of the overview map.

In general, the most common layout geometry was a spherical cap layout (81.5%), followed by spherical (13.2%) and planar (5.3%) layouts. Planar layouts are only found in the initial naïve layout and the final general layout tasks (Fig. 4). During the search, comparison, and route-planning

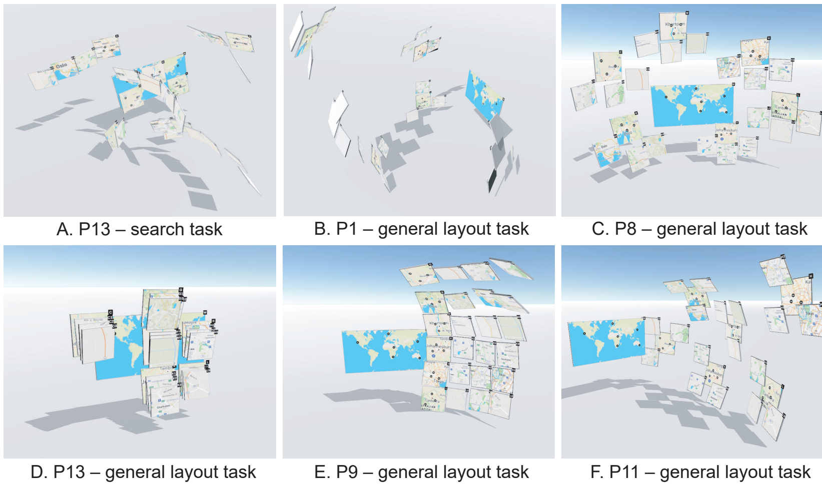

Fig. 3. Examples of layouts and their respective categorisations. Geometry patterns: spherical (B), spherical cap (C, E, F), planar (D), and unidentified (A). Overview-detail relationships: central-window (C), occluding (D), coordinated (E, F), and unidentified (A, B).

tasks, all three participants who started with a naïve planar layout (P5, P8, and P16) finished with a spherical cap layout in all of these tasks and the general layout task. In contrast, the three participants who initially created a spherical layout (P4, P1 and P15) changed to a spherical cap layout less often. However, two of these participants ended off with a spherical cap layout in the general layout task.

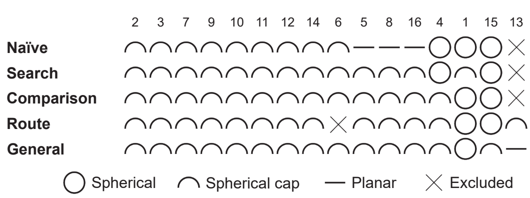

Fig. 4. Layout geometry for all participants (1 to 16) for all tasks. The numbers indicate the participant ID.

We used the least squares method [38] to estimate the radius and the centre point of each spherical and spherical cap layout. The average radius of the estimated spheres is 2.07 m (SD = 0.24 m). The average absolute distance of maps from the sphere surface is 0.17 m (SD = 0.22 m), less than one tenth of the average radius, which quantitatively supports that the spherical and spherical cap layouts indeed consist mostly of maps placed roughly on a spherical surface. We visualised the estimated sphere in the 3D space to visually validate the spherical arrangement of the maps. We also calculated the difference between the sphere centre elevation and eye level of each participant. On average the sphere centre was placed 0.61 m (SD = 0.24 m) above the eye level.

4.2 Overview-Detail Relationships

Our categorisation of overview-detail relationships was based on observations of spatial relationships between parent and child views from the hierarchy, i.e. child metropolitan maps relative to the parent world map, and child city-level maps relative to their respective parent metropolitan maps. We expected that participants would tend to preserve locality in these relationships to reflect their mental model of the hierarchy, but to also generally avoid occlusion.

Our analysis yielded three main layout patterns (Fig. 3, C - F). The **central window layout** minimises the occluded area on the overview map. This results in an overview map that is surrounded by detail maps, maximising the visibility of context [44]. The **occluding layout** is an arrangement where detail maps are placed near their respective location on the overview map. This resembles a magnification lens metaphor where the locality of the spatial relationship is at maximum level [44]. The **coordinated layout** typically has detail maps grouped adjacent to the overview map.

We compared the initial naïve layouts to the final general layouts to see if the intermediate tasks affected the perceived optimal arrangements. Layouts from P1 and P6 do not fit into any overview-detail category. Thus, we discarded those layouts in overview-detail relationship analysis. For the naïve layout, most participants (9 out of 16, 56.3%) avoided occlusion by using a central window layout, four participants (26.7%) created an occluding layout, and one participant created a coordinated layout. The general layout task also resulted in a majority of central window layouts (62.5%). Only one participant (6.3%) suggested an occluding layout and three participants (18.8%) preferred a coordinated layout (Fig. 5, left, Final).

4.3 Interaction Strategy

A significant number of participants moved maps as they performed a task instead of creating a stable layout before starting the task. Consequently, final layouts on each of the tasks do not necessarily tell the whole story because they manifest only the end result of a complex interaction process. This behaviour mainly occurred during the comparison task.

We observed similar strategies as Reda et al. [56]: participants using the **layout-preserving pattern** tended not to change their layouts while performing the task. Participants using the **layout-changing pattern** actively altered the arrangement of maps during the tasks. The layout-changing strategy often led to an unstructured layout at the end of the task. For example, participant 13 in the search task was actively moving maps during the task; we discuss this strategy in the Discussion section. This resulted in maps scattered in space at the end of the task (Fig. 3, A).

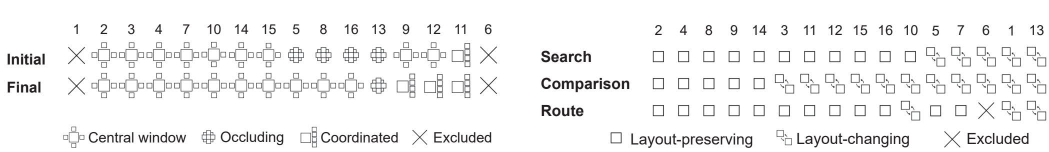

Fig. 5. The overview-detail relationship of naïve layouts and general layouts for all participants (left). The interaction strategies of all participants in search, comparison, and route tasks (right).

The interaction patterns are behaviours that we observed as predominant behaviour throughout the task, ignoring the part when participants prepared the layout. For example, participant 8 began the comparison task by changing the layout and never significantly modified the layout afterwards. In that case, we classified the participant as a layout-preserving pattern. Participant 10 on the other hand, changed the layout by grouping target maps for every task variation in the comparison task. Thus, we considered P10 used the layout-changing pattern.

We excluded the result of participant 6 from the route task due to inability to perform the task. In general, we found that most participants used the layout-preserving approach during the search (68.7%) and route tasks (80%) but not during comparison task (31.2%), as seen in Fig. 5, right.

During the *search task*, five participants changed their layouts, using a "search, group, then count" strategy. They first searched and then placed all target views in arbitrary space around them before finally counting all circles. In the *comparison task*, most participants tended to group target maps before comparing shapes. We also observed that 4 out of 11 participants using the layout-changing pattern were actively re-arranging maps even when all target maps were already grouped. For example, after grouping target maps, some participants sorted maps according to the comparison patterns. P3 stated: "*I am sorting these maps vertically according to the shape of the first leaves. Rectangles are on the second row, and stars are on the first row. Then I sort these maps by the colour of the first leaves. In this way, I can quickly find similar patterns.*"

Most participants had a stable layout during the *route task*. The main strategy was to scan the chain of maps between the two target points. Participants with occluded layouts initially expanded the layout to reduce occlusion. They also moved detail maps closer to the overview map. There were three participants who actively arranged maps in each task variation. Their strategy was to find target maps and cluster them.

5 DISCUSSION

The aim of our study was to understand how people arrange their space to think and manage immersive multiview map layouts. Our analysis focused on the spatial arrangements that participants created, in which specific patterns of layouts emerged. We also focused on how participants handled overview-detail relationships within the map hierarchy, and on the interaction strategies we observed. In the following we discuss the results regarding these themes.

5.1 Layout Geometry

We observed a tendency toward spherical cap arrangements of views, an interesting outcome given the prevalence of flat whiteboards and large display screens in the real world. Moreover, while some fully spherical arrangements were created, there was a greater tendency towards spherical cap layouts, presumably because these allow users to observe the overall layout more easily while reducing the need for turning one's body. The proportion of spherical cap layouts increased throughout the tasks, from over half in the initial naïve layout to all but 2 in the final general layout (Fig. 4). This was in part due to the design of our task set, which allowed users to progress through a series of tasks of increasing difficulty. The final general layout task allowed participants to apply what they had learned through the previous tasks, which resulted in different strategies from the naïve strategies used in the initial layout.

These observations are similar to findings by Batch et al. [9], where during the analysis phase most participants created egocentric arrangements of data visualisations. However, their study differed from ours in that it focused on abstract data visualisations (e.g. scatterplots) rather than multiscale map visualisations. Also, their implementation used a direct grabbing metaphor similar to ours for manipulating the 3D visualisations, but did not include additional features for rotating objects or moving them toward or away from the user. Whereas participants in the study by Batch et al. did not make full use of the 360-degree space, we found that some of our participants did leverage the 360-degree spherical surface space (10 out of 76 layouts from three participants, or 13.2%). Our findings also contrast those of Lisle et al. [46], where a spherical arrangement was made for a text-based analysis task. While some participants in our study also created spherical arrangements of maps, these were in most cases discarded in favour of spherical-cap arrangements. However, the study by Lisle et al. used only a single participant, so further study of a similar task may show more variation. Our findings are more in alignment with those of Liu et al. [47], who found a preference for a semi-circular arrangement over flat and full-circle arrangements for small multiple visualisations. However, our study differs in two principal ways from the study by Liu et al.: *i)* they used pre-configured layouts, and *ii)* the semi-circular and full-circle layouts in their study were cylindrical, as opposed to the spherical cap and spherical layouts in our study.

We took several steps to reduce bias of our study design and minimise influence on the layouts produced by participants, such as providing simple and easy to use controls, including depth and pivot manipulations to freely place maps, placing minified maps initially on the virtual map stick, and setting the study in an empty environment, free of external reference points. Nevertheless, our interface design choices may have influenced the predominance of spherical cap layouts. In particular, the hand-held controller acted as the pivot point for map positioning, so views remained oriented toward the participant by default. Thus, creating a planar layout would have required users to either walk around or align the orientation of detail maps with the world map using the controller, both of which require time and effort. It could be the case that participants found the map rotation feature too cumbersome to use, however, this feature used the same directional pad as the depth-displacement feature, which was used frequently as shown by the large average sphere radius. Thus, the spherical arrangements may have resulted simply from a participant preference for 'billboarding' of maps (orienting them to face the user) to increase information visibility, similar to multiple screen setup on the desktop workspace, e.g. [2]. Nevertheless, the result could be different if an alternative approach for positioning and orienting maps were used; for instance placing a constraint on translation in alignment with the overview might result in more planar layouts.

5.2 Overview-detail Relationship

The results of our study indicate a strong tendency to use a *central window* layout (56.3% for the naïve layout task, 62.5% for the general layout task). The central window layout is the result of emphasising locality preservation and occlusion avoidance. This finding is similar to the detail views placement approach defined in previous work on overview and detail geographical maps by Lekschas et al. [44], in which detail views are placed close to their origin while trying not to occlude the overview map at the centre. Further, our study showed that this *central window* pattern can be repeated at a hierarchical level; our hierarchical views have three levels (world, metropolitan, and city) whereas the insets in the study of Lekschas had two levels. We observed that most participants preserved locality at multiple scales (i.e. world-metropolitan and metropolitan-city).

Our study results also extend the previous finding to 3D VR space with larger maps. As a result, the size of detail maps was greater than the insets used by Lekschas et al. (ratio in surface area relative to the overview map of about 1:10 vs 1:300) because the limits of our VR setup do not allow the same fine resolution as on a 2D display. Although upcoming high-resolution VR display will allow smaller views, immersive space allows interaction beyond the screen real-estate limits of 2D displays.

5.3 Interaction Strategy

We observed that participants did not always create a single layout during the given tasks, but in some cases would continue to move the maps around throughout the task, which we initially did not expect. We found this interaction was predominant in the comparison task, where most participants rearranged the target maps next to each other, and therefore altered their original layout. Presumably, these interactions reduced visual working memory by minimising the distance between the views containing the target patterns. Our finding aligns with other studies on visual exploration [3, 56] where users were found to cluster and group relevant information in proximity. This behaviour could indicate a cognitive offloading strategy [49] where participants externalised their thinking through maps sorting and grouping.

Conversely, most participants kept their layouts stable during the search and route-planning tasks. The main strategy was to ensure that occlusion was minimised so that the participants could scan all maps for the search and route-planning tasks. On the other hand, a few participants did not keep the layout stable during these tasks. Their strategy was to group target maps and separate them from non-target maps.

5.4 Design Implications

Based on what we learned from the study, we identify several key implications to the design of multiview maps in immersive space.

Spherical cap layouts maximise map visibility. Participants who created spherical cap layouts reported in the interviews that they wanted to create a compact layout that allowed them to see all maps at once. Given the relatively large size of the maps, placing all of them in a plane configuration would create a distorted perspective of maps on the edge of the layout. A spherical cap arrangement allows the maps to be viewed from the ideal angle, and minimises the need for users to rotate their body or walk around the environment.

Central window arrangements are recursive. Generally, participants placed child maps along the borders of their parent. This strategy was applied at both hierarchical levels i.e. world map–metropolitan maps, and metropolitan map–city maps. Thus, automatic layout algorithms should consider such recursive central window arrangements.

Automate predefined layouts and manual layout authoring. A variety of predefined layouts can provide initial support for different task types. Whereas most participants did not significantly change their layouts in search and route-planning tasks, most participants did actively move maps in the comparison task. Therefore, it is important for immersive multiview maps system to also support manual map layout adjustments. These manual features should include options for moving both individual views and groups, for instance to move all maps in a hierarchical group or to perform a quick inspection without altering the layout. Smooth transition from one layout to another is also important to maintain the user's mental model of the system, as well as the ability to revert to predefined or bookmarked layouts.

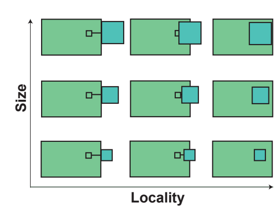

Fig. 6. The 2D locality-size continuum, represented as a matrix for demonstration purposes.

There is a trade-off between view locality and size. Having multiple large maps is the main selling factor to explore multiscale geovisualisation in immersive space. While increasing the size of the detail maps on a 2D screen means sacrificing the size of other maps [37] or increasing occlusion, immersive displays allow layouts to be distributed more widely in 3D space around the user. We see that the size and locality can be modelled as a 2D continuum, providing multiple possible overview-detail relationships (Fig. 6). The locality axis reflects the physical distance from a child map to its origin on the parent, while the size axis reflects the relative size of the child view. The size of the map is further related to the degree of information abstraction in the detail view; larger maps allow deeper zooms and increased visibility of details.

6 REVISITING THE EXPLORATORY IMPLEMENTATION

We updated our initial exploratory implementation in light of the design implications we drew from our observation study. We ran a user walkthrough to gain additional insight to better understand user preferences using a more feature-rich implementation.

6.1 Implementation

We implemented an automated layout manager that allows users to interactively control the *locality* and the *size* of maps for planar, spherical cap, and spherical layout geometries (Fig. 7). We use a tree structure and depth-first search post-order traversal to arrange hierarchical groups. Size and locality parameters are applied recursively on each of the hierarchical groups, which means adjustment of size or locality affects all maps.

At a high level, the model interpolates user-controlled values of *size* and *locality* (ranges from 0 - 1). We set the minimum size to 25 cm × 25 cm and the maximum size to 80 cm × 80 cm for the parent maps and 70 cm × 70 cm for the child maps. For the map view position, we interpolate the edge position – the non-overlapping position where the distance to the parent is minimum – and the origin on the parent after the final size is calculated.

To create a spherical layout geometry, we first create a planar layout and then transform Euclidean map positions to spherical positions with a given radius and centre point. For the spherical layout, we use the locality value to determine the angular size of the layouts. A $locality = 0$ yields a spherical layout whereas a $locality > 0$ creates a spherical cap layout. We also added spacing in z-axis between map levels to increase visibility. Animated transitions were also used when users update the layout.

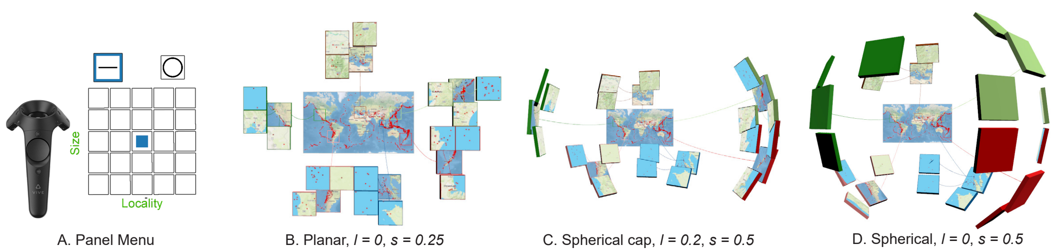

Fig. 7. Panel menu to select flat or spherical layouts and choose different combinations of locality l and map size s (A). Three examples of automated layouts created with the different parameter settings are shown (B–D). A spherical cap layout is created by choosing a spherical layout with non-zero locality value. Note that each of the layouts shown is captured from a different viewpoint distance.

To facilitate manual layout manipulation, we implemented the following interaction modes (demonstration videos are provided in the Supplementary Material SP3):

- *Transient summoning* brings a map to the user's personal space (0.5–1 m from viewpoint) [32] on trigger press and moves it back to the original position when the trigger is released.
- *Permanent summoning* does not move the map back to the original position.
- *Throw-to-reset-position* moves the map back to the default layout position.
- *Continuous distance and rotation* adjusts the position of a grabbed map along the laser pointer direction and the yaw rotation.
- *Group positioning* adjusts position and rotation of a parent-children map group.

6.2 User Walkthrough

Setup – For the qualitative user walkthrough we set the invisible sphere radius to 2 m. The centre point of the sphere was placed in front of the overview map centre. We provided the user a panel menu for switching between planar and spherical layouts, as well as selecting a locality–size setting from a matrix arrangement (Fig. 7 A). The locality–size setting controlled the size of the maps, as well as their vertical and horizontal distribution. For the spherical and the spherical cap layouts, the layout algorithm limited the vertical distribution of maps to a maximum of 150° to prevent excessive neck movement [51]. To avoid overlapping maps, the layout algorithm used a minimum vertical distribution angle of 56°. The layout algorithm also used the locality-size setting to distribute the maps horizontally between 112° (spherical cap layout) and 360° (spherical layout). To indicate overview-detail relationships between maps, we added visual linking [19] and hierarchical colour coding (Fig. 7, B–D). We used the same number of maps as in the previous study but with a different set of static locations. An earthquake dataset[2] was visualised across all maps.

Participants and procedure – We showed our updated implementation to 8 participants (5 male, 3 female) with expertise in information visualisation and/or VR. Half of the participants had participated in the first study. We used interviews and video recordings as the main data collection methods. First, we walked the participants through the new implementation. Then, we asked participants to perform simple tasks such as changing layout geometries, adjusting locality-size parameters, and placing maps around them. Finally, participants were asked questions to encourage exploration such as, "*Find earthquake epicentres located near an airport*" and "*Find maps that show earthquake epicentres that are part of a larger cluster*". We did not validate participants answers because the aim of the exploration was to give participants usage experience of the layout and map manipulation features. The walkthrough of about 60 minutes concluded with a semi-structured interview. The full question set is included with the Supplementary Material SP4.

Preferred layouts – While in the first observation study we found that the spherical cap layout was preferred, this evaluation showed varying preferences: three of eight participants preferred planar layouts, three preferred spherical layouts, and only two preferred spherical cap layouts. Participants perceived the spherical cap layout and the spherical layout as making effective use of space and reducing overlaps. For the *planar* layout, participants appreciated that the full layout was within their field of view so that all maps could be seen from a single perspective. For the *spherical cap* layout, participants liked that all maps were close to their viewpoint and they just needed to turn their head or body slightly to see all maps. One participant did not like *spherical* layout because it required her to turn around to see all maps, it was hard to see the origins of the maps, and made her felt a little claustrophobic.

Preferred map size and locality – Four of the eight participants preferred the largest map size. Participants commented that there was not one ideal size for all tasks. Larger maps were considered as a good choice to use the available VR space (P5), as well as being readable (P4) and good for close inspection (P1, P7). All participants who preferred a planar layout also choose a non-occluding layout.

Interaction patterns – The strategy for solving tasks was generally similar across all participants. Participants tended to repeatedly bring maps to their personal space using permanent summoning and then arrange summoned maps in a new configuration within reaching distance. Surprisingly, this strategy was also used with spherical and spherical cap layouts when all maps were relatively close to participants. This strategy is similar to what we observed in the first study where participants tried to create a clear separation between target maps and non-target maps (Section 5.3). This strategy could also explain why some participants prefer planar layouts; participants could easily summon maps on the edge of the layouts without having to walk closer. Despite the one-hour VR session, all participants but one reported no fatigue or mild fatigue. None of the participants walked around much during the exploration.

Feedback on interactions – In general, participants appreciated the transient and permanent *summoning* interactions. The interviews revealed that transient summoning was valued for a quick inspection, but participants preferred bringing maps to their personal space permanently. The *throw-to-reset-position* and *group positioning* interaction received positive comments while *continuous distance and rotation* was found to be difficult to perform mainly due to the need to press the grip button and touch pad at the same time.

[2]https://www.kaggle.com/usgs/earthquake-database

7 CONCLUSION AND FUTURE WORK

In this paper, we explored the design of multiview map visualisation in virtual reality through observational studies. The results of our first study allowed us to identify distinctive patterns of layout *geometry* (spherical, spherical cap, planar), *overview-detail relationship* (central window, occluding, coordinated) and *interaction strategy*. Our main findings are that participants tend to prefer and arrange multiview maps in a spherical cap layout around them and that they tend to rearrange the views during tasks. Additionally, we identified a two-dimensional locality-size continuum and argue that immersive 3D environments allow more flexibility than 2D displays by providing space for participants to place large size views without the need to occlude or reduce the size of other views.

In a follow up user walkthrough, we observed how participants arranged maps with an interactive implementation that provided both automated and manual control of layouts, including adjustment of the layout geometry and the locality-size continuum. Results with 8 participants indicated that participants' layout preferences were influenced by introduced automatic layout and interaction techniques. In contrast to the preference for spherical cap layouts observed in the first study, participants used a wider variety of layouts. It may be that participants prefer the familiarity of planar layouts but were reluctant to create these using the limited manual tools of the first study. More advanced features such as automated layouts, summoning maps, and throw to reset position may mitigate the disadvantages of planar layouts, such as far reaching distance and distorted oblique views of maps near the layout edge. However, we observed layout-changing patterns in the walkthrough similar to those observed in the comparison task of the first study.

We acknowledge some limitations of our studies. Although the tasks in the first study resemble actual map interaction objectives, potential variations may have arisen if different tasks were used. For instance, due to experiment simplicity purposes, we did not include a *ranking* task which could potentially have yielded different map arrangements. We also acknowledge that the findings of our studies may be limited by a relatively small sample size of participants recruited from our local university.

We envision that automated layouts can be adapted to AR where virtual and physical objects are blended. However, there are several challenges that need to be addressed. For instance, how multiview layouts should be arranged in a setting with a limited physical room size and crowded furniture. Furthermore, future studies in AR should look into integrating virtual map views with physical real-world constraints, for instance sticking large maps on a wall, placing 3D maps horizontally on a table-top, or combining an overview map on a large-screen display with virtual detail maps.

In an AR setting, the use of freehand interaction has been proposed [12, 73], however, gesture based input is prone to usability issues such as fatigue and discoverability. Therefore, integration of freehand gestures for complex tasks beyond simple navigation tasks [5, 65] is worth exploring in the future. For example, how view manipulation (e.g. resize, position, rotation) and content manipulation (e.g. pan, zoom, resizing, details on demand) can be effectively performed.

From our experience while developing the exploratory implementation, we found that integrating complex view and content manipulations with multiple, potentially overlapping views is not trivial. Further research should explore the best way to ensure seamless and "endless" interactions during the exploration by adhering fluid interaction design guidelines, e.g. direct manipulation, clear conceptual model, and minimised mode switching [23]. Ultimately, an "in-the-wild" evaluation of a fully functioning immersive multiview map system with experts is a reasonable future research aim.

REFERENCES

[1] R. Alghofaili, M. S. Solah, H. Huang, Y. Sawahata, M. Pomplun, and L. Yu. 2019. Optimizing Visual Element Placement via Visual Attention Analysis. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. 464–473.

[2] Christopher Andrews, Alex Endert, and Chris North. 2010. Space to think: large high-resolution displays for sensemaking. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 55–64.

[3] Christopher Andrews and Chris North. 2012. Analyst's Workspace: An embodied sensemaking environment for large, high-resolution displays. In *2012 IEEE Conference on Visual Analytics Science and Technology (VAST)*. IEEE, 123–131.

[4] Gennady Andrienko, Natalia Andrienko, Urska Demsar, Doris Dransch, Jason Dykes, Sara Irina Fabrikant, Mikael Jern, Menno-Jan Kraak, Heidrun Schumann, and Christian Tominski. 2010. Space, time and visual analytics. *International Journal of Geographical Information Science* 24, 10 (2010), 1577–1600.

[5] Christopher R Austin, Barrett Ens, Kadek Ananta Satriadi, and Bernhard Jenny. 2020. Elicitation study investigating hand and foot gesture interaction for immersive maps in augmented reality. *Cartography and Geographic Information Science* 47, 3 (2020), 214–228.

[6] Robert Ball and Chris North. 2005. Effects of tiled high-resolution display on basic visualization and navigation tasks. In *CHI'05 Extended Abstracts on Human Factors in Computing Systems*. ACM, 1196–1199.

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

[8] Robert Ball, Michael Varghese, Andrew Sabri, E Dana Cox, Chris Fierer, Matthew Peterson, Bill Carstensen, and Chris North. 2005. Evaluating the benefits of tiled displays for navigating maps. In *IASTED International Conference on Human-Computer Interaction*. 66–71.

[9] A. Batch, A. Cunningham, M. Cordeil, N. Elmqvist, T. Dwyer, B. H. Thomas, and K. Marriott. 2020. 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.

[10] Benjamin B Bederson and James D Hollan. 1994. Pad++: a zooming graphical interface for exploring alternate interface physics. In *Proceedings of the 7th Annual ACM Symposium on User Interface Software and Technology*. ACM, 17–26.

[11] Christophe Bortolaso, Matthew Oskamp, Greg Phillips, Carl Gutwin, and TC Graham. 2014. The effect of view techniques on collaboration and awareness in tabletop map-based tasks. In *Proceedings of the Ninth ACM International Conference on Interactive Tabletops and Surfaces*. ACM, 79–88.

[12] Doug A. Bowman, Ryan P. McMahan, and Eric D. Ragan. 2012. Questioning Naturalism in 3D User Interfaces. *Commun. ACM* 55, 9 (Sept. 2012), 78–88.

[13] Felix Brodkorb, Arjan Kuijper, Gennady Andrienko, Natalia Andrienko, and Tatiana Von Landesberger. 2016. Overview with details for exploring geo-located graphs on maps. *Information Visualization* 15, 3 (2016), 214–237.

[14] Simon Butscher, Kasper Hornbæk, and Harald Reiterer. 2014. SpaceFold and PhysicLenses: simultaneous multifocus navigation on touch surfaces. In *Proceedings of the 2014 International Working Conference on Advanced Visual Interfaces*. ACM, 209–216.

[15] Tom Chandler, Maxime Cordeil, Tobias Czauderna, Tim Dwyer, Jaroslaw Glowacki, Cagatay Goncu, Matthias Klapperstueck, Karsten Klein, Kim Marriott, Falk Schreiber, et al. 2015. Immersive analytics. In *2015 Big Data Visual Analytics (BDVA)*. IEEE, 1–8.

[16] Andy Cockburn, Amy Karlson, and Benjamin B Bederson. 2009. A review of overview+ detail, zooming, and focus+ context interfaces. *ACM Computing Surveys (CSUR)* 41, 1, Article 2 (2009), 31 pages.

[17] Andy Cockburn and Bruce McKenzie. 2001. 3D or not 3D? Evaluating the effect of the third dimension in a document management system. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 434–441.

[18] Andy Cockburn and Bruce McKenzie. 2002. Evaluating the effectiveness of spatial memory in 2D and 3D physical and virtual environments. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 203–210.

[19] Christopher Collins and Sheelagh Carpendale. 2007. VisLink: Revealing relationships amongst visualizations. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1192–1199.

[20] Maxime Cordeil, Andrew Cunningham, Tim Dwyer, Bruce H Thomas, and Kim Marriott. 2017. ImAxes: Immersive axes as embodied affordances for interactive multivariate data visualisation. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*. ACM, 71–83.

[21] Mary Czerwinski, Desney S Tan, and George G Robertson. 2002. Women take a wider view. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 195–202.

[22] Niklas Elmqvist, Nathalie Henry, Yann Riche, and Jean-Daniel Fekete. 2008. Melange: space folding for multi-focus interaction. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1333–1342.

[23] Niklas Elmqvist, Andrew Vande Moere, Hans-Christian Jetter, Daniel Cernea, Harald Reiterer, and TJ Jankun-Kelly. 2011. Fluid interaction for information visualization. *Information Visualization* 10, 4 (2011), 327–340.

[24] Barrett Ens, Juan David Hincapié-Ramos, and Pourang Irani. 2014. Ethereal planes: a design framework for 2D information space in 3D mixed reality environments. In *Proceedings of the 2nd ACM Symposium on Spatial User Interaction*. ACM, 2–12.

[25] Barrett Ens, Eyal Ofek, Neil Bruce, and Pourang Irani. 2015. Spatial Constancy of Surface-Embedded Layouts Across Multiple Environments. In *Proceedings of the 3rd ACM Symposium on Spatial User Interaction* (Los Angeles, California, USA) *(SUI '15)*. ACM, New York, NY, USA, 65–68.

[26] Barrett M Ens, Rory Finnegan, and Pourang P Irani. 2014. The personal cockpit: a spatial interface for effective task switching on head-worn displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 3171–3180.

[27] Andreas Fender, David Lindlbauer, Philipp Herholz, Marc Alexa, and Jörg Müller. 2017. HeatSpace: Automatic Placement of Displays by Empirical Analysis of User Behavior. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology* (Québec City, QC, Canada) *(UIST '17)*. ACM, New York, NY, USA, 611–621.

[28] George W Furnas. 1986. Generalized fisheye views. *SIGCHI Bulletin* 17, 4 (1986), 16–23.

[29] George W Furnas and Benjamin B Bederson. 1995. Space-scale diagrams: Understanding multiscale interfaces. In *CHI*, Vol. 95. 234–241.

[30] Sohaib Ghani, N Henry Riche, and Niklas Elmqvist. 2011. Dynamic Insets for Context-Aware Graph Navigation. In *Computer Graphics Forum*, Vol. 30. Wiley Online Library, 861–870.

[31] Yves Guiard and Michel Beaudouin-Lafon. 2004. Target acquisition in multiscale electronic worlds. *International Journal of Human-Computer Studies* 61, 6 (2004), 875–905.

[32] Edward Twitchell Hall. 1966. *The Hidden Dimension*. Garden City, NY: Doubleday.

[33] Kasper Hornbæk and Erik Frøkjær. 2003. Reading patterns and usability in visualizations of electronic documents. *ACM Transactions on Computer-Human Interaction (TOCHI)* 10, 2 (2003), 119–149.

[34] Alexandra Ion, Y-L Betty Chang, Michael Haller, Mark Hancock, and Stacey D Scott. 2013. Canyon: providing location awareness of multiple moving objects in a detail view on large displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 3149–3158.

[35] Yvonne Jansen, Jonas Schjerlund, and Kasper Hornbæk. 2019. Effects of Locomotion and Visual Overview on Spatial Memory when Interacting with Wall Displays. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–12.

[36] Waqas Javed and Niklas Elmqvist. 2012. Exploring the design space of composite visualization. In *2012 IEEE Pacific Visualization Symposium*. IEEE, 1–8.

[37] Waqas Javed, Sohaib Ghani, and Niklas Elmqvist. 2012. Polyzoom: multiscale and multifocus exploration in 2d visual spaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 287–296.

[38] Charles Jekel. 2016. Obtaining Non-linear Orthotropic Material Models for PVC-Coated Polyester via Inverse Bubble Inflation. Thesis, Stellenbosch University. https://hdl.handle.net/10019.1/98627

[39] Susanne Jul and George W Furnas. 1998. Critical zones in desert fog: aids to multiscale navigation. In *Proceedings of the 11th Annual ACM Symposium on User Interface Software and Technology*. ACM, 97–106.

[40] Pushpak Karnick, David Cline, Stefan Jeschke, Anshuman Razdan, and Peter Wonka. 2009. Route visualization using detail lenses. *IEEE Transactions on Visualization and Computer Graphics* 16, 2 (2009), 235–247.

[41] Matthias Klapperstueck, Tobias Czauderna, Cagatay Goncu, Jaroslaw Glowacki, Tim Dwyer, Falk Schreiber, and Kim Marriott. 2018. ContextuWall: Multi-site collaboration using display walls. *Journal of Visual Languages & Computing* 46 (2018), 35–42.

[42] Heidi Lam and Tamara Munzner. 2010. A guide to visual multi-level interface design from synthesis of empirical study evidence. *Synthesis Lectures on Visualization* 1, 1 (2010), 1–117.

[43] Nina Siu-Ngan Lam and Dale A Quattrochi. 1992. On the issues of scale, resolution, and fractal analysis in the mapping sciences. *The Professional Geographer* 44, 1 (1992), 88–98.

[44] Fritz Lekschas, Michael Behrisch, Benjamin Bach, Peter Kerpedjiev, Nils Gehlenborg, and Hanspeter Pfister. 2019. Pattern-Driven Navigation in 2D Multiscale Visualizations with Scalable Insets. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (2019), 611–621.

[45] Lars Lischke, Sven Mayer, Jan Hoffmann, Philipp Kratzer, Stephan Roth, Katrin Wolf, and Paweł Woźniak. 2017. Interaction techniques for window management on large high-resolution displays. In *Proceedings of the 16th International Conference on Mobile and Ubiquitous Multimedia*. 241–247.

[46] Lee Lisle, Xiaoyu Chen, JK Edward Gitre, Chris North, and Doug A Bowman. 2020. Evaluating the Benefits of the Immersive Space to Think. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*. IEEE, 331–337.

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

[48] Zhicheng Liu and Jeffrey Heer. 2014. The effects of interactive latency on exploratory visual analysis. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014), 2122–2131.
[49] Zhicheng Liu and John Stasko. 2010. Mental models, visual reasoning and interaction in information visualization: A top-down perspective. *IEEE Transactions on Visualization and Computer Graphics* 16, 6 (2010), 999–1008.
[50] Paul Lubos, Gerd Bruder, Oscar Ariza, and Frank Steinicke. 2016. Touching the Sphere: Leveraging Joint-Centered Kinespheres for Spatial User Interaction. In *Proceedings of the 2016 Symposium on Spatial User Interaction* (Tokyo, Japan) *(SUI '16)*. ACM, New York, NY, USA, 13–22.
[51] National Aeronautics and Space Administration (NASA). 1995. NASA-STD-3000: Man-Systems Integration Standards, Revision B, July 1995, Volume I. https://msis.jsc.nasa.gov/
[52] Tao Ni, Doug A Bowman, and Jian Chen. 2006. Increased display size and resolution improve task performance in information-rich virtual environments. In *Proceedings of Graphics Interface 2006*. Canadian Information Processing Society, 139–146.
[53] Matthew Plumlee and Colin Ware. 2002. Zooming, multiple windows, and visual working memory. In *Proceedings of the Working Conference on Advanced Visual Interfaces*. ACM, 59–68.
[54] Matthew D Plumlee and Colin Ware. 2006. Zooming versus multiple window interfaces: Cognitive costs of visual comparisons. *ACM Transactions on Computer-Human Interaction (TOCHI)* 13, 2 (2006), 179–209.
[55] Quang Quach and Bernhard Jenny. 2020. Immersive visualization with bar graphics. *Cartography and Geographic Information Science* (2020).
[56] Khairi Reda, Andrew E Johnson, Michael E Papka, and Jason Leigh. 2015. Effects of display size and resolution on user behavior and insight acquisition in visual exploration. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 2759–2768.
[57] George Robertson, Mary Czerwinski, Kevin Larson, Daniel C Robbins, David Thiel, and Maarten Van Dantzich. 1998. Data mountain: using spatial memory for document management. In *Proceedings of the 11th Annual ACM Symposium on User Interface Software and Technology*. 153–162.
[58] George Robertson, Maarten Van Dantzich, Daniel Robbins, Mary Czerwinski, Ken Hinckley, Kirsten Risden, David Thiel, and Vadim Gorokhovsky. 2000. The Task Gallery: a 3D window manager. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 494–501.
[59] George G Robertson, Jock D Mackinlay, and SK Card. 1991. The perspective wall: Detail and context smoothly integrated. In *Proceedings of ACM CHI*, Vol. 91. 173–179.
[60] Chris Rooney, Alex Endert, Jean-Daniel Fekete, Kasper Hornbæk, and Chris North. 2013. Powerwall: int. workshop on interactive, ultra-high-resolution displays. In *CHI'13 Extended Abstracts on Human Factors in Computing Systems*. 3227–3230.
[61] Chris Rooney and Roy A Ruddle. 2015. HiReD: a high-resolution multi-window visualisation environment for cluster-driven displays. In *Proceedings of the 7th ACM SIGCHI Symposium on Engineering Interactive Computing Systems*. 2–11.
[62] Robert E Roth. 2013. An empirically-derived taxonomy of interaction primitives for interactive cartography and geovisualization. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2356–2365.
[63] Roy A Ruddle, Rhys G Thomas, Rebecca S Randell, Phil Quirke, and Darren Treanor. 2015. Performance and interaction behaviour during visual search on large, high-resolution displays. *Information Visualization* 14, 2 (2015), 137–147.
[64] Vít Rusnák, Caroline Appert, Olivier Chapuis, and Emmanuel Pietriga. 2018. Designing coherent gesture sets for multi-scale navigation on tabletops. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 1–12.
[65] K. A. Satriadi, B. Ens, M. Cordeil, B. Jenny, T. Czauderna, and W. Willett. 2019. Augmented Reality Map Navigation with Freehand Gestures. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. 593–603.
[66] Lauren Shupp, Robert Ball, Beth Yost, John Booker, and Chris North. 2006. Evaluation of viewport size and curvature of large, high-resolution displays. In *Proceedings of Graphics Interface 2006*. Canadian Information Processing Society, 123–130.
[67] Monica Tavanti and Mats Lind. 2001. 2D vs 3D, implications on spatial memory. In *IEEE Symposium on Information Visualization, 2001. INFOVIS 2001*. IEEE, 139–145.
[68] Lucia Terrenghi, Aaron Quigley, and Alan Dix. 2009. A taxonomy for and analysis of multi-person-display ecosystems. *Personal and Ubiquitous Computing* 13, 8 (2009), 583–598.
[69] Sabine Timpf. 1998. Hierarchical structures in map series. Thesis, Department of Geoinformation, Technical University Vienna.
[70] Jorge A Wagner Filho, Wolfgang Stuerzlinger, and Luciana Nedel. 2019. Evaluating an immersive space-time cube geovisualization for intuitive trajectory data exploration. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (2019), 514–524.

[71] Fangzhou Wang, Yang Li, Daisuke Sakamoto, and Takeo Igarashi. 2014. Hierarchical route maps for efficient navigation. In *Proceedings of the 19th International Conference on Intelligent User Interfaces*. ACM, 169–178.

[72] Michelle Q Wang Baldonado, Allison Woodruff, and Allan Kuchinsky. 2000. Guidelines for using multiple views in information visualization. In *Proceedings of the Working Conference on Advanced Visual Interfaces*. ACM, 110–119.

[73] Daniel Wigdor and Dennis Wixon. 2011. *Brave NUI World: Designing Natural User Interfaces for Touch and Gesture*. Morgan Kaufmann.

[74] Yalong Yang, Maxime Cordeil, Johanna Beyer, Tim Dwyer, Kim Marriott, and Hanspeter Pfister. 2020. Embodied Navigation in Immersive Abstract Data Visualization: Is Overview+ Detail or Zooming Better for 3D Scatterplots? *IEEE Transactions on Visualization and Computer Graphics* (2020).

[75] Yalong Yang, Tim Dwyer, Bernhard Jenny, Kim Marriott, Maxime Cordeil, and Haohui Chen. 2018. Origin-destination flow maps in immersive environments. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2018), 693–703.

[76] Yalong Yang, Tim Dwyer, Kimbal Marriott, Bernhard Jenny, and Sarah Goodwin. 2020. Tilt Map: Interactive Transitions Between Choropleth Map, Prism Map and Bar Chart in Immersive Environments. *IEEE Transactions on Visualization and Computer Graphics* (2020).

[77] Shumin Zhai, Paul Milgram, and William Buxton. 1996. The influence of muscle groups on performance of multiple degree-of-freedom input. In *CHI*, Vol. 96. 308–315.