MIRIA: A Mixed Reality Toolkit for the In-Situ Visualization and Analysis of Spatio-Temporal Interaction Data

Wolfgang Büschel
bueschel@acm.org
Interactive Media Lab Dresden
Technische Universität Dresden
Dresden, Germany

Anke Lehmann
anke.lehmann@tu-dresden.de
Interactive Media Lab Dresden
Technische Universität Dresden
Dresden, Germany

Raimund Dachselt*†
dachselt@acm.org
Interactive Media Lab Dresden
Technische Universität Dresden
Dresden, Germany

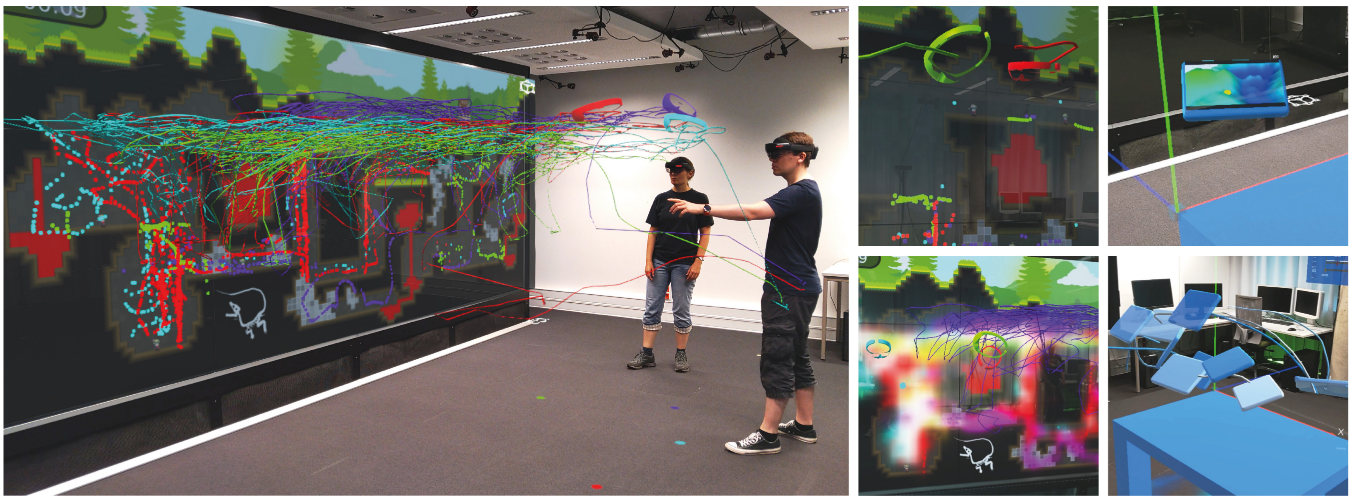

Figure 1: Our MIRIA toolkit supports the co-located, in-situ analysis of spatial interaction data by multiple users in Augmented Reality. It provides 3D visualizations, e.g., trajectories and trails, and 2D visualizations, e.g., scatterplots and heatmaps. MIRIA also supports 3D models, videos, and pictures placed in space, providing additional context to the data.

ABSTRACT

In this paper, we present MIRIA, a Mixed Reality Interaction Analysis toolkit designed to support the in-situ visual analysis of user interaction in mixed reality and multi-display environments. So far, there are few options to effectively explore and analyze interaction patterns in such novel computing systems. With MIRIA, we address this gap by supporting the analysis of user movement, spatial interaction, and event data by multiple, co-located users directly in the original environment. Based on our own experiences and an analysis of the typical data, tasks, and visualizations used in existing approaches, we identify requirements for our system. We report on the design and prototypical implementation of MIRIA, which is informed by these requirements and offers various visualizations such as 3D movement trajectories, position heatmaps, and scatterplots. To demonstrate the value of MIRIA for real-world analysis tasks, we conducted expert feedback sessions using several use cases with authentic study data.

CCS CONCEPTS

• **Human-centered computing** → **Mixed / augmented reality**; **Visualization toolkits**; • **Information systems** → *Data analytics.*

KEYWORDS

interaction analysis, immersive analytics, in-situ analysis, in-situ visualization, augmented reality, human-computer interaction, visualization

ACM Reference Format:
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 *CHI Conference on Human Factors in Computing Systems (CHI '21), May 8–13, 2021, Yokohama, Japan.* ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3411764.3445651

*Also with, Cluster of Excellence Physics of Life, Technische Universität Dresden.
†Also with, Centre for Tactile Internet with Human-in-the-Loop (CeTI), Technische Universität Dresden.

ACM ISBN 978-1-4503-8096-6/21/05.
https://doi.org/10.1145/3411764.3445651

1 INTRODUCTION

In recent years, immersive and multi-display environments receive increased interest both from practitioners and researchers, gaining traction besides traditional desktop or mobile usage scenarios. They make use of a variety of setups such as large interactive display walls or mixed reality head-mounted displays, sometimes in combination with spatially tracked mobile devices or additional, stationary displays. In visualization research, for example, Immersive Analytics [54] combines the use of mixed reality technologies with natural and embodied user interaction to support the visual analysis of data in immersive environments.

In an effort to gain insight into how people use such novel setups, user studies are conducted that frequently result in a large body of heterogeneous, spatio-temporal data of movements, user interactions, and other events that need to be visualized and analyzed to gain the desired insights. However, today's analysis tools (e.g., [8, 74]) are mostly 2D, single-user, and confined to desktops, distancing the analysts from the often rich and complex environments that might have a crucial impact on the users' behaviors. For example, patterns in the data may not be obvious during analysis from a fixed point of view or hard to explain without the original environmental context. Furthermore, classic tools often do not adequately support the various forms of multi-modal interaction employed in advanced computing environments [10], which combine input channels such as user movement, mid air gestures, pen & touch, and gaze.

We believe that visualizing interaction data *in-situ*, i.e., directly in the original environment, can support the analysis of spatial interaction in multi-display and Augmented Reality (AR) environments. To this end, we designed a toolkit, called *MIRIA* (Mixed Reality Interaction Analysis toolkit), that allows this form of in-situ data exploration and analysis by embedding AR visualizations of spatial interaction data into the physical locations were it was originally recorded.

We support the visual analysis of different interaction data as it is typically recorded in user studies: On one hand, we visualize 3D trajectories based on tracked, mostly three-dimensional, position data. This includes the movement of people in the environment, paths of mobile devices and controllers in spatial interaction, and virtual camera paths. On the other hand, we provide configurable 2D views, situated in the 3D environment, that can be used to visualize both 2D projections of spatial data, e.g., user positions on the floor, as well as 2D events on a surface or display, e.g., touches on a display wall. We also support aggregated views such as heatmaps or scatterplots.

The three main scientific contributions of this paper, and its general structure, are as follows:

Our first contribution is the concept of in-situ mixed reality analysis for spatio-temporal interaction data recorded in Augmented Reality and multi-display environments, including its requirements and challenges. Based on a systematic review of the related work and our own experiences in researching spatial and multi-modal user interaction, we thoroughly analyzed existing methods for the analysis of spatial user interaction. We specifically investigated typical, domain-specific analysis tasks, the most common data types and sources, and which visualizations are already used. From this review, we derive requirements and typical challenges of in-situ interaction visualization and analysis. Informed by these requirements, we present the concept of our MIRIA toolkit, highlighting how position and event data can be visualized and how we envision users to interact with such a system.

Our second contribution is to show the technical feasibility of our concepts by presenting a working prototypic system. We describe the setup, important design choices, and implementation details of our MIRIA toolkit. In contrast to many existing solutions, our toolkit is multi-user capable and thus allows for a co-located, collaborative visual analysis, e.g., pairing experts from human-computer interaction (HCI) and different target domains. The MIRIA toolkit runs self-contained on one or more Microsoft HoloLens Augmented Reality head-mounted displays and can easily be deployed in a variety of different environments, independent of any instrumentation such as an external tracking system or a central data server.

Our third contribution is a preliminary evaluation illustrating how our approach might be advantageously used in the collaborative visual analysis of study data. We do so by reporting on practical walkthroughs of our system with experts using real data from four user studies, two of which we discuss in more detail. These example case studies cover different technical setups (mixed reality, a large display wall, and a multi-display environment), various interaction modalities (touch, tangibles, distant pointing, spatial input), and both single- and multi-user scenarios.

2 BACKGROUND & RELATED WORK

Our work is intersecting several fields of research. In the following, we will first look into classic, mostly desktop-based, *analysis of user interaction* and, briefly, movement data in other contexts. In contrast to most of the existing systems, our approach is centered on the exploration of immersive, in-situ visualizations of spatial interaction data. Because of this, we then also discuss related work in the developing field of *Immersive Analytics*.

2.1 Analysis of User Interaction & Movement Data

Mostly, desktop-based or web-based analysis tools are used to explore and analyze user interaction in HCI user studies. Techniques for the analysis of the interaction with classic graphical user interfaces include, e.g., heatmaps of gaze or click data [33, 55]. However, for the analysis of spatial interaction or user behavior in virtual environments, specialized tools are necessary. For instance, *GIAnT* [74] allows to analyze users' locomotion and territoriality in front of a large multi-touch display wall. It provides 2D visualizations like heatmaps and scatterplots to show user movements, interactions and gaze data on the wall display, and captured video streams along with basic statistics like distance from the display or touch frequency. Brudy et al. [8] provide an analysis tool to study group interaction like f-formations of people and devices. Their tool allows to visualize people's position and movement over periods of time with trajectories, video playback, and heatmaps. Furthermore, it supports search queries to identify proxemic zones or attention grouping. *VisTACO* by Tang et al. [71] is a tool to identify patterns of spatial behavior during collaboration on interactive tabletop surfaces. With the desktop-based tool, the recorded interaction sequences of multiple users at a tabletop can be visualized.

There are also tools for the analysis of specific sensors and modalities. For example, *Kinect Analysis* [59] targets Microsoft's Kinect sensor to track and visualize body movements, and *GestureAnalyzer* [42] focuses on the analysis of gesture data. Further tools to visualize time-based user or group interaction data are *EXCITE* [53], *VICPAM* [57], *ChronoViz* [30], and *Panoramic* [80]. *VU-Flow* [19] is an earlier example for the analysis of movement patterns in virtual environments, supporting 2D trajectory plots, heatmaps, and flow maps. Also, there are commercial solutions for interaction analysis in Virtual Reality (VR) and Augmented Reality (AR) applications, such as *Cognitive3D*[1].

Similarly, there have been efforts towards the visual analysis of player behavior in computer games, both by game developers and researchers. Often, such applications make use of heatmaps, line of sight visualizations, and movement trajectories (e.g., [25, 37, 44, 56]). There are also examples for multivariate graph visualizations of player movement or progress (e.g., [58, 78, 79]) and for location-based games [22]. Typically, and in contrast to our approach, these game play analyses are done in external, desktop-based tools. Spatial exploration of interaction traces or an in-situ representation of the data are rare (e.g., [47]).

While existing classic tools provide powerful (mostly) 2D visualizations to analyze user interactions or movements, often the spatial context of user positions and movements is missing. For instance, questions such as "Why are there gaps in the trajectory plots?", "Why do users always initiate interaction from a certain position?", or "Why are certain postures adopted?" are hard to answer with these tools. We believe that the exploration of these phenomena can benefit from an *in-situ* analysis by clearly showing spatial relations between the data and the environment, e.g., movement around physical obstacles like tables or chairs, in a 1:1 scale. Such a system may also help to analyze resulting occlusion, for example in co-located collaboration, similar to the tools presented by Fender et al. for display [28] and virtual content placement [27].

Besides user interaction, spatial data from a multitude of other domains is also visually analyzed. These domains include visualizing movement in contexts such as museums [49, 83] or sports [81], traffic monitoring and surveillance [35, 72], the migration of animals [46], and flight and naval vessel paths [3, 64, 75]. A survey by Andrienko & Andrienko [2] gives an overview of visual analytics for movement data. Most of these systems are based on classic desktop setups. The spatial context is typically shown as maps or by compositing the visualizations with video recordings. Also, in most of the domains mentioned above, the data is two-dimensional or does not require a 3D representation. Therefore, mostly 2D projections are used. In contrast, in our domain the height of users or the 3D-spatial movements of tracked devices are often important. Our data also differs in that user studies typically include relatively few trajectories with high complexity, compared to many of the examples given above, in which many simple trajectories need to be considered.

2.2 Immersive Analytics

Immersive Analytics (IA) [18, 54, 67] makes use of novel display technologies such as mixed reality headsets or large display walls together with spatial or embodied interaction [10] to facilitate the immersive analysis of data. A recent and extensive survey on IA was written by Fonnet and Prié [29]. Our work is related to IA in two ways: On one hand, our system can be categorized as IA itself, both in the technologies that we use, as well as the major design principles that we employ. On the other hand, we contribute to the field of IA by presenting an approach that allows to analyze interaction in a wide range of existing and future Augmented Reality Immersive Analytics systems.

IA suggests itself for collaborative settings [6, 17, 52], allowing user-specific information to be presented via head-mounted displays (HMDs) [41, 70] or to extend (interactive) display surfaces with further AR information visualizations [15, 63]. Another potential advantage of IA is the possibility to use the whole body for interaction in an immersive Visual Analytics environment [31, 45] and to support situated visualizations, either directly attached to physical objects with relation to the data [26] or using them as landmarks that provide a frame of reference [12, 15].

Several recent examples underline the feasibility of IA systems for the exploration of spatio-temporal data: In *Fly with the flock* [46], Klein et al. examined the visualization of bird tracking data in different technical mixed reality setups. Filho et al. [77] presented a space-time cube visualization for the immersive VR visualization of trajectory data. Following the *VirtualDesk* [76] metaphor, the trajectories were shown above a 2D map rendered on a virtual desk environment. They found a lower perceived mental workload for their system compared to a conventional setting, indicating the usefulness of immersion for this type of trajectory visualization. Similarly, Yang et al. [82] examined different forms of origin-destination flow maps, finding evidence that a 3D representation can be beneficial. Batch et al. [5] presented one of the first extensive, mixed-methods studies on Immersive Analytics, testing the *ImAxes* system [21] with expert domain users. Among other measurements, they collected tracking data and analyzed it to find out how users moved in VR and how/where they arranged views such as scatterplot matrices or parallel coordinate plots. For collaborative analysis of multidimensional data, Butscher et al. [15] use an AR environment with an interactive tabletop to support fluid interaction and provide a set of guidelines for such tools.

In the last years, VR and AR frameworks for immersive data analysis have been developed, enabling scientists to explore their multidimensional or spatio-temporal data, sometimes collaboratively. Besides general frameworks (e.g., [14, 20, 63, 66]), there are also some specializations like FiberClay [38], where users can visualize multidimensional data as 3D trajectories with abstract data attributes in VR (e.g., air traffic data).

Recently, Agarwal et al. [1] presented an analysis of the design space for visualizing user actions in mixed reality. This is similar to our analysis of prior work in section 3.1 but covers mostly 2D visualizations. Closely related to our work, Kloiber et al. [47] enable the analysis of user motion in the same VR environment in which the data was recorded, focusing on 3D trajectories and visualization of key events along the timeline. In comparison, our work focuses on real and AR environments. Very relevant to our work is also *MRAT* by Nebeling et al. [60], a mixed reality analytics toolkit that has been presented recently. *MRAT* can be tightly integrated into Unity applications for Microsoft's HoloLens and allows the

[1] See https://cognitive3d.com/

visualization of user data, including support for data collection and pre-processing. The authors present the example use case of a crisis simulation and triage training application.

In contrast to previous work, we focus less on application integration & evaluation and more on the analysis of logged study data in its physical environment. Also, we present various combinable in-situ visualizations not supported in *MRAT* [60] or by Kloiber et al. [47], including a set of 2D visualizations that can be placed freely in the AR environment. Furthermore, in comparison to VR systems such as [38, 47], using Augmented Reality allows us to preserve the often important real-world context of a wide range of AR and multi-display environments. A more detailed analysis of the differences between AR, VR, and desktop-based interaction analysis systems can be found in Table 2.

3 ANALYSIS OF SPATIO-TEMPORAL USER INTERACTION

In this section we examine the specific research questions and requirements for the in-situ visual analysis of spatial user interaction in more detail and report on the data, tasks, and visualization techniques used in prior work. Based on this, we then describe requirements for the mixed reality analysis of spatio-temporal user interaction.

3.1 Use Cases and Analysis of Existing Systems

During user studies in novel, immersive, or multi-display environments, researchers capture a lot of heterogeneous temporal and spatial user-based data. This includes users' movements, logged interaction events or sequences of performed interactions, and the usage of interactive surfaces, including touch and pen interaction. Examples for such setups in the domain of Information Visualization include the visual data analysis on a large, interactive display wall [4, 48] and AR visualizations [13, 41, 52]. In examining their systems, researchers are often interested in varying spatial and time-related effects and events. Their observed or identified behavioral patterns have an impact on the design of user interfaces (UI) and the arrangement or use of technical equipment and devices. For instance, possible research interests include:

Utilization of space: How much did the users have to walk around for the examined tasks? Where in space are interactions performed and how large are spatial interaction volumes? This influences the general arrangement and setup of interactive environments as well as technical decisions about tracking volume or resolution.

Interaction on and around surfaces: Where do users touch an interactive display? Can we identify specific zones? Does interaction differ between displays, e.g., due to size and arrangement? These questions impact the UI design (placement of UI objects/application parts, user-oriented menu parts) and the use of technical equipment, e.g., display arrangement.

Social interaction between users: What are locomotion and movement paths in multi-user systems? Is there any interference between users? Can we detect proxemic zones [34]? This has an impact on workflows and the application of supporting tools.

Awareness of collaborators' actions: How did users spatially relate to each other? Could the collaborators and their actions be perceived? How closely did the users work together? These questions impact the UI design for collaborative tools and workflows, suggest suitable interaction modalities, and inform about possible (social) conflicts.

For further analysis, we examined 25 research papers from various fields dealing with spatial analysis, e.g., mixed reality, proxemic interaction [34], computer games analysis, and video surveillance. This is not a single, confined field of research, as such it is impossible to cover all relevant papers. Instead, our goal was to provide a good cross section of relevant research from different areas. To this end, we took a sample of papers, drawn from major conferences such as ACM CHI, IEEE VIS, or ACM UIST, smaller specialized conferences such as ACM ISS, as well as papers cited by them. The selected papers cover both toolkits (e.g., [7, 8, 47, 60]) as well as studies of specific aspects (e.g., [39, 51, 65]) or use cases (e.g., [25, 46, 52]) for the analysis of spatial data. Although related to spatial data, we do not consider classic GeoVis in depth, as it is typically concerned with larger scales, nor do we consider general visualization frameworks that do not specifically support spatial interaction analysis. We were interested in the typical analysis tasks, the data, and which visualizations are used for the analysis. For an

Table 1: Overview of reviewed papers and most important supported data types and visualizations.

	Data				2D Vis							3D Vis			
Reviewed Paper	Position	Orientation	Event, Activity	Images, Video	Trajectory	Heatmap	Scatterplot	Event Timeline	Line Chart	Images, Video	Statistics [1]	Trajectory	Marker, Object	Density Map	Scatterplot
Study Embodied Interaction [4]	x	x	x		x					x	x				
Social Diffusion Patterns [7]	x		x		x		x		x		x			x	
EagleView [8]	x	x	x	x	x	x		x		x					
PAMOCAT [9]	x	x		x				x	x	x		x	x		
GhostAR [16]	x	x										x	x		
VU-Flow [19]	x	x		x	x	x				x					
Social Interaction [23]	x^2	x^2		x			x			x	x				
HouseFly [24]	x			x				x		x		x		x	
Spat. User Behavior in Games [25]	x		x			x									
Study Is Moving Improving? [39]	x	x	x		x		x				x				
Study Mouse vs. Touch [40]	x	x	x	x			x		x		x				
Fly with the flock [46]	x	x^3		x					x	x		x			
IA of User Motion in VR [47]	x	x	x^4									x	x		
Study Classification Task [51]	x		x		x		x				x				
MR Remote GeoSpatial Vis [52]	x												x		x
EXCITE [53]	x	x	x	x				x		x					
VICPAM [57]			x	x				x		x					
MRAT [60]	x	x	x	x			x	x			x		x		x
VEEVVIE [61]	x				x	x	x				x				
3D Attention Volumes [62]	x	x		x						x		x		x	
Study Tabletop Territoriality [65]	x		x^2	x			x			x	x				
Browsing Videos [69]	x	x^3		x						x		x	x		
VisTACO [71]	x		x		x			x							
Stacked-based Trajectory [72]	x			x	x			x		x	x	x			
GIAnT [74]	x	x	x	x		x	x	x	x	x	x				

[1] e.g., metrics, tables, barcharts, boxplots
[2] data from annotated video
[3] orientation data derived from tracked points [46] or 2D video [69]
[4] teleportation events

overview of the papers and some of the findings discussed below, see Table 1.

In general, most of the systems are desktop based, with only few supporting an in-situ analysis (e.g., [16, 46, 47, 52, 60]). All examined systems allow for general exploration of spatial interaction or movement data, aiming at understanding locomotion or usage patterns. Many also support filtering (e.g., [57, 71]) or querying (e.g., [8, 53]). Some systems (e.g., [19, 24, 47, 61]) specifically support comparing movement data from multiple entities. Based on this and an overview in [10], we believe that we need at least to support the following typical analysis tasks: data and view specification tasks (i.e., encode/visualize, filter, derive, reconfigure), view manipulation tasks (i.e., select, navigate/explore, organize), and process/provenance tasks (i.e., annotate).

The data usually consists of positions and orientations (e.g., from users and mobile devices), sometimes with added trajectory parameters such as speed [9, 19, 46, 47] or tortuosity [72]. This tracking data was usually recorded as 2D or 3D positions and their orientation, if possible, based on the used tracking or capture system. In addition, and when applicable, event data is also visualized, including study data such as task completion times, information about user activities (e.g., user touches, application events), and speech/text. Event data can be very heterogeneous, like occurrences of an event (e.g., task started or finished), points (e.g., touches on display surface), or 2D vectors (e.g., a touch trail). Most of the data is recorded in a temporal sequence with timestamps. Many of the systems also use video recordings, for example to extract position information [23, 24, 69] or activities [65], as well as to support the analysis process later on (e.g., [8, 24, 40, 57, 74]).

In the examined systems, the following visualization techniques are used to analyze the captured data: 2D visualizations typically include simple heatmaps (e.g., [8, 19, 25, 61, 74]), 2D trajectories/paths (e.g., [19, 61, 71, 72]), and scatterplots (e.g., [7, 60, 74]) of positional data. In some cases, flow maps [7, 19] or visualizations of viewing directions [8, 74] are used. Usually, some (event) timeline visualization is also available (e.g., [57, 60, 71, 74]). These visualizations are sometimes combined with abstract visualizations [61, 72, 74] such as line charts, parallel coordinate plots, or bar charts. Only a few systems support 3D visualizations at all. If so, they show entity positions (and sometimes orientations) [9, 16, 47, 60], 3D trajectories [24, 46, 47, 69, 72], or 3D density maps [7, 24, 62]. Sometimes, time instead of height is mapped to the vertical axis, typically for space-time-trajectories.

3.2 Requirements

From the above findings, we conclude that a clever combination of basic 2D and 3D visualizations in an immersive, in-situ environment could help to support a range of basic visual analysis tasks. In the following, we list functional requirements derived from these findings. We also briefly explain how these requirements may be addressed in a system such as MIRIA by suggesting initial visualization and interaction concepts.

(R1) Visualization of Position & Movement Data: Among the most important data when studying user behavior in immersive or multi-display environments are the users' positions and movements over time. Here we have to consider movement trajectories of individual users, areas where people stand, and their viewing or gaze direction.

Based on the literature, we propose to use 3D plots of space-time trajectories to visualize tracked object and user movement paths (Figure 2, 2). Directly visualizing the spatial object position over time in 3D allows the analysts to see the data in relation to its context and helps to show movement patterns. The current time step in the data should be marked (e.g., with a simple glyph or a visualization of a user and their hands, as in [47]). Further attributes of the object can be mapped on color or line thickness (e.g., movement speed, see Figure 4, left). With all of these approaches, the complexity of the visualization is a tradeoff between the amount of information that can be encoded and the resulting level of visual clutter.

(R2) Visualization of Event Data: In addition to the spatial tracking data, event-based data is also typically recorded during a user session. This data includes user interactions (like touch, pen, speech, or other input events) and application events (like mode switches or task completions), which are useful for pattern observation and identification. Order and co-occurrence of these events is often crucial for the data analysis, showing possible patterns and dependencies.

Similar to classic tools, we propose to present such events in a timeline (Figure 2, 6), with time being represented on the horizontal axis and events being shown on the vertical axis (e.g., as glyphs or sorted stacked bands). However, not only the temporal aspects of events are important, but also *where* they happen. In addition, we suggest to also show these events in-situ, indicating location and orientation of interaction events (e.g., with 3D markers, as in [60]). Thus, the analyst can explore relations between preferred interaction distances, how the physical environment affects user interaction, and how the spatial arrangement of devices influences interaction locations.

Where interaction events are performed on surfaces of interactive displays, their position on that surface is obviously relevant for data analysis. We propose to show these events with its two-dimensional relation (e.g., as scatterplots).

(R3) Visualization of Study Context & Stimuli: We believe that providing as much context of a user study as possible, including information about the study environment, tasks, or stimuli, is an important requirement for its analysis. Analysts often use videos or images to document the study setup and application state and to track which user actions happened when and why during the study. This data can consist of single snapshots (e.g., screenshots or photos) or videos. We propose to show video recordings or images in-situ, freely positioned in the environment, e.g., on a wall, to prevent occlusion of other data. For, e.g., screencasts, we suggest coupling them to the actual displays used during the recording. In the case of mobile devices, this means that the video could dynamically move with the virtual device representation during playback; on stationary displays, the screencast can be shown as an overlay. Additional options such as time-lapse or image strips are helpful, so that the analyst can quickly navigate to interesting points in time.

Furthermore, 3D models can be used in the in-situ analysis to show important virtual content of the original study application, e.g., visualizing its tasks or original data. They also allow to accurately reconstruct the study environment in cases where the location has changed (e.g., if furniture or study equipment has been removed, see Figure 5).

(R4) Filtering: Filtering is an essential part of the analysis process for evaluating and comparing data. Accordingly, it should be supported on multiple levels. At its most basic level, analysts should be able to show only data of selected sessions and study conditions. Then, further filtering can be carried out by choosing which tracked objects or users should be shown. In addition, time-based filtering is vital: the analyst specifies a certain time period in the visualizations to analyze in detail and discuss it with their collaborators.

Finally, we propose to also support location-based filtering. The analyst may want to analyze actions in detail that were performed on specific locations. For example, they may select all data points in their vicinity. In practice, the mentioned filtering approaches may be combined, but require sufficient verbal communication between the collaboration partners during the analysis process.

(R5) Annotations: Annotations are an important part of visual data analysis. In HCI studies, a typical task is video coding, where analysts label observed situations or behavior. Therefore, we propose that annotations and tagging should be supported by an in-situ analysis tool, e.g., by placing notes in space. In addition, user-defined tags may be used to filter and cluster specific observations or to highlight them during the analysis. Furthermore, annotations could not only be coupled to a point in time but also to specific locations. Finally, import and export functionality for such tags and annotations helps analysts to transition between in-situ analysis and classic desktop setups.

(R6) Flexibility and User Preferences: For the in-situ mixed reality analysis that we aim to support with MIRIA, the placement of the visualizations must be consistent with the original location where the interaction was performed by the users during the study. Thus, the analysts need to be able to manually rearrange the origin of the virtual scene in the application, registering it with the coordinate space of the recorded data. Accordingly, the visualizations of interactions on surfaces (e.g., scatterplots of touches, 2D trajectories of mouse traces, image planes of screencasts) should be virtually projected onto the physical displays (or their virtual 3D representatives) in space (Figure 2, 6). Visualizations with clustered or aggregated data (e.g., heatmaps of visited positions) should be placed close to the visualization of spatial data, e.g., on the floor for trajectories (Figure 2, 4). However, a suitable placement of visualizations depends on the investigated HCI study and its recorded data, which can vary between studies (for examples, see section 5). We suggest that an initial configuration may be automatically inferred from the recorded data, but we propose to also support configuration files for initial setups and to allow manual rearrangement during runtime.

4 THE MIRIA TOOLKIT

Based on the requirements and informed by the literature review and our own experiences in building and evaluating mixed reality and multi-display environments, we propose a concept to facilitate the visual analysis of spatio-temporal user interaction data. In the following, we first describe this concept in more detail. Afterwards, we present the implementation of our MIRIA toolkit, describing the technical setup, the MIRIA pipeline (e.g., data import, data processing, and data exploration), and the implemented visualization and interaction concepts with respect to the listed requirements.

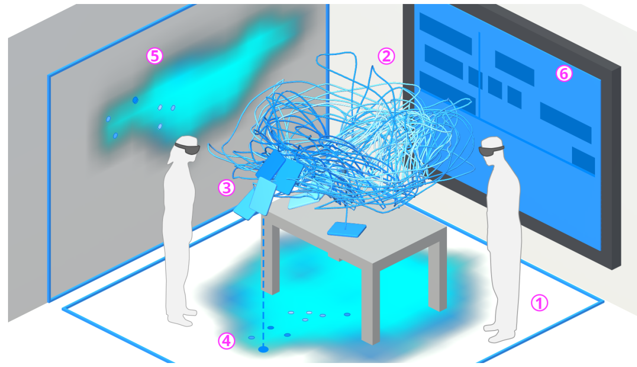

Figure 2: The MIRIA concept. 1: Co-located, in-situ analysis of spatial interaction data by multiple users; 2: 3D trajectories of logged tracking time-series; 3: 3D objects, here tablets, placed according to the currently selected point in time; 4: 2D view containers can be placed freely and show visualizations such as heatmaps or point plots; 5: 2D visualizations can be combined, here a heatmap and a plot of the current tablet positions; 6: 2D views can also show a timeline, images, or videos and can even be attached to physical objects. AR content is colored blue, real objects gray.

4.1 General Concept

With MIRIA, we aim to enable the collaborative exploration and analysis of study data by co-located analysts in mixed reality (Figure 2, 1). While we also recognize the importance of VR systems for Immersive Analytics, we specifically focus on the visual analysis of data from studies in AR or multi-display environments, thus explicitly targeting environments with physical real-world context. For the general setup, we use HMDs for an immersive AR experience, supporting in-situ analysis (Figure 2). In our analysis toolkit, we combine both 3D and 2D visualizations in a single mixed reality environment. This allows us to use the advantages of visual cues such as stereoscopy or motion to improve perception of spatial 3D data (Figure 2, 2 & 3). At the same time, 2D views (Figure 2, 4-6) are used to either visualize natively two-dimensional data (e.g., touches on an interactive surface) or 2D projections of 3D data (e.g., user positions on the floor).

The central paradigm of our approach is to enable an in-situ analysis. We support this in several ways: First, by choosing AR over VR technology, we make it possible to perceive both the data and the originally used, physical environment, blended and simultaneously. Second, we support embedding 3D models of scene geometry that is not available during analysis time but was part of the environment when the data was captured. Examples may be as simple as tables or chairs but could also include special equipment or even virtual content presented to the participants in an AR study. With this, we address that lab settings can change over time and original locations might even become completely inaccessible for the analysis phase. Third, 2D visualizations take the form of virtual, rectangular containers. These view containers can be placed in the scene. We do not only support free placement but also allow to attach views to relevant surfaces or objects in the environment (Figure 2, 4-6). As mentioned above, this enables us, e.g., to visualize movement on the ground or touch positions on a display wall, similar to click maps used in web analytics. This way, we support relating the data to the devices or locations that they are conceptually linked to. This is not limited to a static placement. Instead, visualizations can also be coupled to the tracked, dynamic location of a mobile device that was moved during the study, showing, e.g., how touch behavior changed depending on the device's location.

As video analysis is often used in classic systems, we also propose to support the playback of video recordings and displaying pictures in MIRIA. Footage of any classical observation camera can be placed in the room, showcasing details that may not be apparent from the more abstract data logs. More importantly, we expand on this by also supporting display recordings or screenshots. Here, again, we see added value in showing such recordings of user and system behavior in-situ, at the (even dynamic) position of the actual displays used during the recording. Finally, we consider support for multiple co-located analysts to be an important corner stone of our concepts. In contrast to most existing systems, we thus support analysis in a shared mixed reality space, e.g., allowing experts from different domains to work together.

4.2 Implementation Overview

We implemented most parts of our MIRIA concept in a working system.[2] This implementation serves two goals: First, it demonstrates the technical feasibility of our concepts with today's technology. Second, the MIRIA toolkit can be used by, e.g., HCI researchers for the visual exploration and analysis of user interaction data.

The MIRIA toolkit targets AR headsets, specifically the Microsoft HoloLens v2, and uses the Unity 3D engine and Microsoft's Mixed Reality Toolkit for Unity. We developed MIRIA to be multi-user capable, with one HoloLens device acting as a server that all other clients connect to. Thus, no further infrastructure is necessary. This allows us to use MIRIA in various environments with only little needed setup.

Our framework supports different virtual views placed in the AR environment: The *visualization views*, which visualize the recorded data sets (e.g., 3D trajectory view, heatmap) and the *application views*, which host the user interface for system control and visualization settings, and provide information about the currently loaded study data.

The underlying visualization workflow of MIRIA is depicted in Figure 3: First, the study data is imported, with additional meta data being stored in a configuration file. After import, the data is preprocessed. It is then displayed in visualization views placed in space. The analysts can then use the additional application views to filter the data or reconfigure the visualizations.

[2]Source available at https://github.com/imldresden/miria, see also the project website at https://imld.de/miria for more details.

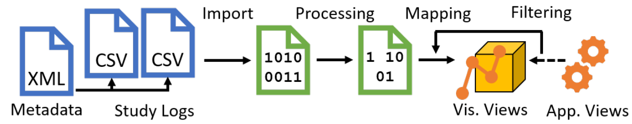

Figure 3: MIRIA pipeline. A configuration file links to the log files. After import, the data is filtered and then shown in 2D and 3D visualization views that are controlled via application views.

4.3 Metadata Description & Data Import

To support a wide range of applications and addressing requirement **R6**, we took special care to make MIRIA as flexible as possible regarding study data and setups. The captured tracking and event data (e.g., timestamps, spatial user & device positions, and interactions such as touch events) is imported from CSV files. A configuration XML file provides all relevant metadata about a study. It contains the definitions of study conditions and sessions, information about tracked and static scene entities such as type or id, the axes definitions of the original coordinate system, initial positions for 2D views, and defines the mapping between the defined objects and their logged data in the CSV files. Additionally, recorded video data and pictures can also be loaded, and the XML configuration file specifies the scene positions of these media objects. This allows analysts to display videos where they were originally taken, such as a screencast of a mobile device used during the study.

At startup of the application, all XML files, representing one study each, are parsed. We then show a list of all studies in the application, allowing users to select and load a study of their choice. Currently, all files are kept on the AR headsets and are loaded locally. In principle, however, they could also be streamed from a server.

4.4 Data Processing

In order to cope with the memory and rendering limitations of the HoloLens, a data preprocessing is used after the import step to filter the displayed data. The specific preprocessing is visualization-dependent but in general, we currently use two strategies: First, we use temporal downsampling to reduce the often high-frequency tracking data down to 15 Hz. Second, we remove tracking points with a very small distance (< 3 cm) to the last point, which are, in our experience, likely to be noise in the tracking. With these strategies and with the data used in our studies, we can typically reduce the number of data points to be drawn by about 90%. These values can easily be adapted depending on the use cases and future technology.

4.5 Visualization Views

Currently, our MIRIA implementation supports *3D trajectory plots*, *3D trails*, *2D heatmaps*, *2D scatterplots*, *2D point plots*, *media views*, and a *2D event timeline*. The 3D visualizations are directly rendered into the AR environment, while all 2D visualizations are drawn on rectangular, planar surfaces situated in 3D space. The size, position, and orientation of these containers or *2D views* can be configured

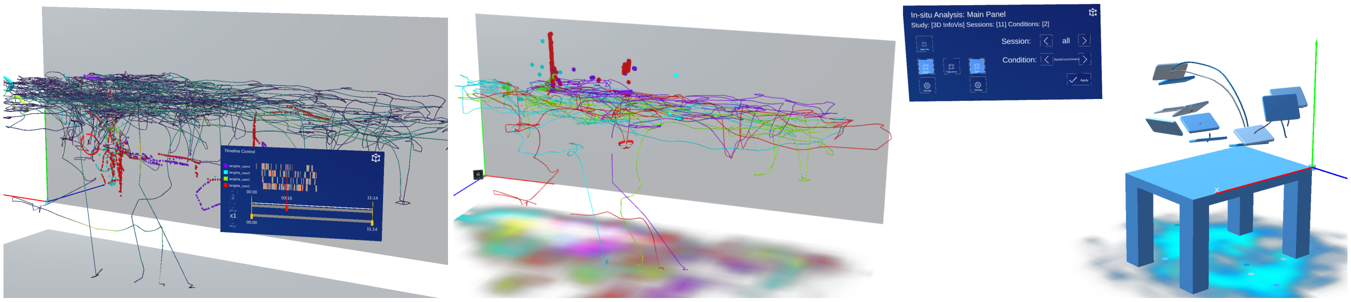

Figure 4: Screenshots of MIRIA in Unity. Left: *3D trajectories* of several users, colored to encode speed. The event timeline in the playback panel is also visible. Center: *3D trajectories* of multiple users, a *2D scatterplot* of touch positions, and a *2D heatmap* of user positions projected to the floor. Right: *3D trail visualization* of multiple tablets, tracked above a table. The *study control panel* is also visible.

through the meta data. Thus, analysts can easily place them according to the study setup or depending on the planned analysis task (Figure 4). During runtime, the 2D views can also be repositioned by simply dragging and moving them. Each 2D view can hold multiple different 2D visualizations at the same time, allowing analysts to combine several visualizations. For example, a 2D view placed on the floor may display both a heatmap and the current 2D positions of the tracked objects simultaneously, giving an overview of the data's distribution, as well as showing current positions (Figure 4, right).

4.5.1 3D Visualizations. In our current MIRIA implementation, *3D trajectories* (Figure 4, left & center) are used to visualize spatial movements of tracked objects (e.g., subjects or mobile devices), addressing requirement **R1**. We use connected, three-dimensional tube segments to visualize the objects' movement over time and employ simple shading and shadows to support the users' depth perception. The trajectories are color-coded to match the object's colors defined in the configuration file. In our current implementation, if multiple sessions are shown at the same time, the trajectories and markers are differentiated by manipulating the color saturation of each visualized session (Figure 2, 2). Alternatively, as the speed of tracked objects may be of special interest, we offer the option to encode this speed directly on the trajectories (Figure 4, left). For this, we currently use the *Viridis* colormap [68]. Using the corresponding settings button, users can open a configuration menu where they can select which trajectories to show and, for each of them, whether local speed or object color should be displayed (**R4**, **R5**). As an alternative to showing the complete paths of tracked objects, we also implemented *3D trails*, which only show the last few seconds from the currently selected time code (Figure 4, right). Apart from this dynamic, time-based filtering, they are functionally equivalent to the 3D trajectories described above.

The current data point on a trajectory or trail can be marked by a 3D object (e.g., a tablet or stylized HoloLens), definable in the configuration file. These *3D models* can be activated or deactivated independently of the trajectories. In addition to the use as dynamic position markers (**R1**), we also support an arbitrary number of static objects. As mentioned in section 4.1, these models can serve to recreate changed parts in the environment (see, e.g., the virtual table in Figure 5) or show additional context about the study tasks (**R3**).

4.5.2 2D Visualizations. The current positions of tracked objects or of interaction events in relation to a 2D plane (e.g., touches on a display) are visualized with a *2D point plot* (Figure 2, 2). It shows the positions as small, colored circles and lends itself to be combined with, e.g., 3D trajectories or a heatmap.

We also implemented *2D scatterplots* to visualize the entirety of such interaction events in two spatial dimensions (**R2**). For instance, we use scatterplots to show all touch events up to the current point in time, color-encoded per user (Figure 6, left & center). Often, there is no further semantic information, e.g., differentiating a touch down/up, a drag or hold action. Therefore, we show all events as single, unconnected points. Consequently, no continuous drag action is visualized in the current implementation.

In MIRIA we also support *2D heatmaps* (Figure 4, center). They visualize the frequency and density of positions (e.g., touches on an interactive surface, users' position in front or next to a display) in an aggregated form. To render a heatmap, we project all sample or event positions to the target plane where the view is located and accumulate their number per grid cell. The values are then normalized, mapped to color values, and applied to the corresponding texel in the heatmap texture. To improve visual quality, we also apply a moderate gaussian blur. For the color mapping, we use the object color for hue and saturation and the sample density for the value in the HSV color space. This results in a series of differently toned monochromatic colormaps that we combine by alpha blending, very similar to the process used in [74].

In an *event timeline*, we show the time and duration of events (**R2**). The objects for which events are available are stacked vertically, the current state/event is color-coded with time being mapped to the horizontal axis (Figure 4, left). In addition to the option of displaying it on any 2D view, we currently show this timeline visualization directly on the *playback panel* (compare 4.6.2).

Finally, *media views* are used to show video and image data captured during the study, allowing us to present the analysts additional information about study stimuli or application contexts (**R3**). Like the other 2D visualizations, media views can be assigned to any 2D view. These can be coupled to any static scene object (e.g., corresponding to a video camera position during the user study or to a fixed wall-display) or tracked, movable object (e.g., a tablet) (**R6**). As mentioned earlier, the media files are referenced in the XML file. Here, the analysts can also define the time offset to the logging data and which part of a video should be shown.

4.6 System & Timeline Control

After having introduced the various visualization views of MIRIA, we describe how users can position the virtual content in the physical environment, the different application views used to control the system, and how we currently support filtering of the data.

4.6.1 Data Alignment. When the first user creates an analysis session, the scene's origin is placed near them and visualized with a simple 3D marker. This position is also shared with all other analysts joining the session. However, at any time, users can move and rotate the scene simply by grabbing it, thus allowing them to match the underlying coordinate system of the logged tracking data. This cannot be automated, because, even for one given environment, the coordinate system could vary considerably between studies.

4.6.2 Application Views. The *application views* include all views that are used to setup and control the MIRIA application and the analysis process. For brevity, we only cover the two most important views here, the *study control panel* and the *playback panel* (Figure 4, left & right). Others include a *session panel* to join an existing analysis session or the *settings panels*, mentioned earlier, where visualizations can be configured (**R6**) on the fly.

The *study control panel* shows the available user studies based on the pre-loaded XML configuration files (see section 4.3) and allows loading the desired user study data. Furthermore, after loading a study, it shows information like the currently selected study conditions and participant. It also allows the users to activate or deactivate the 3D visualizations (2D visualizations are directly controlled from their 2D view panels) and to switch interactively between sessions/participants and conditions in the current data set.

A central part of our MIRIA tool is the *playback panel* (Figure 4, left). Similar to a media player's control interface, it shows the current point in time, also marked on a progress bar, and the total duration of the selected session. Buttons allow analysts to start and stop playback of the time-dependent data. As a default, playback happens at real time, but we also support playback speeds between 0.25- and 8-times normal speed. We also allow users to quickly jump to other points in time, either by dragging the slider or clicking anywhere on the progress bar.

If multiple sessions are shown at once, then all data is aligned at the first samples' time codes. All time-dependent, non-aggregated views are synchronized. This means, videos or screencasts, if available for the session, are also played synchronously. As mentioned above, we also show the *event timeline* in this panel, aligned with the progress bar.

4.6.3 Filtering. Our current MIRIA implementation supports several methods to filter the visualized data, addressing requirement **R4**. First, users can select individual or all sessions and conditions of a study for visualization. We also support simple filtering of data objects by enabling or disabling them in the settings panel of any visualization. Furthermore, we support a time-based filtering of data objects. To this end, the analyst can set an in and out marker in the playback panel to define the desired time period. Only samples within this time period are considered for visualization. The filtering is applied to all currently used visualizations and all objects/entities. Support for filtering based on location, e.g., close to an object or inside a defined volume, is planned for later versions.

4.6.4 Collaboration & Annotations. Our decision for a co-located, in-situ AR system instead of VR or remote collaboration means that the analysis of study data by multiple users is naturally supported by verbal communication and non-verbal cues such as pointing. Thus, we did not include specific software features to support collaboration. However, we found that live information about the tracking and synchronization status between the devices can help to ensure users that the digital content is correctly placed. Therefore, we added a small indicator that marks the current position and orientation of each user. Any tracking error would lead this indicator to freeze or be at an offset from its user's position. In our current MIRIA implementation, we provide some initial support for tags or annotations (**R5**). Users can place markers in space, indicating, e.g., interesting points in the data set. These markers are synchronized between users.

Besides extending the annotation possibilities, we plan to add other features later. In particular, we would also like to include advanced filtering methods, such as spatial filtering for trajectories (e.g., [36]) or filtering by example [43]. The current implementation can support the visual analysis of study data and we believe it clearly shows the feasibility of our concepts. In the next section, we thus describe how experts can use the different MIRIA visualizations during an in-situ data analysis.

5 ANALYSIS WORKFLOW WITH PRACTICAL EXAMPLES AND EXPERT FEEDBACK

As a preliminary evaluation and to illustrate the usefulness of our approach for real-world analysis tasks, we used several example use cases with authentic data and combined different validation approaches: Our implementation serves as a *technical validation* of our concepts and shows that typical HCI study data can be visualized with MIRIA on current AR devices. Additionally, *expert walkthroughs* illustrate the use of MIRIA with practical examples and provided us with helpful insights and feedback. Due to the ongoing pandemic, a full evaluation with outside participants was unfortunately not possible at this time. Instead, we invited four colleagues (three male, one female) from our institute who are HCI or visualization experts, are familiar with some of the example user studies, but did not directly work on the MIRIA implementation or this paper. With each of these experts (E1-E4), we performed a collaborative in-situ data analysis as a guided walkthrough (about one hour per session), in which one of the authors explained the software and helped to control it. For these expert feedback sessions, we used two user studies from previously published papers [12, 73], which are described in detailed below. Furthermore, we briefly report on walkthroughs of two additional example studies by two of the authors, highlighting only some key points from these scenarios. Each practical study example had a different focus on observing the use of spatial interaction or physical space during interaction. We have chosen these HCI study examples to show the diverse aspects

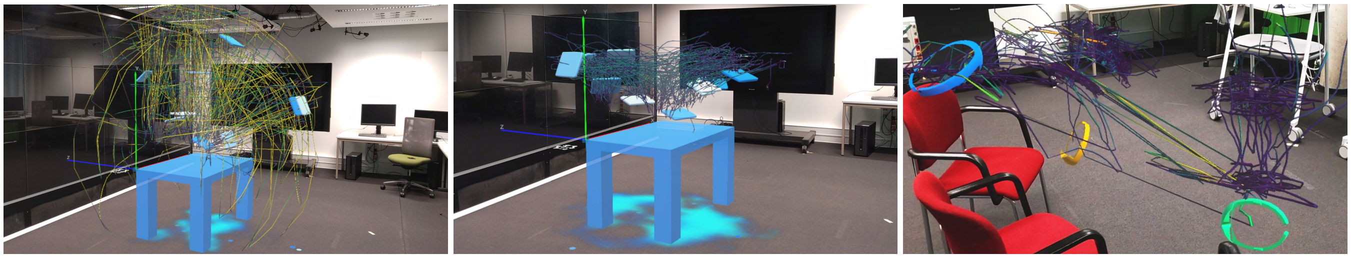

Figure 5: Application cases presented in section 5. Left: In the *spatial interaction for 3D visualization* use case, the baseline condition "touch camera", when visualized in 3D space, shows larger, faster movements impossible with a camera bound to a physical prop. Center: The same task, with the camera spatially controlled by moving a tablet. Right: In the *collaboration on multiple displays* use case, tracking errors in the analyzed data were visible. MIRIA helped to attribute them to problems with the tracking markers.

of our toolkit in the visual exploration of data from single user 3D data analysis to collaborative multiplayer games. In the following, we present these walkthroughs, summarize the expert feedback, and finish with a discussion of our MIRIA toolkit.

5.1 Tablet-based Spatial Interaction for 3D Visualizations

Our first example is a user study by Büschel et al. [12] about the use of spatial interaction for the exploration of 3D visualizations. Specifically, the paper compares the performance of users exploring an AR visualization virtually placed on a real table with a tracked tablet, where the baseline condition was a classic mobile 3D visualization with touch controls on a tablet. In both cases, the data consists of the 3D position and the forward vector of the virtual camera (view direction) logged in the original application. In case of the spatial interaction condition, this virtual camera corresponds to the tracking data of the tablet, obtained from an infrared tracking system. From the study data, we chose a subset of eleven subjects with two conditions (spatial and touch camera) and one task (the navigation task described in [12]), resulting in more than 155,000 samples.

To simulate the physical table present during the original study, we included a simple 3D mesh as a static scene object in the configuration file. Upon loading the study data, we manually adjusted the position of the scene, so that the virtual table was placed on the ground. As our experts wanted to get a better overview of all the data, they opted to show all subjects' data at the same time (Figure 5, left & center). We then demonstrated the available visualizations. For example, we assigned both a position *heatmap* and a *2D point plot* to the 2D view, allowing the experts to better perceive the current tablet position in relation to its distribution over the whole task duration. For instance, two experts (E3, E4) mentioned that the *heatmap* showed that all study subjects started at one position and some subjects walked around the table. The 2D view was partially blocked by the table, which the experts then opted to turn on and off, depending on what visualization they focused on. The *3D trajectories* allowed the experts to clearly see the typical patterns of the touch-controlled camera in the data (e.g., trajectories reaching the ceiling were mentioned by E2 and E3). In the original study paper [12], the importance of visualizing such interaction data in-situ was already stressed, with the physical surroundings putting the data into perspective. Consequently, the experts (E2-E4) could immediately see that the typical distance of the camera to the table in this mode was larger than would be possible in the AR condition of the study (Figure 5, left). In contrast, the trajectories in the AR condition were more densely distributed (Figure 5, center). In addition, color-encoding speed also showed that camera movements were faster in the touch condition. During playback, it became clear that, similar to the touch condition, participants were trying to get an overview by taking a step back and, presumably after finding the target location, quickly focused on that spot.

5.2 Multi-player Game

For the second example, we wanted to especially focus on interaction data by multiple users. We chose an observation study by von Zadow et al. [73] about users' movement and awareness during a fast-paced, collaborative multi-player game on a wall-sized display. In this game, four players had to guide miners lost in a cave back home, under time pressure and by clever use of individual tools on the wall's surface. For each of the players, the available data consists of the tracking logs (head position and orientation, again captured by an IR tracking system), the assigned tool, and touch positions on the display wall. We limited ourselves to the data of two of the groups, leading to ca. 9,800 touch positions and more than 262,000 tracking samples.

We had access to the display wall used during the study [73], allowing us to visualize and explore the study data with MIRIA in an authentic environment. In addition, we showed a screenshot of the game as an AR overlay on the display wall during analysis, providing a strong contextual cue, which would be hardly obtainable in classic analysis tools (Figure 1 & Figure 6, left & center).

The amount of data in this use case led to visual clutter. MIRIA allowed us to counteract this by filtering out the initial planning phase during the analysis, concentrating on the actual gameplay, and also by switching on and off trajectories as needed. In addition to the *3D trajectories* of the players, we chose to visualize their touch positions on a 2D view placed directly on the display wall. This enabled the experts to see touches during playback in the

Figure 6: Further application cases presented in section 5. Left & center: In the *multi-player game* use case, players had to move in front of a display wall, interacting with it by touch. Right: In the *distant 3D pointing techniques* use case, both a tablet and the user were tracked in front of a display wall.

context of the game (see, e.g., the colored dots indicating a touch in Figure 6, center). Here, one expert (E3) pointed out that they could easily recognize which assigned game tool was used by the visual behavior pattern. The game forces the players to interact with large parts of the display wall. This is immediately visible in the data. By looking at the touch positions on the wall and the corresponding trajectories in the context of the physical environment, the experts also found that players had to bend over or kneel in some cases, as seen in the lower parts of the trajectories in Figure 1 and confirmed in the recorded study videos. An observation that, without a system such as MIRIA, may make a time-consuming video analysis necessary. Visualizing the data over time with increased playback speed, we also found that the players would work in very close proximity but would also split up at times, confirming the findings presented in [73]. Furthermore, the provided visualization and application views of MIRIA helped our experts to verify the recorded data. For instance, the *event data* view showed users' actions up to five minutes (the duration of a game level), but the *3D trajectory* view shows user movements up to ten minutes, because the users were waiting for the next level to start (e.g., mentioned by E3). In addition, MIRIA's visualization of the trajectories in the environment revealed that for several minutes after the game ended, the HMDs were resting on a table.

5.3 Further Application Examples

We also tested our concepts and implementation with additional study data, briefly summarized in the following.

Collaboration on Multiple Displays. In this user study [32], groups of three people were observed during a collaborative, co-located web search task in a multi-display environment. We had access to the tracking data of people's positions and observation videos for three sessions, with a session length of about 30 minutes. In this case, there was a concrete analysis problem: an unexplained error with parts of the recorded tracking data. In contrast to prior systems, MIRIA's in-situ analysis allowed us to check the line of sight between the tracking cameras and the user positions along the 3D trajectories. By using the AR visualization in combination with the physical environment, we were thus able to confirm that the tracking problem was not caused by occlusion of the tracking cameras by the used large displays, but probably by problems with the markers on the caps used to track the participant (Figure 5, right). In addition, during analysis setup, we were able to reconstruct the general study arrangement of physical objects (a table, chairs, mobile large displays) in the lab quickly and easily using the *3D trajectories* of several sessions and photos of the study as reference. Furthermore, the easy switching between *3D trajectories* and *3D trails* of individual subjects helped to have a closer look at short data snippets and enabled us to easily pinpoint the brief phases of tracking loss in the long session.

Distant 3D Pointing Techniques. In an unpublished in-house study, several distant techniques for smartphone-based pointing at a display wall were examined (Figure 6, right). The data includes both the tracked data for users as well as the mobile devices, allowing us to visualize them together in one scene. By looking at the *2D heatmap* projected on the floor, we could see the global pattern of "swarming out" from the starting location in each study trial. In addition, the *3D trajectories* also showed a preference for one of the techniques to point from further afar. Finally, by showing the *3D trails* of both the users and the phones at the same time, the spatial relation between user and phone became apparent, allowing to see, e.g., how far away the users were holding the phone, whether they tilted it, and how user behavior changed when moving closer to the wall.

5.4 Expert Feedback and Tool Improvements

The presented in-situ data analysis with MIRIA was judged positively by all of our expert participants (E1-E4). All experts were highly interested and engaged during the sessions, the physical navigation was perceived helpful (E2, E3, E4) and as working well during data exploration (E2, E3). The experts mentioned the playback mode as most useful from the analysis point of view. A few of them also positively mentioned setting the playback speed (E1, E3, E4). In addition to 3D trajectories, showing the 3D trails for several sessions at once was pointed out as helpful (E3). All experts appreciated the efforts to visualize the study context (e.g., screencasts on physical displays, captured videos) for the data analysis (E1-E4). Furthermore, in some cases the 3D trajectory visualization helped to mentally reconstruct the physical setup of the study environment.

The expert feedback also addressed a few critical points, e.g., the limited field of view (E2, E4) and the fatigue when wearing the HoloLens over a longer period was noted (E2). Better support for annotations and the configuration of analysis sessions for reuse was also commented on (E3). The simple annotation feature mentioned in section 4.6.4 is a first step in this direction. Furthermore, experts requested the possibility to interactively synchronize the timelines of multiple sessions (E4) or to support jump marks (E3). To improve the 3D trajectory visualization, experts suggested to label the trajectories, as well as the possibility to configure the color-encoded attribute (E1, E3, E4).

5.5 Discussion & Limitations

The previous example studies show the versatility of our current implementation and the MIRIA concept in general. In particular, the variety of the studies underlines how central *spatial aspects* of user interaction are for many different immersive and multi-display applications; and thus, how important the visual exploration of this inherently three-dimensional data is. In several cases, the *in-situ analysis*, directly in the original environment, allowed analysts get a better sense of the scale of movement paths or their relation to physical objects. This was also supported by the stereoscopic rendering and would have been harder with traditional tools (e.g., [8, 53, 74]). In this regard, the *ease of deployment* and flexibility of MIRIA in various in-door environments, not needing a server or any room instrumentation, is beneficial. Manual calibration of the visualization in respect to the physical space turned out to be fast and precise enough for our use case.

Furthermore, being able to *work collaboratively* in the same space made for an engaging experience and supported us in looking at different areas at the same time. While this could potentially also be done in, e.g., a web-based analysis system, the immediate and natural exchange about each user's current position and region of interest proved helpful. We would like to point out that the *high level of user engagement* that we typically saw, could make a novel system like ours also interesting for reporting findings to others or for in-situ data story telling.

Besides the presented features and demonstrated applicability of our toolkit, it has also some limitations. One issue is that, despite its advantages, being "in the data" at times prevents an outside-in view that may be needed to get an initial overview of the data. We propose to allow analysts to temporarily scale down the whole AR scene. This way, an immersive visual analysis of the 3D interaction data would still be possible, while seamlessly switching between an in-situ view aligned with the physical environment and an overview that is easily accessible and could even be virtually placed on, e.g., a table. We discussed this with the invited experts and conclude that this strategy could be useful depending on the use cases. In addition, while the technical setup of our toolkit is easy, manually writing a configuration file for a dataset can be cumbersome, due to the flexibility of MIRIA. More configuration options in the application itself could address this problem. We plan to include these additions in future versions of the toolkit, together with the other currently unimplemented features pointed out in section 4, most notably extended annotation and filtering support. With such features, more complex user interaction with the MIRIA toolkit will be necessary, e.g., to specify filter parameters or input text. Here, the additional use of a mobile device, as examined in, e.g., [11, 84] may be promising for interface controls. Recently, Lilija et al. [50] proposed a direct manipulation approach to navigate time in VR recordings that might also be interesting for our use case.

On the technical side, MIRIA typically reaches the HoloLens 2 target frame rate of 60 fps with study data as used in our preliminary evaluation. As mentioned above, these data sets had up to several 100,000 samples (before any filtering in MIRIA) and represent diverse study logs of typical size and duration. However, the performance of the hardware currently still limits both the amount of data that can be shown while maintaining interactive framerates as well as the rendering quality, especially when showing multiple visualizations at once. In addition, perception issues such as color fidelity and resolution have to be considered. Importantly, the narrow field of view can make it harder to get an overview of data distributed over a larger area, forcing the user to move their head. While such considerations are not in the scope of this paper, we believe that the hardware limitations will continue to be addressed by manufacturers in future iterations.

Several of the advantages and limitations of MIRIA mentioned above are inherent to AR applications. Table 2 shows a brief comparison of MIRIA to similar systems and serves to highlight some of the differences to VR and desktop-based solutions. Most importantly, AR systems such as MIRIA directly combine the physical environment with virtual data representations and support face-to-face communication, while VR systems make it easier to (re-)create arbitrarily sized virtual environments. Thus, MIRIA is not meant to replace existing systems but rather serves as an alternative specifically designed for contexts that depend on physical artifacts or environmental cues.

6 CONCLUSION & FUTURE WORK

In this paper, we presented MIRIA, a concept and exemplary toolkit for the immersive visual analysis of spatial user interaction data. By analyzing prior, related systems and based on our own experiences in the field of human-computer interaction for mixed reality and multi-display environments, we proposed a series of visualization concepts for this novel type of analysis system. Our central idea is to support researchers and practitioners by providing a mixed reality solution for the in-situ visual analysis of interaction and event data. In our approach, 3D views such as trajectories or trails can be integrated with videos or images and 2D visualizations like heatmaps, scatterplots, or event timelines, all of which are virtually placed in the physical world, preserving valuable contextual and environmental cues. We provided an overview of our current implementation, detailing both its general structure and the individual visualizations implemented so far. Finally, we demonstrated the usefulness of our approach in an initial evaluation. To this end, we used a combination of expert feedback sessions and our own walkthroughs to validate MIRIA in several real-world application examples with real study data and presented their setups and findings. In the future, we plan to investigate extensions of the current interaction and visualization concepts, to continue developing the MIRIA toolkit, and to further evaluate our system in a formal user study. Among the planned additions are more powerful filtering tools, a way to resume analysis sessions and export results, as well as more specific tools to support higher level analysis goals, e.g., visualizing common interaction zones, lines of sight, and occlusion.

Table 2: Comparison of MIRIA to selected systems as described in their publications. See also Table 1 for additional information about the data and visualizations.

System	Environment	Analysis Use Case	Advantages	Disadvantages
MIRIA	AR	Interaction in AR & multi-display environments, pre-recorded room-scale data, supports collaboration by multiple analysts, also supports projected views for 2D visualizations	Support for wide range of immersive and non-immersive analysis scenarios, direct availability of physical artifacts and environmental features, easy face-to-face communication facilitates co-located collaboration	Performance limited by available AR hardware, access to original physical environment important
MRAT [60]	AR	Interaction in AR environments, supports data recording in Unity apps, focus on event data, uses additional tablet for 2D visualizations and filtering		
IA of User Motion in VR [47]	VR	Interaction in VR environments, focus on user motion	Higher rendering quality possible compared to AR, flexibility due to purely virtual environment	Recreation of real environments difficult, collaboration can suffer from HMD usage
GIAnT [74]	Desktop	Interaction in front of large display walls, pre-recorded data, typically room-scale data	Easy setup, high performance	Non-immersive, limited flexibility, 2D visualizations only

ACKNOWLEDGMENTS

We wish to thank Ricardo Langner for the fruitful discussions about our concepts and his help in preparing this paper. We also wish to thank the anonymous reviewers for their detailed and helpful feedback. This work was supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC-2068 – 390729961 – Cluster of Excellence Physics of Life of TU Dresden, DFG grant 389792660 as part of TRR 248 (see https://perspicuous-computing.science) and DFG grant DA 1319/11-1 (CollabWall).

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. In *Vision, Modeling, and Visualization*, Jens Krüger, Matthias Niessner, and Jörg Stückler (Eds.). The Eurographics Association. https://doi.org/10.2312/vmv.20201194

[2] Gennady Andrienko, Natalia Andrienko, Peter Bak, Daniel Keim, and Stefan Wrobel. 2013. *Visual Analytics of Movement*. Springer. https://doi.org/10.1007/978-3-642-37583-5

[3] Natalia Andrienko, Gennady Andrienko, Jose Manuel Cordero Garcia, and David Scarlatti. 2019. Analysis of Flight Variability: a Systematic Approach. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (Jan 2019), 54–64. https://doi.org/10.1109/TVCG.2018.2864811

[4] Robert Ball and Chris North. 2007. Realizing embodied interaction for visual analytics through large displays. *Computers & Graphics* 31, 3 (2007), 380 – 400. https://doi.org/10.1016/j.cag.2007.01.029

[5] Andrea Batch, Andrew Cunningham, Maxime Cordeil, Niklas Elmqvist, Tim Dwyer, Bruce H. Thomas, and Kim Marriott. 2020. There Is No Spoon: Evaluating Performance, Space Use, and Presence with Expert Domain Users in Immersive Analytics. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (Jan 2020), 536–546. https://doi.org/10.1109/TVCG.2019.2934803

[6] Mark Billinghurst, Maxime Cordeil, Anastasia Bezerianos, and Todd Margolis. 2018. *Collaborative Immersive Analytics*. Springer International Publishing, Cham, 221–257. https://doi.org/10.1007/978-3-030-01388-2_8

[7] Katy Börner and Shashikant Penumarthy. 2003. Social Diffusion Patterns in Three-Dimensional Virtual Worlds. *Information Visualization* 2, 3 (Sept. 2003), 182–198. https://doi.org/10.1057/palgrave.ivs.9500050

[8] 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 (ISS '18)*. ACM, New York, NY, USA, 61–72. https://doi.org/10.1145/3279778.3279795

[9] Bernhard Brüning, Christian Schnier, Karola Pitsch, and Sven Wachsmuth. 2012. Integrating PAMOCAT in the Research Cycle: Linking Motion Capturing and Conversation Analysis. In *Proceedings of the 14th ACM International Conference on Multimodal Interaction (ICMI '12)*. ACM, New York, NY, USA, 201–208. https://doi.org/10.1145/2388676.2388716

[10] Wolfgang Büschel, Jian Chen, Raimund Dachselt, Steven Drucker, Tim Dwyer, Carsten Görg, Tobias Isenberg, Andreas Kerren, Chris North, and Wolfgang Stuerzliner. 2018. *Interaction for Immersive Analytics*. Lecture Notes in Computer Science, Vol. 11190. Springer International Publishing, Cham, 95–138. https://doi.org/10.1007/978-3-030-01388-2_4

[11] Wolfgang Büschel, Annett Mitschick, Thomas Meyer, and Raimund Dachselt. 2019. Investigating Smartphone-Based Pan and Zoom in 3D Data Spaces in Augmented Reality. In *Proceedings of the 21st International Conference on Human-Computer Interaction with Mobile Devices and Services* (Taipei, Taiwan) *(MobileHCI '19)*. ACM, New York, NY, USA, Article 2, 13 pages. https://doi.org/10.1145/3338286.3340113

[12] Wolfgang Büschel, Patrick Reipschläger, Ricardo Langner, and Raimund Dachselt. 2017. Investigating the Use of Spatial Interaction for 3D Data Visualization on Mobile Devices. In *Proceedings of the 2017 ACM International Conference on Interactive Surfaces and Spaces* (Brighton, United Kingdom) *(ISS '17)*. ACM, New York, NY, USA, 62–71. https://doi.org/10.1145/3132272.3134125

[13] Wolfgang Büschel, Stefan Vogt, and Raimund Dachselt. 2019. Augmented Reality Graph Visualizations. *IEEE Computer Graphics and Applications* 39, 3 (2019), 29–40. https://doi.org/10.1109/MCG.2019.2897927

[14] Peter W. S. Butcher, Nigel W. John, and Panagiotis D. Ritsos. 2020. VRIA: A Web-based Framework for Creating Immersive Analytics Experiences. *IEEE Transactions on Visualization and Computer Graphics* (2020), 1–1. https://doi.org/10.1109/TVCG.2020.2965109

[15] Simon Butscher, Sebastian Hubenschmid, Jens Müller, Johannes Fuchs, and Harald Reiterer. 2018. Clusters, Trends, and Outliers: How Immersive Technologies Can Facilitate the Collaborative Analysis of Multidimensional Data. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) *(CHI '18)*. ACM, New York, NY, USA, 1–12. https://doi.org/10.1145/3173574.3173664

[16] Yuanzhi Cao, Tianyi Wang, Xun Qian, Pawan S. Rao, Manav Wadhawan, Ke Huo, and Karthik Ramani. 2019. GhostAR: A Time-space Editor for Embodied Authoring of Human-Robot Collaborative Task with Augmented Reality. In *Proceedings of the 32Nd Annual ACM Symposium on User Interface Software and Technology* (New Orleans, LA, USA) *(UIST '19)*. ACM, New York, NY, USA, 521–534. https://doi.org/10.1145/3332165.3347902

[17] Marco Cavallo, Mishal Dholakia, Matous Havlena, Kenneth Ocheltree, and Mark Podlaseck. 2019. Immersive Insights: A Hybrid Analytics System ForCollaborative Exploratory Data Analysis. In *25th ACM Symposium on Virtual Reality Software and Technology* (Parramatta, NSW, Australia) *(VRST '19)*. ACM, New York, NY, USA, Article 9, 12 pages. https://doi.org/10.1145/3359996.3364242

[18] Tom Chandler, Maxime Cordeil, Tobias Czauderna, Tim Dwyer, Jaroslaw Glowacki, Cagatay Goncu, Matthias Klapperstueck, Karsten Klein, Falk Schreiber, and Elliot Wilson. 2015. Immersive Analytics. In *2015 big data visual analytics (BDVA 2015) : Hobart, Australia, 22-25 September 2015*. IEEE, Piscataway, NJ, 73–80. https://doi.org/10.1109/BDVA.2015.7314296

[19] Luca Chittaro, Roberto Ranon, and Lucio Ieronutti. 2006. VU-Flow: A Visualization Tool for Analyzing Navigation in Virtual Environments. *IEEE Transactions on Visualization and Computer Graphics* 12, 6 (2006), 1475–1485. https://doi.org/10.1109/TVCG.2006.109

[20] Maxime Cordeil, Andrew Cunningham, Benjamin Bach, Christophe Hurter, Bruce H. Thomas, Kim Marriott, and Tim Dwyer. 2019. IATK: An Immersive Analytics Toolkit. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE, 200–209. https://doi.org/10.1109/VR.2019.8797978

[21] Maxime Cordeil, Andrew Cunningham, Tim Dwyer, Bruce H. Thomas, and Kim Marriott. 2017. ImAxes: Immersive Axes as Embodied Affordances for Interactive Multivariate Data Visualisation. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology* (Québec City, QC, Canada) *(UIST '17)*. ACM, New York, NY, USA, 71–83. https://doi.org/10.1145/3126594.3126613

[22] Paul Coulton, Will Bamford, Keith Cheverst, and Omer Rashid. 2008. 3D Space-Time Visualization of Player Behaviour in Pervasive Location-Based Games. *Int. J. Comput. Games Technol.* 2008 (Jan. 2008), 2:1–2:5. https://doi.org/10.1155/2008/192153

[23] Marco Cristani, Loris Bazzani, Giulia Paggetti, Andrea Fossati, Diego Tosato, Alessio Del Bue, Gloria Menegaz, and Vittorio Murino. 2011. Social interaction discovery by statistical analysis of F-formations. In *Proceedings of the British Machine Vision Conference*. BMVA Press, 23.1–23.12.

[24] Philip DeCamp, George Shaw, Rony Kubat, and Deb Roy. 2010. An Immersive System for Browsing and Visualizing Surveillance Video. In *Proceedings of the 18th ACM International Conference on Multimedia* (Firenze, Italy) *(MM '10)*. ACM, New York, NY, USA, 371–380. https://doi.org/10.1145/1873951.1874002

[25] Anders Drachen and Alessandro Canossa. 2009. Analyzing Spatial User Behavior in Computer Games Using Geographic Information Systems. In *Proceedings of the 13th International MindTrek Conference: Everyday Life in the Ubiquitous Era*. ACM, 182–189. https://doi.org/10.1145/1621841.1621875

[26] Neven ElSayed, Bruce Thomas, Kim Marriott, Julia Piantadosi, and Ross Smith. 2015. Situated Analytics. In *2015 Big Data Visual Analytics (BDVA)*. IEEE, 1–8. https://doi.org/10.1109/BDVA.2015.7314302

[27] Andreas Fender, Philipp Herholz, Marc Alexa, and Jörg Müller. 2018. OptiSpace: Automated Placement of Interactive 3D Projection Mapping Content. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) *(CHI '18)*. ACM, New York, NY, USA, 1–11. https://doi.org/10.1145/3173574.3173843

[28] Andreas Fender, David Lindlbauer, Philipp Herholz, Marc Alexa, and Jörg Müller. 2017. HeatSpace: Automatic Placement of Displays by Empirical Analysis of User Behavior *(UIST '17)*. ACM, New York, NY, USA, 611–621. https://doi.org/10.1145/3126594.3126621

[29] Adrien Fonnet and Yannick Prié. 2019. Survey of Immersive Analytics. *IEEE Transactions on Visualization and Computer Graphics* (2019), 1–1. https://doi.org/10.1109/TVCG.2019.2929033

[30] Adam Fouse, Nadir Weibel, Edwin Hutchins, and James D. Hollan. 2011. ChronoViz: a system for supporting navigation of time-coded data. In *Proceedings of the International Conference on Human Factors in Computing Systems, CHI 2011, Extended Abstracts Volume, Vancouver, BC, Canada, May 7-12, 2011*. ACM, 299–304. https://doi.org/10.1145/1979742.1979706

[31] Bruno Fruchard, Arnaud Prouzeau, Olivier Chapuis, and Eric Lecolinet. 2019. Leveraging Body Interactions to Support Immersive Analytics. In *CHI Workshop on Immersive Analytics* (Glasgow, United Kingdom). 10 pages. https://hal.archives-ouvertes.fr/hal-02095993

[32] Robert Fuhrmann, Anke Lehmann, Annett Mitschick, Ricardo Langner, and Raimund Dachselt. 2020. CoFind: A Browser Plugin for Investigating Co-Located Collaborative Web Search. In *Proceedings of the Conference on Mensch Und Computer* (Magdeburg, Germany) *(MuC '20)*. ACM, New York, NY, USA, 425–429. https://doi.org/10.1145/3404983.3410012

[33] Vanessa Georges, François Courtemanche, Sylvain Senecal, Thierry Baccino, Marc Fredette, and Pierre-Majorique Leger. 2016. UX Heatmaps: Mapping User Experience on Visual Interfaces. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 4850–4860. https://doi.org/10.1145/2858036.2858271

[34] Saul Greenberg, Nicolai Marquardt, Till Ballendat, Rob Diaz-Marino, and Miaosen Wang. 2011. Proxemic Interactions: The New Ubicomp? *interactions* 18, 1 (Jan. 2011), 42–50. https://doi.org/10.1145/1897239.1897250

[35] Yuejun Guo, Qing Xu, Xiu Li, Xiaoxiao Luo, and Mateu Sbert. 2014. A New Scheme for Trajectory Visualization. In *2014 18th International Conference on Information Visualisation*. IEEE, 40–45. https://doi.org/10.1109/IV.2014.18

[36] 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)*. 729–737. https://doi.org/10.1109/VR46266.2020.00096

[37] Nate Hoobler, Greg Humphreys, and Maneesh Agrawala. 2004. Visualizing Competitive Behaviors in Multi-User Virtual Environments. In *Proceedings of the Conference on Visualization '04 (VIS '04)*. IEEE, Washington, DC, USA, 163–170. https://doi.org/10.1109/VISUAL.2004.120

[38] 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 (Jan 2019), 704–714. https://doi.org/10.1109/TVCG.2018.2865191

[39] Mikkel R. Jakobsen and Kasper Hornbæk. 2015. Is Moving Improving?: Some Effects of Locomotion in Wall-Display Interaction. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) *(CHI '15)*. ACM, New York, NY, USA, 4169–4178. https://doi.org/10.1145/2702123.2702312

[40] Mikkel R. Jakobsen and Kasper Hornbæk. 2016. Negotiating for Space?: Collaborative Work Using a Wall Display with Mouse and Touch Input. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) *(CHI '16)*. ACM, New York, NY, USA, 2050–2061. https://doi.org/10.1145/2858036.2858158

[41] Raphaël James, Anastasia Bezerianos, Olivier Chapuis, Maxime Cordeil, Tim Dwyer, and Arnaud Prouzeau. 2020. Personal+Context navigation: combining AR and shared displays in Network Path-following. In *Proceedings of Graphics Interface (GI '20)*. CHCCS/SCDHM, Toronto, Canada. https://hal.archives-ouvertes.fr/hal-02860063

[42] Sujin Jang, Niklas Elmqvist, and Karthik Ramani. 2014. GestureAnalyzer: Visual Analytics for Pattern Analysis of Mid-Air Hand Gestures. In *Proceedings of the 2nd ACM Symposium on Spatial User Interaction* (Honolulu, Hawaii, USA) *(SUI '14)*. ACM, New York, NY, USA, 30–39. https://doi.org/10.1145/2659766.2659772

[43] Franziska Kahlert and Stefan Gumhold. 2020. Partial Matching of Trajectories with Particle Orientation for Exploratory Trajectory Visualization. In *Vision, Modeling, and Visualization*, Jens Krüger, Matthias Niessner, and Jörg Stückler (Eds.). The Eurographics Association. https://doi.org/10.2312/vmv.20201193

[44] Daniel Kepplinger, Günter Wallner, Simone Kriglstein, and Michael Lankes. 2020. See, Feel, Move: Player Behaviour Analysis through Combined Visualization of Gaze, Emotions, and Movement. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '20)*. ACM, New York, NY, USA, 1–14. https://doi.org/10.1145/3313831.3376401

[45] Ulrike Kister, Patrick Reipschläger, Fabrice Matulic, and Raimund Dachselt. 2015. BodyLenses: Embodied Magic Lenses and Personal Territories for Wall Displays. In *Proceedings of the 2015 International Conference on Interactive Tabletops & Surfaces* (Madeira, Portugal) *(ITS '15)*. ACM, New York, NY, USA, 117–126. https://doi.org/10.1145/2817721.2817726

[46] 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

[47] 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* (2020). https://doi.org/10.1007/s00371-020-01942-1

[48] Ricardo Langner, Ulrike Kister, and Raimund Dachselt. 2019. Multiple Coordinated Views at Large Displays for Multiple Users: Empirical Findings on User Behavior, Movements, and Distances. *IEEE Transactions on Visualization and Computer Graphics* 25 (1 2019), 608–618. Issue 1. https://doi.org/10.1109/TVCG.2018.2865235

[49] J. Lanir, P. Bak, and T. Kuflik. 2014. Visualizing Proximity-Based Spatiotemporal Behavior of Museum Visitors Using Tangram Diagrams. *Computer Graphics Forum* 33, 3 (2014), 261–270. https://doi.org/10.1111/cgf.12382

[50] 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)*. ACM, New York, NY, USA, 1–11. https://doi.org/10.1145/3313831.3376604

[51] Can Liu, Olivier Chapuis, Michel Beaudouin-Lafon, Eric Lecolinet, and Wendy E. Mackay. 2014. Effects of Display Size and Navigation Type on a Classification Task. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Toronto, Ontario, Canada) *(CHI '14)*. ACM, New York, NY, USA, 4147–4156. https://doi.org/10.1145/2556288.2557020

[52] Tahir Mahmood, Willis Fulmer, Neelesh Mungoli, Jian Huang, and Aidong Lu. 2019. Improving Information Sharing and Collaborative Analysis for Remote GeoSpatial Visualization Using Mixed Reality. In *2019 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, 236–247. https://doi.org/10.1109/ISMAR.2019.00021

[53] Nicolai Marquardt, Frederico Schardong, and Anthony Tang. 2015. EXCITE: EXploring Collaborative Interaction in Tracked Environments. In *Human-Computer Interaction – INTERACT 2015*, Julio Abascal, Simone Barbosa, Mirko Fetter, Tom Gross, Philippe Palanque, and Marco Winckler (Eds.). Lecture Notes in Computer Science, Vol. 9297. Springer International Publishing, 89–97. https://doi.org/10.1007/978-3-319-22668-2_8

[54] 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.

[55] Justin Matejka, Tovi Grossman, and George Fitzmaurice. 2013. Patina: Dynamic Heatmaps for Visualizing Application Usage. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 3227–3236. https://doi.org/10.1145/2470654.2466442

[56] Pejman Mirza-Babaei, Günter Wallner, Graham McAllister, and Lennart E. Nacke. 2014. Unified Visualization of Quantitative and Qualitative Playtesting Data.

In *Proceedings of the Extended Abstracts of the 32Nd Annual ACM Conference on Human Factors in Computing Systems (CHI EA '14)*. ACM, New York, NY, USA, 1363–1368. https://doi.org/10.1145/2559206.2581224

[57] Roshanak Zilouchian Moghaddam and Brian Bailey. 2011. VICPAM: A Visualization Tool for Examining Interaction Data in Multiple Display Environments. In *Human Interface and the Management of Information. Interacting with Information*, Michael J. Smith and Gavriel Salvendy (Eds.). Lecture Notes in Computer Science, Vol. 6771. Springer Berlin Heidelberg, 278–287. https://doi.org/10.1007/978-3-642-21793-7_32

[58] Dinara Moura, Magy Seif el-Nasr, and Christopher D. Shaw. 2011. Visualizing and Understanding Players' Behavior in Video Games: Discovering Patterns and Supporting Aggregation and Comparison. In *Proceedings of the 2011 ACM SIGGRAPH Symposium on Video Games (Sandbox '11)*. ACM, New York, NY, USA, 11–15. https://doi.org/10.1145/2018556.2018559

[59] Michael Nebeling, David Ott, and Moira C. Norrie. 2015. Kinect Analysis: A System for Recording, Analysing and Sharing Multimodal Interaction Elicitation Studies. In *Proceedings of the 7th ACM SIGCHI Symposium on Engineering Interactive Computing Systems* (Duisburg, Germany) *(EICS '15)*. ACM, New York, NY, USA, 142–151. https://doi.org/10.1145/2774225.2774846

[60] Michael Nebeling, Maximilian Speicher, Xizi Wang, Shwetha Rajaram, Brian D. Hall, Zijian Xie, Alexander R. E. Raistrick, Michelle Aebersold, Edward G. Happ, Jiayin Wang, Yanan Sun, Lotus Zhang, Leah E. Ramsier, and Rhea Kulkarni. 2020. MRAT: The Mixed Reality Analytics Toolkit. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '20)*. ACM, New York, NY, USA, 1–12. https://doi.org/10.1145/3313831.3376330

[61] Charilaos Papadopoulos, Ievgeniia Gutenko, and Arie E. Kaufman. 2016. VEEVVIE: Visual Explorer for Empirical Visualization, VR and Interaction Experiments. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (Jan 2016), 111–120. https://doi.org/10.1109/TVCG.2015.2467954

[62] Thies Pfeiffer. 2012. Measuring and Visualizing Attention in Space with 3D Attention Volumes. In *Proceedings of the Symposium on Eye Tracking Research and Applications (ETRA '12)*. ACM, New York, NY, USA, 29–36. https://doi.org/10.1145/2168556.2168560

[63] Patrick Reipschläger, Tamara Flemisch, and Raimund Dachselt. 2020. Personal Augmented Reality for Information Visualization on Large Interactive Displays. *IEEE Transactions on Visualization and Computer Graphics* (2020). https://doi.org/10.1109/TVCG.2020.3030460

[64] Roeland Scheepens, Niels Willems, Huub van de Wetering, and Jarke J. van Wijk. 2011. Interactive Visualization of Multivariate Trajectory Data with Density Maps. In *2011 IEEE Pacific Visualization Symposium*. IEEE, 147–154. https://doi.org/10.1109/PACIFICVIS.2011.5742384

[65] Stacey D. Scott, M. Sheelagh T. Carpendale, and Kori M. Inkpen. 2004. Territoriality in Collaborative Tabletop Workspaces. In *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work* (Chicago, Illinois, USA) *(CSCW '04)*. ACM, New York, NY, USA, 294–303. https://doi.org/10.1145/1031607.1031655

[66] Ronell Sicat, Jiabao Li, Junyoung Choi, Maxime Cordeil, Won-Ki Jeong, Benjamin Bach, and Hanspeter Pfister. 2019. DXR: A Toolkit for Building Immersive Data Visualizations. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (Jan 2019), 715–725. https://doi.org/10.1109/TVCG.2018.2865152

[67] Richard Skarbez, Nicholas F. Polys, J. Todd Ogle, Chris North, and Doug A. Bowman. 2019. Immersive Analytics: Theory and Research Agenda. *Frontiers in Robotics and AI* 6 (2019), 82. https://doi.org/10.3389/frobt.2019.00082

[68] Nathaniel Smith and Stéfan van der Walt. 2015. A better default colormap for Matplotlib. *SciPy2015* (2015). URL: https://www.youtube.com/watch?v=xAoljeRJ3lU.

[69] Yuki Sugita, Keita Higuchi, Ryo Yonetani, Rie Kamikubo, and Yoichi Sato. 2018. Browsing Group First-Person Videos with 3D Visualization. In *Proceedings of the 2018 ACM International Conference on Interactive Surfaces and Spaces* (Tokyo, Japan) *(ISS '18)*. ACM, New York, NY, USA, 55–60. https://doi.org/10.1145/3279778.3279783

[70] Tianchen Sun, Yucong Ye, Issei Fujishiro, and Kwan-Liu Ma. 2019. Collaborative Visual Analysis with Multi-level Information Sharing Using a Wall-Size Display and See-Through HMDs. In *2019 IEEE Pacific Visualization Symposium (PacificVis)*. IEEE, 11–20. https://doi.org/10.1109/PacificVis.2019.00010

[71] Anthony Tang, Michel Pahud, Sheelagh Carpendale, and Bill Buxton. 2010. VisTACO: Visualizing Tabletop Collaboration. In *ACM International Conference on Interactive Tabletops and Surfaces (ITS '10)*. ACM, New York, NY, USA, 29–38. https://doi.org/10.1145/1936652.1936659

[72] Christian Tominski, Heidrun Schumann, Gennady Andrienko, and Natalia Andrienko. 2012. Stacking-Based Visualization of Trajectory Attribute Data. *IEEE Transactions on Visualization and Computer Graphics* 18, 12 (Dec. 2012), 2565–2574. https://doi.org/10.1109/TVCG.2012.265

[73] Ulrich von Zadow, Daniel Bösel, Duc Dung Dam, Anke Lehmann, Patrick Reipschläger, and Raimund Dachselt. 2016. Miners: Communication and Awareness in Collaborative Gaming at an Interactive Display Wall. In *Proceedings of the 2016 ACM International Conference on Interactive Surfaces and Spaces* (Niagara Falls, ON, Canada). ACM, New York, NY, USA, 235–240. http://dx.doi.org/10.1145/2992154.2992174

[74] 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)*. ACM, New York, NY, USA, 2639–2647. https://doi.org/10.1145/3025453.3026006

[75] Katerina Vrotsou, Halldor Janetzko, Carlo Navarra, Georg Fuchs, David Spretke, Florian Mansmann, Natalia Andrienko, and Gennady Andrienko. 2015. SimpliFly: A Methodology for Simplification and Thematic Enhancement of Trajectories. *IEEE Transactions on Visualization and Computer Graphics* 21, 1 (Jan. 2015), 107–121. https://doi.org/10.1109/TVCG.2014.2337333

[76] Jorge A. Wagner Filho, Carla Freitas, and Luciana Nedel. 2018. VirtualDesk: A Comfortable and Efficient Immersive Information Visualization Approach. *Computer Graphics Forum* 37, 3 (2018), 415–426. https://doi.org/10.1111/cgf.13430

[77] 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 (Jan 2020), 514–524. https://doi.org/10.1109/TVCG.2019.2934415

[78] 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 (CHI '19)*. ACM, New York, NY, USA, 363:1–363:12. https://doi.org/10.1145/3290605.3300593

[79] Günter Wallner and Simone Kriglstein. 2012. A Spatiotemporal Visualization Approach for the Analysis of Gameplay Data. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, New York, NY, USA, 1115–1124. https://doi.org/10.1145/2207676.2208558

[80] Evan Welbourne, Magdalena Balazinska, Gaetano Borriello, and James Fogarty. 2010. Specification and Verification of Complex Location Events with Panoramic. In *Proceedings of the 8th International Conference on Pervasive Computing* (Helsinki, Finland) *(Pervasive'10)*. Springer-Verlag, Berlin, Heidelberg, 57–75. https://doi.org/10.1007/978-3-642-12654-3_4

[81] Yingcai Wu, Xiao Xie, Jiachen Wang, Dazhen Deng, Hongye Liang, Hui Zhang, Shoubin Cheng, and Wei Chen. 2019. ForVizor: Visualizing Spatio-Temporal Team Formations in Soccer. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (Jan 2019), 65–75. https://doi.org/10.1109/TVCG.2018.2865041

[82] Yalong Yang, Tim Dwyer, Bernhard Jenny, Kim Marriott, Maxime Cordeil, and Haohui Chen. 2019. Origin-Destination Flow Maps in Immersive Environments. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (Jan 2019), 693–703. https://doi.org/10.1109/TVCG.2018.2865192

[83] Massimo Zancanaro, Tsvi Kuflik, Zvi Boger, Dina Goren-Bar, and Dan Goldwasser. 2007. Analyzing Museum Visitors' Behavior Patterns. In *User Modeling 2007 (Lecture Notes in Computer Science)*, Cristina Conati, Kathleen McCoy, and Georgios Paliouras (Eds.). Springer Berlin Heidelberg, 238–246.

[84] Fengyuan Zhu and Tovi Grossman. 2020. BISHARE: Exploring Bidirectional Interactions Between Smartphones and Head-Mounted Augmented Reality. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '20)*. ACM, New York, NY, USA, 1–14. https://doi.org/10.1145/3313831.3376233