

Shared Surfaces and Spaces: Collaborative Data Visualisation in a Co-located Immersive Environment

Benjamin Lee, Xiaoyun Hu, Maxime Cordeil, Arnaud Prouzeau, Bernhard Jenny and Tim Dwyer




Fig. 1. Three participants collaboratively visualising and analysing multidimensional data in the FIESTA system, as seen in reality (top left), the virtual environment from the point of view of an observer (bottom left) and a top-down view of the scene from which we analyse their behaviour (right).

Abstract— Immersive technologies offer new opportunities to support collaborative visual data analysis by providing each collaborator a personal, high-resolution view of a flexible shared visualisation space through a head mounted display. However, most prior studies of collaborative immersive analytics have focused on how groups interact with surface interfaces such as tabletops and wall displays. This paper reports on a study in which teams of three co-located participants are given flexible visualisation authoring tools to allow a great deal of control in how they structure their shared workspace. They do so using a prototype system we call *FIESTA*: the Free-roaming Immersive Environment to Support Team-based Analysis. Unlike traditional visualisation tools, FIESTA allows users to freely position authoring interfaces and visualisation artefacts anywhere in the virtual environment, either on virtual surfaces or suspended within the interaction space. Our participants solved visual analytics tasks on a multivariate data set, doing so individually and collaboratively by creating a large number of 2D and 3D visualisations. Their behaviours suggest that the usage of surfaces is coupled with the type of visualisation used, often using walls to organise 2D visualisations, but positioning 3D visualisations in the space around them. Outside of tightly-coupled collaboration, participants followed social protocols and did not interact with visualisations that did not belong to them even if outside of its owner's personal workspace.

Index Terms—Immersive analytics, collaboration, virtual reality, qualitative study, multivariate data

1 INTRODUCTION

Rapid advances in virtual and augmented reality (VR and AR) technologies offer exciting possibilities for data visualisation. In the future, the places where we work together—especially for exploring and understanding data—will likely differ from our current desktops and meeting rooms. It is, however, difficult to predict precisely what these places for collaborative immersive analytics will look like.

In lab environments, researchers have invested significant efforts in

- Benjamin Lee, Xiaoyun Hu, Maxime Cordeil, Arnaud Prouzeau, Bernhard Jenny, and Tim Dwyer are with Monash University. Email: {benjamin.lee1,annika.hu1,max.cordeil,arnaud.prouzeau,bernie.jenny,tim.dwyer}@monash.edu.

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxx

building collaborative visualisation environments that offer unconventional arrangements of displays, such as wall-sized displays, interactive tabletops, monitors on adjustable stands, and displays projected onto surfaces. More recently, there have been efforts to use AR to extend these displays. An example for individual use is DesignAR, which demonstrates the use of a Microsoft HoloLens with an interactive surface for the creation of 3D models [45]. Specifically for collaborative visualisation, Augmented Reality above the Tabletop (ART) uses immersive parallel coordinates visualisations floating above an interactive tabletop which acts as the interaction surface [6]. The DataSpace environment [8] and its deployed Immersive Insights system [9] combines AR with 15 large high-resolution displays mounted on robot arms and a central projection table. Each of these setups and others like them (as described in Sec. 2) present interesting examples of specific combinations of display hardware and interaction techniques, but they are a subset of many possible configurations of collaborative environments. It has been shown [44] that the collaborative behaviour that occurs in

such environments is dependent on the arrangement of displays.

In this paper we aim to explore how groups collaborate when given complete freedom to organise their workspace, while investigating how users naturally make use of the available surfaces during their analysis. To do so, we built FIESTA: the Free-roaming Immersive Environment to Support Team-based Analysis [36]. FIESTA was designed to allow users to create, manipulate and share data visualisations while freely moving in a physically co-located VR environment. It emulates a conventional room with virtual walls and a table, which users may use during their analysis process. Our main contributions are:

- Our FIESTA prototype, allowing multiple users to create data visualisations in a shared virtual environment – Sec. 3.
- A user study involving the use of FIESTA to solve collaborative immersive analytics tasks consisting of 10 groups, each with three participants (30 participants total) – Sec. 4 and Sec. 5. To the best of our knowledge, this is the first study observing groups of more than two in a VR data visualisation task.
- Our findings from this user study regarding the collaborative use of surfaces and spaces for data visualisation and implications for future development of collaborative immersive analytics – Sec. 6.

Our observations show that groups were able to perform visual analytics tasks in collaboration. Their use of the surfaces in the environment was influenced by the types of visualisations they created, pinning 2D visualisations on the walls, but placing 3D visualisations egocentrically in the space around them—oftentimes ignoring the table. In general, we observed that the room was equally divided between participants into territories, with participants respecting ownership of objects during individual work. During tightly-coupled collaboration however, participants were willing to use each others’ objects regardless of territory. Finally, participants presented findings to an audience in a variety of ways, but often neglected to account for issues specific to 3D visualisation such as viewing angles and occlusion.

2 RELATED WORK

2.1 Collaborative Immersive Analytics

Collaborative immersive analytics has been identified by Billingham et al. as an emerging field of research at the intersection of InfoVis, CSCW and HCI [4]. One of the first immersive platforms for collaborative data visualisation was the CAVE Automatic Virtual Environment (CAVE) [13]. It consisted of a room with screens projected on each wall that immerses users inside the data using stereoscopic technology. Several similar systems have since been developed with improved resolution and tracking systems [14, 19]. Stereoscopic information visualisation has been shown to improve spatial perception, complex scene understanding, and memory [40]. Marai et al. presented different use cases for the CAVE2, highlighting its benefits for collaborative data analysis [38]. To improve this aspect, Amatriain et al. presented AlloSphere, a spherical immersive room with a diameter of 10 meters that can accommodate 30 people [1]. However, in addition to their high price, these rooms can only provide appropriate stereoscopic vision to a single tracked user. Cordeil et al. [12] studied immersive collaborative visualisation of 3D graphs between pairs of participants, comparing performance of a CAVE2 with connected VR head-mounted displays (HMDs). They found that performance was highly accurate in both environments, meaning pairs of analysts do not necessarily require a CAVE-style environment for collaborative immersive analytics.

Using VR HMDs, Nguyen et al. proposed a preliminary collaborative immersive analytics system in a co-located setting, where pairs of users visualise a star plot in the centre of a virtual room [42]. Similar work by the Institute for the Future showcased the Organizational Network in Virtual Reality (ONVR) prototype, which allows groups of remote users to collaboratively interact with a network visualisation in 3D space [22]. To the best of our knowledge, these are the only systems involving the use of HMDs for collaborative immersive analytics with no external interfaces or displays. However, both solely focus on specific visualisations, limiting groups to a singular shared visualisation. Other collaborative systems have sought to integrate AR/VR HMDs

with external displays such as tabletops and large displays, oftentimes using these surfaces as shared interaction and work spaces. For example, Butscher et al. combined a touch sensitive tabletop and AR to allow users to interact with and manipulate a 3D parallel coordinates visualisation [6]. Cavallo et al. proposed DataSpace, a collaborative room-scale immersive environment composed of AR HMDs and many configurable large displays arranged in a CAVE-like fashion [8]. They subsequently presented Immersive Insights, a hybrid analytics system built on the DataSpace environment which integrates touch sensitive displays along the edges of the environment, a touch-enabled projection table in the centre of the room, and an AR data visualisation floating above it for collaborative data analysis [9]. In our work, we aim to understand how larger groups perform collaborative immersive analytics tasks in an unconstrained immersive environment, whereby users can move freely and are not restricted to tabletops or large displays.

2.2 Territoriality in collaborative contexts

When collaborating in shared environments, people oftentimes divide the workspace into separate territories. Scott et al. observed participants working together on a physical tabletop and identified three types of territories which were implicitly created: *personal*, *shared*, and *storage* territories [47]. These territories are dynamic, changing depending on the needs of the activity. When individuals are required to move, Tang et al. found that these territories constantly shift in response to their physical positions [49], as people tend to establish a personal space in the area directly in front of them [32, 50]. In contrast to tabletops, high-resolution wall displays offer a much larger space for collaborators to share. Jakobsen and Hornbæk observed that territories were transient, with groups fully sharing the screen without any explicit negotiation [30]. However, instances where participants interacted with an area on the wall display that was directly in front of someone else only occurred during tightly-coupled work. On the other hand, Liu et al. found that pairs of participants work at different areas of the wall display, but only when given facilities to communicate and assist each other across distances [37]. While territoriality has been studied for tabletops and wall displays, how users move and manage these territories is constrained by the surface. In this work, we investigate a fully unconstrained environment where users can work in the open space around them, and freely utilise any available nearby surfaces.

2.3 View and window management

AR and VR researchers have explored the manner in which views, in our case free-floating data visualisations, can be arranged and managed to evoke spatial memory and cognition. Robertson et al. used an art gallery metaphor to place windows in a 3D space, with the main task window on a ‘stage’ and others along the walls of the gallery [46]. An earlier work by Feiner et al. showcased an AR system for placing virtual windows in the 3D world around the user [20]. They defined three types of windows: *surround-fixed windows* positioned along a sphere surrounding the user, *display-fixed windows* positioned directly on the heads up display, and *world-fixed windows* fixed to movable objects in the 3D world. Ens et al. used this concept of body- and world-fixed windows in The Personal Cockpit, with AR windows automatically arranged in a curved body-fixed or flat world-fixed layout [17]. Extending this, Ens et al. demonstrated how body-fixed layouts can automatically transition into world-fixed layouts while preserving relative spatial consistency in diverse environments [18]. While much work has explored how view management can be automated [3, 33], there is no agreed upon windowing metaphor for the arrangement of these views [39], and little work exploring how users naturally organise these views when performing data visualisation tasks in immersive environments. Most relevant to our work is that by Cordeil et al. on ImAxes [11] and the subsequent evaluation by Batch et al. [2]. They found that users would arrange their views egocentrically in a body-fixed like fashion while working [17, 20], switching to a linear or semi-circular layout for presentation. However, it is uncertain if the introduction of multiple users would influence this behaviour.

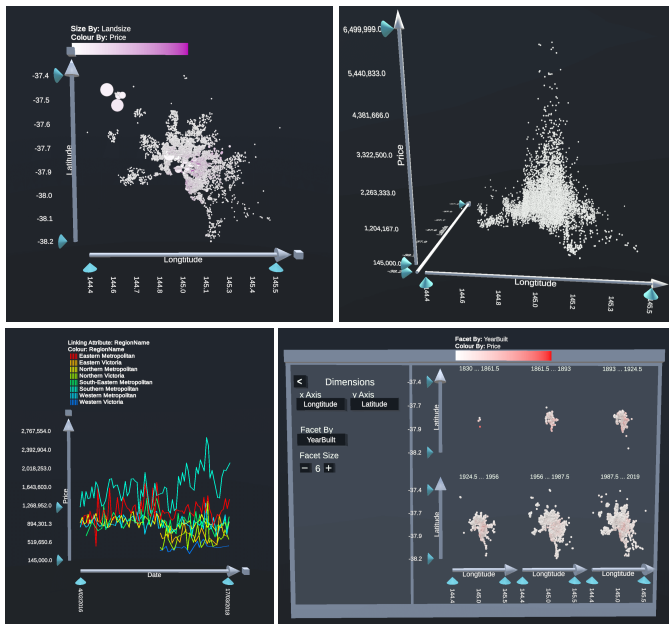


Fig. 2. Visualisation styles available with FIESTA: 2D scatterplot (top left), 3D scatterplot (top right), time series (bottom left), faceted 2D scatterplots on the panel (bottom right).

3 A PROTOTYPE FREE-ROAMING COLLABORATIVE IMMERSIVE ANALYTICS SYSTEM

We designed and developed FIESTA to study how groups collaboratively explore data in a free-roaming immersive analytics environment [36]. The source code is made publicly available on a GitHub repository¹. We developed FIESTA with three fundamental design requirements in mind, which we describe in turn in this section.

3.1 Requirement 1: Baseline Data Visualisation System

Given the novelty of collaborative immersive analytics, we implement a baseline level of functionality expected of any data visualisation system [27]. We describe these in terms of the key characteristics and functions of FIESTA.

Visualisations. FIESTA was built around a customised version of the Immersive Analytics Toolkit (IATK) [10]. We use IATK because it is optimised for rendering hundreds of thousands of data points at high frame rates, which is important when multiple users can produce dozens of visualisations each. Users can assemble three main types of 2D and 3D visualisations: scatterplots, faceted scatterplots, and time series (Fig. 2). For scatterplots, each data record is represented by a circular billboard. For time series, lines are drawn between billboarded points. Visualisations are static objects that float in 3D space.

Interaction. We choose user interface elements that are standard across desktop and VR environments to minimise learning time. We adapt 3D UI interaction metaphors [35] for data visualisation authoring and manipulation: grasping techniques involving direct contact with UI elements at close range, and ranged pointing techniques involving distant interaction with UI elements using a laser pointer.

Authoring. To create visualisations, users interact with flat panels reminiscent of physical touch panel displays (Fig. 2, bottom right). A panel provides a $1.6\text{ m} \times 1.1\text{ m}$ interface that can be moved by grasping and releasing it at the desired location, where it remains suspended in space. One panel is provided for each user. The panel is split between a user interface and an accompanying visualisation. The user interface consists of navigable pages which expose the following functions: binding data dimensions to glyph size, colour, x, y and (optionally) z position; defining data attributes to build time series;




Fig. 3. FIESTA uses a tear-out metaphor to duplicate visualisations from panels, allowing them to be freely positioned either in space or on the surfaces of the virtual room.

faceting by dimension and adjusting number of facets; and defining colour channels with gradients for continuous attributes and discrete palettes for categorical attributes. Visualisations update in real time when their properties are changed. 2D visualisations are positioned parallel to the panel, with 3D visualisations protruding outwards from it. All operations utilise conventional user interface elements, including buttons, dropdown menus, sliders, and HSV colour pickers. These are activated by performing a direct grasping or ranged laser pointer action on the interface element.

Tearing out visualisations. After a visualisation is edited on a panel, it can be cloned using a ‘tearing out’ action, grasping it and pulling it away from the panel (Fig. 3). The copy can then be freely positioned parallel to the panel, with 3D visualisations protruding outwards from it. All operations utilise conventional user interface elements, including buttons, dropdown menus, sliders, and HSV colour pickers. These are activated by performing a direct grasping or ranged laser pointer action on the interface element.

Tearing out visualisations. After a visualisation is edited on a panel, it can be cloned using a ‘tearing out’ action, grasping it and pulling it away from the panel (Fig. 3). The copy can then be freely positioned parallel to the panel, with 3D visualisations protruding outwards from it. Torn out visualisations can be adjusted through grasping interactions on widgets along their axes: resizing the visualisation along a dimension; and rescaling minimum and/or maximum axis ranges to adjust the visible domain of the dimension. Visualisations can be dropped onto a panel for further editing, or destroyed with a ‘throwing’ action aimed at the ground.

Pointer, brushing, and annotation tools. A set of additional tools are available via a spin menu [23] above the user’s offhand which the user can enable: a private brush to make selections that are only visible to the user, a shared brush to make selections which are shown to all users, and a ranged grab which can be used to pull distant visualisations towards the user. Brushing and selection of data records is linked across all visualisations (Fig. 4, left). A pointer appears when a tool is enabled, and is activated whenever the controller’s trigger is pulled. In addition, a details-on-demand tool can be enabled by holding down the touchpad on the user’s dominant hand. This shows details of the nearest data record to the pointer (Fig. 4, right) in an infobox connected by a leader line to the glyph. The pointer can also be used to interact with a panel’s interface, either using point and click (for buttons) or point and drag (for sliders and pickers). Lastly, we include basic annotation tools, accessible through floating marker and eraser objects which users can pick up and use. These annotations ‘attach’ to visualisations when drawn on top of their bounding boxes.

3.2 Requirement 2: Free Roaming Shared Environment

Our goal was to emulate a physically co-located collaborative environment using VR. We use a tetherless VR setup to prevent tripping and enable free roaming around the environment. Users see each other as




Fig. 4. FIESTA supports additional data visualisation functionalities: private and shared brushing modes which are linked across all visualisations (left); details on demand to easily inspect data records (right).

¹<https://github.com/benjaminchlee/FIESTA>

virtual avatars aligned to their real-world positions. We use the Oculus Avatar SDK which provides embodied head, torso, and hand models, supporting presence, deixis (i.e. pointing, finger gestures), and gaze direction [26]. Each user is uniquely identified by a floating nameplate and avatar colour. The same colour is also used for shared brush selections. This allows users to see the actions of others to support collaborative tasks and information sharing, as well as to avoid physical collisions. In addition, we support keyboard and mouse controlled virtual avatars that are visible to VR users, which can be used as a remote audience member (described further in Sec. 4).

Sun *et al.* suggest private workspaces increase the need for communication in order for collaborators to synchronise as a group [48]. To mitigate this and to promote workspace awareness [25], we make everything visible to all users at all times, including pointers, details-on-demand tooltips, etc. The exception is the private brush tool, as shared brushing can disturb other users [43]. While each user is given their own panel at the start of the session, FIESTA allows users to interact with any object in the environment, supporting simultaneous interaction of panels and visualisations with no explicit ownership system.

3.3 Requirement 3: Room and Surface Affordances

To explore how users naturally utilise surface affordances in collaborative immersive analytics, we establish a room metaphor similar to prior work [8, 13, 42] using a standard four-wall room layout. These walls imitate basic interaction with wall-sized displays, allowing users to freely ‘pin’ visualisations on the vertical surface. Likewise, we incorporate a tabletop metaphor through a square table in the middle of the room (1 m × 1 m, 1 m tall). The table allows for visualisations to ‘rest’ on top of it similar to prior work [6, 8, 31]. These pinning functions are performed by releasing a visualisation close to the surface, or by physically throwing a visualisation at it. Virtual surfaces are not aligned with any physical objects meaning users can simply walk through them, but the walls act as a boundary to prevent collisions with the actual walls in the room. Because visualisations can be suspended in space, the use of surfaces is inherently optional. However, since past work places great emphasis on interaction with surfaces (e.g. [34]) we felt it was important to provide them and to study their effect.

3.4 Technical Implementation and Deployment

FIESTA was developed using Unity3D. While we have confirmed it to work with remote collaborators via the internet, we chose to deploy the system on three VR ‘backpacks’ with HTC Vive Pros in a 16 m² space (4 m × 4 m). The VR backpacks are Intel Core i7-7820HK (2.9 GHz, 4 cores) PCs with a Nvidia GeForce GTX 1070 (8 GB) GPU, and 16 GB of RAM VR One PCs from MSI [41]. We bundle and secure any cables to enable truly free-roaming VR with no tethers or tripping hazards. Four SteamVR 2.0 base-stations were positioned at each corner of the room to minimise tracking issues. The backpack PCs were managed remotely by a fourth desktop PC running an Intel Core i7-7800X CPU (3.5 GHz, 6 cores), Nvidia GeForce GTX 1080 (8 GB) GPU, and 32 GB of RAM. This desktop was positioned outside of the VR space. Networking was done using the Photon Unity Networking engine. All computers were connected to a server running on the desktop PC on a dedicated WLAN to minimise external interference and latency. A server update rate of 30 Hz was used to balance between smooth synchronisation of objects and battery life of the backpacks.

4 USER STUDY

The purpose of this work is to explore how groups behave in a virtual environment like FIESTA to solve visual analytics tasks, and to understand how users utilise surface-based functionalities to accomplish this. To do so, we conducted an exploratory study with 10 groups of three participants, split equally between two study parts with five groups each: Part A and Part B. In Part A, we limit participants to 2D visualisations and with only the walls as 2D surfaces. In Part B, we further explore the surface metaphor by including the virtual table, as well as enabling 3D visualisations. Otherwise, study conditions were kept as similar as possible, which we describe in this section.

4.1 Dataset and Study Design

We used a housing auction result dataset with 8,400 data records collected from the Melbourne, Australia region. This was chosen as it is simple and recognisable to most of our participants, and has a combination of temporal, spatial, categorical, and numerical dimensions. The study focused around two types of data analysis tasks:

Directed Tasks (DT). Groups were given specific questions to answer. This was to both observe how groups achieve a specified goal, and to give some guidance to better familiarise them with FIESTA and the dataset. Three of these questions were given, DT1: “Where is the most expensive house located and why is it so expensive?”; DT2: “Have house sizes gone up or down over time?”; DT3: “A family of four is moving to Melbourne with the following requirements... Where would you recommend them to live?” Groups were expected to give an answer with consensus.

Free Exploration Task (FET). Groups were given 15 minutes to freely explore the same dataset to find any insights they deemed interesting. This was to observe how they worked together towards a more open objective. They then had to present these findings to a virtual embodied audience avatar controlled by keyboard and mouse by the experimenter. The avatar was positioned at the direction of the group, and was controlled as passively as possible (i.e. minimal movement, no follow-up questions asked).

4.2 Procedure

Each session was approximately 90 minutes from start to finish, and were conducted as follows.

Introduction and Training (20 min). Each group was given a brief introduction to the study, followed by a demographics questionnaire. Participants were then trained on how to use FIESTA by the experimenter who was present in the virtual environment as a fourth user, allowing them to ‘follow along’ the instructions and become accustomed to having others around them. In Part B, they were also taught how the table and 3D visualisations worked.

Directed Tasks (30 min). Groups were given directed task one after the other. The task was written on a specific virtual wall and was read out loud. At the beginning, authoring panels were instantiated in intentionally awkward positions to force participants to reposition them manually, but these panels (and all visualisations) were not forcefully reset or removed between tasks. Verbal time warnings were given, but much leeway was given for those still actively working on an answer.

Break (10 min). Groups were then given a 10 minute break. During this time, they answered a mid-study questionnaire about their collaborative strategies and behaviours. Participants were not permitted to discuss the questionnaire amongst each other whilst answering it.

Free Exploration Task (25 min). After putting on the VR equipment again, groups were given the free exploration task. Instructions and warnings were the same as for the directed tasks.

Wrap Up (5 min). Participants were then given a final post-study questionnaire, which was similar to the mid-study questionnaire but with additional questions about immersion in the environment.

4.3 Participants

We recruited 30 participants (7 female, 23 male) aged between 18 and 64, with all but one participant having a background in computer science. 74% were familiar with data visualisation and 74% with VR technology (self-reported 3 or higher out of 5). The participants were recruited via word of mouth and were a mix of academics and post-graduate students from our university. All had normal or corrected-to-normal vision and had no strong susceptibility to motion sickness. All but two participants reported being familiar with the other members in their group.

4.4 Data Collected

During the study, a fourth PC was used to collect the spatial positions of each participants heads, hands, and of all objects in the virtual room. It was also used to collect the start and end times of all interactions, such as panel button clicks and visualisation interactions. Finally, it was used to video record a top-down view of the virtual room. The

Task	Duration [min]	Movement [m/min]	Actions Performed [per min]		Collab. Duration [min]		
			Grasping	Ranged	Discussion	Group Work	
Part A	DT1	6.4	6.7	4	1.8	2.0	0.9
	DT2	7.0	6.3	3.5	2.2	3.0	0.6
	DT3	13	5.8	2.6	2.8	3.5	2.6
	FET	20.2	6.2	2.8	3.1	6.5	2.4
	Total	46.7	6.2	3	2.7	3.8	1.6
Part B	DT1	5.9	5.1	3	1	1.2	0.5
	DT2	5.4	5.2	1.7	1	0.8	1.1
	DT3	12.3	4.6	1.4	1	1.1	4.3
	FET	20	5.3	1.7	1.7	5.8	4
	Total	43.5	5.1	1.8	1.3	2.2	2.5

Table 1. Averages of task duration, participant movement, actions performed on panels and visualisations, and collaboration time collected from trace data and collaboration video coding.

questionnaire contained a mix of Likert-scale and short-answer based questions. Likert-scales were on a scale of 1 (disagree) to 5 (agree).

4.5 Analysis Process

Three coders initially coded the videos with an open coding methodology looking for collaborative activity and patterns, specifically identifying periods of discussion and tightly-coupled collaboration. We also processed the log data to analyse participant movement and actions, as well as the placement and properties of objects in the environment. We then more closely viewed the video footage to extract a deeper qualitative contextual understanding of individual and collaborative behaviours. See supplemental material for further details².

5 FINDINGS AND OBSERVATIONS

We first present general findings about the study and how users utilised FIESTA, followed by our main study results structured around a *representative group workflow* described later. For the rest of this paper, groups are prefixed with a G and participants with a P. G1–G5 and P1–P15 are Part A, G6–G10 and P16–P30 are Part B.

5.1 General Results

Table 1 shows the averages of task times, lateral head movement, and number of grasping and ranged actions performed on both panels and visualisations. Groups in Part B took less time than those in Part A, but this was not strictly controlled during the study. Part A performed many more actions per minute than Part B, particularly due to their increased use of the panel to modify visualisation properties (4000 total actions on panels in Part A, 2037 in Part B). Participants had mixed preferences of these grasping vs ranged actions: 17 predominantly used grasping actions, 10 mainly used ranged actions, and three used a mix of the two. To our surprise, only three groups (G2, G3, G4) made substantial use of the brushing feature (longer than three minutes continuous brushing), with other groups only using it either briefly or not at all. In contrast, the details-on-demand tool was frequently used by all. While intended for inspection of specific points such as outliers, many used it to ‘scan’ over the visualisation, rapidly inspecting many points in quick succession. This was also used to point at specific parts of visualisations.

In Part B, a considerable number of participants created more 3D than 2D visualisations (Fig. 5). This suggests a willingness to use the z-axis to visualise a third dimension instead of other visual channels such as size or colour. However, this was partly the result of 3D visualisations being easier to manipulate and rotate when torn out than when still on the panel. When comparing the number of creation actions per minute between study parts in Fig. 6, those in Part B noticeably created more visualisations than those in Part A, but conversely performed far less actions on the panel to modify visualisation properties (Table 1).

Participants responded positively when asked about their experience with FIESTA. They felt immersed in the virtual environment (27 participants reported a 4 or 5), felt comfortable using the VR setup (25 reported a 4 or 5), and that they felt their group-mates were really there with them (25 reported a 4 or 5). Taking into account setup times and training, participants spent roughly 80 minutes in VR, with the longest continuous duration being 50 minutes. Despite this, the majority of




Fig. 5. Proportion of 2D versus 3D visualisations torn out per each participant for all tasks in Part B, further split by if an additional channel (size, colour) was present in the visualisation.

participants reported little to no cybersickness throughout the study (4 participants reported a 3, the rest 1 or 2), but a few complained that the heat of the VR backpacks had started to become uncomfortable. No participant commented on any issues regarding usability or latency, however this was not explicitly asked in the questionnaire.

5.2 Representative Group Workflow

We structure our main study observations around a *representative group workflow*. That is, we identify major stages in the groups’ workflows and describe each with a narrative (*italic text*) of notable behaviours (with generalised personas Alice, Bob, and Carol). Below each stage, we describe how groups followed or deviated from it, as well as any alternate behaviours we observed.

[Stage 1] Setting up the workspace. *Alice, Bob, and Carol, are working together to find insights in the dataset akin to the free exploration task. They each place their panels along a separate wall without any verbal coordination.*

Placement of panels and movement of participants throughout all tasks can be seen in Fig. 7. Many groups did not negotiate where to place panels or how to divide the workspace, but always resulted in roughly equally sized areas for each participant. Two groups (G8, G9) positioned themselves side-by-side to more easily see each other, with the other groups taking separate parts of the room. Groups that took separate parts of the room placed their panels against a wall or the table, only moving them for specific reasons such as to move closer to their group members (G9 in Fig.7 as an example). In almost all cases, panel placement was intrinsically linked to where its owner comfortably moved and worked, essentially defining their personal territory. This is noticeable in Fig. 7 as participant movement was typically in front of their panel. Outside of tightly-coupled collaboration, no participant entered the territory of another.

[Stage 2] Collaboration strategy. *Alice suggests formulating a strategy. The group decides to divide the work: Alice looks at the most expensive properties in the data, Bob at temporal trends, and Carol at spatial characteristics.*

In most cases, groups began exploring the data on their own panels without explicitly discussing strategy. Even so, they ended up using




Fig. 6. Boxplots of actions per minute for certain interaction types for each participant across all tasks. Vertical lines represent quartiles, points represent individual participants.

²<https://sites.google.com/monash.edu/shared-surfaces-and-spaces>

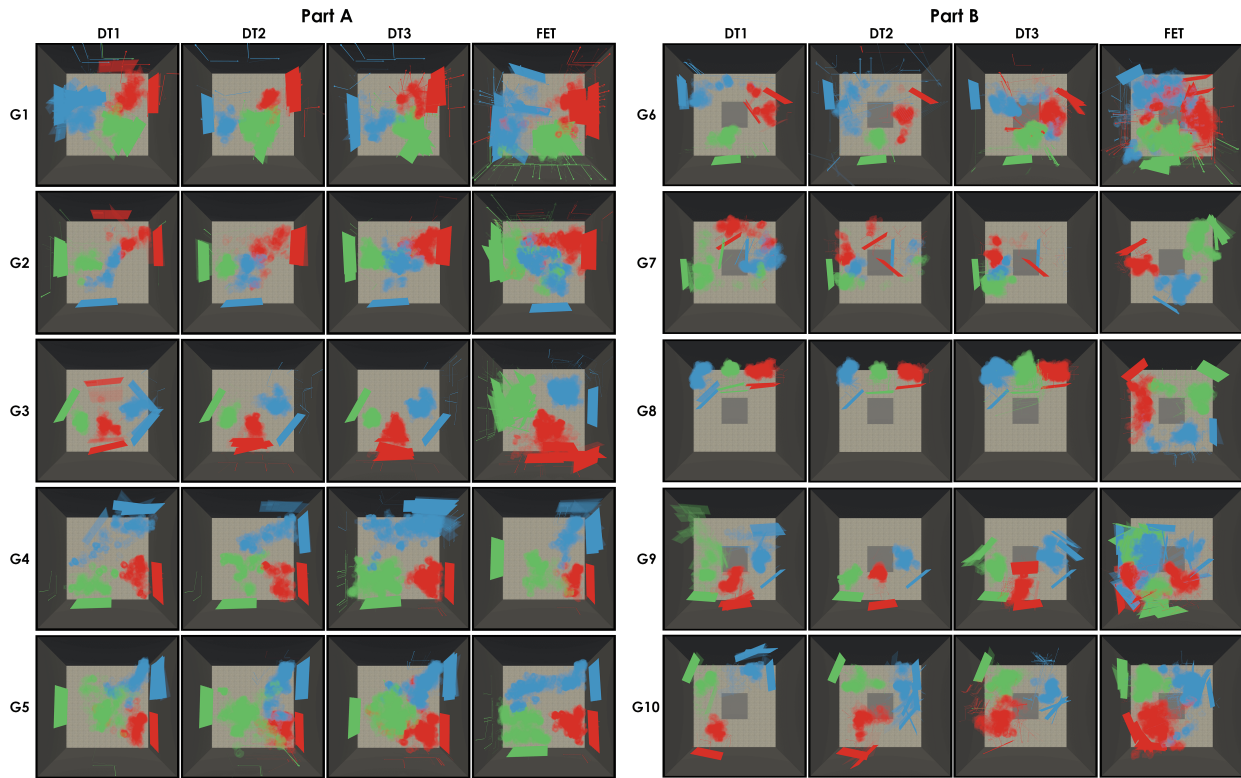


Fig. 7. Bird's-eye view heatmaps of each participant and their objects sampled every 5 seconds using a perspective projection. Large dark grey border is the wall, grey centre square in Part B is the table. Brighter areas indicate more time spent there. Red, green, and blue each represent a unique participant in that group, visualisations and panels use the same colour as their owner.

mixed-focus collaboration styles, seeking assistance and/or sharing findings throughout the task. Three groups (G3, G4, G8) progressed toward divide and conquer strategies, particularly during the free exploration task. They would negotiate aspects and dimensions of the data they each wanted to explore before splitting up, occasionally checking their results with each other. Three other groups (G6, G7, G9) organically shifted to tightly-coupled collaboration for a few tasks in the study, working on a panel or visualisation together for a significant period of time. Overall, participants would usually break off to work individually on their own panels, but would remain engaged with any ongoing discussion and seamlessly rejoin when needed.

[Stage 3] Working individually. *They each work independently in separate areas. Alice isolates outliers along a Price dimension with the axis scalers and cycles through dimensions on the other. She tears out 'completed' 2D scatterplots and pins them neatly on the adjacent wall. Bob tears out a 2D scatterplot of Date \times LandSize, widening it to around 2 m along its x-axis to visually analyse it in front of his panel. Carol tears out multiple 3D scatterplots of Longitude \times Latitude with*




Fig. 8. Actual instances of planar used by P6 (left) and egocentric layouts used by P29 (right) used by participants.

various z-axis dimensions, rotating each directly in front of her and leaving previous ones in the open space around her.

Most participants used two distinct visualisation layouts while working: a planar layout with visualisations in a grid-like fashion, as per Alice (8 in Part A, 1 in Part B), or an egocentric layout with visualisations in a spherical-like fashion around them, as per Carol (5 in Part A, 12 in Part B). Examples of both are shown in Fig. 8. Planar layouts would oftentimes use the wall for 2D visualisations in Part A, exemplified by the higher attach action count in Part A in Fig. 6. Note that some participants with egocentric layouts changed to planar layouts for presentation (described in Stage 8). We also observed other styles, such as P5 who only placed visualisations in a central 'shared' area, and P17 and P18 who would tear out and analyse a single large 2D visualisation, as per Bob. To pan across large visualisations, participants either walked back and forth (P17), or grabbed and moved the visualisation sideways, clutching as many times as necessary (P18, P28) in order to minimise body movement. It was generally not necessary to move for smaller visualisations. We saw no instances of participants orbiting around 3D visualisations when working independently, instead rotating them on the spot by clutching (hence the visibly higher grasp counts for Part B in Fig. 6).

[Stage 4] Transitions to tightly-coupled collaboration. *Carol finds something interesting, and announces to her group that she has something to show them.*

Participants often communicated regardless of physical location to ask for assistance and share findings. In some cases, this lead to tightly-coupled collaboration around a visualisation or a panel (Fig. 9 left). Sometimes this collaboration was instigated more by action: moving visualisations in front of others (P14, P16), or placing visualisations in a group territory (P1, P4, P16). This group territory was only present in three groups (G1, G2, G6) and was typically in the centre of the room, acting as a space to analyse visualisations together. This space was never explicitly negotiated or defined. In other groups, participants would either look from their own personal territory (G3, G4, G8, G10),




Fig. 9. Actual instances of tightly-coupled collaboration: G7 working on the same panel together (left), G6 working on a 3D visualisation on the centre table together (right).

or enter each others' territories only during tightly-coupled collaboration (G5, G7, G10). The openly shared nature of FIESTA allowed for some spontaneous moments of collaboration as well, such as P12 assisting in identifying specific data points from across the room using the pointer, or P25 noticing and offering to help a struggling team member. When asked if VR helped or hindered collaboration, seven stated it was useful to see each others' work in real-time to improve workspace awareness, and another seven stating that it was easy to share findings with one another. Overall, 21 participants rated a 3 or more for being aware of what their team members were doing at all times.

[Stage 5] Surface and 3D visualisation usage. *Carol places a 3D scatterplot of Longitude \times Latitude \times YearBuilt on the table, with the others joining her on opposite sides.*

Contrary to our intentions, only G6 placed 3D visualisations on top of the table with the snapping functionality for group work (Fig. 9, right). Other groups would either place visualisations at head height above the table (G8, G10) or not use the table at all (G7, G9). Despite this, G6 would sometimes place 3D visualisations floating in space, which was usually a result of the table already being occupied. The group would also sometimes struggle with placing the visualisations exactly on top of the table, often needing to make multiple adjustments due to the lack of gravity. In contrast, participants frequently placed 2D visualisations along the walls, as described earlier for planar layouts, but never placed them on the table. Tightly-coupled collaboration with 2D visualisation was always done side-by-side. For 3D visualisations, groups positioned themselves around the visualisation (G6, G7, G8, G10) more often than side-by-side (G6, G9). However this positioning was oftentimes a natural consequence of their original starting positions. Unlike working individually, participants working together with 3D visualisations refrained from rotating them unless contextually appropriate, such as explaining things to others (P23, P28), rotating to help find a specific data point together (P26, P27), or to collectively view it side-on rather than top-down (G6). There was only one instance where a participant, P20, moved around a 3D visualisation to view it from the same perspective as another.

[Stage 6] Spatial awareness of virtual objects. *They notice a conical shape of recent properties extending outwards over time from the central business district. They walk 'through' the virtual table and 'into' the visualisation to reach the city centre, using the details-on-demand tool to investigate historical properties.*

We did not instruct participants to avoid walking through the table (unlike the VRTable in Kraus *et al.* [31]), but eight of 15 were observed to deliberately avoid it. One group of three (G6) did walk through the table in a similar manner to the presented scenario, standing inside of a large 3D visualisation for closer inspection (similar to the VRRoom [31]) rather than rescaling the visible ranges along visualisation axes. However, this occurred both on top of and away from the table, using large visualisations ($\approx 1.7 \text{ m} \times 1.7 \text{ m}$) that were placed in the corners of the room. The remaining four of 15 participants paid little heed to the table, occasionally walking through it either on accident or when walking from one end of the room to the other. While most participants did not use the table, some placed their panels against it as if it were a wall (G8 in Fig. 7 as an example). In contrast, participants never

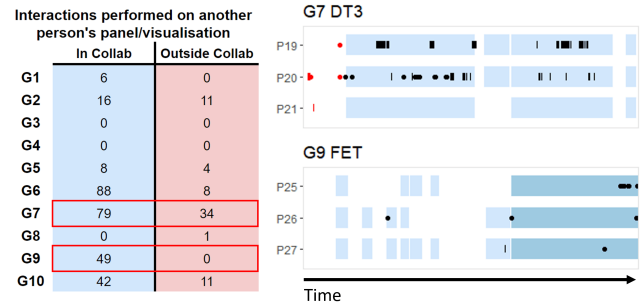


Fig. 10. Total counts of interactions performed on objects created by someone else, in and outside of tightly-coupled collaboration (left); and selected examples showing when these occurred in relation to tightly-coupled collaboration and presentation phases (right). Light blue phases are tightly-coupled collaboration, darker blue is presentation in FET. Black dots and bars are discrete and continuous interactions performed during collaboration, red dots and bars are outside of collaboration.

walked through walls (as it acted as the room boundary) and all but three avoided moving through panels. In two cases (G8, G9) this was pushed to the extreme, as they placed panels side-by-side in such a manner where it cordoned off two-thirds of the room, inadvertently restricting their available space.

[Stage 7] Behaviours and protocols during collaboration. *Alice decides to work independently again, pinning 2D visualisations as before. Bob draws Carol's attention to his 2D visualisation, stating there is no trend between Date and LandSize. Carol suggests that he should change Date to YearBuilt. Bob throws the visualisation onto an empty wall, and quickly creates a new YearBuilt \times LandSize scatterplot. Observing the slight positive trend, Carol brushes the largest properties to highlight them on the 3D scatterplot. While this selection also shows on Alice's visualisations, she makes no fuss about it.*

Those in Part B were much more willing to interact with panels and visualisations belonging to others than in Part A. This was more likely to occur during tightly-coupled collaboration (Fig. 9), such as when working on the same panel together (G7, G9) or to adjust visualisation size of a central 3D visualisation (G6). For occurrences outside of tightly-coupled collaboration, it was either during presentation in FET, or at the start and end of the task when groups manually reset the room of all visualisations (Fig. 10). Some participants, particularly those in Part A, would instead suggest or direct the owner to perform actions for them, for example P6: "Along the longitude latitude graph, colour by result on the map..." P4: "...let me just do it for you...", but three participants were willing to physically reach over someone else to modify a panel/visualisation. When not manipulating visualisation properties, participants made constant use of the pointer (accessed through details on demand) during their discussion to highlight specific parts of visualisations. Ultimately, objects would belong to the participant who created it, with others avoiding interacting with said objects except during tightly-coupled collaboration—and in certain circumstances such as those stated in Stage 5. Only in G5 was this 'ownership' transferred: one instance was implicit with P15 effectively taking control of P13's visualisation during group work, the other explicit with P15 offering P13 to use his panel.

Excluding talking over each other, all interruptions were accidents resulting from the shared nature of the system. For example, P6 being surprised when P5 starts to brush unannounced (as per Alice); P9 accidentally erasing P8's selection when performing opposing actions simultaneously; and P28 creating a very wide 3D visualisations which protrudes into P30's workspace. However, in all cases participants made little to no fuss. In fact, we noticed some would change their behaviour to minimise/avoid conflict, such as P9 only using private brushing after accidentally erasing someone else's selection in a previous task.

[Stage 8] Presentation. *For presentation, they decide to organise their visualisations on a shared wall, clustering their visualisations by author. Carol however places some of her 3D visualisations close to the wall, with the main one left on the table. With the remotely controlled*




Fig. 11. Actual examples of presentation styles in FET: P22 (in red) explaining insights on 3D visualisations without accounting for the audience’s perspective, with his back turned towards the audience (in black) (left); P25 using the marker to highlight a trend on a visualisation while exhibiting proper body language to the audience (right).

audience present, they take turns presenting their findings. Alice grabs a 2D visualisation off the wall and brings it closer to the audience, describing it while maintaining eye-contact and using gestures for each visualisation. Bob instead stands a distance from his large 2D visualisations, using a laser pointer to point during his presentation and periodically turning to the audience. Carol moves to each of her 3D visualisations, using the marker to highlight trends and clusters of interest. However, she pays no attention to the audience, and neglects to account for accurate viewing angles.

Groups presented an average of 7.5 findings ($M = 8$, $SD = 2.46$) to the audience during FET. Two groups (G1, G9) took time to lay out their findings neatly along a shared set of walls. Groups that did not do this generally already had visualisations grouped in each workspace—particularly for those with planar layouts. Of the 17 who used egocentric layouts, six switched to planar layouts for presentation, nine only had a few visualisations and mainly used their panels, while the remaining two left their (mostly 3D) visualisations in the centre of the room. G6 was a major exception with many visualisations scattered around the area—a result of their analysis involving many large 3D visualisations. As a result, P16 lead his group’s presentation for the majority of their findings. All other groups took turns presenting (excluding P1, P2, P5 who did not present). While no presentation had any overarching story, some focused on a specific theme as a result of their division of work, such as P9 focusing on most expensive properties and P22 with comparison between houses, units, and apartments.

While presenting, 16 used gestural hand movements (as per Alice), six used the laser pointer (as per Bob), four used a mix of the two, and one relied on the marker (as per Carol). While the majority in Part B used gestures (3 in Part A, 13 in Part B), this may be because the pointer is disabled when pointing at 3D visualisations. We also observed that 15 participants had poor awareness of the audience, paying no attention to it while presenting (Fig. 11 left). In contrast, 12 made sure to make eye-contact and engage the audience akin to real life (Fig. 11 right). Participants did not appear to be influenced by the presentation styles of others, such as P8 presenting in a very charismatic tone, and P9 following with a very static/rigid tone. While three participants grabbed visualisations off walls to bring them closer to the audience, none made sure to accommodate for the audience’s perspective for 3D visualisations, despite it being controlled as passively as possible. In fact, two participants (P22, P24) reported findings using a 3D visualisation as ‘two 2D scatterplots in one’ akin to the rolling dice metaphor [16], although neither rotated them to properly face the audience.

6 DISCUSSION

Our study revealed how groups behave in the FIESTA collaborative immersive analytics environment. We also observed users solve visual analytics tasks in novel ways made possible by the immersive environment. Based on our results, we highlight and discuss important takeaways which can help guide future work in this area.

6.1 System and Space Usage

Collaborative visual analytics is possible in immersive environments. Overall, our participants were capable of performing visual

analytics (Sec. 5.2). They explored a data set, authored visualisations, discovered insights, organised visualisations in the space around them, and presented their findings to others— doing so both independently and collaboratively through mixed-focus collaboration depending on the given context [24] (Stage 2). They were engaged with the tasks in VR for 45 minutes on average (plus 15 minutes training), and made use of most of the functionalities provided. However, FIESTA is a research prototype, and to observe more complex workflows functions such as textual queries and filters, as well as more advanced collaborative support such as privacy management or token-based protocols would be needed. We see this as a logical next step in collaborative immersive analytics to understand how groups utilise such features. Furthermore, we would like to observe how users use FIESTA in AR as compared to VR, particularly when integrating with real-world surfaces and proper face-to-face collaboration.

Users willingly made use of 3D visualisations, but these require UI designs that support their use. Although we did not emphasise the use of 3D visualisations during training, many of our participants chose to use and experiment with them (Sec. 5.1). Some groups naturally came up with interesting ways of viewing these 3D visualisations, such as standing inside of room-scale 3D visualisations for an egocentric view (within the data) up close to specific points [15, 29, 31], or looking from an axis-aligned viewpoint to view two dimensions, and then rotating 90 degrees to view the other two dimensions (similar to a manual *Scatterdice* [16]) (Stage 5). Participants created many more 3D than 2D visualisations (Fig. 5) as a result of the authoring interface we provided. When exploring the data through iterative visual design (e.g. changing axes dimensions or aesthetics), participants had to tear out 3D visualisations from the panel to inspect them from different angles. This resulted in the creation of redundant visualisations which may have impeded their flow, as compared to 2D visualisations which could be iterated through at a much faster rate (Sec. 5.1). While some of these behaviours are a consequence of the panel-based interface, they demonstrate the need for consistent UI designs that equally support both 2D and 3D visualisations—particularly when using 2D WIMP metaphors. That said, while there was evident interest, engagement, and experimentation with 3D visualisation by our participants, it is difficult to say from our study design whether the use of 3D visualisations had a positive impact on analytic performance. Further study would be required to see if this translates to improved task performance and if the interest persists as people become more experienced with the system.

The presence of 3D visualisations fundamentally changes view management and organisation. With only 2D visualisations available in Part A, half of participants placed them neatly on the wall while working (Stage 3). This careful placement required a little more physical effort than the egocentric layouts used by other participants and in related work [2, 11], but made their work easily visible and presentable to others at all times. In contrast, when 3D visualisations were available in Part B, almost all participants used egocentric layouts. This may be due to convenience, the wall affordance not being as suitable for 3D visualisations, or due to perspective issues with 3D making it impractical to accommodate others’ viewpoints at all times. The introduction of 3D also affected how 2D visualisations were organised, as the same egocentric layout was used for 2D visualisations as well. Overall, these are novel observations that suggest view management and organisation—not just analytic performance—can vary greatly based on the types of visualisations present, as the inclusion of 3D visualisations had a profound impact on how participants worked. Future collaborative immersive analytics systems may need to provide different placement tools to accommodate 2D versus 3D visualisations.

Users are influenced by environment configuration only if there is a tangible benefit. In contrast to previous work which would situate the task in an ‘endless void’ [2] or tightly couple interaction to a visible surface [6, 9, 21], the use of surfaces in FIESTA is a completely optional part of users’ workflows. As previously stated, participants took advantage of the walls in the environment to neatly organise their 2D visualisations throughout each task, whereas previous studies observed users doing so only for presentation [2]. In contrast, only a single group made careful use of the table (Stage 5), despite it being

a common metaphor in previous work [6, 8, 31]. This was not for a lack of awareness, as most participants did acknowledge the presence of the table (Stage 6). Simply put, our participants generally saw no benefit in using the table in its current state, as it adds an unnecessary constraint (e.g. height in space) in an environment where visualisations can be placed anywhere. This raises new questions as to the value of the tabletop metaphor in collaborative immersive analytics environments, as existing systems using physical tabletops have enforced their use [6, 21]. It may be that the role of tabletops is first and foremost to provide tangible and more comfortable interaction—something which the virtual table does not provide. However, introducing unique functionality may promote its use, such as an alternate authoring interface better suited for 3D visualisations.

6.2 Collaboration

Equally distributed interaction resources promotes parallel work and mixed-focus collaboration. We chose to provide each person their own authoring panel during the study. As a result, we observed that almost all groups utilised mixed-focus collaboration styles [24], either working individually and sharing findings throughout the task, or alternating between tightly-coupled and parallel work to try ideas on their own. Only in three tasks (G7 in DT2 and DT3, G9 in FET) did participants work closely with one another for the entire duration (Stage 2). This is consistent with previous findings that users work more in parallel when given the ability to do so [5]. While future work may evaluate the effects of having single versus multiple authoring interfaces in collaborative immersive environments, a factor to consider is the virtuality of VR/AR. Unlike setups with large physical displays, there is no additional cost for virtual displays and interfaces. While a few groups did shift to a tightly-coupled single display context, the ability for mixed-focus collaboration did appear to be invaluable, allowing groups to work in any manner that best suits them.

Territories in free-roaming immersive environments are spaces where users can work individually, defined by movement patterns and placement of artefacts. Movement of participants and placement of visualisation artefacts appeared to be tightly linked with territorial behaviour. These territories acted as an area where one could comfortably work individually without physical intrusion or disturbance (Fig. 7). Territories were not negotiated, but were typically defined by the initial placement of panels, and participants never entered a territory of another unless for tightly-coupled work (Stage 1). This behaviour is similar to that found on tabletops [32, 49, 50] and wall-sized displays [30], as these territories were oftentimes respected.

Participants did not interact with objects that did not belong to them. We observed that participants kept track of and respected ownership of visualisations and panels (Stage 7), refraining from interacting with artefacts which belonged to others. Instances where participants interacted with someone else's visualisations mainly occurred during tightly-coupled collaboration (Fig. 10). Some groups would direct the owner to perform some action for them, while others would simply reach over and do it themselves (Stage 7). The latter was always when the purpose was known to all [30, 49], such as making a visualisation easier for everyone to see, or when the group was already working closely together for an extended period of time. This behaviour may also have been influenced by participants being extra cautious of collisions in co-located VR however. While personal territories used similarly to that of previous work [47], group territories were seldom used, acting more as an area in the centre of the room to congregate and discuss rather than to share resources due to the virtuality of the environment. Overall, some behaviours are consistent with physical environments, and others unique to co-located immersive environments. Future work may seek to understand how users collaborate when physically remote, but virtually co-located.

Participants favoured workspace awareness over privacy. FIESTA contains numerous design choices to improve workspace awareness [25], such as the completely shared nature of the environment. Contrary to our expectations, many participants commented that it helped improve their awareness of others with only a few reporting it was

distracting, which allowed for mixed-focus collaboration based on each others' actions (Stage 4) to frequently transition to tightly-coupled collaboration [25, 28, 30]. The inclusion panel-based authoring tools in FIESTA showcased unique and interesting behaviours as well. While participants could turn their panels away from each other to achieve privacy, no one intentionally did this, with some groups instead trying to optimise their visibility of each other (Stage 1) or simply gave up privacy altogether by working on the same panel (Stage 2). It is apparent that this flexibility allows groups to balance the level of awareness and privacy they have, and that immersive environments have different considerations for awareness than conventional virtual workspaces [24]—namely that users can see much more of the workspace. With this in mind, future work may explore alternate forms for improving awareness or privacy management (e.g. [7]) for collaborative immersive analytics.

Ensuring proper perspective is challenging during presentation of 3D visualisations. While 2D visualisations can accurately be viewed from a wide range of angles, 3D visualisations are reliant on perspective in shared contexts—particularly during presentation. However, despite an audience embodied by avatars, presenters did not accommodate well for their audiences' viewpoint with 3D visualisations (Stage 8). This was not for a lack of willingness to engage with the audience, as we observed them actively use gestures and eye-contact. To the best of our knowledge, only a study by Batch *et al.* had involved presentation to a third-party in an immersive analytics environment [2]. While participants did account for viewpoint, the audience had the same view as the participant themselves, rather than using a separate embodied audience. Future work may explore how to facilitate this presentation, whether through an embodied audience with agency, or simply by mirroring the presenter's point of view.

7 CONCLUSION AND FUTURE DIRECTIONS

Our FIESTA system and study is the first to test the effect of surfaces and spaces on visualisation tasks performed by groups collaborating in a room-sized immersive environment. We observed novel behaviours that are unique to collaborative immersive analytics:

- Participants reported that the shared VR environment was useful in maintaining workspace awareness and to share findings with each other.
- The authoring panel metaphor allowed them to easily create 2D visualisations and organise them neatly along the virtual walls not only for presentation, but in some cases for exploration and analysis,
- 3D visualisations were frequently created, but were freely suspended in convenient locations in space rather than on top of the virtual table.

These surfaces, while optional for users to use, were inherently predefined. A future direction could be to explore emergent workspace configurations when users are able to arrange their own surfaces. Our study also confirms previous results from collaborative work studies in non-immersive environments:

- Groups organically divided the shared virtual environment into equally sized territories which were used for individual work.
- They transitioned to tightly-coupled collaboration either by observing each others' work or through discussion regardless of physical distance.
- The use of avatars and pointers facilitated collaboration, with deixis allowing participants to work while up close or when far apart.

However, a potential factor in how groups worked was the use of co-located VR, as participants may have been apprehensive to perform certain actions due to the risk of collision. FIESTA already supports and has been recently tested with distributed collaboration, thus a natural progression to our work is to explore how users collaborate in distributed immersive environments. Particularly, many social protocols may not exist when collaborators are physically remote, while offering new possibilities such as shared perspective [12].

ACKNOWLEDGMENTS

Thank you to our study participants and anonymous reviewers for their feedback. This research was supported under the Australian Research Councils Discovery Projects funding scheme (project number DP180100755) and by an Australian Government Research Training Program (RTP) Scholarship.

REFERENCES

- [1] X. Amatriain, J. Kuchera-Morin, T. Hollerer, and S. T. Pope. The allosphere: Immersive multimedia for scientific discovery and artistic exploration. *IEEE MultiMedia*, 16(2):64–75, April 2009. doi: 10.1109/MMUL.2009.35
- [2] A. Batch, A. Cunningham, M. Cordeil, N. Elmqvist, T. Dwyer, B. H. Thomas, and K. Marriott. There is no spoon: Evaluating performance, space use, and presence with expert domain users in immersive analytics. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):536–546, Jan 2020. doi: 10.1109/TVCG.2019.2934803
- [3] B. Bell, S. Feiner, and T. Höllerer. View management for virtual and augmented reality. In *Proceedings of the 14th Annual ACM Symposium on User Interface Software and Technology*, UIST '01, pp. 101–110. ACM, New York, NY, USA, 2001. doi: 10.1145/502348.502363
- [4] M. Billinghurst, M. Cordeil, A. Bezerianos, and T. Margolis. *Collaborative Immersive Analytics*, pp. 221–257. Springer International Publishing, Cham, 2018. doi: 10.1007/978-3-030-01388-2_8
- [5] J. P. Birnholtz, T. Grossman, C. Mak, and R. Balakrishnan. An exploratory study of input configuration and group process in a negotiation task using a large display. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '07, pp. 91–100. ACM, New York, NY, USA, 2007. doi: 10.1145/1240624.1240638
- [6] S. Butscher, S. Hubenschmid, J. Müller, J. Fuchs, and H. Reiterer. 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*, CHI '18, pp. 90:1–90:12. ACM, New York, NY, USA, 2018. doi: 10.1145/3173574.3173664
- [7] A. Butz, C. Beshers, and S. Feiner. Of vampire mirrors and privacy lamps: Privacy management in multi-user augmented environments. In *Proceedings of the 11th Annual ACM Symposium on User Interface Software and Technology*, UIST 98, p. 171172. Association for Computing Machinery, New York, NY, USA, 1998. doi: 10.1145/288392.288598
- [8] M. Cavallo, M. Dholakia, M. Havlena, K. Ocheltree, and M. Podlaseck. Dataspace: A reconfigurable hybrid reality environment for collaborative information analysis, 2019.
- [9] M. Cavallo, M. Dolakia, M. Havlena, K. Ocheltree, and M. Podlaseck. Immersive insights: A hybrid analytics system for collaborative exploratory data analysis. In *25th ACM Symposium on Virtual Reality Software and Technology*, VRST 19. Association for Computing Machinery, New York, NY, USA, 2019. doi: 10.1145/3359996.3364242
- [10] M. Cordeil, A. Cunningham, B. Bach, C. Hurter, B. H. Thomas, K. Marriott, and T. Dwyer. Iatk: An immersive analytics toolkit. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 1–10. IEEE, March 2019.
- [11] M. Cordeil, A. Cunningham, T. Dwyer, B. H. Thomas, and K. Marriott. Imaxes: Immersive axes as embodied affordances for interactive multivariate data visualisation. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*, UIST '17, pp. 71–83. ACM, New York, NY, USA, 2017. doi: 10.1145/3126594.3126613
- [12] M. Cordeil, T. Dwyer, K. Klein, B. Laha, K. Marriott, and B. H. Thomas. Immersive collaborative analysis of network connectivity: Cave-style or head-mounted display? *IEEE Transactions on Visualization and Computer Graphics*, 23(1):441–450, Jan 2017. doi: 10.1109/TVCG.2016.2599107
- [13] C. Cruz-Neira, D. J. Sandin, and T. A. DeFanti. Surround-screen projection-based virtual reality: The design and implementation of the cave. In *Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '93, pp. 135–142. ACM, New York, NY, USA, 1993. doi: 10.1145/166117.166134
- [14] T. A. DeFanti, G. Dawe, D. J. Sandin, J. P. Schulze, P. Otto, J. Girado, F. Kuester, L. Smarr, and R. Rao. The starcave, a third-generation cave and virtual reality optiportal. *Future Generation Computer Systems*, 25(2):169–178, 2009. doi: 10.1016/j.future.2008.07.015
- [15] C. Donalek, S. G. Djorgovski, A. Cioc, A. Wang, J. Zhang, E. Lawler, S. Yeh, A. Mahabal, M. Graham, A. Drake, S. Davidoff, J. S. Norris, and G. Longo. Immersive and collaborative data visualization using virtual reality platforms. In *2014 IEEE International Conference on Big Data (Big Data)*, pp. 609–614, Oct 2014. doi: 10.1109/BigData.2014.7004282
- [16] N. Elmqvist, P. Dragicevic, and J. Fekete. Rolling the dice: Multidimensional visual exploration using scatterplot matrix navigation. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1539–1148, Nov 2008. doi: 10.1109/TVCG.2008.153
- [17] B. Ens, R. Finnegan, and P. P. Irani. The personal cockpit: A spatial interface for effective task switching on head-worn displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '14, pp. 3171–3180. ACM, New York, NY, USA, 2014. doi: 10.1145/2556288.2557058
- [18] B. Ens, E. Ofek, N. Bruce, and P. Irani. *Shared Façades: Surface-Embedded Layout Management for Ad Hoc Collaboration Using Head-Worn Displays*, pp. 153–176. Springer International Publishing, Cham, 2016. doi: 10.1007/978-3-319-45853-3_8
- [19] A. Febretti, A. Nishimoto, T. Thigpen, J. Talandis, L. Long, J. D. Pirtle, T. Peterka, A. Verlo, M. Brown, D. Plepys, D. Sandin, L. Renambot, A. Johnson, and J. Leigh. CAVE2: a hybrid reality environment for immersive simulation and information analysis. In M. Dolinsky and I. E. McDowall, eds., *The Engineering Reality of Virtual Reality 2013*, vol. 8649, pp. 9 – 20. International Society for Optics and Photonics, SPIE, 2013. doi: 10.1117/12.2005484
- [20] S. Feiner, B. MacIntyre, M. Haupt, and E. Solomon. Windows on the world: 2d windows for 3d augmented reality. In *Proceedings of the 6th Annual ACM Symposium on User Interface Software and Technology*, UIST '93, pp. 145–155. ACM, New York, NY, USA, 1993. doi: 10.1145/168642.168657
- [21] J. A. W. Filho, C. M. D. S. Freitas, and L. Nedel. Comfortable immersive analytics with the virtualdesk metaphor. *IEEE Computer Graphics and Applications*, 39(3):41–53, May 2019. doi: 10.1109/MCG.2019.2898856
- [22] I. for the Future (IFTF). Onvr - collaborative vr data visualization prototype. <https://www.youtube.com/watch?v=USiF073trok>, 2018. [Accessed: 2020-01-08].
- [23] D. Gerber and D. Bechmann. The spin menu: a menu system for virtual environments. In *IEEE Proceedings. VR 2005. Virtual Reality, 2005.*, pp. 271–272, March 2005. doi: 10.1109/VR.2005.1492790
- [24] C. Gutwin and S. Greenberg. Design for individuals, design for groups: Tradeoffs between power and workspace awareness. In *Proceedings of the 1998 ACM Conference on Computer Supported Cooperative Work*, CSCW 98, p. 207216. Association for Computing Machinery, New York, NY, USA, 1998. doi: 10.1145/289444.289495
- [25] C. Gutwin and S. Greenberg. A descriptive framework of workspace awareness for real-time groupware. *Comput. Supported Coop. Work*, 11(3):411446, Nov. 2002. doi: 10.1023/A:1021271517844
- [26] J. Heer and M. Agrawala. Design considerations for collaborative visual analytics. *Information visualization*, 7(1):49–62, 2008.
- [27] J. Heer and B. Shneiderman. Interactive dynamics for visual analysis. *Queue*, 10(2):30, 2012.
- [28] P. Isenberg, D. Fisher, S. A. Paul, M. R. Morris, K. Inkpen, and M. Czerwinski. Co-located collaborative visual analytics around a tabletop display. *IEEE Transactions on Visualization and Computer Graphics*, 18(5):689–702, May 2012. doi: 10.1109/TVCG.2011.287
- [29] A. Ivanov, K. Danyluk, C. Jacob, and W. Willett. A walk among the data. *IEEE Computer Graphics and Applications*, 39(3):19–28, May 2019. doi: 10.1109/MCG.2019.2898941
- [30] M. R. Jakobsen and K. Hornbæk. Up close and personal: Collaborative work on a high-resolution multitouch wall display. *ACM Trans. Comput.-Hum. Interact.*, 21(2):11:1–11:34, Feb. 2014. doi: 10.1145/2576099
- [31] M. Kraus, N. Weiler, D. Oelke, J. Kehrer, D. A. Keim, and J. Fuchs. The impact of immersion on cluster identification tasks. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):525–535, Jan 2020. doi: 10.1109/TVCG.2019.2934395
- [32] R. Kruger, S. Carpendale, S. D. Scott, and S. Greenberg. Roles of orientation in tabletop collaboration: Comprehension, coordination and communication. *Comput. Supported Coop. Work*, 13(5-6):501–537, Dec. 2004. doi: 10.1007/s10606-004-5062-8
- [33] E. Kruijff, J. E. Swan, and S. Feiner. Perceptual issues in augmented reality revisited. In *2010 IEEE International Symposium on Mixed and Augmented Reality*, pp. 3–12, Oct 2010. doi: 10.1109/ISMAR.2010.5643530
- [34] A. Kunert, T. Weissker, B. Froehlich, and A. Kulik. Multi-window 3d interaction for collaborative virtual reality. *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–1, 2019. doi: 10.1109/TVCG.2019.2914677
- [35] J. J. LaViola Jr, E. Kruijff, R. P. McMahan, D. Bowman, and I. P. Poupyrev. *3D user interfaces: theory and practice*. Addison-Wesley Professional, 2017.
- [36] B. Lee, M. Cordeil, A. Prouzeau, and T. Dwyer. Fiesta: A free roaming collaborative immersive analytics system. In *Proceedings of the 2019 ACM International Conference on Interactive Surfaces and Spaces*, ISS

- 19, p. 335338. Association for Computing Machinery, New York, NY, USA, 2019. doi: 10.1145/3343055.3360746
- [37] C. Liu, O. Chapuis, M. Beaudouin-Lafon, and E. Lecolinet. Shared interaction on a wall-sized display in a data manipulation task. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, CHI '16, pp. 2075–2086. ACM, New York, NY, USA, 2016. doi: 10.1145/2858036.2858039
- [38] G. E. Marai, A. G. Forbes, and A. Johnson. Interdisciplinary immersive analytics at the electronic visualization laboratory: Lessons learned and upcoming challenges. In *2016 Workshop on Immersive Analytics (IA)*, pp. 54–59, March 2016. doi: 10.1109/IMMERSIVE.2016.7932384
- [39] K. Marriott, J. Chen, M. Hlawatsch, T. Itoh, M. A. Nacenta, G. Reina, and W. Stuerzlinger. *Immersive Analytics: Time to Reconsider the Value of 3D for Information Visualisation*, pp. 25–55. Springer International Publishing, Cham, 2018. doi: 10.1007/978-3-030-01388-2_2
- [40] J. P. McIntire and K. K. Liggett. The (possible) utility of stereoscopic 3d displays for information visualization: The good, the bad, and the ugly. In *2014 IEEE VIS International Workshop on 3DVis (3DVis)*, pp. 1–9, Nov 2014. doi: 10.1109/3DVis.2014.7160093
- [41] MSI. Mr ready. <https://mr.msi.com/Backpacks/vrone/>. Accessed: 2020-07-20.
- [42] H. Nguyen, B. Ward, U. Engelke, B. Thomas, and T. Bednarz. Collaborative data analytics using virtual reality. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 1098–1099, March 2019. doi: 10.1109/VR.2019.8797845
- [43] A. Prouzeau, A. Bezerianos, and O. Chapuis. Evaluating multi-user selection for exploring graph topology on wall-displays. *IEEE Transactions on Visualization and Computer Graphics*, 23(8):1936–1951, Aug 2017. doi: 10.1109/TVCG.2016.2592906
- [44] A. Prouzeau, A. Bezerianos, and O. Chapuis. Trade-offs between a vertical shared display and two desktops in a collaborative path-finding task. In *Proceedings of Graphics Interface, GI '17*. CHCCS, May 2017.
- [45] P. Reipschläger and R. Dachselt. Designar: Immersive 3d-modeling combining augmented reality with interactive displays. In *Proceedings of the 2019 ACM International Conference on Interactive Surfaces and Spaces*, ISS 19, p. 2941. Association for Computing Machinery, New York, NY, USA, 2019. doi: 10.1145/3343055.3359718
- [46] G. Robertson, M. van Dantzich, D. Robbins, M. Czerwinski, K. Hinckley, K. Risdén, D. Thiel, and V. Gorokhovskiy. The task gallery: A 3d window manager. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '00, pp. 494–501. ACM, New York, NY, USA, 2000. doi: 10.1145/332040.332482
- [47] S. D. Scott, M. S. T. Carpendale, and K. M. Inkpen. Territoriality in collaborative tabletop workspaces. In *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work*, CSCW '04, pp. 294–303. ACM, New York, NY, USA, 2004. doi: 10.1145/1031607.1031655
- [48] T. Sun, Y. Ye, I. Fujishiro, and K. Ma. Collaborative visual analysis with multi-level information sharing using a wall-size display and see-through hmds. In *2019 IEEE Pacific Visualization Symposium (PacificVis)*, pp. 11–20, April 2019. doi: 10.1109/PacificVis.2019.00010
- [49] A. Tang, M. Tory, B. Po, P. Neumann, and S. Carpendale. Collaborative coupling over tabletop displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '06, pp. 1181–1190. ACM, New York, NY, USA, 2006. doi: 10.1145/1124772.1124950
- [50] J. C. Tang. Findings from observational studies of collaborative work. *International Journal of Man-Machine Studies*, 34(2):143 – 160, 1991. Special Issue: Computer-supported Cooperative Work and Groupware. Part 1. doi: 10.1016/0020-7373(91)90039-A