PEARL: Physical Environment based Augmented Reality Lenses for In-Situ Human Movement Analysis

Weizhou Luo*
Interactive Media Lab
Technische Universität Dresden
Dresden, Germany

Zhongyuan Yu*
Immersive Experience Lab
Technische Universität Dresden
Dresden, Germany

Rufat Rzayev*
Interactive Media Lab
Technische Universität Dresden
Dresden, Germany

Marc Satkowski*†
Interactive Media Lab
Technische Univsersität Dresden
Dresden, Germany

Stefan Gumhold*†‡
Computer Graphics and Visualization
Technische Universität Dresden
Dresden, Germany

Matthew McGinity*
Immersive Experience Lab
Technische Universität Dresden
Dresden, Germany

Raimund Dachselt*†‡
Interactive Media Lab
Technische Universität Dresden
Dresden, Germany

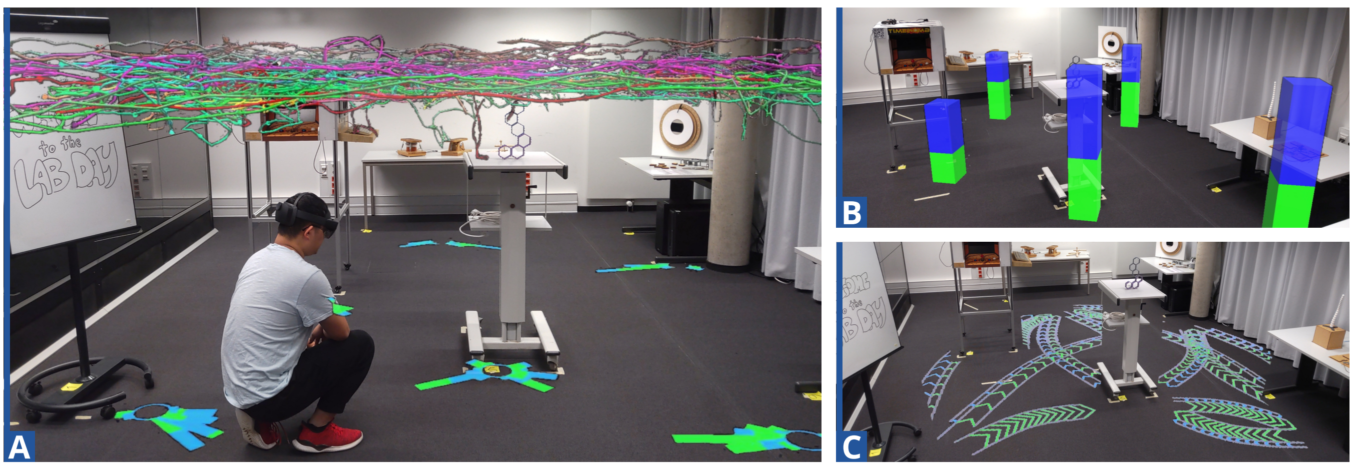

Figure 1: Exemplary PEARL visualizations for human movement data analysis based on *Regions of Interest*. (A) A combination of 3D trajectories showing fine-grained movement recordings and the Approach View summarizing how humans approached associated objects. (B) Superimposed 3D stacked bars embedded directly on the physical objects showing stay duration of two visitor groups. (C) Flow View with two color-coded visitor groups presenting the trend of their movement transitions.

ABSTRACT

This paper presents PEARL, a mixed-reality approach for the analysis of human movement data in situ. As the physical environment shapes human motion and behavior, the analysis of such motion can benefit from the direct inclusion of the environment in the analytical process. We present methods for exploring movement data in relation to surrounding regions of interest, such as objects, furniture, and architectural elements. We introduce concepts for selecting and filtering data through direct interaction with the environment, and a suite of visualizations for revealing aggregated and emergent spatial and temporal relations. More sophisticated analysis is supported through complex queries comprising multiple regions of interest. To illustrate the potential of PEARL, we developed an Augmented Reality-based prototype and conducted

*e-mail: [weizhou.luo, zhongyuan.yu, rufat.rzayev, marc.satkowski, stefan.gumhold, matthew.mcginity, raimund.dachselt]@tu-dresden.de
†Also with Centre for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI) Dresden/Leipzig, Germany
‡Also with Centre for Tactile Internet with Human-in-the-Loop (CeTI) and Cluster of Excellence Physics of Life, Technische Universität Dresden

CHI '23, April 23–28, 2023, Hamburg, Germany

ACM ISBN 978-1-4503-9421-5/23/04.
https://doi.org/10.1145/3544548.3580715

expert review sessions and scenario walkthroughs in a simulated exhibition. Our contribution lays the foundation for leveraging the physical environment in the in-situ analysis of movement data.

CCS CONCEPTS

• **Human-centered computing → Mixed / augmented reality**; **Visual analytics**; *Visualization techniques.*

KEYWORDS

In-situ visualization, augmented/mixed reality, Immersive Analytics, movement data analysis, affordance, physical referents

ACM Reference Format:
Weizhou Luo, Zhongyuan Yu, Rufat Rzayev, Marc Satkowski, Stefan Gumhold, Matthew McGinity, and Raimund Dachselt. 2023. Pearl: Physical Environment based Augmented Reality Lenses for In-Situ Human Movement Analysis. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23), April 23–28, 2023, Hamburg, Germany*. ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3544548.3580715

1 INTRODUCTION

Analyzing human movement data can provide valuable insights into movement behavior and spatial dynamics. Improvements in the affordability and performance of motion-tracking systems allow increasing possibilities to capture and analyze human motion behavior in various contexts [28]. However, due to its complex nature, understanding movement data requires sophisticated tools for visualization and analysis [20]. Researchers have proposed various tools to support the analysis of human movement data in different contexts, including understanding navigational behavior, room occupancy, and flow pattern of museum visitors [13, 38], exploring spatial and group interaction with the environment [10, 70], and analyzing human motion in sports [67] and dances [65].

Emerging mixed reality (MR) solutions provide new opportunities to explore motion data in the exact location where the data was captured. MR systems, such as AvatAR [57] or MIRIA [11], enable visualization and interactive exploration of movement data in the original environment. These kinds of systems provide analysts with awareness of the physical environment, providing context to the motions and actions captured in the data. Moreover, using MR technology in the original space allows analysts to take the perspective of involved humans and contextualize the initial situation.

Nevertheless, existing approaches usually focus on interacting directly with the movement data and largely ignore the interplay between human motion and the physical environment. For example, solutions have been introduced that allow direct interaction with movement data visualization (e.g., trajectories [32, 35]), with the representation of involved humans (e.g., virtual avatars [57]), or through an additional interface (e.g., tablets [35, 57] or tangibles [66]). However, human motion, especially in indoor environments, is often, consciously or unconsciously, affected by the physical objects populating a space and their affordances. A physical object might be an obstacle that has to be avoided to continue the movement, such as a table in the middle of the room, or an object of interest that is intentionally approached to interact with it, like a whiteboard on the wall. Thus, physical objects and their surroundings restrict, shape, or elicit human behavior and motion to varying extents. As such, they can serve as reference points for the analysis of associated human motion and behavior.

In fact, several works have shown the potential of using physical objects as references for movement data visualization, for instance, to present human movement around physical objects (like interactive display [11, 57, 70]), or using virtual objects as references to directly interact with the spatial recordings in virtual reality (VR) [42]. However, to the best of our knowledge, no previous work presented approaches to explore movement data in the original real-world environment by referring to regions of interest (*ROIs*) and providing visualization techniques from their perspective. As the physical environment affects human movement, it is promising to analyze movement data focusing on associated *ROIs* in the original environment. *ROIs* can contain one or more physical objects as referents from the original real-world environment, parts of them, or together with the close-by environment. The *ROIs* can play the role of *lenses* providing interaction and visualizations for movement data from their viewpoint.

In this work, we propose Pearl, **P**hysical **E**nvironment based **A**ugmented **R**eality **L**enses (see Fig. 1). Pearl is an approach for the in-situ exploration and analysis of human motion data using the physical objects in the original space as referents. It provides three functionalities: 1) Defining Lenses by **selecting** *ROIs* from the physical environment for exploring human motion data, 2) **Query-based filtering** the movement data based on the proximity and temporal relation to the selected *ROIs* and building complex queries using logical operations, and 3) **Visualizing** the movement data as detailed (like 3D trajectories) and aggregated embedded visualizations regarding the *ROIs*. In particular, aggregated visualizations provide a novel and unique way to visually summarize the movement data among the *ROIs* in the environment. Moreover, we designed and implemented a prototype of Pearl for Augmented Reality (AR) head-mounted displays (HMDs), including a suite of contextually embedded visualizations with different granularity. To illustrate its potential, we prepared a simulated exhibition room, collected spatial movement data, and designed several use-case scenarios. We conducted expert review sessions with domain experts in the simulated exhibition, systematically analyzed the results, and report the first insights of our approach.

In summary, our main contributions are the following:

(1) An immersive approach for the in-situ analysis of human motion data focusing on leveraging the affordances of physical objects as referents in the situated MR environment.
(2) Concepts for interactive querying and filtering human movement data based on regions of interest and a suite of situated visualizations for analyzing aggregated spatio-temporal data related to the physical environment.
(3) A prototype implementation of our concept, as well as expert review sessions and scenario-based walkthroughs highlighting how our concept can be used to analyze human movement data.

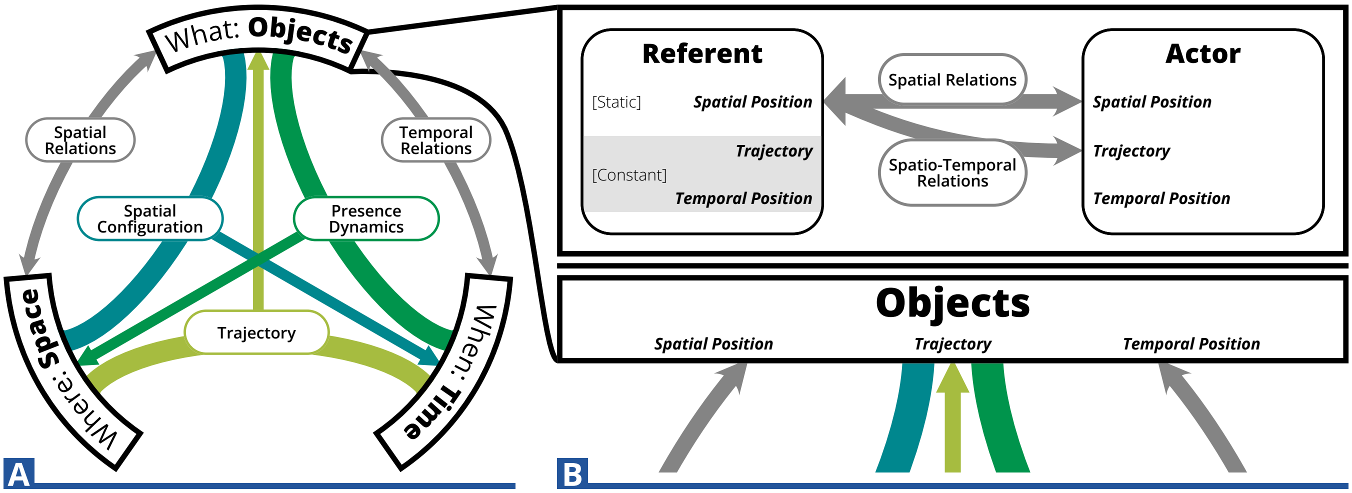

Figure 2: A model describing movement data. (A) shows the three components of spatio-temporal information, as well as their characteristics and relations [3, 54]. (B) highlights two specific Objects: the Referents and the Actors. The Referents differ from general Objects, and more specifically from Actors, by only having a static spatial position that is constant over time (temporal position). Thus, Referents have no trajectories.

2 BACKGROUND AND RELATED WORK

Our work focuses on exploring and analyzing movement data using AR and Immersive Analytics (IA) systems. Below, we outline the definition of spatio-temporal data and describe different perspectives of visual analysis, particularly indicating AR and IA systems.

2.1 Spatio-Temporal and Movement Data

Spatio-temporal data describes changes in space information over time [3, 4]. Its datasets are primarily associated with movement data, describing the *"changes of spatial positions of one or more moving objects"* [4]. Those objects can be inanimate, like drones [75], cars, or other vehicles [27], or animate, like animals [34] or, more relevant to this work, humans [11, 16, 80]. Besides space and time, such datasets can contain further properties, also called thematic attributes [4], for instance, the fuel level of a car or a person's conversation [62].

In general, spatio-temporal data stores information of three components: Where, When, and What, known as the triad scheme [54]. Those components can also be named as Space or Locations, Time, and Objects, respectively (see Fig. 2A) [3, 54]. Analysts can build relations from the elements of these components to answer a range of questions about, for example, specific spatial configurations, presence dynamics, or spatio-temporal positions and trajectories [3, 4] (see Fig. 2A).

Prior works have summarized users' workflows and tasks for analyzing movement data. Specifically, Andrienko et al. [4] classified the movement data analysis tasks as mover-oriented tasks on trajectories, event-oriented tasks about spatial events, space-oriented tasks about places of interest, and time-oriented tasks on specific time units. During exploring spatio-temporal data or model building [5, 19], an analyst can be interested in Space, Time, or Objects, their elements [3], or the type of changes occurring over time [4, 6, 8]. Such changes can be existential changes or changes in spatial properties and thematic attributes. Moreover, an analyst could be interested in the characteristics of the temporal domain, such as a moment, pace, duration, sequence, or frequency [8].

2.2 Visual Analysis of Movement Data

As movement data is highly complex, its visualization can contribute to the analysis procedure [4, 19] by supporting all its stages, including data exploration, cleaning, preprocessing, and querying [51]. The analysis can be driven by different models [3, 54], research questions [19], or challenges [51]. Following, we focus on the model presented in Fig. 2, describing the movement data with its three components, Object, Space, and Time, and their relations, with a particular focus on the human movement analysis.

Displaying color-encoded 2D or 3D trajectories is a common way of representing **Single Object** motion data [31, 79]. For instance, researchers have used trajectories to examine how people occupy space in a social context [41], to support navigation in spatial recordings [42], or to explore human motion data in VR [35]. These visualizations can be augmented with object representations to provide a better understanding of their motion and behavior. For example, several works used 2D avatars [10], 3D simple avatars [11, 35, 42], and 3D detailed avatars [40, 57] for displaying humans. Furthermore, objects are often influenced by other objects [19], making it necessary to support analysis processes related to **Object-Object Relation**. For instance, a direct comparison of co-located objects [62], the characterization of their interaction [40, 41], and additional visualizations [10, 11, 70] can be presented to uncover relations between objects. Moreover, to explore movement data, not only the analysis of single objects but also **Aggregated Objects**, via grouping and clustering, is crucial [3, 51]. For this purpose, elements can be categorized based on characteristics (e.g., demographic information) to summarize patterns or compare differences [1, 62]. It is also possible to categorize by the temporal component of the dataset, which allows for comparing study sessions [31, 35], typical dates [53, 69], or several artificial, aggregated time duration [41, 79].

Another important aspect is the **Object-Space Relation**, the interplay between objects and the environment or the context [64]. Such a relation and, therefore, the usage of space can be presented by a bird-view floor map [2, 57], surveillance video review [10, 31, 70], or a 3D model of the environment [40, 41]. Moreover, as objects can interact with the environment, the analysis of **Object-Space Interaction** with a focus on specific locations in the environment can also be of interest. That interaction can present touch or manipulation [10, 11], attention [2, 57, 70], engagement [10, 40, 42], or other general interaction [1, 31, 41]. The interplay of objects and locations can also be analyzed in groups, named **Aggregated Object-Space Relation**. This can be visualized based on changes in the temporal domain [8], like time duration or frequency (such as in [11, 41, 42]). Moreover, a 2D transition trajectory [41, 63], trail visualization [11, 40], or a synthetic visualization of movement transitions [35, 42] can be used to visualize the flow of objects in the space and how the movement transition between specific locations takes place. This literature research highlights current research aspects and interests with respect to human movement analysis, further informing the design of a visual analytic system. For more details about the resulting literature categories, please refer to supplementary material (Sec. A and Tab. A1).

2.3 Immersive Analytics and Augmented Reality

While previous research has shown how visualizations can support different steps of the movement analysis process [51], most were based on traditional 2D desktop setups. Recently, the research field of Immersive Analytics (IA), where immersive technologies are used to analyze data, has attracted attention and shown its value [36]. IA highlights making *"use of engaging, embodied analysis tools to support data understanding and decision making"* [50]. This analysis process can be enhanced by applying embedded and situated visualizations, where the data representation is tightly integrated or displayed close to the physical referent [21, 22, 37], like seen in the Internet of Things scenarios [25].

In general, visualizations of data in immersive environments have been extensively researched [12, 15, 33], including perceptual problems [18, 23, 46], like the influence of the real-world background on virtual content [59]. While previous work has largely investigated visualization [16, 24, 80], interaction [42, 77], and filtering [30] movement data in VR, research on using AR visualizations for movement data analysis is limited. Recently, MIRIA [11], an in-situ AR analysis toolkit that provides solutions to understand user interaction by visualizing spatial, interaction, and event data in the original environment, has been proposed. It uses 3D trajectories to display movement data and 2D visualizations, such as scatterplots or heatmaps, to show event and interaction data. While the toolkit allows placing additional visualizations on the floor or the walls in the environment, it does not fully integrate the environment for the exploration of the movement data. On the other hand, Reipschläger et al. [57] presented AvatAR, a mixed-reality environment allowing exploring human movement behavior and interaction with the physical environment through movement data visualizations in the real world. AvatAR enables an analyst to understand how a person represented as a full-body 3D virtual avatar touched or gazed at the environment by embedding visualizations on them. While AvatAR allows finding answers to questions focusing on the human's perspective (for instance, "Which part of the display did a particular person touch?"), finding answers to the questions focusing on the perspective of the environment (for example, "Which part of the display was touched the most?") is challenging.

Furthermore, to analyze data situated in the environment, not only the presentation but also the interaction with visualizations is crucial. Particularly, immersive technologies enable a closer coupling between physical objects and their information. This provides great potential to interact with such information intuitively through physical objects or environments, such as direct manipulation [14, 78], gestures [43], or one's own body [39].

2.4 Summary

Related work has shown that AR-based in-situ visualization can support movement data analysis. Considering movement data components, Space, Time, and Objects (see Fig. 2), and an arbitrary number of thematic properties, an analyst can take various perspectives for exploring the data. While existing systems enable the analysis of the movement data, they mainly focus on mover-oriented tasks on trajectories [4]. However, a space-focused analysis, such as Single Space, Space-Space Relation, and Aggregated Space, is also important, as the environmental context is essential for understanding movement behavior and path choices [51].

Subsequently, we see a research gap in approaches allowing in-situ movement data analysis based on the physical objects within the original environment. Such an analysis could benefit from an AR-based immersive system since it enables displaying and freely exploring information in the original environment. We thus aim to provide a solution by utilizing situated physical objects with the considerations of both geometry and temporal aspects of movement data. Besides, our goal is to support a coherent visual analytic workflow including overviewing, filtering, and querying.

3 PEARL: OVERVIEW

We present PEARL - **P**hysical **E**nvironment based **A**ugmented **R**eality **L**enses - an approach enabling analyzing and understanding human movement data in the original space by utilizing the surrounding physical environment. In the following, we describe the terminology used throughout our paper (Sec. 3.1) and the functionalities and requirements for *Lenses* (Sec. 3.2).

3.1 Term Definitions

Objects within the same environment might be in various degrees of relation considering Space and Time (see Fig. 2A). While one Object is moving, it is likely influenced by other Objects within the same environment. In PEARL, we focus on two specific types of Objects (see Fig. 2B), their relationship, and further linked concepts, which are defined as follows and illustrated in Fig. 3:

Actor An *Actor* (see Fig. 2B) is any object capable of self-motion (Mover in [4]). In our work, *Actors* only refer to humans moving through a given environment.

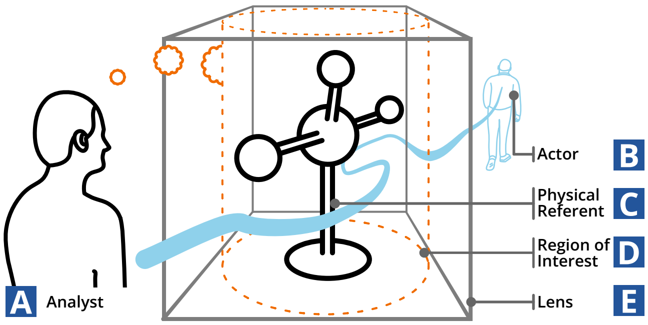

Figure 3: Our definitions of four terms and their relations to an analyst (A). (B) The human-generated movement data of an *Actor* within a given environment. (C) A static physical *Referent* with a set of thematic attributes. (D) The *ROI*, a conceptual space based on the real world that the analyst is interested in. (E) A *Lens*, an approximation to the *ROI*, which is created and defined by the analyst.

Referent A *Referent* is a physical object (see Fig. 2B) that is spatially static (similar to Locations [3]). The *Referents* (assumingly always present) have a single position and an extent in space but no temporal position. However, as they are still objects, they also possess a set of thematic attributes.

Region of Interest (*ROI*) A *ROI* is an arbitrarily imagined space in the real-world environment in which the spatial position and extent are related and based on single or multiple *Referents*. It represents the research interest of an analyst and so can be adjusted with the changing focus of the analysis.

Lens A *Lens* is a simple representation of 3D space and an approximation for a specific *ROI*. *Lenses* are designed for intuitive visualizations, interactions, and spatial queries based on proximity, spanned hulls, or 3D volumes.

The core of PEARL revolves around the use of static physical objects, the *Referents*. Such a focus allows an analyst to directly interact with the physical objects as references for the movement data and to understand the interaction between the *Referents* and the movement of the *Actors*. However, while cases are relevant where human movement is affected by dynamic objects and other people in the environment, those are beyond the scope of this paper.

3.2 Design Considerations for *Lenses*

We present our design choices for the characteristics, features, and functionalities of *Lenses* within PEARL.

***Lens* as *ROI* Representation.** For each *ROI*, a specific *Lens* can be initialized. As the *ROI* is related to a *Referent*, the initial size and position of a *Lens* can be based on the geometry of a *Referent* (via, e.g., Computer Vision techniques). A refinement of this *Lens* through the analyst should be possible, like adjusting spatial extent and position, form, or hull. In case analysts are interested in various parts or the surrounding area of *Referents* (i.e., several *ROIs* on one *Referent*, or *ROI* which does not fully encapsulate a *Referent*), several *Lenses* should be attachable to one *Referent*.

***Lens* as Spatial Filter.** As a *Lens* defines a spatial volume, a spatial query based on the volume of the *Lens* can filter spatio-temporal data. Furthermore, as various *Lenses* can exist simultaneously, their results can be combined to enhance the data exploration further. The resulting dataset can consist of either *Actors*' movement (on an actor level) that passed through the *Lens* or segments of *Actors*' movement (on a data-point level) within the *Lens*.

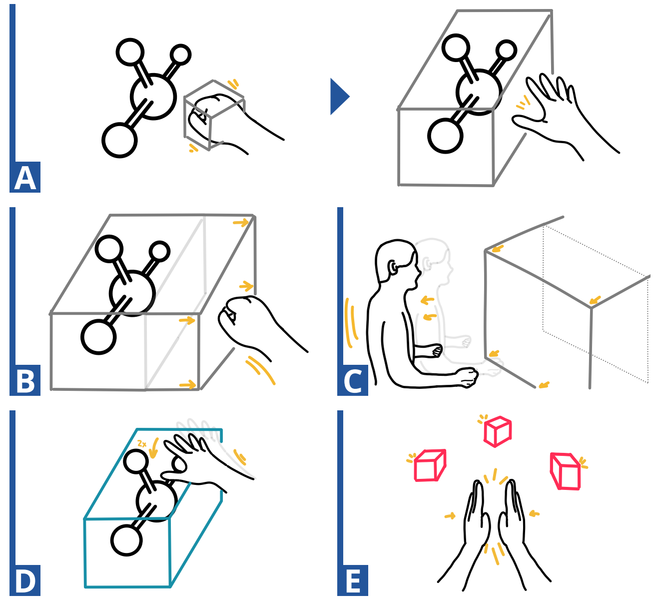

Figure 4: Interactions for creating, adjusting, selecting, and grouping *Lenses*. (A) Holding a hand as a fist and then releasing it creates a *Lens* around a *Referent*. (B) The dimension of a *Lens* can be customized by a drag gesture or (C) by the distance between analysts and *Referents* using a body gesture holding both hands as fists. (D) A *Lens* can be selected by a double tap gesture, highlighting it accordingly. (E) Multiple *Lenses* with filters can be grouped via a clap gesture.

***Lens* for Visualization.** The dataset resulting from the *Lenses* operation can be used as the input for visualizations. Those visualizations include trajectories or avatars focusing on single *Actors*, but in particular, aggregated views combining several *Actors* and their data points. Additionally, involved *Referents* can be used as spatial anchors for placing and structuring additional movement data visualizations to enhance the data exploration within the environmental context. Specifically, visual representations of aggregated information can be integrated into the environment, like embedded visualizations on the floor or the *Referent* itself, situated directly around and close to the *Referent*, or placed on the hull of the *Lens*.

4 PEARL: CONCEPTS

Based on our design considerations and an iterative design process, we present PEARL concepts, which are grouped into three categories accompanied by the information of general functionalities. In particular, we introduce general user interface components and functionalities (Sec. 4.1), describe basic interactions for authoring and managing *Lenses* (Sec. 4.2), introduce filter functions for *Lenses* (Sec. 4.3), and propose visualization techniques related to *Referents* and their *Lenses* (Sec. 4.4).

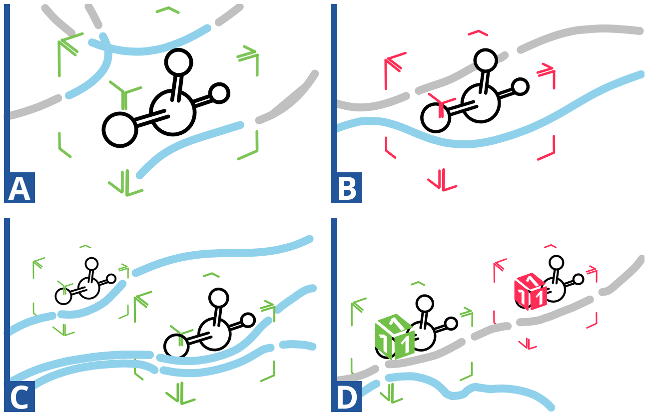

Figure 5: Filters affect trajectory data visualizations on either a data-point (A) or actor level (B–D). (A) A *Positive* filter on the data-point level results in segments of the *Actors*' movement path. (B) A *Negative* filter excludes passing trajectories. (C) illustrate *Positive* filters in an *OR* relation. (D) shows a *Positive* and *Negative* filter in an *AND* relation, indicated by the group number in the left upper corner. Notably, all gray lines represent invisible trajectories.

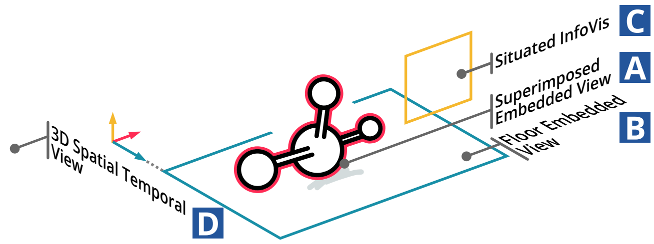

Figure 6: Placement regions of visualizations around the physical object: (A) magnitude information as aggregated visualizations superimposed on the object, (B) the spatial and Object-Object Relation data as aggregated visualizations embedded on the floor, (C) the 2D information visualization next to the object for complex analysis, and (D) detailed motion data as 3D spatial-temporal visualizations in space.

4.1 General Functionality

In addition to the *Lens* as our main interactive element, we propose two menus available to the users, a *Lens*- and a body-anchored menu. The former is placed on the hull of a *Lens*, while the latter is accessible by lifting a hand. The body-anchored menu also allows for global filters on thematic attributes of spatio-temporal data.

4.2 Pearl Selectors: Managing *ROI*

As analysts are often interested in particular *Referents* and, with that, in *ROIs*, defining and managing these regions for data explorations are essential. Pearl enables this functionality through *Lenses* overlapping the *ROIs*. In Pearl, *Lenses* are represented by a cuboid bounding box either as a full wire frame (see Fig. 4A) or only with corner indicators (see Fig. 5). The system can change between both visual states based on the proximity of the analyst and a possible intent to interact with the *Lens* or its *Referent*. To create a *Lens* based on a *Referent* (see Fig. 4A), an analyst has to hold a fist close to a *Referent*, which is highlighted by a preview box attached to the same hand. The transition from the fist to a flat open hand will create a new *Lens* encapsulating the *Referent*. For removing the *Lens*, the same gesture can be inverted.

Common 3D manipulation techniques (see Fig. 4B) can be used to refine the *Lens* to match the *ROI* further. Additionally, an analyst can hold both hands as fists in front of the face of the bounding box to adjust the *Lens* based on the distance traveled by the user (see Fig. 4C). Besides, performing a double tap on the bounding box selects or deselects a *Lens* (see Fig. 4D). When multiple *Lenses* are selected, synchronized operations, like a collective filter configuration, can be executed via the body-anchored menu. Furthermore, to ease managing a growing number of spatially scattered *Lenses* and their properties, other overview UI elements can be used, such as 2D interactive list views or a 2D minimap (similar to [57]).

4.3 Pearl Filter: Refining Data Selections

Lenses can further be used as spatial queries to filter and refine the movement data, in particular, 3D trajectories, as the increasing amount of recorded *Actors* and recorded time leads to the "spaghetti heaps" issue [17] and overwhelms analysts consequently. *Lenses* can be configured for filtering either on an actor or data-point level. In the example of trajectories, this can lead to less cluttered complete trajectories (see Fig. 5B-D) or spatially-confined segments of trajectories (see Fig. 5A), respectively.

A *Lens* as a filter can have either *Positive* or *Negative* polarity. A *Positive* filter includes movement data of all *Actors* or their segments within the selected proximity to a *Referent* defined by the *Lens* (see Fig. 5A). Conversely, a *Negative* filter excludes the movement data of *Actors* passing through the corresponding *Lens* (see Fig. 5B). As analysts might be interested in multiple *ROIs* in the environment, Pearl enables combining filters of several *Lenses* via basic boolean logic operators, *OR* and *AND*, for creating global filter queries on the actor level. If there is more than one filter in the scene, they are combined with *OR* relations by default (see Fig. 5C) for questions like "Who visited at least one of these two *Referents*?". Further, filters on *Lenses* can also be grouped by an *AND* relation (see Fig. 5D) to answer questions like "Who visited both of these two *Referents*?".

The *Lens*-anchored menu or the overview UI element (map or list views) can be used to apply a filter to a *Lens*. Moreover, multiple filters can be grouped either by using the body-anchored UI or by performing a grouping gesture (see Fig. 4E) while the corresponding lenses are selected. Pearl uses the corner indicators surrounding the bounding box to highlight the polarity and filter group state of the *Lens* (see Fig. 5). Therefore, gray, green, and red are used to highlight "no filter", "a positive filter", and "a negative filter", while a filled-out and numbered corner (see Fig. 5D) indicates the affiliation with a filter group.

4.4 Pearl Visualizer: Visualizing Data

Lenses can also be used as data sources to create visualizations. However, over-detailed visualizations can overwhelm analysts due to a growing number of moving objects or *Actors*. Thus, Pearl supports aggregated visualizations structured by the created *Lenses*.

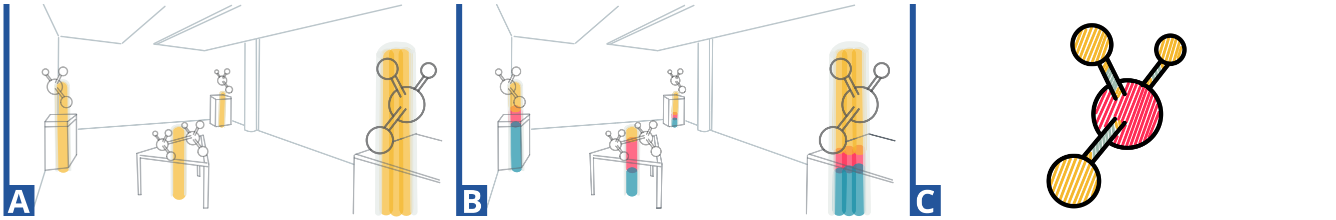

Figure 7: Superimposed embedded visualizations on *Referents* for magnitude information. (A) Embedded 3D bars overlaying *Referents* can visualize magnitude information in situ. (B) Similarly, such visualizations can support comparing *Actor* clusters using 3D stacked bars. (C) The aggregation of touch or gaze data can be shown as embedded heatmaps directly on the *Referents*.

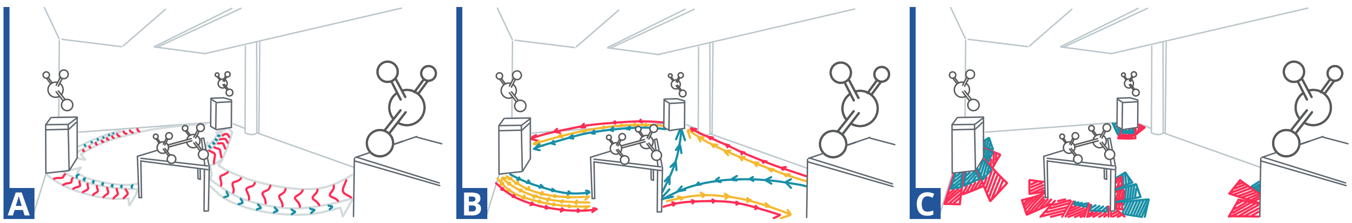

Figure 8: Floor-embedded visualizations for spatial and relational data. (A) Flow View with arrowed flow visualizations shows the direction and quantity of *Actors*' group movement. (B) The navigation strategies of *Actors* can be informed by Sequence View. (C) Approach View with radius bars surrounding the *Referents* shows the number of *Actors* approaching from the respective directions.

Based on spatial relations to the *Lenses* (see Fig. 6), these visualizations can be embedded on a *Referent* (Sec. 4.4.1), embedded in an environment (floor) (Sec. 4.4.2), situated to a *Referent* or placed on a surface of a *Lens* (Sec. 4.4.3), situated within an environment (Sec. 4.4.4), or a combination of above (Sec. 4.4.5).

4.4.1 Referent-Embedded Visualizations. Embedding information directly on top of a *Referent* and with that within a *Lens* is the most direct way to relate a visualization to its physical *Referent*. *Referent*-embedded visualizations are chosen to enhance visual attention [76] and to unify visual representations and data referents. In general, such a direct embedding of visual information can be used to show arbitrary visualizations overlapping the *Referent* or to relate and highlight specific areas of the *Referent*.

Superimposed Visualization. Superimposed visualizations place visual content at the same spatial position as the *Referent* within a *Lens*, thus overlaying virtual information on top of the real world. While such an overlay could cause obfuscation due to the visual competition with *Referents*' presence and natural features, it allows analysts to navigate the environment easily. Accordingly, it becomes important to use suitable simple visual representations for presenting aggregated attributes of the available data. Examples are the presentation of magnitude information (like the visit frequency of a museum) as a 3D bar (see Fig. 7A) or grouped information based on a specific parameter (like the time frames of interest) as a 3D stacked bar (see Fig. 7B). As those visual artifacts are placed throughout the environment, it becomes essential to provide a common frame of reference to enhance the comparison process, such as an absolute vertical scale shared between all representations.

Emphasizing Referent Areas. Fine-grained engagement information, such as touch and gaze attention, can be visualized directly overlapping the *Referents*' geometry. Moreover, the color saturation and hue can be encoded with the data density (see Fig. 7C), as commonly done with heatmap representations.

4.4.2 Environment-Embedded Visualizations. The encoding of space and relationship information of *Referents* can be embedded into *Referents*' surrounding environment. In PEARL, we concentrate on the floor as a surface to show content, as it can help to free up analysts' field of vision [60] and allows for displaying 2D visualizations as seen on 2D map-based movement data analysis tools (such as in [71, 72]). Embedding visualization into the floor enables the design of visual representations for different movement data features, such as the transition, flow, and sequence of *Actors* between *Referents*, the approach direction and the speed of such transitions, and the general stay duration of *Actors*.

Flow View for Aggregated Transitions. As the *Referents* define and guide the movement of *Actors* in the environment, it is crucial to understand the *Referents*' influence on the location transition of *Actors*. Especially the flow of *Actors* within the space can be aggregated and articulated via *Flow View* in the form of, for instance, arrow-directed visualizations as linkages (see Fig. 8A). One such link combines two *Referents* and encodes the number of transitions on the width and the direction with arrows. With this visualization, an analyst can detect not only potential patterns, like a "central hub" (i.e., a physical object connected by many wide incoming or outgoing linkages) but also which properties of the *Referent* resulted in this. Moreover, multiple *Actor* groups can be color-coded to allow for a comparison of navigation strategies.

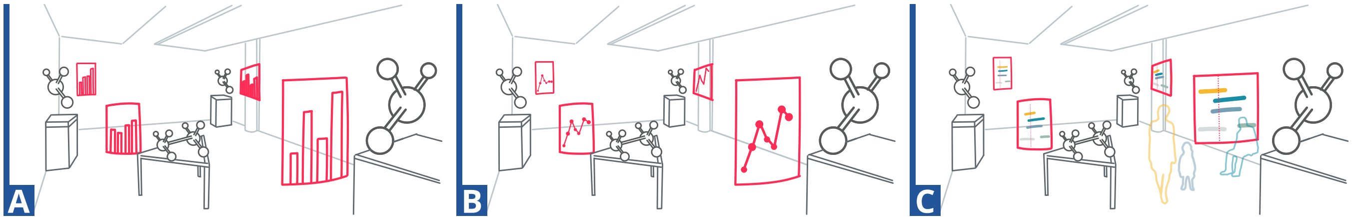

Figure 9: Situated visualizations with *Referents* side-by-side for complex analysis. (A) Situated bar charts next to *ROIs* show the statistics of corresponding *Referents*. (B) Likewise, situated dotted line charts can show the visiting frequency over time in detail. (C) 2D Gantt charts display co-located *Actors*' timelines with 3D animated avatars for concurrent moments.

Sequence View for Individual Transitions. To understand the movement behavior of a single *Actor*, more fine-grained visualization like a trace depicting the sequence of the *Actors*' engagement on *Referents*, is needed. Hence, we use a fishbone-like *Sequence View* (see Fig. 8B) with a similar visual encoding as the *Flow View*, representing one specific transition of a single *Actor*. To better reveal any reoccurring transitions between several *Referents*, single fish-bones between the same *Actors* are bundled together side-by-side. With *Sequence View*, navigation strategies can be analyzed either collectively via comprehending the pattern by the overview of linkages, or individually via following the visiting sequence of an *Actor*.

Pace View for Aggregated Transitions. Aside from the quantity and direction of location transitions, the *Actors*' movement speed between two *Referents* can also be of interest. Thus, we design *Pace View* to show the average speed of *Actors*' movement. It consists of a fixed-width linkage on the floor connecting two involved *Referents*. The traversing speed is encoded by color saturation and hue, resulting in heatmap-like, ribbon-shape visualizations. Thus, analysts are able to have a general overview of the pace of movement among *Referents*.

Approach View for Aggregated Moving Directions. To understand how *Actors* approach a specific *Referent*, we designed an *Approach View* (see Fig. 8C) displayed around the *Lens*. Each *Approach View* presents a number of bars ordered radially around the *Lens*. The height of each bar indicates the number of *Actors* who approached the *Referent* from the corresponding direction. Additionally, such bars can be color-coded and stacked to allow for comparisons of approaching strategies between *Actor* groups.

Heatmap for Aggregated Stay Duration. Moreover, 2D heatmap visualizations can be embedded on the floor to show an overview of aggregated time spent by *Actors* in different locations, thereby revealing relationships between location and *Actors*' engagement. We thus use color hue and saturation to visualize the duration of stay across the tracked space.

4.4.3 Referent-Situated Information Visualization.

Visualizations can be embedded and placed close to a given *Referent* (similar to [25]). In PEARL, situated visualization can be placed directly mapped to the hull of the *Lens* or only anchored to it while always orienting towards the analyst. As situated visualizations likely compete less with the visual features of the *Referent*, they can present more complex information. Those include demographic attributes like education status, age, or occupation, but also *Referent*-related information, like stay duration within a *Lens* or visiting frequency for multiple *Actor* groups. Such information can be presented as, e.g., bar charts (see Fig. 9A) or line charts (see Fig. 9B). Furthermore, auxiliary data that might aid the comprehension, such as video recordings or images of a *Referent*, can be shown within the space in appropriate locations.

4.4.4 3D Spatio-Temporal Visualizations.

PEARL also supports the visualization of spatial movements of *Actors* over time in the environment, such as 3D trajectories of head and hand movement using connected three-dimensional tube segments. These tube segments can be color-coded according to the given *Actor*'s thematic property. Besides, to reduce visual clutter, non-salient 3D trails, displaying only the recent moments of the movement, can be shown. Lastly, 3D animated avatars, including head and two hands models, can be displayed together with 3D trajectories or trails.

4.4.5 Combinative Visualizations.

To analyze movement data in various granularities, PEARL supports displaying a combination of aggregated and detailed visualizations. This also alleviates the challenges of individual visualization techniques. In the case of trajectories, it is hard to understand proximal interactions between *Actors* [11], as the spatial proximity of two trajectories does not indicate spatio-temporal proximity - being together at the same moment in time. Thus, we design an additional situated 2D Gantt chart (situated to a *Referent*) as a local timeline overview, helping identify when such concurrent moments occur (see Fig. 9C). Relevant timelines of *Actors* are presented horizontally as color-coded 2D bars, highlighting the entry, exit time, and duration of stay of *Actors* in relation to *Lens*. An interactive vertical line indicates the current display time point and allows for navigating the time conveniently. Thus, concurrent visiting time frames among *Actors* can be pinpointed via the overlapping areas of bars (namely, parts of *Actors*' timelines share the same horizontal axis). Using this visualization, analysts can dive into the details of the situational context by combining it with, for instance, a 3D trail view to observe the detailed motion of *Actors* and social interactions among *Actors*.

5 INTERACTIVE PROTOTYPE

In this section, we introduce the prototype implementation of PEARL. We describe the technical setup (Sec. 5.1), the overview of our system (Sec. 5.2), and details for the computing and rendering of aggregated and embedded visualizations (Sec. 5.3). Additionally, we describe the recording of room-coupled human movement data (Sec. 5.4) used in the evaluation presented later.

5.1 Technical Setup

We implemented all three proposed PEARL components (Selector, Filter and Visualizer) and their corresponding functionalities. However, as we focused on the core functionalities, the current object detection is based on QR codes, the situated charts (in Fig. 9) and overview UIs (like 2D minimap in Sec. 4.2) are not interactive, and the referent-embedded visualizations presenting fine-grained engagement (like showing touch or gaze density in Fig. 7C) are absent.

We implemented the prototype for Microsoft HoloLens 2 HMD by using Unity 3D, MRTK[1], and available open-source toolkits, including u2Vis [58] and MIRIA [11]. Rendering is performed remotely on a PC with a GeForce RTX 3070 GPU and streamed wirelessly to the HMD. Notably, the goal of the prototype is to prove the feasibility of our concept. Thus, our source code is available on the project page[2] to support further research and development toward a comprehensive human motion analysis solution.

5.2 System Overview

In our prototype, we used a combination of QR code recognition with pre-measured dimensions of the *Referents* to create pre-defined *Lenses*. Other possibilities, like using computer vision techniques, were also considered but rated as beyond the scope of this paper.

In general, logic filters can become arbitrarily complex. However, as we primarily aim to demonstrate the feasibility of our concepts, we opted only to implement a lightweight filtering logic. Thus, we realized one global filter query, which combines *Positive* filters via a single *OR* or *AND* relation, while all *Negative* filters are absolute and connected with the *Positive* filters via an *AND*.

We used MRTK to implement a *Lens*-anchored menu positioned at one face of the *Lens* and a body-anchored menu next to the user's hand for interacting with the system on demand. Additionally, to implement the proposed free-hand gestures for the interaction with *Lenses*, we used MRTK hand models and their joints.

5.3 In-Situ Visualization Computing and Rendering

To generate the visualizations (see Fig. 10), our prototype uses a graph-based data structure, considering the *Lenses* as nodes and the spatio-temporal relations as edges. Based on this, it creates textures for rendering floor-embedded visualizations on the ground. For instance, to create a *Flow View* (see Fig. 1C), the prototype first identifies the positions of all *Lenses* on a texture. Then, it renders the edges as flows from the origin to the destination *Lens* with the width defined by the number of *Actors*' movement paths, the color defined by clusters of *Actors* (such as age groups or gender), and ensures no overlapping of edges by using varied control point locations.

For proposed superimposed visualizations (like Fig. 7A+B) as well as situated visualizations (like Fig. 9A+B), we utilized u2vis [58] that supports the display of common information visualizations for immersive analytics, such as 2D and 3D bar charts, scatter plots, line charts, and pie charts. Moreover, spatio-temporal visualizations, including 3D trajectory, 3D trail, and 3D avatar, are based on the MIRIA toolkit [11], where connected, three-dimensional tube segments represent trajectories. These trajectories are color-coded to match the *Actors*' colors defined by the configuration file. Additionally, our prototype indicates the direction of the trajectories by 3D cones at the sample points. For more details about the implementation, please refer to the supplementary material (Sec. B) and the provided source code.

[1] https://www.github.com/microsoft/MixedRealityToolkit-Unity
[2] https://www.imld.de/PEARL

5.4 Movement Data Recording

To test our prototype and use it in the following expert review, we collected movement data within a simulated exhibition room. We used Oculus Quest 2 VR headsets to record data in any environment without needing elaborate tracking systems. The onboard inside-out tracking system of the Quest 2 provides sufficiently accurate head and hand pose over large areas. At the same time, the video pass-through mixed reality feature allows the tracked participants to view the real-world surroundings and move naturally without fear of collision. A custom application was developed for the Quest 2 devices to manage the recording while the captured data was uploaded via the network. A simulated exhibition room was prepared ($8.6 \times 5.9m = 50.74m^2$) to replicate a typical museum or lab demonstration event for the motion data recordings (see Fig. 1). The exhibition room had six objects of various shapes and heights. These objects can be grouped as (1) Interactive Exhibits (an eye-like music set, a gaming station, and a foot pedal set) and (2) Visual Exhibits (an introduction board, a spine sculpture, and a molecular sculpture) similar to an actual technical museum.

We invited nine university students (4 female, 5 male) to the simulated exhibition room for data recording, which resulted in a total of 27 min spatial movement data. To guide the process and simulate everyday situations of an actual exhibition, we assigned two personas to each participant in a counterbalanced order. Specifically, participants acted either as an in-a-hurry persona visiting only a few exhibits during peak hours or as a casual persona being curious about everything. However, they were also asked to move and act freely. Meanwhile, experimenters briefly introduced the exhibits to simulate the visiting experience. Our supplemental material contains the recorded spatio-temporal data.

6 PRELIMINARY EVALUATION

To evaluate our concepts and the prototype, we conducted a preliminary expert review (as in [47, 68, 73]) consisting of a guided hands-on session and scenario-based cognitive walkthroughs (similar to [57, 58]). We aimed to gain first experiences and discuss the overall PEARL approach with invited experts. Thus, in this section, we first describe two scenarios and then detail the expert review and our findings.

6.1 Walkthroughs and Scenarios

We designed two scenarios considering a simulated exhibition room for review (see Fig. 1). In general, the first scenario (Guided Interactive Walkthrough, Sec. 6.1.1) follows more closely the functionalities of our current implementation (see Sec. 5) and is based on the aforementioned collected movement data. In contrast, the second one (Envisioning Scenario, Sec. 6.1.2) describes a more sophisticated scenario based on our visions of PEARL being utilized in the future.

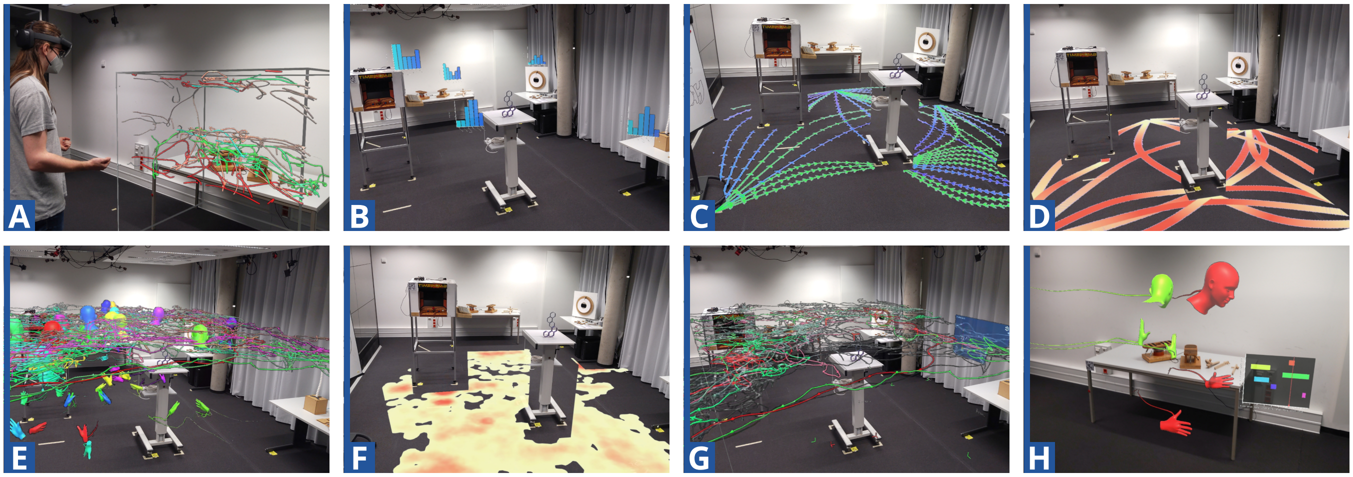

Figure 10: Our prototype was demonstrated in a simulated exhibition room: (A) A body gesture to modify the scope of *Lenses* showing segmented trajectories. (B) Situated 2D bar charts placed next to *Lenses* for comparing visit frequency. (C) Sequence View summarizing *Actors*' location transitions among *Lenses*. (D) Pace View aggregating the movement speed between *Lenses*. (E) 3D trajectories with associated 3D animated avatars. (F) Floor Heatmap showing *Actors*' stay duration. (G) Several actor-level filters with logic operators for 3D trajectories. (H) Co-located View combining a 2D Gantt chart, 3D animated avatars, and trails.

6.1.1 Scenario of the Guided Interactive Walkthrough. A curator, Lizzy, wants to understand exhibition visitors' movements to optimize their navigation and path guidance. After **creating *Lenses*** (see Fig. 10A) based on exhibits, Lizzy first starts comparing the visit frequency based on periods of the day (morning, afternoon, and evening) via the **situated bar charts** (see Fig. 10B) next to the objects. Following this, she reduces the time range to concentrate on the relatively hectic rush hours.

She uses the **Sequence View** (see Fig. 10C) to get a general overview of how people travel through the exhibition. Lizzy notices a pattern that people moved back and forth between an introduction board and the gaming station, as revealed by multiple in- and outwards facing links of the same color. This is further confirmed by the combination of **Approach View** and **3D trajectories** (see Fig. 1A), given that the approach direction is mostly from the information board, whilst the same color trajectories appear very often. Inspecting the introduction board, she sees that it also contains detailed descriptions of the gaming station, which could have been placed closer for better reference.

After noting this, she continues exploring the dataset via the **Pace View** (see Fig. 10D). She notices that people moved slower on the way to the exhibit on the corner. However, the **Flow View** (see Fig. 1C) does not show an increasing amount of visitors there, indicating that it is not due to long queuing. She then filters based on this exhibit to only focus on segments of movements. Combining with the **3D Avatar View** (see Fig. 10E), she finally finds that people were uncertain whether they were allowed to interact with this exhibit due to the lack of signage, resulting in a longer waiting time.

6.1.2 Envisioning Scenario. Sam wants to analyze visitors' movement in the now-ending exhibition to optimize the layout of the upcoming one. He first **creates several lenses** around exhibits to reduce the size of the data and then performs **body gestures** (see Fig. 10A) to alter the dimensions of auto-generated lenses to fit his *ROIs* better. As Sam wants to focus first on interactive exhibits (such as the gaming station) instead of visual ones, he applies ***Positive* and *Negative* filters** on both categories, respectively.

Sam would like to understand how long visitors of different ages engaged with exhibits. For that, he uses **superimposed 3D stacked bars** (see Fig. 1B) that plot stay durations against different age groups (children, teenagers, and adults) for selected exhibits. Surprisingly, he finds that adults visited the gaming station longer than children. At the same time, the **2D heatmap** (see Fig. 10F) shows that this exhibit had the most bystanders, as many visitors stayed relatively far away. To further understand why this system was not so popular with children, Sam activates the **3D avatar view** (see Fig. 10E). He can now identify that only a few children could reach the gaming machine and control it on their own. Combined with the **3D trajectories** (see Fig. 10G) of the children, it becomes apparent that many children took several steps back to view the exhibit. Moreover, looking closer at adults' movements as **3D avatars with trails** (see Fig. 10H) nearby the gaming station, it seems that adults often pointed to it while engaging with the exhibit, seemingly explaining something. Thus, it becomes clear that the accessibility of this exhibit is limited for children.

6.2 Participants

We invited four experts (2 female, 2 male; $M = 44\ years, SD = 10.23\ years$), independent from our team. The experts (E1 - E4) had diverse backgrounds: At the time of the study, E1 had 7 years of experience as a museum curator and 10 years of experience teaching museology at a university; E2 was a university professor of computer science with 13 years of experience designing interactive and visualization computer systems; E3 had 10 years of experience analyzing human movement data, especially via trajectories; and E4 had 16 years of experience organizing and curating exhibitions in combination with 10 years of experience with interactive system design. In addition, E1 and E3 were only slightly familiar with immersive technology (AR or VR), while E2 and E4 were very familiar.

6.3 Procedure and Tasks

Each of our study sessions, based on the principles of semi-structured qualitative studies [7], consisted of the following parts: (1) a brief introduction accompanied by a short video demonstration and presentation of slides about Pearl, (2) a guided interactive walkthrough based on system features (see Sec. 6.1.1), (3) an open-ended discussion of two envisioning scenarios based on our prototype and concepts (see Sec. 6.1.2), (4) and a post-study interview. Participants were also asked to think aloud during the whole session. During each session, the participants were audio-recorded, and an experimenter occasionally took notes. Each session lasted around 90 min. We provide our study materials, including the guided walkthrough and the envisioning scenarios, and the thematically grouped interview answers in the supplementary material and on our project page[3].

6.4 Findings

We summarized all feedback from the experts, including observations, think-aloud comments, and post-study interviews. That information was then used as a basis for thematic analysis [9], followed by a cross-validation of two authors. In this section, we present our findings as the following themes: a general impression of our approach (Sec. 6.4.1), feedback for Pearl functions (Sec. 6.4.2), discussions of perceptional issues (Sec. 6.4.3), and reflections on analysis workflow and perspective (Sec. 6.4.4).

6.4.1 General Impression. The experts agreed with the general mindset of Pearl, as *"movement is also always correlated with [the] environment"* (E3). Further, they were excited about the potential of Pearl, especially domain experts (E1, E3-4) expressed a strong willingness to integrate it into their existing workflow, as it can be *"a great chance for museums"* (E1). The experts rated the use of AR for data analysis as beneficial since it feels more immersive (E2, E4) and intuitive (E3), and enables data sensemaking and reasoning as *"you have your own body to represent, compare, and relate"* (E4). Moreover, experts suggested potential use cases with Pearl, including the evaluation of the sufficiency of the exhibition setup (E2) and the visitors' engagement format and duration (E1), as well as the design of the physical environment, such as the navigation guidance (E4) and the layout iteration (E1, E4). However, while appreciating the potential, two experts (E1, E4) were also concerned about the data collection and its potential influence on the visitors' behaviors.

6.4.2 Pearl Feature Feedback. Focusing on physical referents with the *Lenses* as a representation of *ROIs* was generally liked (E1-2, E4). For instance, E4 mentioned that it was helpful to find *"if people have enough space to watch exhibits"* (E4). Experts also appreciated the potential flexibility offered by *"different [analysis] resolutions"* (E3), such as having several *Lenses* for multiple parts of the same exhibit. Furthermore, most experts (E2-4) could easily and intuitively interact with *Lenses* using gestures. While having difficulty using gestures at the beginning of the hands-on session, E1 could quickly learn and successfully perform them.

Both the actor-level (E1, E4) and data-point level filtering (E2-3) based on the *Lenses* were considered useful. For example, E1 found it *"interesting to see the difference between visitors who visited vs. did not visit the Introduction Board"*. Moreover, the reduction of data complexity was deemed beneficial: *"otherwise, how should I analyze the data in such a huge space"* (E3). For the sake of study duration, while we only presented two filters (one positive data-point level and one negative actor level) in our study, the experts highlighted the benefits of logic operations: *"If there are two objects that need to be visited together, [I want to see] if my idea really works"* (E4).

Experts were excited about the presented visualizations and their relation to *Lenses*: *"It is good that visualizations are updating based on the lenses"* (E4). The novelty and potential usefulness of floor-embedded aggregated visualizations, like Sequence View (E1-4), Flow View (E1, E3-4), Approach View (E1-4), and Pace View (E2-3), were highly appreciated. Experts explained that *"floor encoding is particularly helpful to reduce the stress of raw data"* (E2) and *"they are so useful to find the [patterns]"* (E3). In addition, experts also mentioned the wish to combine single visualization techniques, like Superimposed View and Sequence View, as *"it could be interesting to [...] know how many [visitors] and the ways they came"* (E4). Experts also indicated the need for transitioning between different visualizations, as those *"can supplement each other"* (E3). For that, other novel gestures, like *"pushing traces (trajectories) to the floor to get clearer [aggreagted] visualizations"* (E2), were proposed.

6.4.3 Perceptual Issues. During the review, experts pointed out some perceptual issues. For instance, E2 commented that attention switches between floor-embedded visualizations and referents *"have a trade-off because you focus on the floor and lose the focus [of exhibits]"*. Also, the need for having an overview of the visualizations was mentioned (E3-4), considering viewing perspective (inside-out vs. outside-in) and the HoloLens' field of view (FoV) (E4). For superimposed visualizations, experts also mentioned difficulties in precisely comparing different 3D bars due to the viewing perspective, *"it is hard to tell or be sure [about the volume]"* (E3).

Furthermore, experts suggested improvements in the visual design of *Lenses*. For example, the limited FoV of the HoloLens leads to visibility issues as *"lenses should be smaller"* (E1). On the other hand, experts also wanted additional visual cues to avoid missing information due to the improper configurations, like *"some previews [...] to show this [lens] has been actually visited, or some hints to suggest increasing the size of a lens"* (E2).

6.4.4 Human Movement Analysis Workflow and Perspective. Based on the hands-on review experience, two experts (E3-4) further detailed their envisioning workflow with Pearl: starting with aggregated visualizations for an overview (such as the floor-embedded visualization) and then continuing with detailed visualizations (such as 3D Animated Avatars or 3D Trajectories). Moreover, E3 mentioned: *"it is really good that you could have different starting points [...]. Perhaps, I could change my procedure because the system allows*

[3]Project page: https://www.imld.de/PEARL

me to do so". However, it was also highlighted that such a workflow would depend on the research questions and the number of recorded visitors. Notably, E3 (the human movement analysis and trajectory expert) said: *"I [usually] think about people (trajectories) [while exploring movement data]; perhaps, I [will] have a different perspective from now on"*. E1 also highlighted the potential benefit over his existing workflow: *"right now [it] requires the exhibitor to stay in the room, but with [Pearl], you can review [and] compare [the movement data] anytime you want"*.

7 DISCUSSION AND FUTURE WORK

In this section, we discuss our approach and design decisions in the context of our overall vision of immersive analytics and in-situ visualizations. We first reflect on lessons learned (the first four topics) and then describe future steps (the last two topics). In particular, we highlight some specific points mentioned by experts.

Power of Physical Referents for Human Movement Analysis. While the evaluation remains preliminary and more general conclusions would require further investigations, we could see the exciting potential of Pearl for human movement analysis by referring to physical objects that *Actors* engaged with. In particular, we recognize that Pearl could help break down the inherent complex dataset by reducing, summarizing, or aggregating it. Specifically, as experts (E2-3) mentioned, functions like filters are useful for narrowing down the exploratory space for movement data analysis. Moreover, physical referents enable to further summarize movement data to fine-grained aggregations, such as floor-embedded visualizations that were particularly favored by experts (E1-4). This brings new opportunities for introducing innovative visualization techniques beyond generic aggregated visualizations like heatmaps. Lastly, analyzing data in situ, in general, can facilitate data sensemaking and reasoning. As E4 suggested, the presence of physical references helps interpret the original context. Future MR systems could explicitly indicate such context by labeling *ROIs* into categories based on the affordance of physical referents.

Physical Object Detection and ROIs Definition. Pearl can be used in various physical environments with objects that have diverse geometries and shapes. Thus, defining and adjusting *Lenses* for appropriate *ROIs* can become demanding. We currently implemented a simple cuboid bounding box as the first step. In the future, a mixed method of computing automation and user inputs could be promising. Specifically, advanced Computer Vision can facilitate this process [26, 56] and predict the dimension of captured objects on the fly [45, 52]. As a result, suitable dimensions of *ROIs* could be recommended based on individual objects. However, it is currently unclear and up for future research if such high-fidelity *Lenses* are even needed, as no expert in our study had any concern regarding the basic cuboid shape in the current implementation of Pearl. Nevertheless, experts (E1-2, E4) indicated the limitation of the used HMD regarding the visibility of *Lenses*. Besides, *Lenses* can be further contextualized given the semantics of objects. For instance, a vertical whiteboard or an exhibition shelf likely has *ROIs* in the front instead of the backside. Moreover, remote (like raycasting) or proxy interaction [55] possibilities could be considered for larger objects, as direct manipulation might no longer be suitable.

Presence and Mobility of Physical Objects. Embedding AR visualization in the original environment can keep users' actual perception of physical objects and help analysts take the involved humans' perspective, which enables contextualizing *ROIs* and understanding the initial situation. However, the physical environment and their situated objects are subject to change regarding their position, dimensions, or presence, resulting in a mismatch between the recorded situation and the existing environment. To tackle that, it is possible to show a virtual object (like a table) in the analytic scene for replicating the original environment [11] and to manipulate this virtual representation for interacting with spatio-temporal data [42]. According to the expert walkthrough of Büschel et al. [11], this helped users to understand the initial situation despite the absence of the physical object. Thus, it seems a promising approach to integrate virtual replacements of physical objects to maintain situation awareness. At the same time, as our participants mentioned, this would likewise help *"to compare different setups"* (E1) and *"to train practitioners [through] previous exhibitions"* (E3). Lastly, while we limited our current system to static physical objects, it is also imaginable to include and analyze the movement of physical objects besides humans.

Visualization Placement with the Physical Referents. We designed four visualization classes (see Sec. 4.4 and Fig. 6) to coordinate the data referents and associated visualizations. In fact, though studies have shown that users tend to place virtual content next to physical objects [48, 49], it is unclear how to best arrange a layout of virtual content in relation to physical referents in mixed reality. Moreover, researchers have highlighted the importance of layout for Immersive Analytics. For instance, Lin et al. [44] suggested embedded, situated, and boundary placement for virtual labels to out-of-view referents for situated visual searching. Similarly, Satriadi et al. [61] proposed above, side-by-side, around, and overlay placement on augmented globes. Additionally, the placement of virtual content on the ceiling and floor in close proximity to the object is imaginable [60]. However, a far distance between visualizations and physical referents might lead to an undesirable attention switch and extra cognitive burden, as E2 commented when examining floor-embedded visualizations. Thus, future research could consider how the layout of visualizations to the physical data referents affects visual analytic tasks in mixed reality to inform the layout design. This becomes especially important when combining several visualizations, as mentioned by E2.

Analysis within a Large In-Situ Space. Human movement can be analyzed not only in a single room but also in a larger space with several interconnected rooms, like a typical museum. As experts suggested, this can apply to future application scenarios, like *"architecture design, public space design, [and] interior design"* (E4). Hence, the system might be required to support managing the analysis process when *ROIs* are out of reach or out of view. As the experts (E3-4) indicated, an overview functionality of visualizations would support the analysis in these cases. Future work can explore possible solutions, such as using a mobile device offering an additional overview interface, like the overview floor map for several rooms [57]. Similarly, DistanciAR [74] introduced several remote modes on a mobile tablet to manage out-of-view workspaces. Moreover, the Pearl filter can also be utilized to restrict the exploratory scope to a manageable space by, for instance, setting up filters on the passages or doors.

Scalability and Design For Large Datasets. While our prototype has shown the feasibility of the proposed approach, the system design could be further extended for large datasets. Specifically, the computation performance of in-situ visualization can be improved by the space-partitioning data structure, whilst the rendering of the floor-embedded visualizations can be moved from CPU to GPU by a compute shader. From the system design perspective, future applications could also facilitate combining heterogeneous filters, including spatial queries based on location, global filters based on *Actors*, and temporal filters based on time, to narrow down the data scope. Besides, to mitigate the intersections of edges of floor-embedded visualizations, methods like varying and limiting the control points with a Voronoi-based method or force-directed edge bundling techniques [29] can be applied. In addition, advancements in user interaction capabilities, such as highlighting edges of floor-embedded visualizations while hovering via raycasting, can also be considered to improve visibility and wayfinding.

8 CONCLUSION

We presented Pearl, a physical environment based approach for the in-situ analysis of human movement data in mixed-reality environments. In this work, we proposed an immersive analytics approach by utilizing physical referents as regions of interest. Based on this, we introduced various concepts to support the exploration of human motion data via selecting, filtering, and visualizing information. Those concepts were further supported by our prototype implementation and evaluated by expert review sessions. With that, we demonstrated how Pearl can facilitate the analysis of movement data as well as our vision of future in-situ analytics workflows. Building upon this, future work can apply our approach to real-world use cases and further empower domain experts.

Our work aims to enrich the repertoire of immersive in-situ analysis strategies. Ultimately, we hope to inspire the exploration of this exciting new perspective of designing and applying future immersive analytics, which leverages physical referents for data exploration.

ACKNOWLEDGMENTS

We thank Wolfgang Büschel and Annett Mitschick for their support in this paper, our participants in the study, and the anonymous reviewers for their constructive feedback and suggestions. This work was supported by the German Research Foundation (DFG, Deutsche Forschungsgemeinschaft) under Germany's Excellence Strategy – EXC-2068 – 390729961 – Cluster of Excellence "Physics of Life" and EXC 2050/1 – Project ID 390696704 – Cluster of Excellence "Centre for Tactile Internet with Human-in-the-Loop" (CeTI) of TU Dresden, DFG grant 389792660 as part of TRR 248 (see https://perspicuous-computing.science), and by the German Federal Ministry of Education and Research (BMBF, SCADS22B) and the Saxon State Ministry for Science, Culture and Tourism (SMWK) by funding the competence center for Big Data and AI "ScaDS.AI Dresden/Leipzig".

REFERENCES

[1] Shivam Agarwal, Jonas Auda, Stefan Schneegaß, and Fabian Beck. 2020. A Design and Application Space for Visualizing User Sessions of Virtual and Mixed Reality Environments. (2020). https://doi.org/10.2312/vmv.20201194

[2] Hamed S. Alavi, Himanshu Verma, Jakub Mlynar, and Denis Lalanne. 2018. The Hide and Seek of Workspace: Towards Human-Centric Sustainable Architecture. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) *(CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–12. https://doi.org/10.1145/3173574.3173649

[3] Gennady Andrienko, Natalia Andrienko, Peter Bak, Daniel Keim, Slava Kisilevich, and Stefan Wrobel. 2011. A Conceptual Framework and Taxonomy of Techniques for Analyzing Movement. *Journal of Visual Languages & Computing* 22, 3 (2011), 213–232. https://doi.org/10.1016/j.jvlc.2011.02.003 Special Issue on Visual Analytics and Visual Semantics.

[4] Gennady Andrienko, Natalia Andrienko, Peter Bak, Daniel Keim, and Stefan Wrobel. 2013. *Visual analytics of movement.* Springer Science & Business Media. https://doi.org/10.1007/978-3-642-37583-5

[5] Natalia Andrienko and Gennady Andrienko. 2020. Spatio-Temporal Visual Analytics: A Vision for 2020s. *Journal of Spatial Information Science* 20 (2020), 87–95. https://doi.org/10.5311/JOSIS.2020.20.661

[6] Natalia Andrienko, Gennady Andrienko, and Peter Gatalsky. 2003. Exploratory Spatio-Temporal Visualization: An Analytical Review. *Journal of Visual Languages & Computing* 14, 6 (2003), 503–541. https://doi.org/10.1016/S1045-926X(03)00046-6

[7] Ann Blandford. 2013. *Semi-structured qualitative studies.* Interaction Design Foundation, Aarhus, Denmark. http://www.interactiondesign.org/encyclopedia/semi-structured_qualitative_studies.html

[8] Connie Blok. 2000. Monitoring Change: Characteristics of Dynamic Geo-spatial Phenomena for Visual Exploration. In *Spatial Cognition II: Integrating Abstract Theories, Empirical Studies, Formal Methods, and Practical Applications*, Christian Freksa, Christopher Habel, Wilfried Brauer, and Karl F. Wender (Eds.). Springer, Berlin, Heidelberg, 16–30. https://doi.org/10.1007/3-540-45460-8_2

[9] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology* 3, 2 (2006), 77–101. https://doi.org/10.1191/1478088706qp063oa

[10] Frederik Brudy, Suppachai Suwanwatcharachat, Wenyu Zhang, Steven Houben, and Nicolai Marquardt. 2018. EagleView: A Video Analysis Tool for Visualising and Querying Spatial Interactions of People and Devices. In *Proceedings of the 2018 ACM International Conference on Interactive Surfaces and Spaces* (Tokyo, Japan) *(ISS '18)*. Association for Computing Machinery, New York, NY, USA, 61–72. https://doi.org/10.1145/3279778.3279795

[11] Wolfgang Büschel, Anke Lehmann, and Raimund Dachselt. 2021. MIRIA: A Mixed Reality Toolkit for the In-Situ Visualization and Analysis of Spatio-Temporal Interaction Data. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 470, 15 pages. https://doi.org/10.1145/3411764.3445651

[12] Wolfgang Büschel, Stefan Vogt, and Raimund Dachselt. 2018. Investigating Link Attributes of Graph Visualizations in Mobile Augmented Reality. *Proceedings of the CHI 2018 Workshop on Data Visualization on Mobile Devices. MobileVis* 18 (2018).

[13] Pietro Centorrino, Alessandro Corbetta, Emiliano Cristiani, and Elia Onofri. 2021. Managing crowded museums: Visitors flow measurement, analysis, modeling, and optimization. *Journal of Computational Science* 53 (2021), 101357. https://doi.org/10.1016/j.jocs.2021.101357

[14] Han Joo Chae, Youli Chang, Minji Kim, Gwanmo Park, and Jinwook Seo. 2020. ARphy: Managing Photo Collections Using Physical Objects in AR. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems (CHI EA '20)*. Association for Computing Machinery, New York, NY, USA, 1–7. https://doi.org/10.1145/3334480.3382885

[15] Zhutian Chen, Yijia Su, Yifang Wang, Qianwen Wang, Huamin Qu, and Yingcai Wu. 2019. Marvist: Authoring glyph-based visualization in mobile augmented reality. *IEEE Transactions on Visualization and Computer Graphics* 26, 8 (2019), 2645–2658. https://doi.org/10.1109/TVCG.2019.2892415

[16] Xiangtong Chu, Xiao Xie, Shuainan Ye, Haolin Lu, Hongguang Xiao, Zeqing Yuan, Zhutian Chen, Hui Zhang, and Yingcai Wu. 2022. TIVEE: Visual Exploration and Explanation of Badminton Tactics in Immersive Visualizations. *IEEE Transactions on Visualization and Computer Graphics* 28, 1 (2022), 118–128. https://doi.org/10.1109/TVCG.2021.3114861

[17] Bram Custers, Wouter Meulemans, Bettina Speckmann, and Kevin Verbeek. 2021. Coordinated Schematization for Visualizing Mobility Patterns on Networks. In *11th International Conference on Geographic Information Science (GIScience 2021)-Part II (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 208)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 7:1–7:16. https://doi.org/10.4230/LIPIcs.GIScience.2021.II.7

[18] Juan David Hincapié-Ramos, Levko Ivanchuk, Srikanth Kirshnamachari Sridharan, and Pourang Irani. 2014. SmartColor: Real-time Color Correction and

Contrast for Optical See-through Head-Mounted Displays. In *2014 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, 187–194. https://doi.org/10.1109/ISMAR.2014.6948426

[19] Urška Demšar, Jed A. Long, Fernando Benitez-Paez, Vanessa Brum Bastos, Solène Marion, Gina Martin, Sebastijan Sekulić, Kamil Smolak, Beate Zein, and Katarzyna Siła-Nowicka. 2021. Establishing the Integrated Science of Movement: Bringing Together Concepts and Methods from Animal and Human Movement Analysis. *International Journal of Geographical Information Science* 35, 7 (2021), 1273–1308. https://doi.org/10.1080/13658816.2021.1880589

[20] Somayeh Dodge, Robert Weibel, Sean C. Ahearn, Maike Buchin, and Jennifer A. Miller. 2016. Analysis of movement data. *International Journal of Geographical Information Science* 30, 5 (2016), 825–834. https://doi.org/10.1080/13658816.2015.1132424

[21] Neven AM ElSayed, Bruce H Thomas, Kim Marriott, Julia Piantadosi, and Ross T Smith. 2016. Situated Analytics: Demonstrating immersive analytical tools with Augmented Reality. *Journal of Visual Languages & Computing* 36 (2016), 13–23. https://doi.org/10.1016/j.jvlc.2016.07.006

[22] Barrett Ens, Benjamin Bach, Maxime Cordeil, Ulrich Engelke, Marcos Serrano, Wesley Willett, Arnaud Prouzeau, Christoph Anthes, Wolfgang Büschel, Cody Dunne, Tim Dwyer, Jens Grubert, Jason H Haga, Nurit Kirshenbaum, Dylan Kobayashi, Tica Lin, Monsurat Olaosebikan, Fabian Pointecker, David Saffo, Nazmus Saquib, Dieter Schmalstieg, Danielle Albers Szafir, Matthew Whitlock, and Yalong Yang. 2021. Grand Challenges in Immersive Analytics. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 459, 17 pages. https://doi.org/10.1145/3411764.3446866

[23] Austin Erickson, Kangsoo Kim, Gerd Bruder, and Gregory Welch. 2020. A Review of Visual Perception Research in Optical See-Through Augmented Reality. In *ICAT-EGVE 2020 - International Conference on Artificial Reality and Telexistence and Eurographics Symposium on Virtual Environments*, Ferran Argelaguet, Ryan McMahan, and Maki Sugimoto (Eds.). The Eurographics Association. https://doi.org/10.2312/egve.20201256

[24] Jorge A. Wagner Filho, Wolfgang Stuerzlinger, and Luciana Nedel. 2020. Evaluating an Immersive Space-Time Cube Geovisualization for Intuitive Trajectory Data Exploration. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (2020), 514–524. https://doi.org/10.1109/TVCG.2019.2934415

[25] Philipp Fleck, Aimee Sousa Calepso, Sebastian Hubenschmid, Michael Sedlmair, and Dieter Schmalstieg. 2022. RagRug: A Toolkit for Situated Analytics. *IEEE Transactions on Visualization and Computer Graphics* (2022), 1–1. https://doi.org/10.1109/TVCG.2022.3157058

[26] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. 2014. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 580–587. https://doi.org/10.1109/CVPR.2014.81

[27] Timo P. Gros, David Groß, Stefan Gumhold, Jörg Hoffmann, Michaela Klauck, and Marcel Steinmetz. 2021. TraceVis: Towards Visualization for Deep Statistical Model Checking. In *Leveraging Applications of Formal Methods, Verification and Validation: Tools and Trends (Lecture Notes in Computer Science)*, Tiziana Margaria and Bernhard Steffen (Eds.). Springer International Publishing, Cham, 27–46. https://doi.org/10.1007/978-3-030-83723-5_3

[28] Tony Hey, Dennis Gannon, and Jim Pinkelman. 2012. The Future of Data-Intensive Science. *Computer* 45, 05 (2012), 81–82. https://doi.org/10.1109/MC.2012.181

[29] Danny Holten and Jarke J. Van Wijk. 2009. Force-Directed Edge Bundling for Graph Visualization. *Computer Graphics Forum* 28, 3 (2009), 983–990. https://doi.org/10.1111/j.1467-8659.2009.01450.x

[30] François Homps, Yohan Beugin, and Romain Vuillemot. 2020. ReViVD: Exploration and Filtering of Trajectories in an Immersive Environment using 3D Shapes. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE, 729–737. https://doi.org/10.1109/VR46266.2020.00096

[31] Sebastian Hubenschmid, Jonathan Wieland, Daniel Immanuel Fink, Andrea Batch, Johannes Zagermann, Niklas Elmqvist, and Harald Reiterer. 2022. ReLive: Bridging In-Situ and Ex-Situ Visual Analytics for Analyzing Mixed Reality User Studies. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 24, 20 pages. https://doi.org/10.1145/3491102.3517550

[32] Christophe Hurter, Nathalie Henry Riche, Steven M. Drucker, Maxime Cordeil, Richard Alligier, and Romain Vuillemot. 2019. FiberClay: Sculpting Three Dimensional Trajectories to Reveal Structural Insights. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2019), 704–714. https://doi.org/10.1109/TVCG.2018.2865191

[33] Nikhita Joshi, Matthew Lakier, Daniel Vogel, and Jian Zhao. 2022. A Design Framework for Contextual and Embedded Information Visualizations in Spatial Augmented Reality. In *Graphics Interface 2022*. https://openreview.net/forum?id=xdwZoGTiMJ-

[34] Karsten Klein, Björn Sommer, Hieu T. Nim, Andrea Flack, Kamran Safi, Máté Nagy, Stefan P. Feyer, Ying Zhang, Kim Rehberg, Alexej Gluschkow, Michael Quetting, Wolfgang Fiedler, Martin Wikelski, and Falk Schreiber. 2019. Fly with the Flock: Immersive Solutions for Animal Movement Visualization and Analytics. *Journal of The Royal Society Interface* 16, 153 (2019), 20180794. https://doi.org/10.1098/rsif.2018.0794

[35] Simon Kloiber, Volker Settgast, Christoph Schinko, Martin Weinzerl, Johannes Fritz, Tobias Schreck, and Reinhold Preiner. 2020. Immersive Analysis of User Motion in VR Applications. *The Visual Computer* 36, 10–12 (Oct 2020), 1937–1949. https://doi.org/10.1007/s00371-020-01942-1

[36] Matthias Kraus, Johannes Fuchs, Björn Sommer, Karsten Klein, Ulrich Engelke, Daniel Keim, and Falk Schreiber. 2022. Immersive Analytics with Abstract 3D Visualizations: A Survey. In *Computer Graphics Forum*, Vol. 41. Wiley Online Library, 201–229. https://doi.org/10.1111/cgf.14430

[37] Matthias Kraus, Karsten Klein, Johannes Fuchs, Daniel A Keim, Falk Schreiber, and Michael Sedlmair. 2021. The Value of Immersive Visualization. *IEEE Computer Graphics and Applications* 41, 4 (2021), 125–132. https://doi.org/10.1109/MCG.2021.3075258

[38] Joel Lanir, Tsvi Kuflik, Julia Sheidin, Nisan Yavin, Kate Leiderman, and Michael Segal. 2017. Visualizing museum visitors' behavior: Where do they go and what do they do there? *Personal and Ubiquitous Computing* 21, 2 (2017), 313–326. https://doi.org/10.1007/s00779-016-0994-9

[39] Bokyung Lee, Minjoo Cho, Joonhee Min, and Daniel Saakes. 2016. Posing and Acting as Input for Personalizing Furniture. In *Proceedings of the 9th Nordic Conference on Human-Computer Interaction* (Gothenburg, Sweden) *(NordiCHI '16)*. Association for Computing Machinery, New York, NY, USA, Article 44, 10 pages. https://doi.org/10.1145/2971485.2971487

[40] Bokyung Lee, Michael Lee, Jeremy Mogk, Rhys Goldstein, Jacobo Bibliowicz, Frederik Brudy, and Alexander Tessier. 2021. Designing a Multi-Agent Occupant Simulation System to Support Facility Planning and Analysis for COVID-19. In *Designing Interactive Systems Conference 2021* (Virtual Event, USA) *(DIS '21)*. Association for Computing Machinery, New York, NY, USA, 15–30. https://doi.org/10.1145/3461778.3462030

[41] Bokyung Lee, Michael Lee, Pan Zhang, Alexander Tessier, Daniel Saakes, and Azam Khan. 2021. Socio-Spatial Comfort: Using Vision-Based Analysis to Inform User-Centred Human-Building Interactions. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW3, Article 238 (jan 2021), 33 pages. https://doi.org/10.1145/3432937

[42] Klemen Lilija, Henning Pohl, and Kasper Hornbæk. 2020. Who Put That There? Temporal Navigation of Spatial Recordings by Direct Manipulation. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–11. https://doi.org/10.1145/3313831.3376604

[43] Sikun Lin, Hao Fei Cheng, Weikai Li, Zhanpeng Huang, Pan Hui, and Christoph Peylo. 2017. Ubii: Physical World Interaction Through Augmented Reality. *IEEE Transactions on Mobile Computing* 16, 3 (2017), 872–885. https://doi.org/10.1109/TMC.2016.2567378

[44] Tica Lin, Yalong Yang, Johanna Beyer, and Hanspeter Pfister. 2023. Labeling Out-of-View Objects in Immersive Analytics to Support Situated Visual Searching. *IEEE Transactions on Visualization and Computer Graphics* 29, 3 (2023), 1831–1844. https://doi.org/10.1109/TVCG.2021.3133511

[45] Jingyang Liu. 2021. Semantic Mapping: A Semantics-based Approach to Virtual Content Placement for Immersive Environments. In *2021 17th International Conference on Intelligent Environments (IE)*. IEEE, 1–8. https://doi.org/10.1109/IE51775.2021.9486580

[46] Mark A. Livingston, Joseph L. Gabbard, J. Edward Swan, Ciara M. Sibley, and Jane H. Barrow. 2013. *Basic Perception in Head-Worn Augmented Reality Displays*. Springer New York, New York, NY, 35–65. https://doi.org/10.1007/978-1-4614-4205-9_3

[47] Claes Lundstrom, Thomas Rydell, Camilla Forsell, Anders Persson, and Anders Ynnerman. 2011. Multi-Touch Table System for Medical Visualization: Application to Orthopedic Surgery Planning. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (2011), 1775–1784. https://doi.org/10.1109/TVCG.2011.224

[48] Weizhou Luo, Anke Lehmann, Hjalmar Widengren, and Raimund Dachselt. 2022. Where Should We Put It? Layout and Placement Strategies of Documents in Augmented Reality for Collaborative Sensemaking. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 627, 16 pages. https://doi.org/10.1145/3491102.3501946

[49] Weizhou Luo, Anke Lehmann, Yushan Yang, and Raimund Dachselt. 2021. Investigating Document Layout and Placement Strategies for Collaborative Sensemaking in Augmented Reality. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI EA '21)*. Association for Computing Machinery, New York, NY, USA, Article 456, 7 pages. https://doi.org/10.1145/3411763.3451588

[50] Kim Marriott, Falk Schreiber, Tim Dwyer, Karsten Klein, Nathalie Henry Riche, Takayuki Itoh, Wolfgang Stuerzlinger, and Bruce H. Thomas. 2018. *Immersive Analytics*. Vol. 11190. Springer International Publishing, Cham. https://doi.org/10.1007/978-3-030-01388-2

[51] Harvey J. Miller, Somayeh Dodge, Jennifer Miller, and Gil Bohrer. 2019. Towards an Integrated Science of Movement: Converging Research on Animal Movement Ecology and Human Mobility Science. *International Journal of Geographical Information Science* 33, 5 (2019), 855–876. https://doi.org/10.1080/13658816.2018.

1564317

[52] Benjamin Nuernberger, Eyal Ofek, Hrvoje Benko, and Andrew D. Wilson. 2016. SnapToReality: Aligning Augmented Reality to the Real World *(CHI '16)*. Association for Computing Machinery, New York, NY, USA, 1233–1244. https://doi.org/10.1145/2858036.2858250

[53] Michael Oppermann and Tamara Munzner. 2020. Ocupado: Visualizing Location-Based Counts Over Time Across Buildings. *Computer Graphics Forum* 39, 3, 127–138. https://doi.org/10.1111/cgf.13968

[54] Donna J. Peuquet. 1994. It's About Time: A Conceptual Framework for the Representation of Temporal Dynamics in Geographic Information Systems. *Annals of the Association of American Geographers* 84, 3 (1994), 441–461. https://doi.org/10.1111/j.1467-8306.1994.tb01869.x

[55] Henning Pohl, Klemen Lilija, Jess McIntosh, and Kasper Hornbæk. 2021. Poros: Configurable Proxies for Distant Interactions in VR. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 532, 12 pages. https://doi.org/10.1145/3411764.3445685

[56] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. 2016. You Only Look Once: Unified, Real-Time Object Detection. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 779–788. https://doi.org/10.1109/CVPR.2016.91

[57] Patrick Reipschläger, Frederik Brudy, Raimund Dachselt, Justin Matejka, George Fitzmaurice, and Fraser Anderson. 2022. AvatAR: An Immersive Analysis Environment for Human Motion Data Combining Interactive 3D Avatars and Trajectories. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 23, 15 pages. https://doi.org/10.1145/3491102.3517676

[58] Patrick Reipschlager, Tamara Flemisch, and Raimund Dachselt. 2021. Personal Augmented Reality for Information Visualization on Large Interactive Displays. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2021), 1182–1192. https://doi.org/10.1109/TVCG.2020.3030460

[59] Marc Satkowski and Raimund Dachselt. 2021. Investigating the Impact of Real-World Environments on the Perception of 2D Visualizations in Augmented Reality. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 522, 15 pages. https://doi.org/10.1145/3411764.3445330

[60] Marc Satkowski, Rufat Rzayev, Eva Goebel, and Raimund Dachselt. 2022. ABOVE & BELOW: Investigating Ceiling and Floor for Augmented Reality Content Placement. In *2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. 518–527. https://doi.org/10.1109/ISMAR55827.2022.00068

[61] Kadek Ananta Satriadi, Jim Smiley, Barrett Ens, Maxime Cordeil, Tobias Czauderna, Benjamin Lee, Ying Yang, Tim Dwyer, and Bernhard Jenny. 2022. Tangible Globes for Data Visualisation in Augmented Reality. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 505, 16 pages. https://doi.org/10.1145/3491102.3517715

[62] Ben Rydal Shapiro. 2017. Using Space Time Visualization in Learning Environment Design. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems* (Denver, Colorado, USA) *(CHI EA '17)*. Association for Computing Machinery, New York, NY, USA, 178–183. https://doi.org/10.1145/3027063.3048422

[63] Ben Rydal Shapiro and Rogers P. Hall. 2017. Interaction Geography in a Museum. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems* (Denver, Colorado, USA) *(CHI EA '17)*. Association for Computing Machinery, New York, NY, USA, 2076–2083. https://doi.org/10.1145/3027063.3053146

[64] Mohammad Sharif and Ali Asghar Alesheikh. 2018. Context-aware movement analytics: implications, taxonomy, and design framework. *WIREs Data Mining and Knowledge Discovery* 8, 1 (2018), e1233. https://doi.org/10.1002/widm.1233

[65] Henry Shimooka and Kazuo Umemura. 2012. Trajectory study of ballroom dance using millisecond video analysis. In *ISBS-Conference Proceedings Archive*.

[66] Seung Youb Ssin, James A. Walsh, Ross T. Smith, Andrew Cunningham, and Bruce H. Thomas. 2019. GeoGate: Correlating Geo-Temporal Datasets Using an Augmented Reality Space-Time Cube and Tangible Interactions. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. 210–219. https://doi.org/10.1109/VR.2019.8797812

[67] Manuel Stein, Halldor Janetzko, Andreas Lamprecht, Thorsten Breitkreutz, Philipp Zimmermann, Bastian Goldlücke, Tobias Schreck, Gennady Andrienko, Michael Grossniklaus, and Daniel A Keim. 2017. Bring it to the pitch: Combining video and movement data to enhance team sport analysis. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (2017), 13–22. https://doi.org/10.1109/TVCG.2017.2745181

[68] Ryo Suzuki, Rubaiat Habib Kazi, Li-yi Wei, Stephen DiVerdi, Wilmot Li, and Daniel Leithinger. 2020. RealitySketch: Embedding Responsive Graphics and Visualizations in AR through Dynamic Sketching. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology* (Virtual Event, USA) *(UIST '20)*. Association for Computing Machinery, New York, NY, USA, 166–181. https://doi.org/10.1145/3379337.3415892

[69] Himanshu Verma, Hamed S. Alavi, and Denis Lalanne. 2017. Studying Space Use: Bringing HCI Tools to Architectural Projects. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) *(CHI '17)*. Association for Computing Machinery, New York, NY, USA, 3856–3866. https://doi.org/10.1145/3025453.3026055

[70] Ulrich von Zadow and Raimund Dachselt. 2017. GIAnT: Visualizing Group Interaction at Large Wall Displays. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) *(CHI '17)*. Association for Computing Machinery, New York, NY, USA, 2639–2647. https://doi.org/10.1145/3025453.3026006

[71] Günter Wallner, Nour Halabi, and Pejman Mirza-Babaei. 2019. Aggregated Visualization of Playtesting Data. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) *(CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–12. https://doi.org/10.1145/3290605.3300593

[72] Guenter Wallner and Simone Kriglstein. 2016. Visualizations for Retrospective Analysis of Battles in Team-Based Combat Games: A User Study. In *Proceedings of the 2016 Annual Symposium on Computer-Human Interaction in Play* (Austin, Texas, USA) *(CHI PLAY '16)*. Association for Computing Machinery, New York, NY, USA, 22–32. https://doi.org/10.1145/2967934.2968093

[73] Xiyao Wang, Lonni Besançon, David Rousseau, Mickael Sereno, Mehdi Ammi, and Tobias Isenberg. 2020. Towards an Understanding of Augmented Reality Extensions for Existing 3D Data Analysis Tools. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–13. https://doi.org/10.1145/3313831.3376657

[74] Zeyu Wang, Cuong Nguyen, Paul Asente, and Julie Dorsey. 2021. DistanciAR: Authoring Site-Specific Augmented Reality Experiences for Remote Environments. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 411, 12 pages. https://doi.org/10.1145/3411764.3445552

[75] Zhepei Wang, Chao Xu, and Fei Gao. 2022. Robust Trajectory Planning for Spatial-Temporal Multi-Drone Coordination in Large Scenes. (2022), 12182–12188. https://doi.org/10.1109/IROS47612.2022.9982032

[76] Wesley Willett, Yvonne Jansen, and Pierre Dragicevic. 2017. Embedded Data Representations. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (2017), 461–470. https://doi.org/10.1109/TVCG.2016.2598608

[77] Marc Wolter, Irene Tedjo-Palczynski, Bernd Hentschel, and Torsten Kuhlen. 2009. Spatial Input for Temporal Navigation in Scientific Visualizations. *IEEE Computer Graphics and Applications* 29, 6 (2009), 54–64. https://doi.org/10.1109/MCG.2009.125

[78] Shengzhi Wu, Daragh Byrne, and Molly Wright Steenson. 2020. "Megereality": Leveraging Physical Affordances for Multi-Device Gestural Interaction in Augmented Reality. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI EA '20)*. Association for Computing Machinery, New York, NY, USA, 1–4. https://doi.org/10.1145/3334480.3383170

[79] Yuan Yao, Junai Cai, Kexin Du, Yuxuan Hou, and Haipeng Mi. 2022. Establishing Design Consensus toward Next-Generation Retail: Data-Enabled Design Exploration and Participatory Analysis. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 586, 22 pages. https://doi.org/10.1145/3491102.3517637

[80] Shuainan Ye, Zhutian Chen, Xiangtong Chu, Yifan Wang, Siwei Fu, Lejun Shen, Kun Zhou, and Yingcai Wu. 2021. ShuttleSpace: Exploring and Analyzing Movement Trajectory in Immersive Visualization. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2021), 860–869. https://doi.org/10.1109/TVCG.2020.3030392