

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/358970209>

TimeTables: Embodied Exploration of Immersive Spatio-Temporal Data

Preprint · March 2022

CITATIONS

0

READS

395

5 authors, including:



Yidan Zhang

Monash University (Australia)

3 PUBLICATIONS 2 CITATIONS

SEE PROFILE



Barrett Ens

Monash University (Australia)

79 PUBLICATIONS 1,863 CITATIONS

SEE PROFILE



Kadek Ananta Satriadi

Monash University (Australia)

15 PUBLICATIONS 121 CITATIONS

SEE PROFILE



Arnaud Prouzeau

Monash University (Australia)

39 PUBLICATIONS 355 CITATIONS

SEE PROFILE

Some of the authors of this publication are also working on these related projects:



Multi-scale gesture interaction (Micro + Macro) [View project](#)



Context-Aware Smart Energy Recommender (CASER) [View project](#)

TimeTables: Embodied Exploration of Immersive Spatio-Temporal Data

Yidan Zhang*
Monash University

Barrett Ens†
Monash University

Kadek Ananta Satriadi‡
University of South Australia
Monash University

Arnaud Prouzeau§
Inria, Bordeaux, France

Sarah Goodwin¶
Monash University

ABSTRACT

We propose TimeTables, a novel prototype system that aims to support data exploration, using embodiment with space-time cubes in virtual reality. TimeTables uses multiple space-time cubes on virtual tabletops, which users can manipulate by extracting time layers or individual buildings to create new tabletop views. The surrounding environment includes a large space for multiple linked tabletops and a storage wall. TimeTables presents information at different time scales by stretching layers to drill down in time. Users can also jump into tabletops to inspect data from an egocentric perspective. We present a use case scenario of energy consumption displayed on a university campus to demonstrate how our system could support data exploration and analysis over space and time. From our experience and analysis we believe the system has a high potential in assisting spatio-temporal data exploration and analysis.

Index Terms: Human-centered computing—Visualization—Visualization application domains—Information visualization; Human-centered computing—Human computer interaction (HCI)—Interaction paradigms—Virtual reality; Human-centered computing—Human computer interaction (HCI)—Interaction techniques;

1 INTRODUCTION

Immersive Analytics [2, 24] is an emerging field that uses technologies such as Virtual Reality (VR) to support data exploration. Such immersive technologies have attracted attention due to their ability to support intuitive interaction and spatial reasoning by leveraging innate human abilities such as proprioception, kinaesthesia, and spatial memory [23, 25].

A compelling aspect of Immersive Analytics is embodiment, the mapping of data artifacts to 3D virtual constructs that users can directly manipulate, examine, and rearrange [7]. This concept is exemplified in works such as ImAxes [3], where the data analyst can create a wide variety of different data visualisations by simply arranging virtual axes using direct manipulation, eliminating the need for buttons or menus.

We explore embodiment through TimeTables, a prototype system for exploring spatio-temporal data, using space-time cubes (STCs). Inspired by Uplift [8], which places a virtual STC above a tabletop campus map, we use the constructs of tabletop STCs, campus buildings, and time slices as the embodied elements of our interactive analysis process. Using a large immersive space in VR, TimeTables allows users to explore both space and time dimensions of the spatio-temporal data set, by filtering and scaling time through the arrangement of multiple tables containing different STC views.

Our exploratory prototype provides an initial implementation for investigating the use of embodiment for supporting engaging and

intuitive interaction with spatio-temporal data. We explore a set of novel interactions in TimeTables through a design probe, which presents a use case scenario describing the exploration of energy consumption data.

2 RELATED WORK

2.1 Spatio-temporal Visualisation

The exploratory analysis of spatio-temporal data is common in geographic visualisation and it has become increasingly popular for immersive visualisation. The visualisation methods used for spatio-temporal data typically combine with a map representation to display spatial attributes [35]. The data is geo-located on top of the map and could be represented in various forms, such as 2D/3D trajectories [32], bar charts [28], and an entity-relationship graph [4].

A space-time cube [11, 17] is a common technique for visualising spatio-temporal data. STC uses the horizontal plane for representing geographical space and time is mapped to the vertical axis. With its 3D nature, STC could benefit from the immersive environment and are gradually used in immersive analytics [8, 14, 27, 33]. However, most previous research still requires the use of GUI menus or other indirect means to conduct the analysis. For instance, Uplift [8] used interactive tangible widgets, but these were mapped to the input of slider controls. Our work takes a different approach by focusing on direct manipulation of embodied elements of the space-time cube visualisation themselves as a part of the analysis process. We take inspiration from Uplift to visualise energy data on a campus at different scales by using a STC.

2.2 Embodiment in Immersive Analytics

Embodiment is primarily defined as the degree at which the user perceives their physical body coinciding with a virtual one [9]. More broadly, Dourish proposed that besides relating to the user’s body, we can define an embodied system as a system in which the representation of the digital artefacts is designed in a way that they can be acted upon through embodied interaction [7], which means the representation of abstract ideas (e.g. data) could be regarded as objects that can be manipulated. Embodiment and embodied interaction have been adopted in many immersive analytics applications. One prominent work is ImAxes [3], a system for multivariate data exploration in virtual reality where axes are represented as three-dimensional objects that can be combined and arranged using direct manipulation to create various type of immersive charts. Recent works on immersive geographical data visualisation use various manual actions to facilitate embodied exploration [13, 26, 33, 36]. Lee et al. evaluated the collaborative aspect of immersive data exploration where users could freely manipulate and arrange visualisation on a virtual table and walls [20]. In this paper, we take inspiration from this work to design the visual representation of data and the user interaction to facilitate embodied spatio-temporal data exploration.

2.3 Multi-Scale Views

Exploration of complex spatio-temporal data requires multiple visualisations at different levels of data granularity [10] (e.g. yearly and monthly, country and state levels). For 2D desktop visualisation, guidance on how to visualise data using multi-level interfaces exists [18] and techniques such as views stacking has been introduced

*e-mail: yidan.zhang@monash.edu

†e-mail: barrett.ens@monash.edu

‡e-mail: kadek.satriadi@unisa.edu.au

§e-mail: arnaud.prouzeau@inria.fr

¶e-mail: sarah.goodwin@monash.edu

for multi-scale visualisation of temporal [15] and geo-spatial [16] data.

View management is a crucial factor for multiple view techniques [34] and researchers have started to utilise the vast available 3D space in virtual and augmented reality for arranging these multiple views. Existing works in this domain make use of multiple views to visualise geo-spatial data using hierarchical 2D maps [30] and linked flow maps [26]. Other works in immersive visualisation use multiple views to show different data layers on a map [31], multivariate visualisations [12, 22], and unstructured data such as text corpus [21]. While previous work use multiple views to visualise spatial information, the design of TimeTables focuses on using multiple tabletop maps as well as multi-scale views for exploring the temporal aspect of the data.

3 DESIGN CRITERIA

Taking into account the previous literature and our motivations, we identified several design criteria for a system that supports exploration of data over space and time freely. We refer back to these design criteria (C1–5) when describing the system interactions in Section 4.3.

C1. Support drilling down in time: Users should have the ability to keep drilling down the temporal data to a more detailed level of time granularity to support exploration in time. This criteria is motivated by Uplift [8], which has the ability to fluidly shift between views of increasing granularity.

C2. Provide multiple views: Users should have the ability to view multiple views with differing time granularity. This criteria reflects the motivation of Satriadi et al. [30] to use the space around users for placing multi-view maps, however we apply this principle primarily granularity of time rather than space.

C3. Support full use of virtual space: The system should make use of the large available space in the virtual environment. This criteria is inspired by DataHop [12], which takes advantage of a large virtual space to arrange different data views that are created during the analysis process. These can be traversed to assist users in their understanding of the analytical process. We extend this concept by making use of available space above the tabletops [8], and on available wall surfaces [20].

C4. Enable embodied interactions: The system should support embodied interactions for users to directly manipulate the data representations. This criteria is motivated by the highly engaging interactions provided by ImAxes [3]. Ideally, all significant components of the visualisations should be embodied as interactive elements. The interactions should be simple and intuitive so that users can quickly become proficient.

C5. Support multiple physical scales: The system should provide multiple physical scales of data visualisation for viewing data. This criteria is motivated by recent work that has explored the use of physical scale to help establish context and support data-driven storytelling [6, 19]. A tabletop scale should be included as a work space to provide a broad overview to support data analysis. A large scale can be added to fully immerse users into the visualisations and support a ‘more closed to real world’ experience, such as walk on the campus.

4 IMPLEMENTATION

Based on our motivations and design criteria, we developed a prototype system called TimeTables. TimeTables allows users to explore spatial-temporal data freely in the virtual environment.

4.1 Technical set-up

TimeTables is built in Unity3D¹ (version 2020.3.17f1) and displayed with a Windows Mixed Reality VR device. A Windows 10 PC

¹<https://unity.com/>

equipped with an Intel i7 processor and an NVIDIA GeForce GTX 1080 graphics card is used to run the prototype. We use Mixed Reality Toolkit² (MRTK) and Mapbox Unity SDK³ for implementing the interactive components and maps shown in TimeTables respectively.

4.2 Visualisations

The main visualisation concept used in TimeTables is a STC. We place the STC above a virtual table, whose surface contains a 2.5D map with extruded buildings. We chose a tabletop as the default setting because it allows a top-down overview of the map and provides an engaging workspace [8] that the user can walk around to view from different perspectives. A single tabletop is initially created in the virtual environment and multiple tabletops can be added by users during their exploration (see details in Section 4.3).

The tabletop map is an abstract map in dark mode which provides a neutral background for the visualisations above the map. A scale model of the buildings, each with a simplified shape of its real-world counterpart, is placed on the map. The surface details of these building models are abstracted and the heights of the buildings are re-scaled to all be the same height in order to avoid distraction and help the users to link the buildings to the visualisations above them more easily. The visualisations of the building’s temporal data is situated above each respective building and these visualisations reflect the same shape of each corresponding building.

Each time frame comprises of a set of visualisations in a horizontal layer. Each time frame is an increment of time shown with height and the number of layers shown in the visualisation is based on the given time granularity, for example, 12 layers for a monthly view (1 year) and 7 layers for a daily view (1 week). Each time layer is labeled with the corresponding time information.

The data is encoded by color hue, using a continuous sequential color scheme. The color scheme is derived from a color bar which is generated from a continuous Brewer color table (YlOrBr6). We decided to use this color scheme because it is colorblind safe and is commonly used in prior work in immersive and desktop geovisualisation and cartography (e.g. [1, 36]).

4.3 Interactions

TimeTables provides multiple interactions with embodied data constructs (C4) to support exploration and analysis, which we describe below.

4.3.1 Interactive elements

Users can manipulate visualisation components in TimeTables by a ‘grab and place’ interaction. Components can be grabbed using the VR controller, or using a raycast interaction for objects out of arms’ reach. The following components are interactive elements that users can manipulate in the system:

- **Tabletop:** Users can grab the tabletop and move it around the space or onto the storage wall.
- **Building Model:** Users can interact with the building on the tabletop to get the building name and basic information such as space allocation.
- **Layer:** A layer presents the temporal data for each building at a certain time frame (see Figure 1 (2)).
- **Block:** A block is a data representation of a single building at a certain frame of time (see Figure 2 (2)).

²<https://developer.leapmotion.com/mrtk-unity>

³<https://www.mapbox.com/>




Figure 1: Interacting with Tabletops: (1) Users can move a semitransparent selection plane above the tabletop to make a copy of the layer where the plane is located. (2) Move the copied layer to an empty location. (3) When users drop the layer, a tabletop is created to hold the filtered layer.




Figure 2: Interacting with Blocks: (1 - 2) Grabbing a block on the tabletop creates a free block that can be placed on an empty space. (3) A virtual, embodied, pedestal is provided to hold the block at the ideal height.

4.3.2 Time filter

TimeTables enables users to filter a layer or a block. The new filters will be placed at distinct spatial positions, aiming to fully use the virtual space around the user (C3). The system allows users to create many filters of layers and/or blocks (C2).

Filtering a layer (Figure 1). Users can move a semitransparent plane above the tabletop vertically to collide with the targeting layer. A movable layer can be duplicated from the original STC by pointing at the plane and long pressing the trigger for two seconds. Users can grab the generated layer and place it at a preferred location in the virtual environment. After releasing the trigger, a tabletop will be created at the selected position to hold the filtered layer above.

Filtering a block (Figure 2). To filter a block, users can extract a copy of the original block by selecting and dragging the initial block. Similar to filtering a layer, the copied block can be placed anywhere in the environment and a virtual pedestal will be created to hold the block after users release the trigger. The copied block(s) then can be performed as a new original block(s) and create a new filter, especially with the time stretch interaction (discussed in the following section). The copied block(s) and the original block(s) are linked together by lines to help the users trace the relationship in space and time of the block analysis.

4.3.3 Time stretch

The time stretch interaction allows users to manipulate the visualisations of the data in a single building or building layer to increase the level of time granularity (C1) (Figure 3), for instance to view the set of individual months within the selected year. The levels of time granularity in TimeTables include from most general to most detailed (e.g. yearly, monthly, weekly, daily and hourly). In this example, by default the yearly view is shown over the initial tabletop, which allows a filter for a specific year. After a block or a layer in a

specific year is extracted, users can drill down the time to a monthly view, with 12 layers for each month, by selecting the central area of the top surface of the targeting component using the trigger and stretching the target. The monthly view can be further filtered and stretched down to a weekly view (7 layers) and finally all hours within a selected day (24 layers).

4.3.4 Map storage

The wall on the right side of users' initial position is used for storage (C3). Figure 4 shows the process of the map storage interaction. By default a tabletop of a yearly view is placed in the environment and its corresponding 2D map is shown on the wall. As mentioned in Section 4.3.2, a new tabletop is created in the environment after filtering a layer. If there are too many tables in the space or the created table is no longer needed, users can pick up the table and place it on the wall. Once the edge of the table collides with the storage wall, the tabletop visualisation will be transformed into a 2D map with a title showing the time granularity of that table. If users want to investigate that tabletop again later, they can then drag the relevant 2D map from the wall to the virtual space. Once they place the map in a certain location and release the trigger, the 2D map will transform back to the tabletop representation.

4.3.5 Transferring between physical scales

TimeTables allows users to transfer between the visualisation scales of a tabletop and a large map by jumping into and out of the tabletop (C5). Users can select the button on the table to jump into the table to investigate the visualisation from an egocentric perspective (Figure 5 (1-2)), described as the large scale map, and can click on the map under their feet in the large scale to jump back to the tabletop scale.

Within the large scale map, users can open a mini map by tilting the left controller. A semitransparent cube becomes visible together with the mini map. We call this the 'invisibox' feature as it filters




Figure 3: Users can stretch a layer (1) or a block (2) to see the data visualisation in a more detailed time granularity

data in the virtual space aiming to remove unwanted occlusion during the data exploration. Users can grab the invisibox and place it onto the visualisation above a certain building on the mini map. The visualisation on the mini map will be hidden by the box and the corresponding visualisation information on the large map will become invisible synchronously (Figure 5 (2-4)). The invisibox can be re-scaled so that it can hide multiple buildings and help users to explore the visualisations of buildings being occluded.

5 USE CASE

As illustrated in the accompanying video⁴, we present a use case scenario to describe how people can conduct data exploration and analysis tasks using our system. In our scenario, we visualise the historical energy consumption data of 5 popular buildings on a university campus. The dataset includes aggregated hourly, daily, weekly, monthly and yearly data from 2019 to 2020. Amy, a data analyst, would like to investigate the data and try to perform analysis tasks, such as identifying which building has used the most energy in February 2019 and in which month in 2020 the least energy is used in Building A, comparing the energy consumption of Building A in June 2019 and June 2020, and comparing the energy consumption of Building A and Building B in June 2019.

Explore building information. Before Amy explores the dataset, Amy starts with an overview of the campus to see the absolute and relative locations of the buildings as well as retrieve some basic information about the building. To do so, she points at each building model to see the building name. For a closer inspection, Amy picks the model up and the building information is shown next to the building model (Figure 6 (1)). The model will be automatically returned to its original location on the map after the inspection.

Investigate data in the large map view. After inspecting the buildings, Amy narrows down her exploration to the yearly energy consumption data on the default tabletop (Tabletop A). To get a deeper immersion in the visualisation and explore the data closely,




Figure 4: (1) Users can move a tabletop close to the wall (2) When collide with the wall, the tabletop will be stored on the wall as a 2D map. (3) The 2D map can be grabbed back to the space. (4) when users drop the map, it will return to the original tabletop.




Figure 5: (1) Users can select a button to jump into the tabletop. (2) In the large view, users can open a mini map and manipulate the invisibox. (3) The invisibox can hide the visualisation of buildings in the front to solve occlusion. (4) The mini map and invisibox.

Amy presses the **‘jump into the table’** button on the edge of the table and enters the tabletop. She uses the **‘invisibox’** to *investigate one of the buildings occluded by other buildings*. Personally standing on the map provides Amy with a sense of spatial situatedness which helps her with the data understanding and exploration.

Explore monthly data in 2019. After the situated exploration on the map, Amy returns to the tabletop view and decides to *focus on the monthly energy consumption data in 2019*. She **extracts the layer using time filter** and places it in an empty location of the virtual space to form another tabletop (Tabletop B). She **stretches the 2019 layer** and the *monthly data* is shown on Tabletop B (Figure 6 (2)). She knows that February is the start of a teaching semester (for this Australian university) and would like to know if that affects the energy consumption of all 5 buildings, so she **filters the layer for February** as Tabletop C. As Amy only wants to focus on the data in February for now, she **places Tabletop B onto the map storage wall**.

Compare energy consumption between years. Amy found Building A has more energy consumption in February 2019. She

⁴<https://youtu.be/5MHfGxPpYTY>




Figure 6: (1) Pick up building model to show building name and basic information. (2) Investigate monthly data in 2019. (3) Compare energy consumption of Building A in February 2020 with that in February 2019. (3) Explore the energy consumption of Building A in June 2020. (5) Compare energy consumption of two buildings. (6) Compare the hourly data of two buildings from Thursday in week 3 of June 2019.

decides to compare the energy consumption of Building A in February 2019 with that in 2020 to investigate whether the consumption in these two years is similar. To do that, she **extracts** the targeting block from Tabletop A and **stretches** the block for a **monthly view**. She then further **filters** the block for *February* and **compares** it with the corresponding block shown on the Tabletop C (Figure 6 (3)). During the exploration of monthly data of Building A in 2020, Amy found the energy consumption in each month seems much lower than that in 2019. To confirm this, she **drags the stored 2D map of monthly data in 2019 from the wall** and places it at an empty location to show Tabletop B. She **compares the monthly data in 2019 and 2020** and confirms her guess. Amy thinks this result for Building A may be caused by the pandemic and would like to know if lifting lockdown during the pandemic brought the energy consumption back to 2019 levels. So she **extracts the block for June 2020** (the month when lockdown was lifted) (Figure 6 (4)) and **filters the layer of energy data in June** from Tabletop B and puts it somewhere in the space as Tabletop D. After that, she **extracts the block for June 2019** and **compares** it with that for June 2020. She found the energy consumption in June 2020 is still lower than that in 2019. She thinks this may be because despite lockdown ending, there were still fewer people on campus as less students were in the country and not all staff returned to work full time.

Compare energy consumption of two buildings. Amy would now like to compare the energy consumption of two buildings. She knows the purpose for Building A is teaching and Building B is used for sports and exercise from the initial inspection at the very beginning of the exploration. She would like to know if the distribution of energy consumption is the same for these two buildings. Thus she **stretches the layer to a weekly view**, **filters blocks for both buildings** from Tabletop D and drills down to a **daily view** and further an **hourly view** (Figure 6 (5-6)). After **placing the two building side-by-side** for comparison, she finds the energy consumption in Building A is higher during the day and in Building B is higher during the evening hours. This is to be expected as the gym and

sports facilities are open in the evening, whilst most teaching occurs during the day.

This use case scenario shows how our system could support exploration and analysis tasks. It also highlights the potential benefits compared to previous systems, such as showing the links between different time granularity levels to help users to track the relationships of the temporal data, making use of the map storage wall and allowing the transition between the tabletop view and the large map view for both analysis and experience benefits.

6 DISCUSSION

Based on our implementation of TimeTables and user case scenario, we identified some limitations and highlighted several interesting topics for our future research.

First, this scenario walkthrough shows the potential of our system to provide an engaging experience with embodied data interactions for performing data exploration and analysis in the virtual environment. However, we would like to further evaluate the system in a user study with specific analysis tasks to see the performance of our system and the feedback from users.

The use case we demonstrated shows how easy it is to use TimeTables, but that it can also easily create visual clutter when there are too many tabletops or pedestals in the environment. We would potentially improve this based on work such as DataHop [12] and Data mountain [29], which show the potential for leveraging spatial memory in virtual space for keeping track of a large number of views.

Furthermore, in our current implementation and use case, the objects for data representations are easy to manipulate, such as layers and blocks. However, traditional versions of STCs use lines and points that would be more difficult to embody with interactive constructs. Therefore, we would like to explore the generalisation of embodiment in our system to other types of STC data representations.

We've made some improvements on Uplift [8], such as enabling

embodied exploration rather than the use of tangible objects and GUI interaction. We also added the ability for multiple views and enabled continuous exploration and the history of views. Uplift does provide the added benefit of a casual collaborative experience. In our future work we wish to explore the possibilities for collaboration of TimeTables. In addition, we would like to further explore the concept of *embodied provenance* – the use of data embodiment to support tracking of analysis history through physical control of view placement – and implement features to support this concept. These could include the ability to undo/redo, store more information on the wall such as analysis notes, and provide more intuitive links between the tabletop and storage wall views for enabling the user to track the provenance of the data exploration [5].

Our example shows a novel use of multi-view visualisation for drilling down in time. For future research, we could have multiple views in both the space and time dimensions. With a larger map and larger data set (e.g. a city or a nation), we expect to explore the ability to drill down in space as well, potentially including but not limited in national, regional and local views.

We would also like to expand the ability for storytelling. In the current system, we allow users to change their physical scale to see a large egocentric map view. We expect to make this view to be more visual or visceral by including 360 images of the campus to show contextual information about weather, campus use, state of buildings at that time, etc. These 360 images could be useful for the analysis because the presented contextual information would affect the energy consumption. For instance, weather has an influence on energy consumption due to heating, ventilation, and air conditioning (HVAC), or a huge occupancy of one building could also increase the energy usage. This feature would be similar to Google Earth VR but would allow users to explore the context down over time.

7 CONCLUSION

In this paper, we present TimeTables, a novel prototype system based on five design criteria and inspired by prior work. We present novel interactions for exploring embodied spatio-temporal data in a large immersive space. Our use case and walkthrough demonstrates the potential for TimeTables for exploring spatio-temporal campus energy consumption data in an embodied manner. Through this work we have identified a number of potential improvements over previous implementations that inspired our prototype. We have also identified several areas for future work. In particular, we have identified a need for embodied provenance, which would leverage the benefits of embodied data constructs to improve user’s mental model of their path through the data analysis process. We look forward to continuing these explorations in future work.

REFERENCES

- [1] R. Beecham, J. Dykes, W. Meulemans, A. Slingsby, C. Turkay, and J. Wood. Map lineups: effects of spatial structure on graphical inference. *IEEE transactions on visualization and computer graphics*, 23(1):391–400, 2016.
- [2] T. Chandler, M. Cordeil, T. Czauderna, T. Dwyer, J. Glowacki, C. Goncu, M. Klapperstueck, K. Klein, K. Marriott, F. Schreiber, and E. Wilson. Immersive analytics. In *2015 Big Data Visual Analytics (BDVA)*, pp. 1–8, 2015. doi: 10.1109/BDVA.2015.7314296
- [3] M. Cordeil, A. Cunningham, T. Dwyer, B. H. Thomas, and K. Marriott. Imaxes: Immersive axes as embodied affordances for interactive multivariate data visualisation. *UIST ’17*, p. 71–83. Association for Computing Machinery, New York, NY, USA, 2017. doi: 10.1145/3126594.3126613
- [4] A. Cunningham, J. Walsh, and B. Thomas. Immersive visualisation of geo-temporal narratives in law enforcement. In *2018 International Symposium on Big Data Visual and Immersive Analytics (BDVA)*, pp. 1–8. IEEE, 2018.

- [5] Z. Cutler, K. Gadhav, and A. Lex. Trrack: A library for provenance-tracking in web-based visualizations. pp. 116–120, 2020. doi: 10.1109/VIS47514.2020.00030
- [6] K. Danyluk, T. T. Ulusoy, W. Wei, and W. Willett. Touch and beyond: Comparing physical and virtual reality visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 2020.
- [7] P. Dourish. *Where the action is: the foundations of embodied interaction*. MIT press, 2004.
- [8] B. Ens, S. Goodwin, A. Prouzeau, F. Anderson, F. Y. Wang, S. Gratzl, Z. Lucarelli, B. Moyle, J. Smiley, and T. Dwyer. Uplift: A tangible and immersive tabletop system for casual collaborative visual analytics. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1193–1203, 2021. doi: 10.1109/TVCG.2020.3030334
- [9] C. J. Falconer, M. Slater, A. Rovira, J. A. King, P. Gilbert, A. Antley, and C. R. Brewin. Embodying compassion: a virtual reality paradigm for overcoming excessive self-criticism. *PLoS one*, 9(11):e111933, 2014.
- [10] S. Goodwin, J. Dykes, A. Slingsby, and C. Turkay. Visualizing multiple variables across scale and geography. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):599–608, 2015.
- [11] T. Hägerstrand et al. What about people in regional science? *Papers in Regional Science*, 24(1):7–24, 1970.
- [12] D. Hayatpur, H. Xia, and D. Wigdor. Datahop: Spatial data exploration in virtual reality. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*, pp. 818–828, 2020.
- [13] 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, 2018.
- [14] M. Jahnke, E. P. Boguckaa, and M. Turchenkoa. Space-time cube and mixed reality—new approaches to support understanding historical landscape changes. In *Proceedings of the ICA*, vol. 2. Copernicus GmbH, 2019.
- [15] W. Javed and N. Elmqvist. Stack zooming for multi-focus interaction in time-series data visualization. In *2010 IEEE Pacific Visualization Symposium (PacificVis)*, pp. 33–40. IEEE, 2010.
- [16] W. Javed, S. Ghani, and N. Elmqvist. Polyzoom: multiscale and multifocus exploration in 2d visual spaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 287–296, 2012.
- [17] M.-J. Kraak. The space-time cube revisited from a geovisualization perspective. In *Proc. 21st International Cartographic Conference*, pp. 1988–1996. Citeseer, 2003.
- [18] H. Lam and T. Munzner. A guide to visual multi-level interface design from synthesis of empirical study evidence. *Synthesis Lectures on Visualization*, 1(1):1–117, 2010.
- [19] B. Lee, D. Brown, B. Lee, C. Hurter, S. Drucker, and T. Dwyer. Data visceralization: Enabling deeper understanding of data using virtual reality. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1095–1105, 2020.
- [20] B. Lee, X. Hu, M. Cordeil, A. Prouzeau, B. Jenny, and T. Dwyer. Shared surfaces and spaces: Collaborative data visualisation in a co-located immersive environment. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1171–1181, 2020.
- [21] L. Lisle, K. Davidson, E. J. Gitre, C. North, and D. A. Bowman. Sensemaking strategies with immersive space to think. In *2021 IEEE Virtual Reality and 3D User Interfaces (VR)*, pp. 529–537. IEEE, 2021.
- [22] J. Liu, A. Prouzeau, B. Ens, and T. Dwyer. Design and evaluation of interactive small multiples data visualisation in immersive spaces. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 588–597. IEEE, 2020.
- [23] K. Marriott, J. Chen, M. Hlawatsch, T. Itoh, M. A. Nacenta, G. Reina, and W. Stuerzlinger. *Immersive Analytics: Time to Reconsider the Value of 3D for Information Visualisation*, pp. 25–55. Springer International Publishing, Cham, 2018. doi: 10.1007/978-3-030-01388-2_2
- [24] K. Marriott, F. Schreiber, T. Dwyer, K. Klein, N. H. Riche, T. Itoh, W. Stuerzlinger, and B. H. Thomas. *Immersive Analytics*, vol. 11190. Springer, 2018.
- [25] J. McCormack, J. C. Roberts, B. Bach, C. D. S. Freitas, T. Itoh, C. Hurter, and K. Marriott. *Multisensory Immersive Analytics*, pp.

- 57–94. Springer International Publishing, Cham, 2018. doi: 10.1007/978-3-030-01388-2_3
- [26] R. Newbury, K. A. Satriadi, J. Bolton, J. Liu, M. Cordeil, A. Prouzeau, and B. Jenny. Embodied gesture interaction for immersive maps. *Cartography and Geographic Information Science*, pp. 1–15, 2021.
- [27] K. Okada, M. Yoshida, T. Itoh, T. Czauderna, and K. Stephens. Vr system for spatio-temporal visualization of tweet data. In *2018 22nd International Conference Information Visualisation (IV)*, pp. 91–95. IEEE, 2018.
- [28] M. Ready, T. Dwyer, and J. H. Haga. Immersive visualisation of big data for river disaster management. In *Workshop on Immersive Analytics: Exploring Future Visualization and Interaction Technologies for Data Analytics, Phoenix, AZ*, 2018.
- [29] G. Robertson, M. Czerwinski, K. Larson, D. C. Robbins, D. Thiel, and M. Van Dantich. Data mountain: using spatial memory for document management. In *Proceedings of the 11th annual ACM symposium on User interface software and technology*, pp. 153–162, 1998.
- [30] K. A. Satriadi, B. Ens, M. Cordeil, T. Czauderna, and B. Jenny. Maps around me: 3d multiview layouts in immersive spaces. *Proceedings of the ACM on Human-Computer Interaction*, 4(ISS):1–20, 2020.
- [31] M. Spur, V. Tourre, E. David, G. Moreau, and P. Le Callet. Map-stack: Exploring multilayered geospatial data in virtual reality. In *11th International Conference on Information Visualization Theory and Applications*, pp. 88–99. SCITEPRESS-Science and Technology Publications, 2020.
- [32] C. Tominski, H. Schumann, G. Andrienko, and N. Andrienko. Stacking-based visualization of trajectory attribute data. *IEEE Transactions on visualization and Computer Graphics*, 18(12):2565–2574, 2012.
- [33] J. A. Wagner Filho, W. Stuerzlinger, and L. Nedel. Evaluating an immersive space-time cube geovisualization for intuitive trajectory data exploration. *IEEE transactions on visualization and computer graphics*, 26(1):514–524, 2019.
- [34] M. Q. Wang Baldonado, A. Woodruff, and A. Kuchinsky. Guidelines for using multiple views in information visualization. In *Proceedings of the working conference on Advanced visual interfaces*, pp. 110–119, 2000.
- [35] L. Yang, Z. Ma, L. Zhu, and L. Liu. Research on the visualization of spatio-temporal data. In *IOP Conference Series: Earth and Environmental Science*, vol. 234, p. 012013. IOP Publishing, 2019.
- [36] Y. Yang, T. Dwyer, K. Marriott, B. Jenny, and S. Goodwin. Tilt map: Interactive transitions between choropleth map, prism map and bar chart in immersive environments. *IEEE Transactions on Visualization and Computer Graphics*, 27(12):4507–4519, 2021. doi: 10.1109/TVCG.2020.3004137