

Reimagining *TaxiVis* through an Immersive Space-Time Cube metaphor and reflecting on potential benefits of Immersive Analytics for urban data exploration

Jorge Wagner*
Federal University of
Rio Grande do Sul

Claudio T. Silva†
New York University

Wolfgang Stuerzlinger‡
Simon Fraser University

Luciana Nedeľ§
Federal University of
Rio Grande do Sul



Figure 1: What would *TaxiVis*, a landmark Urban Visualization system from ten years ago, look like if it were to be created within the context of an Immersive Analytics interface? In our reimagination, taxi trips are visualized over space and time (center left), and the visual query language extends to time through 3D prisms (center right). Bi-manual interactions allow brushing and pairwise comparisons (left, center), while embedded time series views bring associated data into the main visualization (right). ●

ABSTRACT

Current visualization research has identified the potential of more immersive settings for data exploration, leveraging VR and AR technologies. To explore how a traditional visualization system could be adapted into an immersive framework, and how it could benefit from this, we decided to revisit a landmark paper presented ten years ago at IEEE VIS. *TaxiVis*, by Ferreira et al., enabled interactive spatio-temporal querying of a large dataset of taxi trips in New York City. Here, we reimagine how *TaxiVis*’ functionalities could be implemented and extended in a 3D immersive environment. Among the unique features we identify as being enabled by the *Immersive TaxiVis* prototype are alternative uses of the additional visual dimension, a fully visual 3D spatio-temporal query framework, and the opportunity to explore the data at different scales and frames of reference. By revisiting the case studies from the original paper, we demonstrate workflows that can benefit from this immersive perspective. Through reporting on our experience, and on the vision and reasoning behind our design decisions, we hope to contribute to the debate on how conventional and immersive visualization paradigms can complement each other and on how the exploration of urban datasets can be facilitated in the coming years.

Index Terms: Human-centered computing—Visualization—Visualization systems and tools; Human-centered computing—Human

computer interaction (HCI)—Interaction paradigms—Virtual reality.

1 INTRODUCTION

When Ferreira et al. introduced *TaxiVis* in 2013 [23], they amazed data experts with the potential provided by a clever visualization tool combining efficient data storage and intuitive real-time visual querying [60]. Over the following decade, their work became highly influential, receiving the 10-year VAST *Test of Time* Award at IEEE VIS 2023 [29]. Since then, several research and commercial visualization systems for the exploration of big spatio-temporal origin-destination (OD) urban datasets have been presented, sometimes adopting similar visual querying metaphors [27].

In this work, our goal is to reimagine a tool like *TaxiVis* through the lens of *Immersive Analytics* (IA), a novel visualization paradigm that has similarly gained relevance in the past decade [24, 33, 41]. Our idea is to reflect, simultaneously, on (1) how a traditional visualization tool could be implemented in an immersive 3D virtual environment without losing its core functionality and (2) how this scenario could potentially bring in new functionality (what we refer to as *immersive extensions*). We also want to shed light onto potential benefits of IA for the exploration of urban datasets, a relatively under-explored topic despite the 3D nature of cities.

The background to our work is the growing research on IA [20], which has shown that the application of immersive AR/VR displays and interactions can improve data exploration activities, for example by offering a new visual dimension [63], by supporting different interaction metaphors [3, 11] and perspectives [62], or by facilitating collaboration [21, 36]. Increasing interest within the VIS, HCI, and AR/VR communities [24, 33] suggest a growing acceptability of such approaches, and a potential partial shift into this direction within visualization research and industry [65]. The idea is most frequently to complement established techniques, not to replace them [51, 59]. While most work at this point is exploratory and not intended for

*e-mail: jorge.wagner@ufrgs.br

†e-mail: csilva@nyu.edu

‡e-mail: w.s@sfu.ca

§e-mail: nedel@inf.ufrgs.br

immediate adoption, it is conceivable that such approaches will become more popular over the next years, either due to the release of new immersive devices, due to the discovery of new objective benefits, or simply due to growing consumer and corporate interest.

In this context, and given the complexity and multidimensional nature of urban data together with the strong requirement for collaboration between people from different backgrounds [21], we believe that urban data will be a natural application area for IA. Large datasets obtained from cities capture their behavior [18] and integrate both spatial and abstract information. Alternative interface paradigms such as IA may help analysts, researchers, and citizens leverage such data to better understand human mobility and other urban phenomena in the future. Despite the known drawbacks of desktop-based 3D visualizations, the spatial nature of urban spaces has already led to the development of multiple such tools [22], though their adaptation to immersive spaces has seen limited study [8, 72]. While some work has targeted immersive movement visualizations [28, 67], these have mostly not focused on urban contexts. This might be partially related to the difficulty of the necessary integration of 2D and 3D views and corresponding filters in this domain, one of the issues we seek to address.

Our vision is that, in the medium-term future, immersive tools will support urban experts in the execution of specific tasks and also enhance remote collaboration. In the long term, immersive AR/VR tools will seamlessly complement existing tools in the workplace, including support for the real-time analysis of data (for example, in a municipal control room during a large urban emergency).

In this work, we focus specifically on the immersive visualization of origin-destination data, prevalent in the urban context. Besides taxis, multiple other urban systems can provide this type of data, such as ride-sharing apps [27], bike-sharing stations [46], transportation cards [70], or mobile phone antennas [49]. We illustrate our ideas through a series of examples employing the taxi trips dataset from New York City (NYC) [17], as originally visualized by *TaxiVis*.

To reflect on the potential benefits of immersive visualizations for the exploration of this dataset, we iteratively developed a single-user proof-of-concept prototype. This process helped us envision a series of *immersive extensions* enabled by IA (Sect. 3), some of which we implemented and others we discuss as future work. These immersive extensions include extending the *TaxiVis*' visual query model to encompass the temporal dimension through a Space-Time Cube (STC) metaphor. The STC view, aided by on-demand 2D projections and embedded time series views, makes temporal patterns in the data more evident and enables the use of visual spatio-temporal filters and brushes through 3D prisms. Other extensions envisioned and already implemented in our prototype include the use of complementary ego- and exocentric immersive perspectives and uses of bi-manual interactions for brushing and comparisons. We also reflect on how to adapt key capabilities of the classic *TaxiVis* system into IA, for example, through the use of embedded views and controls (Sect. 4).

In summary, we present as contributions:

1. An extensive discussion of the potential benefits and extensions enabled by Immersive Analytics for the urban domain, as well as how an existing visualization tool could potentially be migrated to an immersive paradigm.
2. A single-user prototypical system demonstrating most of the ideas we discuss and how they can be implemented, openly available to inform the development of similar applications.¹
3. A series of examples that revisit case studies from the original paper and demonstrate the use of our prototype system for the exploration of large real-world urban datasets with OD data.
4. From an engineering point of view, a demonstration of how the core ideas behind the *Immersive Analytics Toolkit* (IATK) [10]

¹<https://github.com/jorge-wagner/ImmersiveTaxiVis>

can be adapted to construct new custom applications scalable to reasonably large datasets.

Further, we also make an effort to guide interested readers to additional relevant work within the IA literature.

2 BACKGROUND AND RELATED WORK

Our research is at the intersection of Urban Data Visualization [15] and Immersive Analytics [41], looking at the benefits stemming from merging these two areas. Within these two areas, we look at relevant work associated with the visualization of urban data, especially origin-destination data, which inspired and informed our design rationale for our reimagined *TaxiVis*.

2.1 *TaxiVis* and subsequent systems

Urban Visualization applications offer analysts visual tools to explore datasets generated by urban spaces, and have gained increased relevance with the progressive release of large open urban datasets. Being able to explore such datasets is becoming highly relevant to better inform expert decisions [18].

TaxiVis [23] sought to address some of the main challenges in this area at its time of publication, in particular, the difficulty in interactively querying the spatio-temporal urban data. To this end, *TaxiVis* proposed an optimized storage solution and a visual query language to intuitively specify origin-destination queries by drawing or selecting regions on a map. The system also allowed the construction of more complex queries by combining a series of atomic queries and offered a series of coordinated visualizations (see Fig. 2). The core ideas of *TaxiVis* became highly influential in the field, resulting in a 10-year *Test of Time* recognition awarded at IEEE VIS 2023 [29].

Subsequent systems have targeted the efficient exploration and uncovering of spatio-temporal patterns in other related large origin-destination datasets, such as bike-sharing data from New York City [46, 54] and subway trips data from Singapore [69, 71].

Kepler.gl, a web-based open-source geospatial visualization system introduced by Uber in 2018 [27] is currently a state-of-the-art tool for the visualization of large origin-destination datasets. It allows the visualization of either raw point or aggregate data through several possible 2D and 3D encodings. Similarly to *TaxiVis*, *Kepler.gl* supports polygon-based visual spatial querying of the data, and also allows interactive origin-destination brushing.

In this work, we reimagine how many of the features introduced by these systems could be extended to an immersive paradigm.

2.2 Uses of 3D in conventional urban visualization

Given the inherently 3D nature of urban spaces, the use of 3D visualization approaches in this domain precedes the advent of immersive visualizations. In *Urbane*, for example, Ferreira et al. employed 3D city representations in a desktop-based system to visually depict the effect of planned construction on landmark visibility and sunlight exposure [22]. Mota et al. [45] compared different approaches to visualize time-varying data embedded in or alongside 3D city models, identifying that plot-based visualizations were more accurate than color-based ones. Moreira et al. [44] introduced *UrbanTK*, a grammar-based toolkit that allows integrating multiple urban datasets into a 3D city model. In our work, we identify better understanding and interacting with the 3D topography of cities as one of the extensions enabled by IA for such applications.

When visualizing spatio-temporal urban datasets, however, another key use of the third dimension is not linked to the 3D representation of cities. As originally proposed by Hägerstrand [26], and revisited by Kraak [31], the third axis can be used to depict the progression of time, resulting in a *Space-Time Cube* (STC). This representation emphasizes the perception of patterns over space and time simultaneously, and has been widely applied in different areas.

Lock et al. [39], for example, constructed a 3D STC of one year of bus trips in the city of Sydney, making the occurrence of major



Figure 2: Screenshot of the original *TaxiVis* system illustrating the application of temporal constraints (A) and of visually-defined spatial queries (B). Coordinated views depict associated trip data for each query over time (D). Source: Ferreira et al. [23] (©2013 IEEE).

delays very evident by unusual point densities. They called this representation a *Performance Point Cloud*. In *iVizTRANS*, Yu et al. [69] similarly used an STC view to represent *tap-in tap-out* smart card data from the transport system in Singapore. They rendered each trip between two selected regions as a line connecting origin and destination, highlighting temporal commuting patterns. Cheng et al. [9] employed 3D isosurface volumes in STC views to represent the incidence of traffic congestion over time.

Gonçalves et al. [25], Ding et al. [16], and Yang et al. [66] all used STC views to visualize, over time, taxi trip *pick-up* and *drop-off* data, for New York, Shanghai, and Shenzhen, respectively. In our work, we integrate similar STC point cloud views of taxi data in the immersive paradigm, where, according to previous findings [64], STC exploration leads to lower user mental workload. We also extend the STC metaphor by introducing embedded 2D plots and a 3D spatio-temporal visual query language adapted from *TaxiVis*.

2.3 Immersive Analytics of movement data and initial attempts at immersive urban analytics

Started around 2015 in its current form [5], Immersive Analytics (IA) is a relatively novel research area at the intersection of Visualization, AR/VR, and HCI. It focuses on facilitating data analysis and understanding through the application of novel visualization and interaction paradigms. The area has been gaining growing prominence in recent years [24, 33], and current research challenges [20] include supporting the visualization of large datasets [10, 28], conducting evaluations with experts and in-the-wild [3, 21], and supporting collaboration [36, 51].

Of particular interest to our current project are the applications of IA to movement data and urban data. Notable works in this area include Yang et al.’s evaluation of immersive origin-destination flow maps at a global level [67]. They found that a 3D globe with raised flows was faster and more accurate than other (flat) alternatives. Wagner et al. adapted the previously-mentioned Space-Time Cube metaphor to immersive environments, and, in a comparative study against a desktop-based baseline, found that immersion decreases user workload [64]. They also identified complementary benefits for

egocentric and exocentric immersive perspectives of the cube [62]. In *GeoGate*, Ssin et al. combined an AR STC view with a tabletop display, using a ring-shaped tangible controller to select focal regions on the map [57]. With *FiberClay*, Hurter et al. extended IA to much larger datasets, combining a custom rendering engine and a series of interaction metaphors to allow the interactive querying of millions of aircraft trajectories. In particular, their system allowed “sculpting” queries by manually adding or subtracting selections, and bi-manually brushing combinations of origin and destination areas. All these efforts influenced the design of our *Immersive TaxiVis*.

Other work has also presented initial attempts at applying IA to urban data visualization. Chen et al. [8] presented an early exploration of the design space of immersive urban analytics, proposing, in particular, a typology of visual integration of 3D physical and 2D abstract data, distinguishing between linked views, embedded views, and mixed views. They also proposed an exploded views method to overcome occlusion in urban visualizations [7].

In recent years, a series of prototypes leveraging different technologies have been designed to integrate data into urban representations. In the *CityScope* interactive simulation tool for urban design, Alonso et al. explored a combination of a tangible *LEGO*-based city model overlaid with projected data and a conventional display for 2D coordinated views [1]. In *HoloCity*, Lock et al. presented an AR prototype visualizing real-time and aggregate city data, including transportation and social media data, overlaid onto a city model [38].

Several works adopted VR approaches. In *UrbanVR*, Zhang et al. implemented an immersive system to support the assessment of the impact of new planned buildings [72], similarly to the aforementioned *Urbane*. In *ViBe* (Virtual Berlin), Al Bondakji et al. presented a prototypical 3D environment for sustainable city planning through the visualization of open source datasets of demographic, environmental, and land use data [40]. Spur et al. demonstrated two different metaphors to visualize layers of urban data in VR: stacking several layers at different planes [56] or distorting the data layers into a city-scale *Urban Datasphere* [55]. Chen et al. proposed a related distortion-based *focus+context* approach, by folding the more distant parts of the map upwards with *UrbanRama* [6].

Further, Cunningham et al. extended the *ImAxes* system [11] to enable embodied interaction with geospatial data, using energy consumption data as an example [12] – similarly to *ImAxes*, different ways of attaching different data axes then results in different 3D spatial visualizations of the data. Elvezio et al. demonstrated a collaborative prototype for the situated visualization of live urban data such as service requests obtained from public city APIs [19], overlaid onto a city model to indicate their origins.

In this work, our primary focus is not on conceiving new applications for immersive urban analytics, but on imagining how a conventional tool could be adapted and extended into the IA paradigm. To this end, all examples above informed our proposed designs.

3 ENVISIONED IMMERSIVE EXTENSIONS

In this section, we detail how we envision that IA could potentially extend the functionality of *TaxiVis*, introducing a series of *immersive extensions* enabled by IA (underlined throughout this section). In the next section, we will focus on how it could also adapt features originally designed for a desktop interaction paradigm.

We start with an expansion to the *TaxiVis* visual query model by integrating the temporal component into the visual representation, following the Space-Time Cube (STC) metaphor [2, 64]. Further extensions correspond to alternative uses of the third dimension for: immersive interaction, immersive perspectives, immersive collaboration, and immersive engagement. This conceptual vision was iteratively informed by the development of our *Immersive TaxiVis* proof-of-concept prototype, which we start introducing in this section to illustrate the discussed concepts. Extensions already implemented into this prototype are marked with check marks (✓).

We implemented this prototype using the Unity game engine, the *Immersive Analytics Toolkit (IATK)* [10] for data rendering, the *Mixed Reality Toolkit (MRTK)* for interaction support, and the *Bing Maps SDK* for mapping resources. Please refer to the Supplementary Materials for more information on how we adapted the core ideas of *IATK* to implement spatio-temporal queries.

3.1 Extended visual query model (or: *TaxiVis* meets the Space-Time Cube)

The additional visual dimension enabled by an immersive visualization approach, i.e., the third dimension, can be utilized in many different ways when working with urban data, as we will discuss later (Sect. 3.2). Our main idea to leverage this additional dimension for the exploration of urban origin-destination data is to extend the visualization to encompass the temporal component ✓, through a Space-Time Cube (STC) metaphor. In an STC, the vertical axis corresponds to time (which advances upwards in all examples in this paper, although the opposite is also possible [64]), enabling a clear visualization of patterns over space and time. While the original *TaxiVis* system offered the possibility of visualizing several time steps side-by-side, an STC could be seen as a potentially infinite number of time steps in the same view. Our examples in Sect. 5 illustrate how this can be advantageous, for example, to make temporal patterns more evident. As discussed earlier, this approach has already been explored non-immersively in recent work [16, 25, 39, 66]. Prior work also demonstrated that an immersive STC can minimize the interaction challenges previously associated with this visualization [64].

In an STC for origin-destination data, data positions can be either represented individually, as a point cloud [39], or connected to each other, forming space-time paths [25, 42]. We enable both such visualization forms in our prototype. Following the *TaxiVis* convention, we color *pick-up* points in blue and *drop-off* points in red. Users are able to adjust point sizes through a menu slider. Smaller data points lead to better visualization of patterns in the full point cloud, while bigger data points can help emphasize a selected subset. To avoid clutter, showing connecting lines between pairs is more convenient when associated with specific origin-destination filters, such as in *iVizTRANS* [69]. Connected trajectories are also well suited for data with multiple origin-destination pairs for a same individual—in this case, they can also be colored according to thematic attributes, for example, the means of transportation or the trip purpose [42]. Showing aggregate views of the STC data is also possible through a density heat map [14], but currently not implemented. We also support on-demand 2D projections of inspected data points both onto the map and onto the time walls to allow the exploration of each component separately and reduce the effort needed to understand the locations and times of mid-air trajectories.

A key contribution of *TaxiVis* was its **visual query model**, through which queries could be visually composed by combining spatial, temporal, and attribute constraints, and iteratively refined by direct manipulation of the results. One could manually draw spatial queries on the map view (Fig. 2–B), while specifying temporal queries in a different module of the interface, the *time selection widget* (Fig. 2–A). For *Immersive TaxiVis*, we propose extending the visual query model ✓ so that a user can query both spatial and temporal components simultaneously by creating and adjusting visual representations in the STC view. To this end, instead of 2D query polygons, we support query prisms (i.e., extruded polygons) in 3D space, whose height corresponds to their temporal span, as seen in Fig. 1.

In our prototype, we support all kinds of queries proposed by the original *TaxiVis* model, with the extension of handling the temporal component as inseparable from the spatial one.

Atomic queries. As in *TaxiVis*, atomic queries can be created by either freely drawing a region on the map, a time span on the time walls, or by selecting a specific neighborhood on the choropleth map. A new query specified on the map targets initially the whole

time span of the dataset, and vice-versa. Besides being visually represented by the query prisms, queries are also visually projected on the map and on the time walls, to facilitate the user’s understanding. Wall projections are also labeled to indicate their lower and upper limits. Users can freely translate and scale the query prisms to modify queries, and they can also directly manipulate the query projections to modify their limits. The time granularity of the time walls is dynamically adjusted according to how much the user zooms in or out on the data, to facilitate selecting periods of interest. As in *TaxiVis*, each atomic query is assigned a new query color, which is used to color the projection contours. This color is used to refer to the query in associated visualizations.

Atomic queries can indicate origin-only (taxi pick-ups, in the case of taxi data) or destination-only constraints (taxi drop-offs), or their combination (either pick-ups or drop-offs). Following *TaxiVis*’s encoding, these are indicated by blue, red, or green areas or volumes, respectively. To avoid overloading the main menu (see Sect. 4), we defined that new queries correspond to both origin and destination (i.e., green volumes) by default. When a new query is created, a 2D panel of buttons (see Fig. 1–center right) appears with it, which enables query modifications. This panel is dynamically positioned close to the user to facilitate interactions, with a leading line linking it to its query. Through these buttons, the user can modify the kind of constraint it represents. To avoid overloading the interface, these buttons are automatically hidden when the user exits the query editing mode, and the query prism is then also made more translucent (or invisible, depending on user preference). Each query is also connected to a floating 2D panel summarizing statistics about the selected trips.

Directional queries. One of the most semantically useful query types in *TaxiVis* might be the directional one, where the user specifies both origin and destination constraints. Such a query can be used, for example, to select taxis that departed from an airport and went to a given area in Manhattan. In our implementation, the user can select one of the atomic query buttons (“*Directional Constraint*”) to directionally link a query to another existing or new one. Such a link is visually indicated by an arrow line rendered over the map surface, as illustrated in Fig. 1–left. A new directional query is immediately assigned a single color and a single set of buttons (where one of the options, naturally, is to revert it to its two atomic queries).

Recurrent queries. In *TaxiVis*, recurrent temporal selections (e.g., every Monday, every day from 10 to 11AM) could be filtered through the time selection widgets. As we extend the visual representation to integrate time, we decided to represent queries associated with recurrent selections as sliced prisms, clearly indicating the currently applied selections, as shown in Fig. 1–center right. Such selections can be applied either to one individual query or to all existing queries. The selection widget is designed similarly to the original interface.

Merged queries. Multiple atomic queries can be merged into single *merged query* (or *complex query*) by logical disjunction (“or”), so that the results of multiple alternative constraints are aggregated in the associated visualizations. This action can be performed by selecting the “*Merge*” button on a query panel, and then selecting a new query to be merged into it. Merged queries are assigned a single color and a single set of query option buttons.

3.2 Other potential uses of the third dimension

As discussed above, our main use for the third dimension enabled by IA is to integrate the temporal component into the visualization and thus render all time steps on top of each other through an STC.

Another obvious use of 3D in the urban context, which we also support (see Fig. 1–center), is correlating the data with the urban topography ✓. Depending on the dataset, features such as flat or hilly areas can be of interest due to their impact on different forms of mobility, as can be the density of tall buildings in a given area. Some desktop-based urban visualization systems, such as *Urbane* [22],

already leverage 3D views to visualize building shadows and landmarks visibility. We believe that the combination of the stereoscopic display and more suitable 3D interaction metaphors will increase the utility of the 3D views when compared to conventional systems.

Other possible uses of 3D which are not included in our prototype would include layering query results into a limited number of planes or time steps to facilitate comparisons, thus resulting in a *Stratified STC*, which would be comparable to the side-by-side views offered by *TaxiVis*. Depending on the dataset and user needs, we could also have a *Cyclic STC* which resets the data height after a certain amount of time (depending on the desired granularity, such as day or hour) in order to compare mobility patterns over time. This latter option would be more convenient for smaller datasets.

Finally, instead of visualizing individual data points, the 3D space can also be used to render alternative 3D geospatial visualizations, such as 3D heatmaps, as studied by Kraus et al. [32], or *Prism Maps*, as studied by Yang et al. [68].

3.3 Uses of immersive and bi-manual interactions

Our system supports interactions with either controllers or hand tracking in the *Oculus Quest 2* headset. Both modes support at-a-distance ray-based interactions, and the hand-tracked mode also supports direct touch interactions with menus. This combination [61] enables easy interaction with both 2D and 3D interface components.

The fact that an immersed user conventionally has access to two interaction devices (controllers or hands) instead of one also enables data exploration by simultaneous origin-destination brushing ✓, i.e., dynamically specifying both origin and destination constraints by pointing to different locations or time periods with each hand (see Fig. 1–left), as proposed originally by *FiberClay* [28]. This could also be seen as an extension of the single desktop-based time lens proposed by *Trajectory Wall* [58] or the brushing feature available in *Kepler.gl* [27]. Two-handed brushing can also be used for quick pairwise comparisons ✓ between two simultaneous selections. We support this through prism brushes (Fig. 3–right) and through our *choropleth mode*, where users can quickly compare stacked color-coded prisms for two different neighborhoods, allowing a detailed comparison of traffic patterns over time (Fig. 3–left). This latter option is similar to the visualization of energy consumption over time in *Uplift* [21] and *TimeTables* [73]. We aggregate time dynamically according to the temporal zoom level applied to the STC. Pairwise comparisons are also available through the *cutting planes* (one controlled by each hand) and *finger inspection lenses* (tapping data points to obtain details on demand) features.




Figure 3: Extensions enable pairwise comparisons between neighborhoods with bi-manual interaction, either through choropleth prism stacks (left) or two-handed brushes and embedded plots (right).

3.4 Complementary uses of immersive perspectives

We are convinced that a tabletop metaphor (as seen in Fig. 6) is the ideal standard mode for most urban data exploration scenarios. This

metaphor evokes the familiar concept of an architectural scale model and allows easy user operation, replacing most navigation with manipulation actions. Also, it supports a group of collaborating users around the “table”, such as seen in *Uplift* [21] and *CityScope* [1] (where real tabletop screens were used). Depending on user preference, it also supports both seated and standing use. However, as demonstrated by Wagner et al. [62], both exocentric and egocentric immersive perspectives can be preferable depending on the task and goal. An interesting extension enabled by IA is the ability to inspect urban data at different perspectives and scales ✓, including egocentrically in room-scale when desired (see Fig. 4–left).

Another interesting approach enabled by VR in urban spaces is the inspection of data at real city-scale ✓. The idea is that the user would be able to *dive-in* and *dive-out* of this real-scale egocentric view to inspect the actual location where some phenomena of interest transpired. Given the currently limited resolution of photo-realistic 3D tiles such as *Bing’s* at the street level, we opted to implement this functionality by using 360 degrees images from *Google StreetView*, an approach already adopted in some commercial applications such as *Google Earth VR* and *Woorld*. By positioning a 3D widget, the user can select any position on the map, and the closest possible image is then loaded as a scene *skybox*, preserving the STC view as context (see Fig. 4–right). In future work, we intend to replace *StreetView* data with a larger dataset of millions of images of NYC over time to allow coverage of the temporal dimension as well [43].




Figure 4: Immersion allows the integration of complementary frames of reference, such as egocentric room-scale (left) and egocentric view of 360-degree street images (right).

3.5 Future uses of immersive collaboration

Enhanced support for collaboration is a key IA research challenge [20] with multiple potential applications to immersive urban analytics. As demonstrated by *Uplift*, IA may be particularly useful to support knowledge sharing between users with different backgrounds [21]. At the present time, supporting such collaboration has not been the focus of our proof-of-concept prototype, but we identify the following possible designs as of particular interest: (1) remote collaboration enhanced by avatars and body cues [36], (2) access to both private and shared views [50], (3) multi-scale collaboration (i.e., between users immersed at different scales to investigate different aspects of the data) [47], and (4) asymmetric collaboration (i.e., between immersed and non-immersed users to leverage the strengths of each paradigm) [51, 59].

3.6 Future uses of immersive engagement

Besides the potential advantages already discussed, immersive interfaces can be considered engaging and playful to casual users due to their novelty. This can be leveraged to support and motivate casual data exploration by general citizens [48]. In this context, recent work has explored immersive urban simulations to support participatory planning [30, 52] and urban noise understanding [4].

One possible future way of integrating a more playful metaphor into *Immersive TaxiVis* could be simulating the movement of taxis in the city at any given time by showing 3D taxi models, which could be implemented either at a city- or desk-scale. This feature would then be comparable to the animations used in *TaxiVis* to illustrate the probable route taken by the visualized trips. Such an interface would also qualify as *data visceralization*, i.e., immersive experiences that facilitate intuitive understanding of quantities [34].

4 ADAPTING DESKTOP FEATURES TO IA

Beyond extending *TaxiVis*' functionality, as discussed in the previous section, we consider it essential to adapt the original system's functionality effectively to an immersive context. In this section, we discuss the design rationale for our *immersive adaptations* of elements from *TaxiVis*' original WIMP (Windows, icons, menus, pointer) interface. To this end, we leveraged, in particular, a combination of embedded views, coordinated views, and contextual menus, as discussed below.

4.1 Supporting interactions with space and time

To support the core functionality of *TaxiVis*, our first priority was to allow efficient interactions with space and time. Besides the strong stereopsis depth cues provided by VR headsets, we integrated additional features to facilitate comparisons, particularly in the temporal (vertical) dimension. Inspection lines indicate the date and time currently pointed at by each hand. All points corresponding to the temporal slices in focus are highlighted in a darker color (slices adjust dynamically to the level of temporal zoom applied). Translucent cutting planes are also available on demand, and users can inspect any data point with their fingers through *inspection lenses*.

To avoid overwhelming the user and leading to unintended interactions, we divided the system into two main interaction modes: *exploration mode*, where the user is free to drag and zoom into space and time, and *query mode*, where the user is free to create new spatio-temporal queries or edit existing ones.

4.2 Uses of embedded views and controls

A significant part of *TaxiVis*' original interface is dedicated to temporal constraints (Fig. 2–A) and visualization of time series data (Fig. 2–D) for the identification of patterns over time. When distributing the taxi data over space and time simultaneously, our design rationale was to also integrate as much temporal information as possible into this Space-Time Cube view.

In this context, we implemented the ability to embed all attribute time series views into the STC walls, as shown in Fig. 1–right, Fig. 6, and Fig. 7. As in *TaxiVis*, time series can display data for the whole dataset or per query, and can be switched between line graphs and histogram views depending on user preference. These views dynamically move and scale with the STC time walls.

Similarly, the controls for *regular selection* of continuous time intervals are also fully embedded into the time walls, as sphere widgets on the right side of the STC (see Fig. 1–right). However, we opted to preserve the *recurrent selection* widget as a separate panel, as we considered the user interface would become unnecessarily more complex if integrated into the STC walls. As discussed in the next section, we also preserved a separate 2D panel for time series views in case this view is more familiar and preferable to the user.

Besides embedding the temporal data, we also overlay a choropleth view of the data on the 2D base map of the STC, similarly to *TaxiVis*. As mentioned, this view can be used to create new neighborhood queries and for pairwise comparisons between neighborhoods.

4.3 Uses of coordinated views

The coordinated views offered in *TaxiVis* through different windows for attribute exploration are also available through separate 2D panels in our prototype (Fig. 5). These include the *attribute constraints*

panel that allows the analysis of value distributions and application of minimum and maximum constraints, and the *attributes over time* panel that displays time series in a more conventional format.

Leveraging the 3D space of the immersive environment, the coordinated panels sit by default at the sides of the user in the exocentric desk-scale mode, but they are fully movable and can be positioned next to or inside the STC. In room-scale mode, they are originally positioned within an open side of the STC walls, but users can bring them inside the STC if preferred (as in Fig. 5–right).

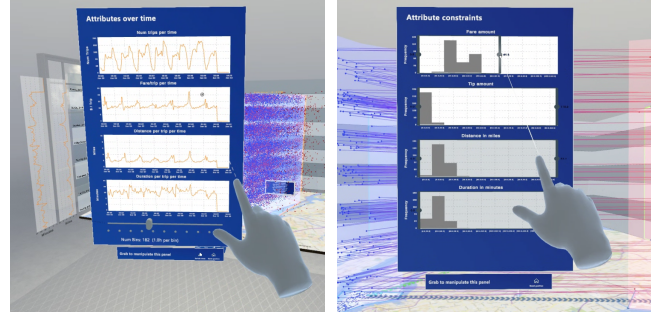


Figure 5: Movable 2D panels complement the interface, offering similar attribute exploration functionality as in *TaxiVis*.

4.4 Uses of contextual menus

TaxiVis' toolbar (Fig. 2–C) provides buttons for the creation and combination of different kinds of queries and for activating different features. Some of these buttons work in combination with conventional mouse-keyboard metaphors, such as holding a key for multiple selections, or with keyboard shortcuts. In our proposed implementation, we also have a main menu of buttons, which can be attached to the user's hands (revealed with a palm-up gesture), to the hand controller, or to the desk edge (see Fig. 1–left). Through this menu, users can enter the query editing mode, and activate several features we propose, such as displaying links, displaying cutting planes, displaying point projections, and changing perspectives.

However, taking advantage of the 3D space, we also propose to decompose this menu into contextual menus as much as possible. We believe this avoids overloading the interface and makes relevant buttons easier to find when needed. In this context, map settings (such as activating satellite images or 3D buildings) are attached to the desk, while query modification buttons are attached to each specific query (see Fig. 1–center right). To avoid having menus out of user reach, the query menus, which are only visible in the query editing mode, are positioned closer to the user's head.

5 CONCEPT DEMONSTRATION

In this section, we present and discuss a series of examples that showcase how our adapted and extended design works for exploring the NYC taxi dataset. To reflect on how an immersive reimagination could reproduce and extend *TaxiVis*'s workflows, we revisit case studies from the original paper [23]. As discussed earlier, our idea is not to replace *TaxiVis*, which works remarkably well for what it was designed for, but to reflect on the potential benefits opened up by IA.

We use taxi data from 2010-2015 [17] in these examples, as more recent data no longer includes detailed latitude-longitude information due to privacy concerns. The analysis of newer data with approximate locations is still possible but would not be as directly comparable to *TaxiVis*. Given the scalability challenges we discuss in Sect. 6.3, in these examples, we work with randomly sampled subsets of data containing up to a hundred thousand taxi trips each.

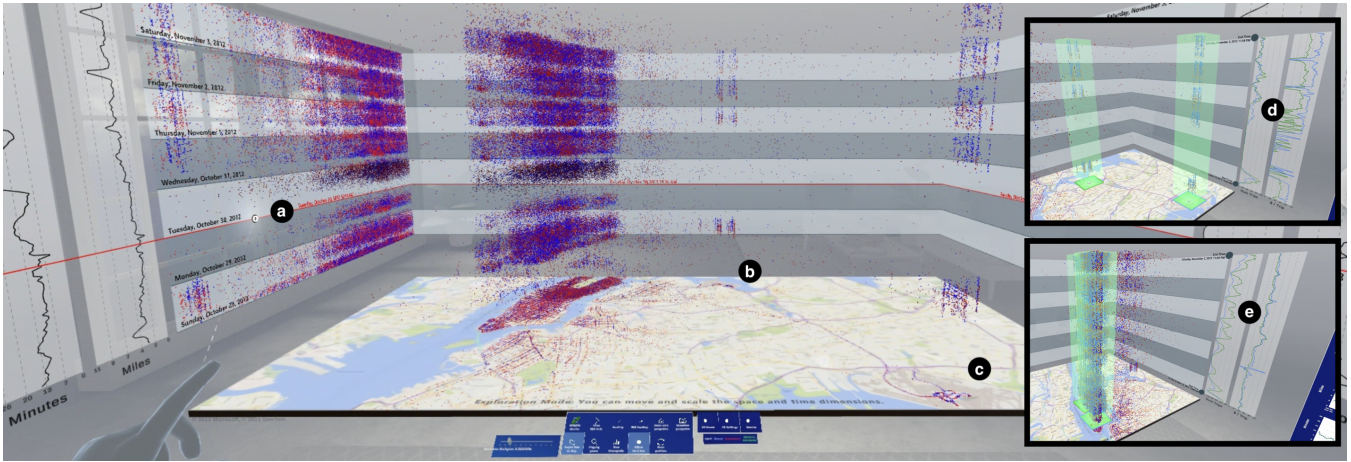


Figure 6: Overview of taxi activity in the week of Hurricane Sandy in October 2012. The STC point cloud and its projections (a) evidence the sudden drop in activity from Monday to Tuesday and the gradual recovery in Lower Manhattan over the next days. The insets highlight the use of interactive brushing and embedded plots to compare activity at different airports (top) and in different parts of Manhattan (bottom).

5.1 Using the self-evidence of temporal patterns for an improved analysis of behavior over time

The workflow for the identification of unexpected temporal behaviors could start by analyzing a time series distribution of trips over time, using the coordinated 2D panel or the embedded plots (for example, in Fig. 1–right, the unusually low activity at the STC’s top corresponds to Memorial Day). While this would be typically followed by an inspection of periods of interest through the *time-space exploration mechanism* in *TaxiVis*, i.e., creating a series of small multiple visualizations with a user-defined step size, in our immersive system, the STC point cloud presents a clear picture of the evolution of an event over time, helping users to detect and inspect both gradual or sudden variations of behavior over time. Some features, such as the ability to manually adjust point sizes, allow the user to make patterns more evident. Further, bi-manual cutting planes and brushes also help in identifying periods of interest. On-demand 2D projections on the map and walls (Fig. 6–a) also contribute to the analysis when combined with mid-air brushes and queries.

A good example is the case of major weather events that disrupt life in NYC. Like Ferreira et al., we analyze the cases of Hurricanes Irene and Sandy, which severely impacted the city in August 2011 and October 2012, respectively. Fig. 1–center left depicts the evolution of Irene, while Fig. 6 depicts Sandy. The STC view allows a clear comparison between the patterns for each storm, evidencing, for example, how Irene led to a more extreme interruption of taxi services. While the same conclusion was reached by Ferreira et al., in our system, it is more straightforward to arrive at this insight.

The STC also helps highlight location-specific temporal patterns. One example is the clear difference between activity at the airports LaGuardia (Fig. 6–b) and JFK (Fig. 6–c) following Sandy. Clearly, taxis return to JFK on Oct. 31st, while they only return to LaGuardia on Nov. 1st, what is also confirmed by the embedded plots when brushing these areas (Fig. 6–d). News articles from the time confirm that LaGuardia flights resumed later due to significant flooding. Another example for the use of comparative brushes or queries is an analysis of how activity was restored in different areas after the storm. Fig. 6–e illustrates a comparison of Midtown to Lower Manhattan, with the latter presenting a slower recovery due to a power outage. Moving into a room-scale view of the data and going inside the point cloud is also useful to inspect such patterns.

Reflecting on future use scenarios, and assuming the use of real-time data, the Space-Time Cube view could also aid an analysis by providing clear contextual information to current activity.

5.2 Using query prisms and alternative perspectives to support the analysis of transportation hubs

Taxi data can also shed light on how people move in and out of the city, by focusing on taxi trips to and from transportation hubs such as airports and train stations. In *Immersive TaxiVis*, it is possible to investigate these hubs with the help of pointer brushes (indicating either regions or origin-destination pairs) or query prisms positioned around these locations. The height of the prisms can be adjusted to specify a period of interest, and recurrent selections can focus on specific parts of the day, for example (Fig. 1–center right).

Following Ferreira et al., we compare trips departing from Lower Manhattan to LaGuardia and JFK airports over the first week of May 2011. Fig. 7 shows the corresponding directional queries (a and b) and the patterns revealed in the embedded time series plots. Note how the number of trips varies over the week (with a large spike on Friday afternoon, Fig. 7–c) and how the trip duration in minutes changes according to the day of the week and time of the day (Fig. 7–d). Meanwhile, the trip distances naturally remain constant, as do the fares per trip (Fig. 7–e), since trips from Manhattan to airports have predetermined prices. As in this scenario fewer trips are selected, connecting lines (Fig. 7–f) can be used to emphasize their patterns. The opposite query is also possible, selecting trips starting at a transportation hub and observing the most common destinations in the city, as shown in Fig. 4–left with Grand Central Station. For this, both the on-demand 2D map projections and 2D choropleth mode are useful.

Switching to the egocentric room-scale perspective allows the analyst to step onto the map and walk around to observe points of interest more closely (Fig. 4–left). In some cases, the egocentric 360° images from *Google StreetView* can also aid in the exploration of potential causes for large numbers of taxi destinations in a specific area (or the lack thereof). Fig. 4–right, for example, reveals a hotel in Manhattan—not surprisingly, the picture itself captures a large number of yellow taxis. In a collaborative exploration scenario, it is possible to imagine two analysts doing such inspections together or exploring different perspectives simultaneously.

5.3 Using comparative brushing and embedded attribute plots to investigate activity in different city regions

One of the ideas behind *TaxiVis* is using taxis as urban sensors which can help to reveal broader urban behaviors. One such possible analysis is to compare taxi activity in different regions of the city. Although the same analysis is also possible with visual queries, Fig. 3

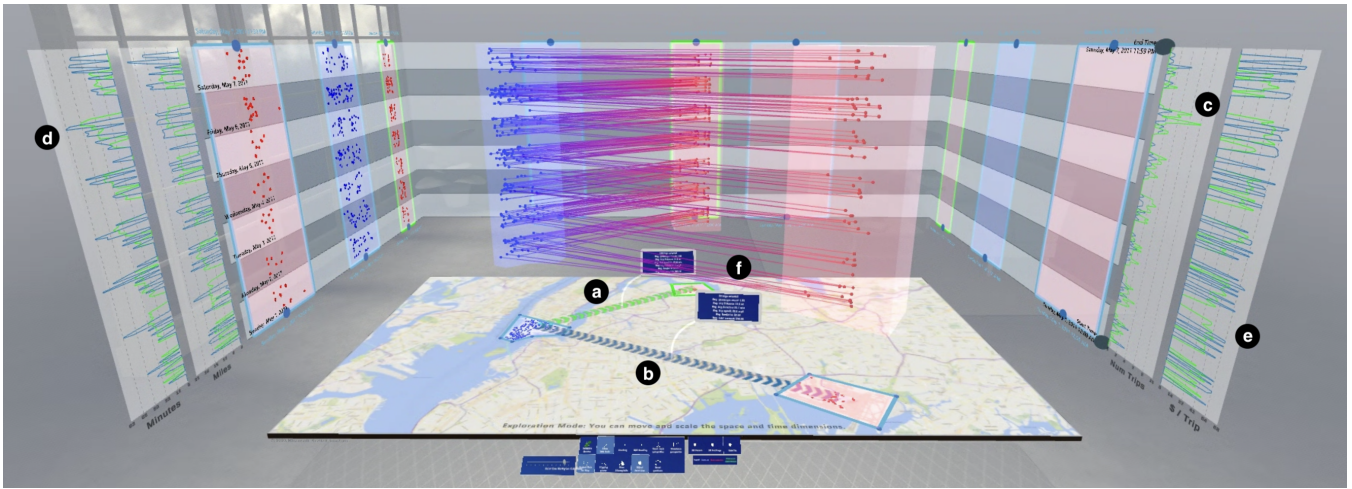


Figure 7: Directional queries selecting trips departing from Lower Manhattan to LaGuardia (a) and JFK (b) airports. The embedded plots indicate how trip numbers (c) and duration (d) vary during one week. On-demand origin-destination connecting lines (f) help emphasize patterns.

shows a comparative exploratory analysis between different neighborhoods using both pointer brushes (whose selections immediately affect all coordinated plots) and the choropleth mode brushing that builds prism stacks evidencing periods of higher activity.

Another example, Fig. 1—center right, demonstrates how the use of recurrent queries helps the user to focus on specific times. In this case, we investigate incoming taxi activity (dropoffs) in Lower Manhattan in the evenings (6PM-12AM). The sliced prisms emphasize the currently selected times, while the point cloud and embedded plots show that activity grows over the week (upwards, in this case), and the on-demand choropleth map specifies the number of trips that originated at other Manhattan neighborhoods.

Building on these examples, we can conceive multiple application scenarios that incorporate other *immersive extensions*. In a collaborative scenario, multiple users could quickly investigate different hypotheses through appropriate extensions. For example, the query duplication button helps to create multiple variations of a query targeting different areas or times. The future integration of private and shared views could thus be useful to enable parallel exploration of different aspects. In a casual exploration scenario, a general citizen could be interested in exploring the activities in their neighborhoods, or compare their personal experiences to the city trends. In a real-time data scenario, query prisms could emphasize the currently applied filters when only some trips are relevant.

6 DISCUSSION

Having designed, implemented, and demonstrated our proof-of-concept, we revisit our original question on the potential benefits of IA for urban data exploration in this section. We discuss our current perspective on benefits and drawbacks of the immersive paradigm, technical challenges, and our research vision for the next years.

6.1 Researchers' perspective: What does the immersive paradigm bring to the table?

Our team includes researchers from both Urban Visualization and IA backgrounds, and with this project, our goal was to pursue a critical reflection, highlighting potential benefits but also noting potential drawbacks of an immersive approach to urban analytics.

The following initial collected perceptions correspond to our opinions based on our experience developing the *Immersive TaxiVis* system, and thus should not be interpreted as a completely unbiased validation. Still, we believe some potential benefits are clear. As illustrated in Sect. 5.1, some analyses can benefit from a better

visualization of the temporal component of the data, which is well supported in a 3D STC view. Without a VR interface, such an analysis would lack sufficient depth cues and rely on (effortful) 3D manipulation through 2D interaction devices, or would involve extremely cluttered 2D visualizations. Using the third dimension also allowed us to better integrate and correlate trip attribute data, juxtaposing embedded time series with queries and brushes, as shown in Fig. 6 and Fig. 7, for example. Prior IA work, such as *ImAxes* [11, 12], also enabled intuitive exploration of combinations of attributes, and this seems to be a promising direction. In the future, allowing direct transformations between 2D and 3D views should also further increase their usefulness [35, 53].

In a world of growing familiarity with spatial interfaces, the use of intuitive metaphors such as sculpting queries [28] and simultaneously brushing origin and destination by two-handed interaction can potentially simplify data exploration activities, conceivably decreasing the system-incurred cognitive load and making tools more accessible to people of different backgrounds. As also demonstrated by *Uplift* [21], such tools can facilitate knowledge sharing between different kinds of stakeholders. In a single-user scenario, the use of a VR HMD is also potentially helpful to enable more focused exploration, removing external distractions. We also foresee the growing use of VR applications for user engagement, particularly in the context of motivating casual data exploration by citizens.

We also consider that our 3D prisms clearly achieve the purpose of indicating the applied space-time filters accurately and enabling direct modifications to them. We see the decomposition of complex interfaces into spatial components connected to 3D objects as another appropriate simplification that leverages humans' abilities to interact with spatial environments. Being able to efficiently interact with the 3D terrain and building models is also an advantage within many urban analysis scenarios where the data is connected to the local topography or urban density. Finally, recent work also clearly indicates the potential for immersive collaboration [36, 51].

6.2 Current limitations of immersive urban analytics

On the other hand, we also identified limitations and concerns, especially for the short-term future. While we believe *Immersive TaxiVis* supports exploratory data analysis well, the most efficient placement of some interface components is still not entirely clear. For example, when experimenting with complex queries combining many atomic ones, we observed that the process could become more complex than in a conventional interface, where keyboard shortcuts facilitate

multiple selections, a functionality that *Immersive TaxiVis* currently does not support. The risk of potentially over-complicating the exploration through the limitations of the user interface is present and must be navigated by the designers, though we expect that best practices will emerge. Conversely, we also see the risk of over-simplifying, seeking to merely reproduce desktop workflows (or even only parts of them). The idea of immersive interfaces should be to extend and complement desktop ones and, in this context, it is not surprising that not all tasks will be simpler under this paradigm.

Other current concerns include the disproportionate difficulty of developing IA systems, particularly ones that can be deployed to real applications. While the often-used Unity platform provides a vast set of resources, including dealing with maps, data rendering, and interaction toolkits, their integration is not straightforward, and some, such as *IATK*, still pose a considerable learning curve for more advanced applications. This challenge is even greater in the case of systems that must scale to very large datasets. Another concern is the limited generalizability of existing prototypes to other datasets, as many approaches, including in our case, are partly customized.

6.3 Implementation challenges and scalability

The main challenges we faced in building our prototype correspond to identifying the most adequate resources for implementing each component, leading to constant rework; learning to extend the functionality of *IATK*; and finding solutions that would scale to large urban datasets without compromising the high frame rate required.

As originally discussed in the *IATK* paper [10], the toolkit is able to sustain good performance up to around two million data points, after which performance drops significantly. Our implementation for spatio-temporal querying and brushing, however, sustains a very satisfactory interactive performance for datasets up to a hundred thousand trips (that is, two hundred thousand points, or six hundred thousand points with all projections enabled). Above that size, query modifications can cause the system to freeze for a few seconds, partially hindering the immersive experience. This is mostly due to the cost of computing query stats for the associated visualizations. As reference: we worked in PC-VR mode, with the Oculus Quest 2 headset directly connected to a desktop PC equipped with an i7-8700 3.2GHz CPU, nVidia GTX 1070 GPU, and 16GB of RAM.

While our current proof-of-concept system already efficiently supports many moderately-sized urban datasets, for example, those obtained from origin-destination surveys [42], others, such as the taxi and bike-sharing datasets from NYC can only be currently imported in subsets, as they contain millions of points per week or month. We intend to continue optimizing query computations to support larger datasets and will assess options to decouple the query computing and rendering components in the future.

6.4 Research vision and the path ahead

6.4.1 For immersive urban analytics

We expect that, through better support for context-switching between immersive and non-immersive approaches, most of the issues discussed above will be either addressed or significantly reduced over the next years, enabling users to benefit from the best of both “worlds”. More capable devices will also allow users to more seamlessly switch between systems to access different interaction modalities, e.g., text entry on a keyboard and 3D manipulation through a controller. More accurate and precise interaction devices, together with enhanced display resolutions, will also support efficient 2D interaction within immersive systems. Moreover, asymmetric collaboration, where one user investigates the data within an immersed perspective while another interacts with a more conventional interface, can also overcome some of the limitations [51, 59].

From an engineering perspective, we expect that novel toolkits will simplify the development of IA applications, for example, through the use of visualization grammars [37]. Web-based urban

visualization toolkits [27, 44] will likely also be increasingly adapted into AR and VR through *WebXR* platforms.

In terms of applications, we expect to see the development of novel approaches to take advantage of more photo-realistic 3D city models, for example, to directly visualize the effects of buildings in the urban space [22, 72]. As immersive applications are already seeing increasing adoption in the context of architecture [13], an extension to urban planning is thus also expected. In this context, the use of immersive systems will also likely be applied to inform and engage citizens in the planning process [30, 52]. As a medium-term goal, the integration of multiple sources of urban data, including real-time data, will certainly be of interest to many applications.

6.4.2 For the *Immersive TaxiVis* system

In future work, our goal is to complement *Immersive TaxiVis* with additional functionality, such as 2D and 3D heat map visualizations, and *immersive extensions*. In particular, we will look at collaborative analysis activities performed by multiple users. We will also further refine the design of our user interface and consider options to improve the system’s performance to support larger urban datasets.

While we see *Immersive TaxiVis* as a proof-of-concept demonstration for the *immersive extensions* and *adaptations* we envisioned, we also consider it a great platform for evaluations with users. In future work, we thus intend to conduct expert and usability evaluations using different urban datasets. While this was beyond the scope of this paper, such evaluations will be essential to better understand how data exploration workflows vary between different paradigms and how they can complement each other. Towards this purpose, we plan to improve the generalizability of our visualizations to support additional datasets, for example, extending views in the attribute exploration panels to also support categorical attributes.

7 CONCLUSION

Above all, this work was a reflection on how the growing adoption of immersive paradigms for data visualization can contribute to the visualization of urban datasets from the point of view of Immersive Analytics and Urban Visualization research. Our reflections were guided by the development of an immersive proof-of-concept prototype for the seminal *TaxiVis* system in urban data visualization. We designed our prototype with the simultaneous intentions of preserving key functionality of the original desktop interface while also complementing that set of features with new ones enabled by *immersive extensions*. With a series of examples related to the NYC taxi dataset, we demonstrated how this system could be used in practice, for example, to more clearly inspect behaviors over time or to inspect locations of interest with complementary egocentric perspectives. We envision possible continuations of this work through expert and usability user evaluations and by reimagining additional well-known 2D data exploration systems as a way of identifying additional potential benefits of IA environments.

ACKNOWLEDGMENTS

Jorge Wagner was supported in this project by the Fulbright Foreign Student Program, by a Microsoft Research PhD Fellowship, and by the Brazilian National Council for Scientific and Technological Development (CNPq). *Claudio Silva* is partially supported by the DARPA PTG program, National Science Foundation CNS-1828576, and NASA. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of DARPA. *Wolfgang Stuerzlinger* acknowledges support from NSERC. *Luciana Nedel* also acknowledges support from CNPq (311259/2022-7, 421962/2023-2). This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) - Finance Code 001.

A SUPPLEMENTARY MATERIALS: NOTES ON THE PROTOTYPE SYSTEM IMPLEMENTATION

We implemented our proof-of-concept prototype using the Unity game engine, the most frequently used platform in IA research, given the multitude of valuable toolkits available for it. In this section, we highlight relevant resources we used and key implementation decisions we took, as well as their present limitations, with the goal of informing subsequent projects.

A.1 Data rendering resources

To overcome Unity’s performance limitations that arise with the creation of a large number of scene objects, we decided early on to adopt the *Immersive Analytics Toolkit (IATK)* [10]. To our knowledge, *IATK* is presently the most efficient form of implementing big data visualization in Unity, being a more accessible alternative than developing a custom rendering engine, as done, for example, by *FiberClay* [28].

IATK’s efficiency is the result of a series of clever design decisions: in *IATK*, data points are not represented as individual scene objects, but as vertices in a single big mesh. Depending on the selected visualization topology, different Unity shaders (i.e., code that runs on the computer’s GPU) render this mesh in the appropriate way. Correspondingly, interactions with data points are not detected by individual colliders, but by the intersection of brushes with mesh vertices, efficiently computed by Unity `Compute Shader` programs and encoded into the RGB pixel values of two-dimensional textures with size $\sqrt{n} \times \sqrt{n}$, where n is the number of points in the dataset.

However, the main limitation of *IATK* is that it was mostly intended for manual configuration of views through the Unity editor interface, with limited and little-documented support for code-based customization. We believe our approach to adapting it to suit our specific needs can serve as an example for future IA applications.

A.2 Interaction and mapping resources

In terms of user interaction, we decided to adopt the *Mixed Reality Toolkit (MRTK)*. This toolkit provides excellent support for easy implementation of interface components, with support for multiple devices of different brands through the *OpenXR* protocol and good support for integration with the *Oculus Package* to support both controller and hand-tracking actions as enabled by the *Oculus Quest 2* HMD, the one we worked with in this project. Future works may also consider other alternatives, such as the *Unity XRI Toolkit*.

To develop our STC visualization, we adopted the *Bing Maps SDK* for Unity to integrate a dynamic mapping API. This SDK provides efficient rendering of 3D terrain meshes and also of 3D buildings for some cities, including New York City, as seen in Fig. 1–center. Several other alternatives are also available and may be preferable depending on the project, including *Cesium for Unity* using photo-realistic 3D data from *Google Maps*, the *Mapbox Maps SDK for Unity*, or the *ArcGIS Maps SDK for Unity*.

A.3 Integrating STC queries into IATK

Our prototype implementation, whose source code is publicly available, demonstrates how the *IATK* code and its core optimizations can be programmatically adapted to construct a query-able STC. We adapted the original *Scatterplot View* from *IATK* to be dynamically attached to the interactions performed on the STC map and time walls. For datasets with a bipartite nature such as taxi data (independent OD pairs), we create two *View* components with *Point* topologies, as well as a *LinkingView*, whose visual OD connections are hidden by default but can be displayed on demand. We also create four additional *View* components to render bidimensional projections on space and time when enabled.

To implement efficient display and querying of millions of data points, *TaxiVis*’ architecture decoupled its rendering and storage components. Replicating this behavior would require dynamically

generating new *IATK* big meshes for each query result, which could be impracticable without hindering interactivity and breaking the immersive experience. Therefore, for our proof-of-concept, we opted to instead heavily adapt the original *BrushingAndLinking IATK* class to compute all queries through Unity shaders and encode query results into three texture images (encoding filtered points, brushed points, and highlighted points). We also significantly extended the *IATK MyComputeShader* to enable filtering by arbitrary 3D prisms, and to enable several such prisms to be combined, through texture post-processing done by a second *ComputeShader* program, depending on the intended query logic. The overall query logic is managed by our *QueryManager* class.

By coupling the rendering and querying components of the system, we gain the ability to efficiently apply spatio-temporal filters that scale up to hundreds of thousands of data points, without having to regenerate the *IATK View* mesh. At the same time, this prevents us from computing queries on millions of data points, something that *TaxiVis* achieved (technically speaking, *TaxiVis* never rendered more than one million points at the same time, but could query and compute stats over larger sets of points in the background). We see this as a reasonable implementation compromise at this time.

Given our performance concerns to maintain a comfortable frame rate for VR exploration, we also worked to reduce the number of view updates when possible. Further technical details about our query implementation approach are beyond the scope of this publication, but are documented in our GitHub repository.

REFERENCES

- [1] L. Alonso, Y. R. Zhang, A. Grignard, A. Noyman, Y. Sakai, M. ElKasha, R. Doorley, and K. Larson. Cityscope: A data-driven interactive simulation tool for urban design. use case volpe. In A. J. Morales, C. Gershenson, D. Braha, A. A. Minai, and Y. Bar-Yam, eds., *Unifying Themes in Complex Systems IX*, pp. 253–261. Springer International Publishing, Cham, 2018. doi: 10.1007/978-3-319-96661-8_27
- [2] B. Bach, P. Dragicevic, D. Archambault, C. Hurter, and S. Carpendale. A review of temporal data visualizations based on space-time cube operations. In R. Borgo, R. Maciejewski, and I. Viola, eds., *Eurographics Conference on Visualization - STARs*. The Eurographics Association, Swansea, UK, 2014. doi: 10.2312/eurovisstar.20141171
- [3] A. Batch, A. Cunningham, M. Cordeil, N. Elmqvist, T. Dwyer, B. H. Thomas, and K. Marriott. There is no spoon: Evaluating performance, space use, and presence with expert domain users in immersive analytics. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):536–546, 2020. doi: 10.1109/TVCG.2019.2934803
- [4] M. Berger and R. Bill. Combining vr visualization and sonification for immersive exploration of urban noise standards. *Multimodal Technologies and Interaction*, 3(2), 2019. doi: 10.3390/mti3020034
- [5] T. Chandler, M. Cordeil, T. Czauderna, T. Dwyer, J. Glowacki, C. Goncu, M. Klapperstueck, K. Klein, K. Marriott, F. Schreiber, et al. Immersive analytics. In *2015 Big Data Visual Analytics (BDVA)*, pp. 1–8. IEEE, September 2015. doi: 10.1109/BDVA.2015.7314296
- [6] S. Chen, F. Miranda, N. Ferreira, M. Lage, H. Doraiswamy, C. Brenner, C. Defanti, M. Koutsoubis, L. Wilson, K. Perlin, and C. Silva. Urbanrama: Navigating cities in virtual reality. *IEEE Transactions on Visualization and Computer Graphics*, 28(12):4685–4699, 2022. doi: 10.1109/TVCG.2021.3099012
- [7] Z. Chen, H. Qu, and Y. Wu. Immersive urban analytics through exploded views. In *Workshop on Immersive Analytics*. IEEE, Phoenix, AZ, USA, 2017.
- [8] Z. Chen, Y. Wang, T. Sun, X. Gao, W. Chen, Z. Pan, H. Qu, and Y. Wu. Exploring the design space of immersive urban analytics. *Visual Informatics*, 1(2):132–142, 2017. doi: 10.1016/j.visinf.2017.11.002
- [9] T. Cheng, G. Tanaksaranond, C. Brunsdon, and J. Haworth. Exploratory visualisation of congestion evolutions on urban transport networks. *Transportation Research Part C: Emerging Technologies*, 36:296–306, 2013. doi: 10.1016/j.trc.2013.09.001
- [10] M. Cordeil, A. Cunningham, B. Bach, C. Hurter, B. H. Thomas, K. Marriott, and T. Dwyer. IATK: An immersive analytics toolkit. In *IEEE*

- Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 200–209. IEEE, Osaka, Japan, 2019. doi: 10.1109/VR.2019.8797978
- [11] M. Cordeil, A. Cunningham, T. Dwyer, B. H. Thomas, and K. Marriott. ImAxes: Immersive axes as embodied affordances for interactive multivariate data visualisation. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*, pp. 71–83. ACM, Québec City, QC, Canada, 2017. doi: 10.1145/3126594.3126613
 - [12] A. Cunningham, J. D. Hart, U. Engelke, M. Adcock, and B. H. Thomas. Towards embodied interaction for geospatial energy sector analytics in immersive environments. In *Proceedings of the Twelfth ACM International Conference on Future Energy Systems, e-Energy '21*, p. 396–400. Association for Computing Machinery, Virtual Event, Italy, 2021. doi: 10.1145/3447555.3466602
 - [13] J. M. Davila Delgado, L. Oyedele, P. Demian, and T. Beach. A research agenda for augmented and virtual reality in architecture, engineering and construction. *Advanced Engineering Informatics*, 45:101122, 2020. doi: 10.1016/j.aei.2020.101122
 - [14] U. Demšar and K. Verrantaus. Space–time density of trajectories: Exploring spatio-temporal patterns in movement data. *International Journal of Geographical Information Science*, 24(10):1527–1542, 2010. doi: 10.1080/13658816.2010.511223
 - [15] Z. Deng, D. Weng, S. Liu, Y. Tian, M. Xu, and Y. Wu. A survey of urban visual analytics: Advances and future directions. *Computational Visual Media*, 9(1):3–39, 2023. doi: 10.1007/s41095-022-0275-7
 - [16] L. Ding, J. M. Krisp, L. Meng, G. Xiao, and A. Keler. Visual exploration of multivariate movement events in space-time cube. In *19th AGILE International Conference on Geographic Information Science*. AGILE, Helsinki, Finland, 2016.
 - [17] B. Donovan and D. Work. New york city taxi trip data (2010–2013), 2016. doi: 10.13012/J8PN93H8
 - [18] H. Doraiswamy, J. Freire, M. Lage, F. Miranda, and C. Silva. Spatio-temporal urban data analysis: A visual analytics perspective. *IEEE Computer Graphics and Applications*, 38(5):26–35, 2018. doi: 10.1109/MCG.2018.053491728
 - [19] C. Elvezio, F. Ling, J.-S. Liu, B. Tversky, and S. Feiner. Collaborative exploration of urban data in virtual and augmented reality. In *ACM SIGGRAPH 2018 Virtual, Augmented, and Mixed Reality*, SIGGRAPH '18. Association for Computing Machinery, Vancouver, BC, Canada, 2018. doi: 10.1145/3226552.3226570
 - [20] B. Ens, B. Bach, M. Cordeil, U. Engelke, M. Serrano, W. Willett, A. Prouzeau, C. Anthes, W. Büschel, C. Dunne, T. Dwyer, J. Grubert, J. H. Haga, N. Kirshenbaum, D. Kobayashi, T. Lin, M. Olaosebikan, F. Pointecker, D. Saffo, N. Saquib, D. Schmalstieg, D. A. Szafr, M. Whitlock, and Y. Yang. Grand challenges in immersive analytics. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI '21. Association for Computing Machinery, Yokohama, Japan, 2021. doi: 10.1145/3411764.3446866
 - [21] 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
 - [22] N. Ferreira, M. Lage, H. Doraiswamy, H. Vo, L. Wilson, H. Werner, M. Park, and C. Silva. Urbane: A 3d framework to support data driven decision making in urban development. In *2015 IEEE Conference on Visual Analytics Science and Technology (VAST)*, pp. 97–104. IEEE, Chicago, IL, USA, 2015. doi: 10.1109/NAIST.2015.7347636
 - [23] N. Ferreira, J. Poco, H. T. Vo, J. Freire, and C. T. Silva. Visual exploration of big spatio-temporal urban data: A study of new york city taxi trips. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2149–2158, 2013. doi: 10.1109/TVCG.2013.226
 - [24] A. Fonet and Y. Prié. Survey of immersive analytics. *IEEE Transactions on Visualization and Computer Graphics*, 27(3):2101–2122, 2021. doi: 10.1109/TVCG.2019.2929033
 - [25] T. Gonçalves, A. P. Afonso, and B. Martins. Improving spatial awareness for human trajectory visualization in space-time cubes. In J. Abascal, S. Barbosa, M. Fetter, T. Gross, P. Palanque, and M. Winckler, eds., *Human-Computer Interaction – INTERACT 2015*, pp. 327–334. Springer International Publishing, Cham, 2015. doi: 10.1007/978-3-319-22723-8_26
 - [26] T. Hägerstrand. What about people in regional science? *Papers in Regional Science*, 24(1):7–24, 1970.
 - [27] S. He. From beautiful maps to actionable insights: Introducing kepler.gl, uber's open source geospatial toolbox. <https://www.uber.com/blog/kepler.gl/>, May 2018. Accessed: 2024-01-31.
 - [28] C. Hurter, N. H. Riche, S. M. Drucker, M. Cordeil, R. Alligier, and R. Vuilleumot. 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
 - [29] IEEE VIS 2023. <https://ieevis.org/year/2023/info/award-s/test-of-time-awards>, September 2023. Accessed: 2024-01-31.
 - [30] J. Keil, M. Weißmann, A. Korte, D. Edler, and F. Dickmann. Measuring physiological responses to visualizations of urban planning scenarios in immersive virtual reality. *KN-Journal of Cartography and Geographic Information*, 73:117–126, 2023. doi: 10.1007/s42489-023-00137-7
 - [31] M.-J. Kraak. The space-time cube revisited from a geovisualization perspective. In *Proceedings of the 21st International Cartographic Conference (ICC 2003)*, pp. 1988–1996. International Cartographic Association, Durban, South Africa, 2003.
 - [32] M. Kraus, K. Angerbauer, J. Buchmüller, D. Schweitzer, D. A. Keim, M. Sedlmair, and J. Fuchs. Assessing 2d and 3d heatmaps for comparative analysis: An empirical study. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI '20, p. 1–14. Association for Computing Machinery, Honolulu, HI, USA, 2020. doi: 10.1145/3313831.3376675
 - [33] M. Kraus, J. Fuchs, B. Sommer, K. Klein, U. Engelke, D. Keim, and F. Schreiber. Immersive analytics with abstract 3d visualizations: A survey. *Computer Graphics Forum*, 41(1):201–229, 2022. doi: 10.1111/cgf.14430
 - [34] 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, 2021. doi: 10.1109/TVCG.2020.3030435
 - [35] B. Lee, M. Cordeil, A. Prouzeau, B. Jenny, and T. Dwyer. A design space for data visualisation transformations between 2d and 3d in mixed-reality environments. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, CHI '22. Association for Computing Machinery, New Orleans, LA, USA, 2022. doi: 10.1145/3491102.3501859
 - [36] 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(02):1171–1181, Feb. 2021. doi: 10.1109/TVCG.2020.3030450
 - [37] B. Lee, A. Satyanarayan, M. Cordeil, A. Prouzeau, B. Jenny, and T. Dwyer. Deimos: A grammar of dynamic embodied immersive visualisation morphs and transitions. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23. Association for Computing Machinery, Hamburg, Germany, 2023. doi: 10.1145/3544548.3580754
 - [38] O. Lock, T. Bednarz, and C. Pettit. Holocity – exploring the use of augmented reality cityscapes for collaborative understanding of high-volume urban sensor data. In *Proceedings of the 17th International Conference on Virtual-Reality Continuum and Its Applications in Industry, VRCAI '19*. Association for Computing Machinery, Brisbane, QLD, Australia, 2019. doi: 10.1145/3359997.3365734
 - [39] O. Lock, T. Bednarz, and C. Pettit. The visual analytics of big, open public transport data – a framework and pipeline for monitoring system performance in greater sydney. *Big Earth Data*, 5(1):134–159, 2021. doi: 10.1080/20964471.2020.1758537
 - [40] L. C. W. Louna Al Bondakji, Anne-Liese Lammich. Vibe (virtual berlin) - immersive interactive 3d urban data visualization. *Blucher Design Proceedings*, 7(1):83–90, 2019. doi: 10.5151/proceedings-ecaadesigradi2019.318
 - [41] K. Marriott, F. Schreiber, T. Dwyer, K. Klein, N. H. Riche, T. Itoh, W. Stuerzlinger, and B. H. Thomas, eds. *Immersive Analytics*, vol. 11190 of *Lecture Notes in Computer Science*. Springer, 2018. doi: 10.1007/978-3-030-01388-2
 - [42] A. Menin, S. Chardonnel, P.-A. Davoine, and L. Nedel. estime: To-

- wards an all-in-one geovisualization environment for daily mobility analysis. In *32nd SIBGRAPI Conference on Graphics, Patterns and Images (SIBGRAPI)*, pp. 39–46. IEEE, Rio de Janeiro, Brazil, 2019. doi: 10.1109/SIBGRAPI.2019.00014
- [43] F. Miranda, M. Hosseini, M. Lage, H. Doraiswamy, G. Dove, and C. T. Silva. Urban mosaic: Visual exploration of streetscapes using large-scale image data. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pp. 1–15. ACM, Honolulu, HI, USA, 2020. doi: 10.1145/3313831.3376399
- [44] G. Moreira, M. Hosseini, M. N. A. Nipu, M. Lage, N. Ferreira, and F. Miranda. The urban toolkit: A grammar-based framework for urban visual analytics. *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–11, 2023. doi: 10.1109/TVCG.2023.3326598
- [45] R. Mota, N. Ferreira, J. Silva, M. Horga, M. Lage, L. Ceferino, U. Alim, E. Sharlin, and F. Miranda. A comparison of spatiotemporal visualizations for 3d urban analytics. *IEEE Transactions on Visualization and Computer Graphics*, 29(01):1277–1287, Jan. 2023. doi: 10.1109/TVCG.2022.3209474
- [46] G. N. Oliveira, J. L. Sotomayor, R. P. Torchelsen, C. T. Silva, and J. L. Comba. Visual analysis of bike-sharing systems. *Computers & Graphics*, 60:119–129, 2016. doi: 10.1016/j.cag.2016.08.005
- [47] T. Piumsomboon, G. A. Lee, A. Irlitti, B. Ens, B. H. Thomas, and M. Billingham. On the shoulder of the giant: A multi-scale mixed reality collaboration with 360 video sharing and tangible interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, p. 1–17. Association for Computing Machinery, Glasgow, Scotland Uk, 2019. doi: 10.1145/3290605.3300458
- [48] Z. Pousman, J. Stasko, and M. Mateas. Casual information visualization: Depictions of data in everyday life. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1145–1152, 2007. doi: 10.1109/TVCG.2007.70541
- [49] J. Reades, F. Calabrese, A. Sevtsuk, and C. Ratti. Cellular census: Explorations in urban data collection. *IEEE Pervasive Computing*, 6(3):30–38, 2007. doi: 10.1109/MPRV.2007.53
- [50] P. Reipschläger, T. Flemisch, and R. Dachsel. Personal augmented reality for information visualization on large interactive displays. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1182–1192, 2021. doi: 10.1109/TVCG.2020.3030460
- [51] D. Saffo, A. Batch, C. Dunne, and N. Elmqvist. Through their eyes and in their shoes: Providing group awareness during collaboration across virtual reality and desktop platforms. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI ’23. Association for Computing Machinery, Hamburg, Germany, 2023. doi: 10.1145/3544548.3581093
- [52] H. Schrom-Feiertag, M. Stubenschrott, G. Regal, T. Matyus, and S. Seer. An interactive and responsive virtual reality environment for participatory urban planning. In *Proceedings of the 11th Annual Symposium on Simulation for Architecture and Urban Design*, SimAUD ’20. Society for Computer Simulation International, Virtual Event, Austria, 2020.
- [53] M. R. Seraji and W. Stuerzlinger. Hybridaxes: An immersive analytics tool with interoperability between 2d and immersive reality modes. In *2022 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*, pp. 155–160. IEEE, Singapore, Singapore, 2022. doi: 10.1109/ISMAR-Adjunct57072.2022.00036
- [54] X. Shi, F. Lv, D. Seng, B. Xing, and J. Chen. Exploring the evolutionary patterns of urban activity areas based on origin-destination data. *IEEE Access*, 7:20416–20431, 2019. doi: 10.1109/ACCESS.2019.2897070
- [55] M. Spur and V. Tourre. Urban DataSphere: Exploring Immersive Multiview Visualizations in Cities. *IEEE VIS*, Sept. 2018. Poster.
- [56] M. Spur, V. Tourre, E. David, G. Moreau, and P. Le Callet. Exploring multiple and coordinated views for multilayered geospatial data in virtual reality. *Information*, 11(9), 2020. doi: 10.3390/info11090425
- [57] S. Y. Ssin, J. A. Walsh, R. T. Smith, A. Cunningham, and B. H. Thomas. GeoGate: Correlating geo-temporal datasets using an augmented reality space-time cube and tangible interactions. In *Proceedings of the 26th IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 210–219. IEEE, Osaka, Japan, 2019. doi: 10.1109/VR.2019.8797812
- [58] 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. doi: 10.1109/TVCG.2012.265
- [59] W. Tong, M. Xia, K. K. Wong, D. A. Bowman, T.-C. Pong, H. Qu, and Y. Yang. Towards an understanding of distributed asymmetric collaborative visualization on problem-solving. In *2023 IEEE Conference Virtual Reality and 3D User Interfaces (VR)*, pp. 387–397. IEEE, Shanghai, China, 2023. doi: 10.1109/VR55154.2023.00054
- [60] H. T. Vo. Visual exploration of big urban data - a case study with taxivis. https://youtu.be/Lp_UdtQMsx4?t=1890, Dec. 2014. Accessed: 2024-01-31.
- [61] J. Wagner, W. Stuerzlinger, and L. Nedel. Comparing and combining virtual hand and virtual ray pointer interactions for data manipulation in immersive analytics. *IEEE Transactions on Visualization and Computer Graphics*, 27(5):2513–2523, 2021. doi: 10.1109/TVCG.2021.3067759
- [62] J. Wagner, W. Stuerzlinger, and L. Nedel. The effect of exploration mode and frame of reference in immersive analytics. *IEEE Transactions on Visualization and Computer Graphics*, 28(9):3252–3264, 2022. doi: 10.1109/TVCG.2021.3060666
- [63] J. A. Wagner Filho, M. F. Rey, C. M. D. S. Freitas, and L. Nedel. Immersive visualization of abstract information: An evaluation on dimensionally-reduced data scatterplots. In *Proceedings of the 25th IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 483–490. IEEE, Reutlingen, Germany, Mar. 2018. doi: 10.1109/VR.2018.8447558
- [64] 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. doi: 10.1109/TVCG.2019.2934415
- [65] K. Wiggers. Data visualization startup Virtualitics lands \$37m investment. <https://techcrunch.com/2023/08/10/data-visualization-startup-virtualitics-lands-37m-investment/>, August 2023. Accessed: 2024-01-31.
- [66] L. Yang, M.-P. Kwan, X. Pan, B. Wan, and S. Zhou. Scalable space-time trajectory cube for path-finding: A study using big taxi trajectory data. *Transportation Research Part B: Methodological*, 101:1–27, 2017. doi: 10.1016/j.trb.2017.03.010
- [67] Y. Yang, T. Dwyer, B. Jenny, K. Marriott, M. Cordeil, and H. Chen. Origin-destination flow maps in immersive environments. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):693–703, 2019. doi: 10.1109/TVCG.2018.2865192
- [68] 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
- [69] L. Yu, W. Wu, X. Li, G. Li, W. S. Ng, S.-K. Ng, Z. Huang, A. Arunan, and H. M. Watt. iviztrans: Interactive visual learning for home and work place detection from massive public transportation data. In *2015 IEEE Conference on Visual Analytics Science and Technology (VAST)*, pp. 49–56. IEEE, Chicago, IL, USA, 2015. doi: 10.1109/VAST.2015.7347630
- [70] W. Zeng, C.-W. Fu, S. M. Arisona, A. Erath, and H. Qu. Visualizing mobility of public transportation system. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1833–1842, 2014. doi: 10.1109/TVCG.2014.2346893
- [71] W. Zeng, C.-W. Fu, S. Müller Arisona, A. Erath, and H. Qu. Visualizing waypoints-constrained origin-destination patterns for massive transportation data. *Computer Graphics Forum*, 35(8):95–107, 2016. doi: 10.1111/cgf.12778
- [72] C. Zhang, W. Zeng, and L. Liu. Urbanvr: An immersive analytics system for context-aware urban design. *Computers & Graphics*, 99:128–138, 2021. doi: 10.1016/j.cag.2021.07.006
- [73] Y. Zhang, B. Ens, K. A. Satriadi, A. Prouzeau, and S. Goodwin. Timetables: Embodied exploration of immersive spatio-temporal data. In *2022 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 599–605. IEEE, Christchurch, New Zealand, 2022. doi: 10.1109/VR51125.2022.00080