

Design and Evaluation of Interactive Small Multiples Data Visualisation in Immersive Spaces

Jiazhou Liu*

Arnaud Prouzeau†

Barrett Ens‡

Tim Dwyer§

Monash University

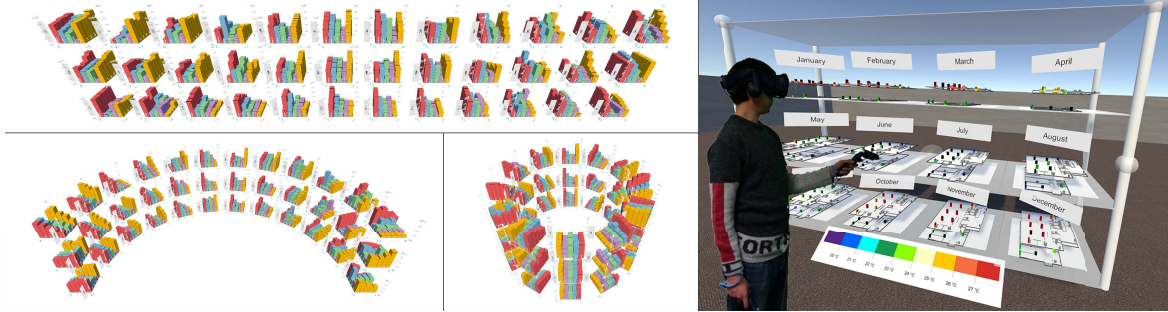


Figure 1: Left: different possibilities for layout curvature; right: small multiples presented in VR using a “shelves” metaphor.

ABSTRACT

We explore the adaptation of 2D small-multiples visualisation on flat screens to 3D immersive spaces. We use a “shelves” metaphor for layout of small multiples and consider a design space across a number of layout and interaction dimensions. We demonstrate the applicability of a prototype system informed by this design space to data sets from different domains. We perform two user studies comparing the effect of the shelf curvature dimension from our design space on users’ ability to perform comparison and trend analysis tasks. Our results suggest that, with fewer multiples, a flat layout is more performant despite the need for participants to walk further. With an increase in the number of multiples, this performance difference disappears due to the time participants had to spend walking. In the latter case, users prefer a semi-circular layout over either a fully surrounding or a flat arrangement.

Index Terms: H.5.1 [Multimedia Information Systems]: Artificial, augmented, and virtual realities; H.5.2 [User Interfaces]: Evaluation/methodology

1 INTRODUCTION

Currently, data visualisation—as with most interactive computing tasks—is overwhelmingly performed using flat screens. For data that is inherently 3D, such as aircraft trajectories or building models, this presents a problem, as rendering 3D data displays on screens is well known to suffer from issues of occlusion, perspective distortion and so on, and generally a loss of information [39]. However, as augmented and virtual reality (AR and VR) headset devices improve in tracking stability, field of view, and resolution, there is a real possibility that traditional screens may be cheaply replaced with wearable headsets that offer immersive display of such 3D data. If headsets do replace screens it presents both an opportunity and a

challenge for data visualisation. It could be a paradigm shift in allowing immersive data visualisation in the context of other activities (i.e. situated analytics [50]). But it will also represent an interaction design challenge in translating everything we have learned about visualisation design on flat screens to the spaces around us. How do we adapt all of the common visualisation idioms and interaction techniques to take advantage of this space?

In this paper, we consider the adaptation of a common visualisation design pattern that is very well studied on screens to immersive interaction spaces; namely *small multiples displays*, wherein a number of different data sets are represented using the same visualisation idiom in a tiled display to support easy comparison. As discussed in our Related Work section, small multiples displays are ubiquitous in many domains, e.g. stock market trading floors, scatterplot matrices, tiled medical images, and so on. However, to our knowledge, layouts for small multiples have not previously been systematically explored or evaluated in 3D immersive environments.

We envision applications for immersive small multiples in domains that rely on exploration of 3D data. For instance, changes in aircraft trajectories above an airport over different time periods may be analysed to reveal patterns about the efficiency of airspace utilisation and risks of collision. Another relevant domain is Building Information Modeling (BIM), which is concerned with the management of a facility’s digital information assets. Building managers may benefit by comparing temperature sensor readings and energy consumption over time to identify trends. Analysts may also be interested in data without a physical spatial embedding or abstract quantitative data, for instance using 3D bar charts to compare wealth and productivity statistics across different populations (e.g. Gapminder [49]).

Based on a design space for immersive small multiples layouts, we develop a prototype implementation with features supporting the three use cases above (air traffic data, building information models, and abstract 3D bar charts; Fig. 2). This implementation explores several interactions for manipulating 3D layouts and interacting with the data visualisations within. To help us understand the benefits of different layouts, we run two comparison studies with different numbers of multiples to compare layout curvature with different data types.

Our contributions include: (1) a design space for layout of and interaction with small multiples in an immersive environment; (2)

*e-mail:Jiazhou.liu@monash.edu

†e-mail:Arnaud.Prouzeau@monash.edu

‡e-mail:Barrett.ens@monash.edu

§e-mail:Tim.Dwyer@monash.edu




Figure 2: (a) Aircraft Trajectories; (b) Building Information Models (BIM); (c) Demographic indicators

a prototype system allowing us to explore layout and interaction designs; (3) two user studies evaluating the effect of introducing curvature into the shelves such that they wrap around the user; and (4) the finding that a flat layout is more efficient than curved with a small number of multiples although it requires more walking. With a high number of multiples, walking hinders the flat layout performance and user preference, fully enclosing circular shelves are particularly disorienting, but half circle is a popular compromise.

2 RELATED WORK

2.1 Small multiples on Conventional Displays

Small multiples are commonly used to perform visual comparisons through a tiled display of charts or models using the same axes and measure system [51]. That is, different data sets are represented using the same encoding [39]. Thus, they provide an overview of the data, but also allow for comparison with minimal interaction and without overloading the visual working memory [41]. In data exploration, small multiples have been shown to provide a broader perspective on the data to avoid missing important information [7]. Compared to other techniques, small multiples have also been shown to improve user performance for global time series tasks requiring the user to consider the entire display width [23].

The visualisations in small multiples are traditionally arranged in a grid with a fixed and predefined order. Liu et al. proposed reordering the grid to bring similar multiples together [33]. Javed et al. used a single column for small multiples of time series to ease temporal comparison [23]. Finally, Meulemans et al. designed an algorithm to break the grid to match the multiples with geographic locations [38]. A hybridisation approach of small multiples combines several visualisation types [36], to provide different perspectives on a graph [5] or to highlight differences between several maps [27].

2.2 Small multiples in immersive displays

On desktop displays there may be insufficient screen space for effective small multiples [22]. Research has focused on large displays for collaborative use of small multiples, e.g.: for software maps [44]; road traffic data [42]; and biological data [13]. Similarly in VR, Johnson et al. propose a system to visualise 3D small multiples on a flat layout [24]. Despite greater size, large displays only support the same flat grid layout available on desktop displays. With VR or AR, other layouts are possible. In FiberClay [20], Hurter et al. visualise small multiples on the ground, providing users an overview of a dataset, with the focus presented directly in front. In Encube [52], Vohl et al. use a circular layout to visualise small multiples, allowing the visualisation of a large number of multiples (up to 80), without increasing the distance between the user and each multiple.

Other research has explored 3D spatial layouts of 2D information displays [8] to support spatial memory [15] or analytic taskwork [9]. For instance, Virtual Shelves [28] distributes app shortcuts in an invisible hemisphere, which users can retrieve using spatial and kinaesthetic memory. Curved virtual “cockpit” [11] or “amphitheatre” [14] display layouts distribute items equidistant from the user, making them easier to view or select [54]. Other layouts embed virtual displays in the physical environment [10, 12], or situate them

in 3D space around desktop monitors [45] mobile devices [18] or smartwatches [17] to facilitate easy context switching.

While many layout possibilities have been demonstrated in immersive environments, no study has been done to validate the performance of small multiples in 3D space. We propose a design space for the such layouts, as well as a practical testbed system, and perform two studies to assess their performance.

2.3 Interaction with Small Multiples

Being able to interact with the small multiples is very important to facilitate visual comparison. A popular technique, which is used more generally in multiple views visualisation, is called brushing and linking [19]. With this technique, when a user selects points in one view, matching records in the other views will be selected. In Cerebral, Barsky et al. extended the concept of linked views to navigation by applying pan and zoom to all the multiples [16]. Finally, in their tool Dream Lens, Matejka et al. [37], allow users to transition from small multiples to a superimposed view of several multiples.

In Encube [52], users can interact with the multiples using a handheld device. It allows them to rotate the multiples either globally or individually. Virtual Reality systems tend to favour direct spatial interaction techniques. In IATK [35], Cordeil et al. propose using controllers to directly brush in a multiple views visualisation. Following this initiative, we apply direct brushing and linking interaction in our prototype.

3 MOTIVATING SCENARIOS

Our design-space investigation is motivated by several real-world use cases of data with a natural 3D embedding (Fig. 2), that is difficult to present in small multiples on flat screens.

Aircraft Trajectories – Airport managers and air traffic analysts routinely analyse aircraft trajectories to assess the maximal capacity of an airport or before opening a new route. The use of 3D visualisations is very important as airplanes move in 3D space and important constraints apply on both latitude/longitude and altitude. In order to identify peak traffic period, it is important for them to be able to compare the traffic hourly, daily or weekly. We are developing immersive visualisation techniques for this data with domain experts from the aviation industries in France and Australia.

Building Information Models (BIM) – With the use of the Building Information Modeling format, Facilities managers have more and more opportunities to visualise data coming from different building sensors (e.g. CO₂, Temperature) overlaid on 3D CAD model of buildings. Seasonality is very important while looking at building data as the outside weather has a big influence on parameters like Air Conditioning and Lighting. It is then important for them to be able to aggregate and visualise their data per hours, days, months, etc. We are trialling small multiples displays with the Buildings and Properties department at our institution, as well as a major commercial supplier of building management systems.

Demographic indicators – When looking at demographic data like population, GDP, spending in different areas, it is important to see both the temporal and spatial evolution (for instance by years and countries). The use of 3D barcharts in small multiples array allows for four dimensions of data to be viewed simultaneously, and potentially for trends involving more than two variables. For instance, in some countries, the population can increase with the GDP, while others will see their GDP decrease when the population increase. Demographic data from the GapMinder website [49] is a popular and reliable baseline dataset in visualisation research.

4 DESIGN SPACE FOR IMMERSIVE SMALL MULTIPLES

While small multiples layouts have been explored in traditional flat-screen implementations, there is no existing design space to describe such layouts in 3D space. We identify 4 design dimensions that

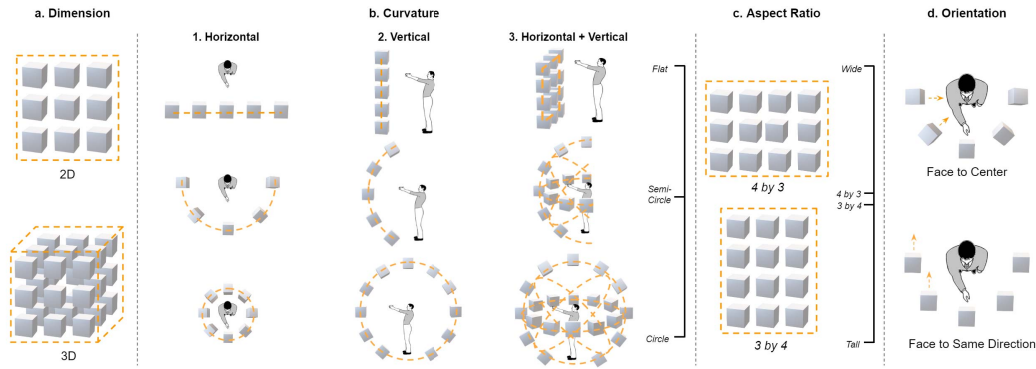


Figure 3: Immersive small multiples layout design space: (a) Dimension, (b) Curvature, (c) Aspect Ratio, and (d) Orientation.

describe many possible layouts of small multiples for 3D data in immersive environments (Fig. 3):

Dimension – We refer to the dimensionality of the grid of small multiples. A 1D display would be a single row, 2D is the traditional grid used on screens, while 3D is a new possibility afforded by immersive environments, adding a depth dimension to the grid. Adding more dimensions allows more multiples to be compacted into a volume but stacking in the depth dimension will introduce occlusions.

Curvature – Allows multiples to wrap around the user, reducing the need for walking. While curving a 1D layout is relatively straightforward, there are several possible ways to curve layouts in higher dimensions (e.g. curving a 2D layout into a cylinder or a sphere).

Aspect ratio – Relates the number of multiples in each orthogonal dimension, e.g., a 2D array of 12 multiples can be arranged in ratios: 4×3 , 3×4 , 2×6 , etc.

Orientation – Refers to the relative orientations of the individual 3D data visualisations. Instinctively, one might align all the layouts to the same forward-facing direction, similar to flat-screen 2D layouts. However, with a curved layout (or potentially with a flat layout), rotating each visualisation to face the user may help them to more easily make comparisons.

Our design uses a *shelves* metaphor [28] to provide cues for interaction. Fig. 1 shows the realisation of this metaphor in our prototype system. The shelf visuals provide affordances for users to understand and orient 3D small multiples and provide clear horizontal and vertical alignment of the small multiples to enhance spatial memory. Interactive elements, such as corner pillars, provide an interface for users to directly manipulate [46] the display layout, for instance changing its *curvature* or *aspect ratio* to best suit the data and task. In the following section we discuss how our implementation allows users to manipulate, brush and filter the data, with visual feedback coordinated across all of the multiples.

5 PROTOTYPE IMPLEMENTATION

We developed a VR prototype to explore the design space for immersive small multiples. It supports both manipulation of the layout and coordinated interaction with the small multiples.

Apparatus – We use an HTC Vive Pro room-scale VR device and the Unity development environment (2017.3.0f3). The prototype runs on a Windows 10 PC with an Intel I7 7800X (3.5GHz) processor and an NVIDIA GeForce GTX 1080 (32GB RAM) graphics card. We leverage VRTK [4] for interactive components, and IATK [35] for rendering for the small multiples data visualisations. The code is available on GitHub [30].

5.1 Interacting with Shelves

In addition to walking around the data and viewing it from different perspectives, users can reconfigure the small multiples layout by “grasping” and manipulating different components of the shelves’ visible form. Affordances for layout operations are revealed to users by visible “handles” on the pillars or shelves, which also provide visual feedback during manipulations. Since the shelf is too large for users to easily reach the pillars, the handles can be manipulated from a distance through a ranged pointing gesture with the Vive controllers. Pointing rays extend from each controller to provide additional visual feedback. Implemented operations include changing the layout aspect ratio, curvature, height, detail level, or shelf position.

Aspect ratio – By grabbing one of the front shelf pillars, the shelving unit can be “stretched” or “compressed” horizontally (Fig. 4–a, b). As the shelf width changes, the multiples automatically rearrange themselves to fit the new aspect ratio, with animated movements between shelf positions.

Curvature – Grabbing both front pillars simultaneously allows the shelf to be “bent”, adjusting the layout curvature (Fig. 4–c, d). The shelves can be adjusted continuously between a straight layout and a half-circle configuration.

Height – Grabbing the top corner of either pillar allows users to “stretch” the shelves vertically. This adjusts the height of the shelving unit, and accordingly the space between shelves, but without changing the aspect ratio.

Detail level – With some data sets, such as BIM data, users may want to adjust the level of semantic detail shown. Pressing a button on the controller increases the separation between the floors of the building model, i.e. creating an “exploded” view.

Shelf position – Grabbing a handle in the centre of the shelving unit allows users to move the entire layout horizontally or vertically. This can be used to get a closer view of distant multiples, for instance the far edges or lower shelves, without walking or crouching.

5.2 Interacting with Data

We implemented several operations for interacting with the 3D data visualisations to allow us to investigate the use of the small multiples layouts with data analytics tasks. These include selection, rotation, brushing, filtering and a ruler tool. These operations (except for selection) are coordinated, so that manipulating any single data visualisation results in the same effect applied to all multiples.

Selection – Users can select one of the small multiples, either by moving the controller near and pressing the trigger button, or by using the controller’s ray from afar (Fig. 4f).




Figure 4: Interactions with the shelf layout (top) and contained data visualisations (bottom): (a,b) adjusting layout aspect ratio by “grabbing” and moving a shelf post, (c,d) adjusting layout curvature by moving both posts, (e) rotating multiples via the ViewCube [25], (f) brushing a single data point, (g) brushing an axis using both controllers, (h) brushing a volume selection on all axes, (i) filtering on the y-axis with cutting planes.

Rotation – To view the data visualisations from different sides, users can press both controller triggers to present a view cube [25] (Fig. 4e). Users can then manipulate the cube rotation, which is reflected across all multiples.

Brushing – Brushing [19] allows users to select one or more data points on a single visualisation, and see the selection linked across all coordinated views. We provide several brushing methods. Users can brush a single data point with a pointer that extends from the controller (Fig. 4f). Bimanual interaction of a pair of sliders on any axis brushes a range in one dimension (Fig. 4g). Finally, a bimanual gesture within a data volume brushes a cube shaped region on all 3 axes at once (Fig. 4h).

Filtering – A pair of cutting planes can be moved along the vertical axes to select a specific range of values (Fig. 4i).

Ruler tool – When touching the vertical axis of a data visualisation the pointer is annotated with a numeric value (Fig. 4i) supporting detailed height comparisons across multiples.

6 USER STUDY 1: LAYOUT COMPARISON

It is unclear how the layout of small multiples in immersive spaces impacts performance of users in a visual comparison task. Therefore, we evaluate three different layouts with two different datasets for such tasks. The design of this experiment has been preregistered on Open Science Framework (OSF) [31].

6.1 Study Design

Task – Our task consists of a visual comparison between pairs of visualisations that are part of a small multiples display. More specifically, participants have to compare the value of two specific data points between two specific multiples in a total of 12 multiples placed in a 2D shelves with grid of 4 columns and 3 rows.

We test three layouts (LAYOUT): *Flat*, *Quarter-Circle* and *Half-Circle*. We focus on horizontal curvature, similar to existing large displays, to prevent combinatorial explosion. We chose not to evaluate a full circle layout, because with just 12 multiples, either the circle will be too small, or the distance between columns too large.

To provide generalizable results, we include two common data sets in our studies: *Bar*, which is a typical representation of multi-dimensional, non-spatial 3D data, and *BIM*, which contains data that has a spatial reference frame such as a floor plan.

Finally, we vary the locations of the multiples to compare within the grid, controlling for distance between the two that need to be

compared. In one condition, *Short*, the two multiples are at a Manhattan distance of 1 or 2, meaning that participants can do the task with both multiples simultaneously in their field of view. In the second one, *Long*, the Manhattan distance were 4 or 5.

Dataset – We use one dataset for each DATASET condition. The data for *Bar* is 12 indicators for 10 years and 10 countries from the GapMinder website [49] (Fig. 1-left). One multiple represents the evolution of one indicator for 10 countries over 10 years. For the *BIM* dataset, we took inside temperature data from a building at our institution. There are 25 temperature sensors in this building. We aggregated the temperature by months for each sensor. Each multiple shows the mean temperature for one month for each sensor (Fig. 1-right). The questions were created manually by the authors.

Design – We used a within-subjects design with 3 factors: 3 LAYOUT (*Flat*, *Quarter-Circle* and *Half-Circle*) \times 2 DATASET (*Bar* and *BIM*) \times 2 COMPARISON DISTANCE (*Short* and *Long*). There were 2 repetitions for each combination, which yielded a total of 288 trials with 12 participants. A latin square design counterbalanced the order of LAYOUT and DATASET. For COMPARISON DISTANCE, participants did *Short* first and then *Long* within each condition.

Based on a pilot study (4 participants), the related work and our design space, we formulate a number of hypotheses. Hypotheses D1 and D2 are related to the difficulty of the task, L1-2 consider effect of layout on performance, M1-2 are related to participants’ movement and P1 relates to participants’ preferences.

We expect participant performance (time and accuracy):

D1 – will be better with *Short* comparison distance than *Long* - because for *Short* they will be more easily able to have both visualisations in view;

D2 – will be better with less abstract data in the visualisation (i.e. *BIM* better than the *Bar*);

L1 – will be better for *Short* distance, *Half-Circle* layout over other layouts - because it should involve least participant movement to see both visualisations simultaneously;

L2 – will be better for *Long* distance, *Flat* layout over other layouts - because it allows participants to step back to see both visualisations simultaneously;

In terms of participant movement, we expect:

M1 – greater distance covered with *Flat* layout than *Quarter-Circle* and *Half-Circle* ones.

M2 – more back and forth movement between the two multiples with the *Quarter-Circle* and *Half-Circle* than *Flat*.

For preference we expect:




Figure 5: The shelf configuration: (a) shelf height, (b) flat shelf length, (c) quarter-circle shelf arc length, (d) half-circle shelf arc length

P1 – *Quarter-Circle* will be preferred over either *Half-Circle* (requires too much rotation) or *Flat* (requires too much movement).

Procedure – After completing the consent form and a demographic questionnaire participants were trained with three interaction techniques: rotation, brushing and filtering. Each study block (i.e. one LAYOUT and one DATASET) followed the same process: the eye tracker was calibrated; followed by two training trials; then four experiment trials; and then a short questionnaire regarding the condition.

At the beginning of each trial, participants were asked to position themselves at a centre position indicated on the ground. Then the question was displayed and they had time to read it. They then had to press a controller trigger to reveal the small multiples and begin the task. By pressing the trigger again, they stopped the task, the small multiples disappeared and they could answer the question by choosing their answer in a menu. They had a maximum of 60 seconds to do the task, after which the small multiples disappeared and they had to choose an answer.

In our initial pilots we found participants spent an inordinate amount of time finding the two visualisations named in the question. We wanted to focus in this study on participants’ ability to compare visualisations at a distance rather than the spatial search task, which is not specific to visualisation. We therefore highlighted the two visualisations for comparison from the beginning of the trial.

Apparatus – We used the apparatus and prototype described in the “Implementation” section. The only difference was that all interactions with the shelf itself were blocked. Each small multiple’s size is $0.4m^3$. The horizontal and vertical offset between each pair of small multiples is 0.15 m (Fig. 5). The shelf height was adjusted dynamically, so that the top edge of the top multiple was aligned with the participant’s eye level (based on pilot studies). The only available interactions were: brushing, filtering and rotation of the small multiples.

Participants – We recruited 12 participants (7 males and 5 females. mean age=27.5 and SD=4.7). 4 participants had already experienced VR 2 used Small Multiples before and 5 had experience with brushing and linking techniques.

Measures – We recorded completion time (i.e. the time between pressing the trigger after reading the question until pressing the trigger to give their answer) and accuracy (whether the participants found the correct answer or not) for each trial. In the analysis, we are actually looking at the percentage of wrong answer for each condition. At the end of each block of trials with the same LAYOUT, preference was measured using a ranking. They were also asked about their strategy to solve the task. Participants’ head was tracked during the entire experiment, which we use to calculate the distance they travelled during trials. Finally, we use an eye-tracker to find the number of times they switched between the two small multiples under comparison.

6.2 Results

Overall task difficulty between COMPARISON DISTANCE and DATASET conditions is shown in Fig. 6. The remaining results were analysed for each DATASET individually, as we are more interested in the nuances with each condition than overall effects.

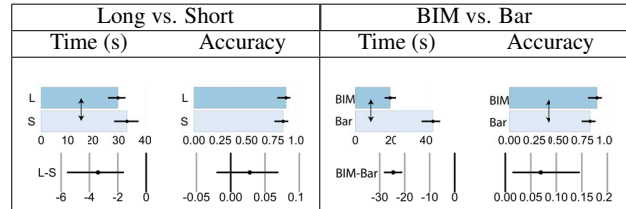


Figure 6: The top bars compare *Long* (L) and *Short* (S) conditions (left), and *BIM* and *Bar* conditions (right). The bottom charts show corresponding 95% CIs for the mean differences. Arrows indicate significant difference between two conditions.

Results for time, accuracy, distance and gaze are shown in Fig. 7, and participant rankings of the different layouts are shown in Fig. 8.

Statistical Method – Following APA recommendations [1], we report our analysis using estimation techniques with confidence intervals and effect sizes (i.e., not using *p*-values) following recent precedents in HCI [3, 53]. Our confidence intervals were computed using BCa bootstrapping, and the term *effect size* here refers to the measured difference of means. Error bars in our charts reporting means are computed using all data for a given condition. When comparing means, we average the data by participants/groups and compare the three conditions globally by computing the CI of the set of differences. A difference is considered as significant when the CI of the difference do not cross 0. In our charts we display the computed CI of the differences. While we make use of estimation techniques, a *p*-value-approach reading of our results can be done by comparing our CIs spacing with common *p*-value spacing as shown by Krzywinski and Altman [26]. No corrections for multiple comparisons were performed [6, 40]. All the results reported in the analysis are significant.

Difficulty – Participants were faster (~ 3 seconds) to answer questions with *Long* than with *Short*, however, this difference may be attributed to a learning effect, since participants always completed the *Short* condition first. We did find any difference between the two conditions in accuracy. Participants were faster (~ 25 seconds) to answer questions on the *BIM* condition than on the *Bar* one. Accuracy is also a little lower (~ 6%) in the *Bar* condition than with the *BIM* one.

6.2.1 Bar Dataset

Time – Overall, participants were faster to complete the task with the *Flat* layout (*Quarter-Circle* was ~ 6 seconds slower, *Half-Circle* ~ 8 seconds slower). When we look at the different COMPARISON DISTANCE, we can see that there is a difference only for the *Long* comparison distance with participants being faster by ~ 9 seconds compared to the *Quarter-Circle* and 14 seconds with the *Half-Circle*.

Accuracy – Overall, participants were less accurate with the *Quarter-Circle* layout (0.3 against 0.1 for the *Flat* and 0.12 for the *Half-Circle*). This difference is present for the *Long* comparison distance, where participants have an error rate of 0.3 with the *Quarter-Circle* layout against one of 0.04 with the two others.

Travel Distance – Overall, we did not see any difference in the distance participants travelled during trials between the different layouts. We do find a difference for the *Long* comparison distance between *Flat* and *Half-Circle* layout (~ 0.7 m with *Half-Circle*).

Gaze Change – Overall, participants did 1 more back-and-forth between the two multiples with the *Flat* layout than with the *Half-Circle* one. This is reflected in both conditions with a similar effect.

Preferences – When we look at the ranking of the LAYOUT, we can see that 5 participants ranked *Quarter-Circle* first, against 4 participants for the *Flat* and 3 for the *Half-Circle*. Six participants ranked the *Flat* and the *Half-Circle* layouts last, against none for the

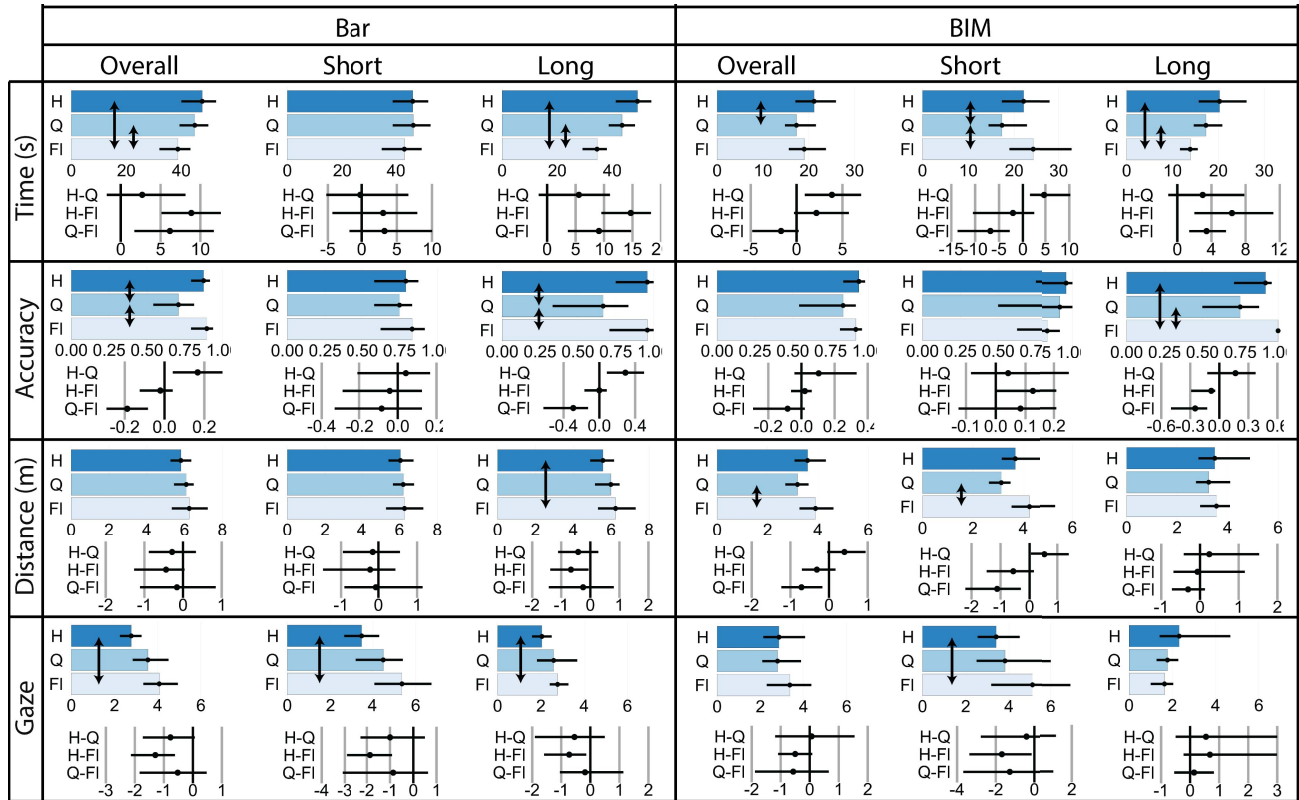


Figure 7: The top chart of each pair shows Means and CIs for all measures for Layout (Flat (FI), Quarter-Circle (Q) and Half-Circle (H)) across conditions. The bottom chart shows 95% CIs for the mean differences between Layouts. Arrows indicate significant differences between conditions.

Quarter-Circle one. In their comments, a few participants said that the *Flat* layout required too much walking (3/12), but one mentioned that it was possible to easily have an overview, without rotating. The rotation with the *Half-Circle* layout was considered an issue by some participants (5/12). Finally, the *Quarter-Circle* was considered as a good compromise between walking and rotating (6/12).

6.2.2 BIM Dataset

Time – The only difference found is between the *Quarter-Circle* and the *Half-Circle* layout (*Quarter-Circle* faster by 4 seconds). If we break by COMPARISON DISTANCE, the *Quarter-Circle* layout was slightly faster for the *Short* comparison distance (~ 6 seconds against the *Flat* and 4 seconds against the *Half-Circle*), while *Flat* was faster for the *Long* comparison distance (~ 6 seconds against the *Half-Circle* and 3 seconds against the *Quarter-Circle*).

Accuracy – Overall, there is no difference between the three LAYOUT. There is also no difference for the *Short* comparison distance. For the *Long* comparison distance, participants have a lower error rate with the *Flat* layout (0.0) than with the *Quarter-Circle* (0.25) and the *Half-Circle* (0.09).

Travelled Distance – Overall, there is a difference between the travelled distance between the *Quarter-Circle* and the *Flat* (~ 0.7m). The same difference can be seen for *Short* comparison distance (~ 1.1m), and additionally, there is a difference between the *Quarter-Circle* and the *Half-Circle* layout (~ 0.5m). No difference can be observed for the *Long* comparison distance.

Gaze Change – Overall, we can not find any difference between the number of back-and-forth switches between the two multiples. We can only see a difference in the *Short* in which participants did almost 2 more switches with the *Flat* than with the *Half-Circle* one.

Preferences – Six participants ranked *Quarter-Circle* and *Half-Circle* layout first, against 2 for *Flat*. Only 1 ranked *Quarter-Circle* layout last, against 5 for *Flat* and 6 for *Half-Circle*. Similar to *Bar*, participants found walking an issue with *Flat* (4/12), but liked the overview that it allowed without rotation (3/12). Rotation in *Half-Circle* was also considered an issue (6/12), and *Quarter-Circle* was a good compromise between walking and rotation (8/12).

6.3 Discussion

The fact that participants were faster with the *Long* condition than with *Short* is interesting (rejecting *D1*). This might be due to several phenomena, including a learning effect due to participants completing *Short* first in each condition, but also due to only 12 small multiples being insufficient to really force much movement. On the other hand, participants were clearly faster in *BIM* condition than in *Bar* (confirming *D2*), which is expected due to the number of candidate data points for comparison.

Regarding performance in *Short*, participants were faster with *Quarter-Circle* for the *BIM* dataset (Rejecting *L1*). No other difference was found regarding time or accuracy. The layout in the *Short* condition probably does not impact performance.

In *Long* condition, participants were faster and more accurate for both datasets with *Flat* layout (supporting *L2*) despite reporting that they found they had to walk more. In fact, analysis of tracked movements revealed only a small increase in movement (there is difference of only 0.7 meters against the *Half-Circle* with the *Bar* dataset) which likely did not significantly affect their completion times. Also we found that participants routinely found positions such that their field of view included both multiples under comparison.

Overall, regarding the distance travelled by participants, there is only a difference with the *BIM* layout (*M1* partially confirmed),




Figure 8: Result of the ranking for the three layouts for the Bar condition (left) and BIM condition (right).

and it is rather small (between 0.5 and 0.7 meters shorter than with the *Flat* layout). This may be due to the low number of multiples which means that participants do not have to walk much to closely inspect all of them. Regarding the number of gaze switches between the two multiples, we can see a global difference only in the *Bar* condition, and it shows that there is more with the *Flat* layout than with the *Half-Circle* one (rejecting M2). Plumlee and Ware explain that the less costly switches are in visual comparison, the more users are going to use them. This is in order to limit the load on their visual working memory. In our case, this seems confirmed by the finding that it is less easy to find a position where users can transition between the two multiples with the *Half-Circle* layout than with *Flat*.

Finally, participants stated that the *Flat* provides a good overview of the data, and that they can easily see all of the multiples at once and keep them in their peripheral view when they focus on one, which is not possible with the two other layouts.

7 USER STUDY 2: LARGE SCALE COMPARISON

The first user study indicated performance advantages of the flat layout, as it provides a broad overview without the need for rotation, despite the need for some walking. We conducted a second user study to determine whether these findings scale to a more extreme design, with a larger layout containing more multiples (from 12 to 36, with 12 columns and 3 rows). To better understand the effects of participant rotation, we also included a full-circle condition (and thus, removed the *Quarter-Circle* layout). Prior work in multi-display environments recommends to never place displays behind the user [48], however we are interested to see if this holds true in an immersive setting where the user is standing, rather than sitting. The design of this experiment has been preregistered on OSF [32].

7.1 Study Design

Tasks – We have two tasks in this follow up study. In both tasks, each multiple presents a 3D bar chart. The first task is the same as in the first study, it is a visual comparison between two multiples. In this study we have only one condition in which the multiples are at a Manhattan distance of 7 or 8.

For the second task, we wanted participants to have to look at all multiples. We decided to go with a task in which participants have to find the maximum value for a specific bars in the small multiples. **Dataset** – We used a dataset inspired by the world indicator dataset used for the *Bar* condition in the previous study. Each multiple represented the value of 5 indicators for 5 countries for a specific year (Figure 1-left). Contrary to the previous experiment, the data were simulated in order to easily create one with a maximum value. The questions were created manually by the authors.

Design – We used a within-subject design which consisted of: 3 LAYOUT (*Flat*, *Half-Circle* and *Full-Circle*) \times 2 TASK (*Comparison* and *Max*). There were 3 repetitions of each combination, which yielded a total of 216 trials for 12 participants. A latin square was used to counterbalance the order of LAYOUT. As we did not intend to compare the two TASK, order was fixed. Participants started with the *Comparison* task.

As this is an exploratory study (we removed a condition to try a new task), we did not have strong hypotheses regarding performance of each LAYOUT. Our goal was to observe the nuances of each and how they compare to each other. However, the metric and the

analysis methods were determined before the study: we explore the effect of LAYOUT performance, movement, gaze and preference.

Procedure – The procedure is similar to Study 1, except the maximum task time is increased to 90 seconds. For the *Max* task, no multiples were highlighted as participants should look at all of them.

Apparatus – We use the same apparatus and prototype as in Study 1. However, we adapted the shelf configuration for the large scale. Each small multiple's size now is $0.3 m^3$. The horizontal and vertical offsets between small multiples are 0.18m and 0.15m, respectively. The shelf height was adjusted dynamically as in study 1.

Participants – We recruited 12 participants (5 male, 7 female; mean age=26.1 with SD=3.9); 4 participants returned from Study 1; 8 had already experienced VR; 5 had used Small Multiples before; and 5 had experience with brushing and linking techniques.

Measures – We use the same statistical method and take the same measure as in Study 1. Using the head tracking data, we calculated the amount of time participants spent walking (over 1.8 km/h), which we called *Walking Time*. Finally, using the eye tracker data, we calculated the distance between the participants and the object they are looking at. We then calculate the amount of time they spent looking at objects that are more than 1 metre away (*Time looking at distant objects*).

7.2 Results

Results regarding time, accuracy, distance, gaze change, walking time and distance between participants and objects are presented in Figure 9, and Figure 10 shows the results of the participant ranking of the different layouts.

7.2.1 Comparison task

Time – We do not observe any difference of completion time between the three LAYOUT.

Accuracy – We do not observe any difference in accuracy between the three LAYOUT. However, there may be a difference between *Full-Circle* and *Flat*, with participants being more accurate with *Flat*, but it is not significant.

Travel Distance – Participants travelled more with *Flat* than with *Half-Circle* ($\sim 2.4m$) than with *Full-Circle* ($\sim 3.2m$).

Gaze Change – We did not find any difference between the number of gaze switches between layouts.

Walking Time – Participants spent more time walking with the *Flat* layout than with *Full-Circle* (5 seconds more) and *Half-Circle* (3 seconds more). They also spent more time walking with *Half-Circle* than with *Full-Circle* (2 seconds more).

Time looking at distant multiples – Participants spent less time looking at distant multiples with *Full-Circle* than with *Flat* (4 seconds less) and *Half-Circle* (5 seconds less).

Preferences – 5 participants ranked *Flat* first, against 4 participants for *Half-Circle* and 3 for *Full-Circle*. Only 1 ranked *Half-Circle* last, against 4 for *Flat* and 7 for *Full-Circle*. In the comments, participants found that walking with *Flat* was an issue (7/12), but some thought that this layout allowed for good overview of the multiples (3/12). For *Half-Circle*, some participants also complained about walking (4/12), and one specifically mentioned that walking “around” was not convenient. Finally, the main issue mentioned for *Full-Circle* was that it was hard to locate the graphs to compare (despite them being highlighted), because they had to do a 180 degrees rotation (8/12).

7.2.2 Max task

Time – We do not observe any difference of completion time between the three LAYOUT.

Accuracy – Participants were slightly more accurate with *Flat* than with *Full-Circle*, with a difference of 0.05. There seems to be a difference between *Flat* and *Half-Circle*, with participants being more accurate with *Flat*, but it is not significant.

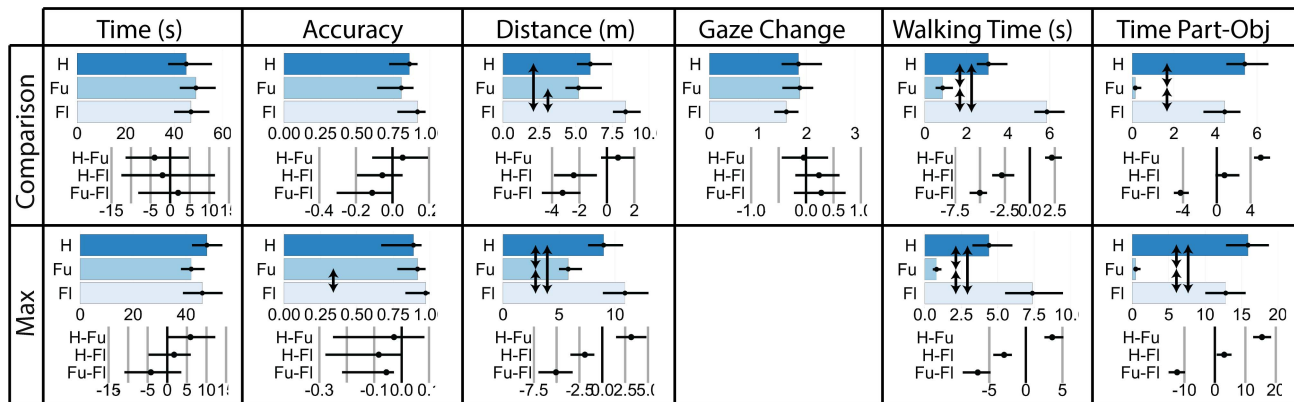


Figure 9: In each cell, the top bars show Means and CIs for all measures for Layout (Flat (FI), Half-Circle (H) and Full-Circle (Fu)) across conditions. The bottom charts show Corresponding 95% CIs for mean difference. Arrows indicate significant difference between two conditions.




Figure 10: Result of the ranking for the three layouts for comparison (left) and max task (right).

Travelled Distance – Participants travelled more distance with *Flat* than with *Half-Circle* (~1.9m) than with *Full-Circle* (~5m). They also walked more with *Half-Circle* than with *Full-Circle* (~3.2m).

Walking Time – Participants spent more time walking with *Flat* than with *Full-Circle* (7 seconds more) or *Half-Circle* (3 seconds more). They also spent more time walking with *Half-Circle* than with *Full-Circle* (4 seconds more).

Time looking at distant multiples – Participants spent less time looking at distant multiples with *Full-Circle* than with *Flat* (12 seconds less) and *Half-Circle* (16 seconds less). They also spent less time with *Half-Circle* than with *Flat* (3 seconds less).

Preferences – 9 participants ranked *Half-Circle* first, against 3 participants for *Flat* and none for *Full-Circle* one. No participant ranked *Half-Circle* layout last, against 1 for *Flat* and 11 for *Full-Circle*. With *Flat* layout, participants liked that they could easily get an overview of the multiples (3/12), but not the fact that they had to walk a lot (5/12); for instance, one participant commented “I wish I could perform less physical walking. Panning the vis would be great in this case.” Some participants thought *Half-Circle* was a good compromise as it provided an overview without too much walking and rotation (4/12). In accordance with that, some participants stated that the *Full-Circle* layout did not provide them with an overview at a glance, and required too much rotation (4/12).

7.3 Discussion

Results from both tasks are very similar. We cannot see significant differences between different layouts regarding time and accuracy. For the *Comparison* task, this could mean that the better performance of *Flat* in the previous experiment is countered here by the greater number of multiples. This explanation is supported by the fact that participants had to walk a greater distance, costing additional time. A similar effect has been observed by Shupp et al. when comparing flat and curved physical displays on search and path tracing tasks [47]. However, their findings suggested that curved display leads to better performance while our results do not provide statistically significant evidence reproducing this result in our VR environment.

So participants, in both tasks, spent more time walking in the *Flat* layout, but also in the *Half-Circle* one. Additionally, they spent more time looking at distant objects, this could mean that they seized the opportunity to step back and get an overview of the multiples, while it is not possible in the *Full-Circle* layout.

Finally, participants preferred, for both tasks, the *Half-Circle* layout. Their comments explained that it was a good compromise between walking and rotation, and that it allows for an overview at a glance by taking a step back. Participants identified that rotation in the *Full-Circle* layout was disorienting and made it harder for them to locate specific multiples. Similar to the *Half-Circle* layout, *Flat* was appreciated for its easy to access overview, but the amount of walking necessary was considered an issue. Interesting future work would be to explore techniques to reduce walking, like panning of the shelf, and VR teleportation. Similar work has been done on wall displays, and they showed that while users tend to favour Virtual Navigation [21] (in our case it would be panning the shelf or teleportation), physical navigation leads to better performance [2, 21, 29] as it improves spatial memory [43]. These studies involved flat wall displays, the influence of curvature on this issue would also be an interesting future direction.

8 CONCLUSION AND FUTURE WORK

Our user studies revealed that the performance of different layouts is dependent on the number of multiples displayed. With a small number, a *Flat* layout is more performant, even if it is not the users’ preferred one, due to the amount of walking required. With a significant increase in the number of multiples, the difference in completion time was less noticeable. However, participants complained about disorientation with *Full-Circle* and that it made locating a specific multiple difficult. It was also an issue that *Full-Circle* made getting an overview at a glance difficult. On the contrary, the *Flat* layout allows users to easily obtain an overview of all the multiples, but requires too much walking. Regarding all these issues, the *Half-Circle* provided a good compromise and was preferred by participants.

There is future opportunity to more thoroughly explore the curvature design space. For instance, vertical curvature has been used to support direct input in joint-centered layouts [34]. It would also be interesting to study the impact of interaction techniques that avoid walking, like VR teleportation or virtual panning of the shelf. However, this may also be less natural, disorienting and may impede any kinesthetic memory effect – but all of these aspects would need to be teased apart in low-level studies.

REFERENCES

- [1] A. P. Association. *Publication Manual of the American Psychological Association*. American Psychological Association (APA), 2009.
- [2] R. Ball, C. North, C. North, and D. A. Bowman. 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*, CHI '07, pp. 191–200. ACM, New York, NY, USA, 2007. doi: 10.1145/1240624.1240656
- [3] L. Besançon, M. Ammi, and T. Isenberg. Pressure-based gain factor control for mobile 3d interaction using locally-coupled devices. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, CHI '17, pp. 1831–1842. ACM, New York, NY, USA, 2017. doi: 10.1145/3025453.3025890
- [4] C.-M. Boddecker. *ExtendRealityLtd/vrtrk*. <https://github.com/ExtendRealityLtd/vrtrk>, 2018. Accessed: 2019-07-08.
- [5] M. Burch and D. Weiskopf. A flip-book of edge-splatted small multiples for visualizing dynamic graphs. In *Proceedings of the 7th International Symposium on Visual Information Communication and Interaction*, VINCI '14, pp. 29:29–29:38. ACM, New York, NY, USA, 2014. doi: 10.1145/2636240.2636839
- [6] G. Cumming. *Understanding the new statistics: Effect sizes, confidence intervals, and meta-analysis*. Routledge, 2013. doi: 10.1111/j.1751-5823.2012.00187.26.x
- [7] S. v. d. Elzen and J. J. v. Wijk. Small multiples, large singles: A new approach for visual data exploration. *Computer Graphics Forum*, 32(3pt2):191–200, 6 2013. doi: 10.1111/cgf.12106
- [8] B. Ens, J. D. Hincapié-Ramos, and P. Irani. 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*, SUI '14, pp. 2–12. ACM, New York, NY, USA, 2014. doi: 10.1145/2659766.2659769
- [9] B. Ens and P. Irani. Spatial analytic interfaces: Spatial user interfaces for in situ visual analytics. *IEEE Computer Graphics and Applications*, 37(2):66–79, Mar 2017. doi: 10.1109/MCG.2016.38
- [10] B. Ens, E. Ofek, N. Bruce, and P. Irani. Spatial constancy of surface-embedded layouts across multiple environments. In *Proceedings of the 3rd ACM Symposium on Spatial User Interaction*, SUI '15, pp. 65–68. ACM, New York, NY, USA, 2015. doi: 10.1145/2788940.2788954
- [11] B. M. Ens, R. Finnegan, and P. P. Irani. 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*, CHI '14, pp. 3171–3180. ACM, New York, NY, USA, 2014. doi: 10.1145/2556288.2557058
- [12] A. Fender, D. Lindbauer, P. Herholz, M. Alexa, and J. Müller. 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*, UIST '17, pp. 611–621. ACM, New York, NY, USA, 2017. doi: 10.1145/3126594.3126621
- [13] Y. Fujiwara, K. Ichikawa, H. Takemura, et al. A multi-application controller for sage-enabled tiled display wall in wide-area distributed computing environments. *Journal of Information Processing Systems*, 7(4):581–594, 2011. doi: 10.3745/JIPS.2011.7.4.581
- [14] B. Gao, B. Kim, J.-I. Kim, and H. Kim. Amphitheater layout with egocentric distance-based item sizing and landmarks for browsing in virtual reality. *International Journal of Human-Computer Interaction*, 35(10):831–845, 2019. doi: 10.1080/10447318.2018.1498654
- [15] B. Gao, H. Kim, B. Kim, and J. Kim. Artificial landmarks to facilitate spatial learning and recalling for curved visual wall layout in virtual reality. In *2018 IEEE International Conference on Big Data and Smart Computing (BigComp)*, pp. 475–482, Jan 2018. doi: 10.1109/BigComp.2018.00076
- [16] J. Gardy, A. Barsky, T. Munzner, and R. Kincaid. Cerebral: Visualizing multiple experimental conditions on a graph with biological context. *IEEE Transactions on Visualization and Computer Graphics*, 14(06):1253–1260, nov 2008. doi: 10.1109/TVCG.2008.117
- [17] J. Grubert, M. Heinisch, A. Quigley, and D. Schmalstieg. Multifidelity: Multi fidelity interaction with displays on and around the body. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '15, pp. 3933–3942. ACM, New York, NY, USA, 2015. doi: 10.1145/2702123.2702331
- [18] K. Hasan, D. Ahlström, J. Kim, and P. Irani. Airpanes: Two-handed around-device interaction for pane switching on smartphones. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, CHI '17, pp. 679–691. ACM, New York, NY, USA, 2017. doi: 10.1145/3025453.3026029
- [19] J. Heer and B. Shneiderman. Interactive dynamics for visual analysis. *Queue*, 10(2):30:30–30:55, Feb. 2012. doi: 10.1145/2133416.2146416
- [20] C. Hurter, N. H. Riche, S. M. Drucker, M. Cordeil, R. Alligier, and R. Vuillemot. Fiberclay: Sculpting three dimensional trajectories to reveal structural insights. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):704–714, Jan 2019. doi: 10.1109/TVCG.2018.2865191
- [21] M. R. Jakobsen and K. Hornbæk. Is moving improving?: Some effects of locomotion in wall-display interaction. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '15, pp. 4169–4178. ACM, New York, NY, USA, 2015. doi: 10.1145/2702123.2702312
- [22] W. Javed and N. Elmqvist. Exploring the design space of composite visualization. In *2012 IEEE Pacific Visualization Symposium*, pp. 1–8, Feb 2012. doi: 10.1109/PacificVis.2012.6183556
- [23] W. Javed, B. McDonnell, and N. Elmqvist. Graphical perception of multiple time series. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):927–934, Nov 2010. doi: 10.1109/TVCG.2010.162
- [24] S. Johnson, D. Orban, H. B. Runesha, L. Meng, B. Juhnke, A. Erdman, F. Samsel, and D. F. Keefe. Bento box: An interactive and zoomable small multiples technique for visualizing 4d simulation ensembles in virtual reality. *Frontiers in Robotics and AI*, 6:61, 2019. doi: 10.3389/frobt.2019.00061
- [25] A. Khan, I. Mordatch, G. Fitzmaurice, J. Matejka, and G. Kurtenbach. Viewcube: A 3d orientation indicator and controller. In *Proceedings of the 2008 Symposium on Interactive 3D Graphics and Games*, I3D '08, pp. 17–25. ACM, New York, NY, USA, 2008. doi: 10.1145/1342250.1342253
- [26] M. Krzywinski and N. Altman. Points of significance: error bars. *Nature Methods*, 10(10):921–922, Sept 2013. doi: 10.1038/nmeth.2659
- [27] O. D. Lampe, J. Kehler, and H. Hauser. Visual Analysis of Multivariate Movement Data using Interactive Difference Views. In R. Koch, A. Kolb, and C. Rezk-Salama, eds., *Vision, Modeling, and Visualization (2010)*. The Eurographics Association, 2010. doi: 10.2312/PE/VMV/VMV10/315-322
- [28] F. C. Y. Li, D. Dearman, and K. N. Truong. Virtual shelves: Interactions with orientation aware devices. In *Proceedings of the 22Nd Annual ACM Symposium on User Interface Software and Technology*, UIST '09, pp. 125–128. ACM, New York, NY, USA, 2009. doi: 10.1145/1622176.1622200
- [29] C. Liu, O. Chapuis, M. Beaudouin-Lafon, E. Lecolinet, and W. E. Mackay. Effects of display size and navigation type on a classification task. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '14, pp. 4147–4156. ACM, New York, NY, USA, 2014. doi: 10.1145/2556288.2557020
- [30] J. Liu. Immersive small multiples project. <https://github.com/joeprincer/ImmersiveSmallMultiples>, 2019.
- [31] J. Liu, A. Prouzeau, B. Ens, and T. Dwyer. Evaluation of various design choices for small multiples visualisations in immersive environments - study 1. <https://osf.io/sbr6u>, May 2019.
- [32] J. Liu, A. Prouzeau, B. Ens, and T. Dwyer. Evaluation of various design choices for small multiples visualisations in immersive environments - study 2. <https://osf.io/2g8d3>, June 2019.
- [33] X. Liu, Y. Hu, S. North, and H.-W. Shen. Correlatedmultiples: Spatially coherent small multiples with constrained multi-dimensional scaling. *Computer Graphics Forum*, 37(1):7–18, 2018. doi: 10.1111/cgf.12526
- [34] P. Lubos, G. Bruder, O. Ariza, and F. Steinicke. Touching the sphere: Leveraging joint-centered kinespheres for spatial user interaction. In *Proceedings of the 2016 Symposium on Spatial User Interaction*, SUI '16, pp. 13–22. ACM, New York, NY, USA, 2016. doi: 10.1145/2983310.2985753
- [35] B. B. C. H. B. T. K. M. T. D. M. Cordeil, A. Cunningham. IatK: An immersive analytics toolkit. In *Proceedings of IEEE VR*, 2019.

- [36] A. MacEachren, D. Xiping, F. Hardisty, Diansheng Guo, and G. Lengerich. Exploring high-d spaces with multiform matrices and small multiples. In *IEEE Symposium on Information Visualization 2003 (IEEE Cat. No.03TH8714)*, pp. 31–38, Oct 2003. doi: 10.1109/INFVIS.2003.1249006
- [37] J. Matejka, M. Glueck, E. Bradner, A. Hashemi, T. Grossman, and G. Fitzmaurice. Dream lens: Exploration and visualization of large-scale generative design datasets. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, p. 369. ACM, Montreal QC, Canada, 2018. doi: 10.1145/3173574.3173943
- [38] W. Meulemans, J. Dykes, A. Slingsby, C. Turkay, and J. Wood. Small multiples with gaps. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):381–390, Jan 2017. doi: 10.1109/TVCG.2016.2598542
- [39] T. Munzner. *Visualization analysis and design*. CRC press, New York, NY, USA, 2014. doi: 10.1201/b17511
- [40] T. V. Perneger. What’s wrong with bonferroni adjustments. *BMJ*, 316(7139):1236–1238, 1998. doi: 10.1136/bmj.316.7139.1236
- [41] M. D. Plumlee and C. Ware. Zooming versus multiple window interfaces: Cognitive costs of visual comparisons. *ACM Trans. Comput.-Hum. Interact.*, 13(2):179–209, June 2006. doi: 10.1145/1165734.1165736
- [42] A. Prouzeau, A. Bezerianos, and O. Chapuis. Towards road traffic management with forecasting on wall displays. In *Proceedings of the 2016 ACM International Conference on Interactive Surfaces and Spaces*, pp. 119–128. ACM, Niagara Falls, Ontario, Canada, 2016. doi: 10.1145/2992154.2992158
- [43] R. Rädle, H.-C. Jetter, S. Butscher, and H. Reiterer. The effect of egocentric body movements on users’ navigation performance and spatial memory in zoomable user interfaces. In *Proceedings of the 2013 ACM International Conference on Interactive Tabletops and Surfaces, ITS ’13*, pp. 23–32. ACM, New York, NY, USA, 2013. doi: 10.1145/2512349.2512811
- [44] W. Scheibel, M. Trapp, and J. Döllner. Interactive revision exploration using small multiples of software maps. In *VISIGRAPP (2: IVAPP)*, pp. 133–140, 2016.
- [45] M. Serrano, B. Ens, X.-D. Yang, and P. Irani. Gluey: Developing a head-worn display interface to unify the interaction experience in distributed display environments. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services, MobileHCI ’15*, pp. 161–171. ACM, New York, NY, USA, 2015. doi: 10.1145/2785830.2785838
- [46] B. Shneiderman. Direct manipulation: A step beyond programming languages (abstract only). In *Proceedings of the Joint Conference on Easier and More Productive Use of Computer Systems. (Part - II): Human Interface and the User Interface - Volume 1981, CHI ’81*, pp. 143–. ACM, New York, NY, USA, 1981. doi: 10.1145/800276.810991
- [47] L. Shupp, R. Ball, B. Yost, J. Booker, and C. North. Evaluation of viewport size and curvature of large, high-resolution displays. In *Proceedings of Graphics Interface 2006*, pp. 123–130. Canadian Information Processing Society, 2006.
- [48] R. E. Su and B. P. Bailey. Put them where? towards guidelines for positioning large displays in interactive workspaces. In M. F. Costabile and F. Paternò, eds., *Human-Computer Interaction - INTERACT 2005*, pp. 337–349. Springer Berlin Heidelberg, Berlin, Heidelberg, 2005.
- [49] G. Team. Gapminder. <https://www.gapminder.org/>, 2019. Accessed: 2019-07-08.
- [50] B. H. Thomas, G. F. Welch, P. Dragicevic, N. Elmqvist, P. Irani, Y. Jansen, D. Schmalstieg, A. Tabard, N. A. M. ElSayed, R. T. Smith, and W. Willett. Situated analytics. In K. Marriott, F. Schreiber, T. Dwyer, K. Klein, N. H. Riche, T. Itoh, W. Suerzlinger, and B. H. Thomas, eds., *Immersive Analytics*, vol. 11190 of *Lecture Notes in Computer Science*, pp. 185–220. Springer. doi: 10.1007/978-3-030-01388-2_7
- [51] E. R. Tufte. *Envisioning information*. Graphics Press, 1990.
- [52] D. Vohl, D. G. Barnes, C. J. Fluke, G. Poudel, N. Georgiou-Karistianis, A. H. Hassan, Y. Benovitski, T. H. Wong, O. L. Kaluza, T. D. Nguyen, et al. Large-scale comparative visualisation of sets of multidimensional data. *PeerJ Computer Science*, 2:e88, 2016. doi: 10.7717/peerj-cs.88
- [53] W. Willett, B. Jenny, T. Isenberg, and P. Dragicevic. Lightweight relief shearing for enhanced terrain perception on interactive maps. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, CHI ’15*, pp. 3563–3572. ACM, New York, NY, USA, 2015. doi: 10.1145/2702123.2702172
- [54] X. Xu, A. Dancu, P. Maes, and S. Nanayakkara. Hand range interface: Information always at hand with a body-centric mid-air input surface. In *Proceedings of the 20th International Conference on Human-Computer Interaction with Mobile Devices and Services, MobileHCI ’18*, pp. 5:1–5:12. ACM, New York, NY, USA, 2018. doi: 10.1145/3229434.3229449