Analyzing User Behaviour Patterns in a Cross-Virtuality Immersive Analytics System

Mohammad Rajabi Seraji, Parastoo Piray, Vahid Zahednejad, and Wolfgang Stuerzlinger

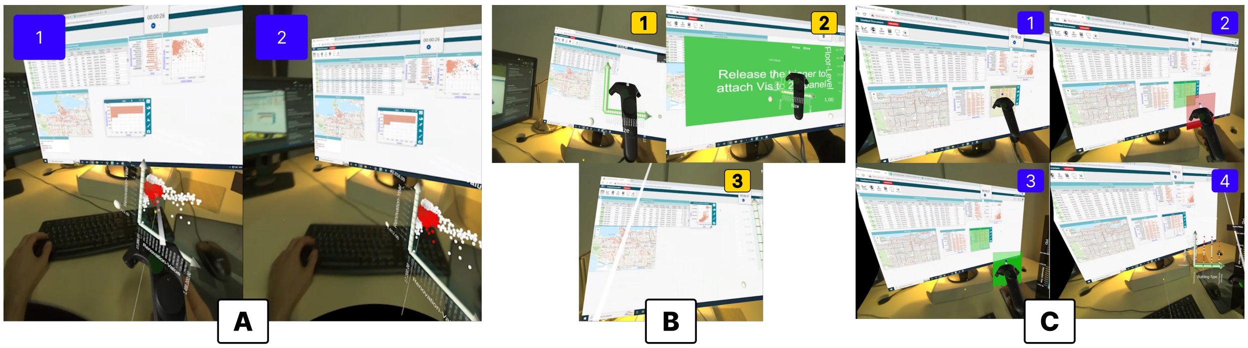

Fig. 1: A) Synchronized brushing and linking between desktop and AR, B) Moving a graph from AR to desktop, and C) vice versa.

Abstract— Recent work in immersive analytics suggests benefits for systems that support work across both 2D and 3D data visualizations, i.e., cross-virtuality analytics systems. Here, we introduce HybridAxes, an immersive visual analytics system that enables users to conduct their analysis either in 2D on desktop monitors or in 3D within an immersive AR environment – while enabling them to seamlessly switch and transfer their graphs between modes. Our user study results show that the cross-virtuality sub-systems in HybridAxes complement each other well in helping the users in their data-understanding journey. We show that users preferred using the AR component for exploring the data, while they used the desktop to work on more detail-intensive tasks. Despite encountering some minor challenges in switching between the two virtuality modes, users consistently rated the whole system as highly engaging, user-friendly, and helpful in streamlining their analytics processes. Finally, we present suggestions for designers of cross-virtuality visual analytics systems and identify avenues for future work.

Index Terms—Immersive Analytics, Cross-virtuality Analytics, Visualization, Human-computer Interaction

1 Introduction

Visual Analytics (VA) focuses on analytical reasoning facilitated by interactive visual interfaces [52]. Desktop VA tools like Tableau or Power BI offer powerful user interfaces for data exploration and analysis but have a steep learning curve [27]. Most current VA tools use the traditional *Windows, Icons, Menus, and Pointer* (WIMP) metaphor [2]. Post-WIMP data exploration tools have demonstrated substantial improvements on a user's data understanding [32,46]. Here we explore a system that bridges both metaphors. One common class of post-WIMP interfaces uses immersive technologies, such as *Virtual, Augmented, Mixed, or Extended Reality* (VR/AR/MR/XR). *Immersive Analytics* (IA) uses engaging, embodied analysis tools to support data understanding and decision-making through immersive technologies, multi-sensory presentation, data physicalization, and natural interfaces [17, 19].

IA was first proposed in the 70's [22], but current affordable VR hardware has re-increased interest in it. IA systems enable users to represent multi-dimensional data in three-dimensional (3D) space, and better exploit human vision capabilities and leverage embodied interaction [6, 23, 51]. As many application areas, such as aerospace, industry, education, and cultural heritage, deal with complex data in domains that can benefit from 3D digital twins, i.e., virtual replicas of reality, IA is of increasing interest [24].

• *Mohammad Rajabi Seraji, Parastoo Piray, Vahid Zahednejad, and Wolfgang Stuerzlinger are with Simon Fraser University. E-mails: {mrajabis | ppa35 | vahid_zahednejad | w.s}@sfu.ca.*

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.xxxx/TVCG.201x.xxxxxxx

Previous work has shown that immersive systems have the potential to improve the sense-making process, and that using a combination of 2D displays and 3D MR environments could be beneficial [9, 30]. Many researchers and practitioners have thus created tools to visualize and analyze data in VR, AR, or XR. Millais et al. [41] and Hubenschmid et al. [30] demonstrated the advantages of using immersion for data exploration in VR. Also, using IA systems can improve a user's ability to find connections in the data [24, 25]. To realize the full potential of IA, some challenges need to be overcome. Snowdon et al. [50] identified important ones as 1) providing users with a shared context for the data, 2) showing them role-specific views of the data, and 3) keeping users aware of what is happening in the system. More recent work [8, 9] adds that 4) providing users with external tools for note-taking and recording insights is a helpful addition to any IA system.

Here, we use the terms 'virtualities', 'modes of virtuality', and 'modes' in an interchangeable manner. All these terms refer to where the user resides on the *reality-virtuality continuum* (RVC) [40]. In this work, we focus only on two such modes of virtuality: immersive AR ("3D mode") and a non-immersive desktop system ("2D mode").

1.1 Motivation and Research Questions

In the real world humans live in 3D spaces and naturally interact with 3D objects. Research has provided objective evidence for the benefits of seeing a virtual 3D data visualization as 3D objects in a real or virtual 3D space and not "just" as a 2D projection [1].

Building on the results of previous studies that suggest benefits for *cross-virtuality analytics* (XVA) systems in the user's sense-making and data-understanding, we investigate such an XVA system that allows the user to simultaneously work at two different points along the RVC, to identify the effect on users and their analysis process, which will also inform the design of future XVA systems.

Only a handful of previous VA systems support interaction in both 2D and virtual environments (VEs) [9, 20, 43]. Even fewer support dynamic visualization creation/modification. But how will users use and perform in a system where they seamlessly interact with their data on both a desktop (non-immersive 2D mode) and in AR (immersive 3D mode) and can instantly switch between these modes? The literature [26, 54] suggests that enabling users to switch between 2D and 3D helps them get a better overview of the data, extends their workspace, and facilitates the data-understanding process, but only for static visualizations. Thus, we need to investigate what happens if users are able to rapidly create graphs while being able to switch between 2D and 3D modes. Our overall goal was thus to investigate: "How does the ability to operate at two different points along the RVC and transitioning between them affect the *interactive* analysis process in IA systems?" Consequently, we targeted the following *research questions* (RQs):

- **RQ1:** How do people use a cross-virtuality IA system? Specifically, what behaviours do they exhibit in each of the two subsystems (AR and desktop)? When do they switch between the different modes of virtuality? And how often do they switch between the two different modes of virtuality?
- **RQ2:** How will the users perform (in terms of task completion time and error rate) in an XVA system, compared to a standalone desktop counterpart?

Here, both "hybrid" and "cross-virtuality" interchangeably refer to a system that can operate in multiple modes of virtuality.

1.2 Contributions

To address the questions mentioned above, we created HybridAxes, an IA prototype. Our contributions are as follows:

- The design of HybridAxes, a mixed reality IA system that builds on previous work and which allows the user to interoperate along two different points along the RVC.
- Evaluating HybridAxes in a mixed-method between-subject study to discover some of the differences between a desktop VA and cross-virtuality IA system.
- Identifying usage patterns for a cross-virtuality IA system through observational sessions and open-ended interviews.
- A list of design considerations for both the desktop and AR side of cross-virtuality IA systems.

To support open science and enable the replication of our work by others, we will open-source our prototype, to enable future IA systems to build on previous ones, within the open-source community.

2 Background and Related Work

Building on the IA work already mentioned in the introduction, we review other relevant work here.

2.1 Interactive Virtual Multi-dimensional Graphs

High-dimensional or complex datasets quickly result in cluttered displays, and interaction is then often required to filter and interactively explore the data, regardless of the visualization technique [17]. To highlight a subset of data points across graphs, focus+context techniques such as brushing and linking of views were introduced. Drag-and-drop to modify the visual encodings is now fairly common in BI tools, such as Tableau. Such interfaces remain constrained by the particular set of visual idioms supported by the tools. Only few systems, e.g., VizInteract [10], avoid these restrictions through new interaction metaphors and methods. Researchers also explored direct interaction with multidimensional graphs in VEs using a mouse [29, 55] or touch [33, 47].

Systems in this category all support the creation and manipulation of graphs in a VE. ImAxes [15] uses the metaphor of an axis to embody data. By pulling them from the virtual axes shelf/panel users can creates a histogram, and then further interact with these virtual axes like physical objects in space and create different graphs by attaching axes differently relative to each other. To move a graph, users grab its substrate (the view area) and drag it. Users can pull an axis of a graph away, or delete a whole graph by simply "throwing it away". Each axis has a range slider for simple filtering. More complex filtering is limited, as users have to construct new graphs to filter existing ones. This means that that users cannot filter more complex graphs without constructing many auxiliary ones. Also, the authors did not explore the effects of using such a filtering system on the overall visual analysis process. ImAxes also provided (partially implemented) menu-based interactions to change the colour or size of data points.

To support casual collaborative VA, Uplift [20] combined a 3D model on a tabletop display with tangible interaction and mid-air AR data visualization. Uplift still relied on menus for filtering and manipulation of the data attributes. U2vis [43] aims to reduce perceptional issues on large displays in data exploration and also facilitate the analysis of dense data sets. While the authors did not validate this claim, they suggested combining their hybrid 2D/AR displays with interactive graph creation. Wagner et al. [21] showed a space-time cube with embodied interaction methods to rotate, scale, and query the visualization. Embodied Axes [13] presented a 3D tangible controller that permits users to perform more precise 3D selections on immersive 3D graphs than with traditional controllers.

Wang et al. [54] extended desktop VA with AR display for particle physicists. In their work, the synchronized desktop and AR systems used similar interfaces, both controlled by the mouse and keyboard. They reported that participants used the AR view to walk around the data, grasp the 3D nature of the data, and create differently-configured mini AR displays to complement their desktop experience. Their users preferred the hybrid system and the consistent input reduced their mental overhead and learning time, albeit at the cost of having to use (relatively) un-intuitive 6-DOF interaction mappings with the mouse.

Still, there is currently a lack of intuitive data filtering methods for IA systems, an important aspect of usable IA systems [17, 19]. There are also few embodied visualization manipulation methods that allow the users to feel both immersed and in control. Also, many aspects of embodied data visualization interactions such as storage, retrieval, undo, redo, and other low-level interactions remain at least under-explored.

2.2 Cross-virtuality Analytics

McIntire et al. [39] outlined that stereoscopic display for information visualization still has limitations. While HMDs may be convenient for exploring spatial or multidimensional data, they can fall short in displaying detailed statistical and abstract information, which is instead better handled by 2D graphs [51]. Thus, research has recently explored XR/MR applications that offer interoperability between desktop and immersive 3D environments [24].

Integrating traditional 2D desktop systems and immersive ones promises to addressing human-centred challenges such as effective interaction and collaborative VA [34]. Frohler et al. [24] summarized such integration as *"Cross-virtuality analytics (XVA)"* and defined it as creating "systems for data visualization and analysis that seamlessly integrate different visual metaphors and devices along the entire RVC to support multiple users with transitional and collaborative interfaces." Others call the same concept hybrid reality or cross-reality analytics. XVA also includes IA systems with interfaces that allow users to interact in multiple modalities (AR, VR, desktop) and transition between them [28]. Recent research identified that cross-virtuality transitions could open new possibilities for interacting with multi-dimensional data across the RVC [45].

Some work demonstrates the potential of dynamically switching between a 2D desktop display and a virtual 3D environment. For instance, DataSpace and Immersive Insights [9] supported collaborative analysis of spatial datasets either as high-resolution tabular information (displayed on 2D screens) or as 3D representations of high-dimensional data (visualized in AR). Their work hinted at the possible advantages of hybrid reality in the sense-making process, but did not provide concrete evidence. Using a gesture-based system, VITA [4] showed that switching graphs between 2D and 3D could be beneficial to the collaborative aspect of immersive excavations. However, they did not explore the effects of such interactions in data visualization scenarios.

Wang et al. [54] showed that using an AR extension of a desktop could help physicists better understand particle collision effects, through a switch from 2D screens to an AR environment. Gesslein et al. [26] explored the effects of extending a small tablet interface to 3D space for spreadsheet editing and data visualization. They reported benefits due to the ability to create new constructs in 3D, while still showing the original 2D spreadsheet on the tablet. Lee et al. introduced FIESTA [35] which allows multiple users to freely move around in a shared room-sized environment and collaboratively gather insights from immersive visualizations. Their studies [37] suggested that surfaces were coupled with the type of graph used. Thus users often used walls to organize 2D graphs but use the space around them in VR for positioning 3D graphs.

The authors of the "Grand Challenges in Immersive Analytics" [19] also identified a need for IA systems with 2D desktop interoperability. They also suggested that transitions between these environments should cause meaningful transformations in data visualizations. Lee et al. [36] introduced a new design space for visualization transformations between 2D and 3D spaces, suggested guidelines for this novel hybrid design space, and that meaningful 2D/3D transformations would benefit IA users and help them in the sense-making process. Yet, none of these guidelines were validated. More recent work [31] explored the core advantages and challenges of such transitional interfaces. Lee et al. studied the effects of spatially situated visualizations in a collaborative VE [35, 37]. They showed promising results for spatial placement, but they did not explore the potential of a hybrid approach, i.e., combining it with an AR or desktop system. More recently, Saffo et al. [48] explored the effect of diverging devices on group awareness in a collaborative asymmetric system across desktop and VR platforms. Within the same domain, Tong et al. [53] showed that a well-designed asymmetric collaboration system could be as effective as a symmetric system.

3 HybridAxes

To answer the questions put forward in subsection 1.1, we needed a functional IA system to perform the user studies. Here, we present the design of HybridAxes: its features, interaction modalities, and supported types of graphs.

3.1 System Design

We focused on creating an hybrid, XVA system that can inter-operate between a (virtual) 2D desktop and an immersive AR environment, to enable users to rapidly create graphs using embodied metaphors [10, 15], but also in the context of their existing real environment. We drew inspiration from Immersive Insights [9], DynSpace [18], VizInteract [10], and Reski et al.'s system [44]. Further, cross-virtuality experiences that involve multiple platforms simultaneously can easily feel inconsistent and disconnected across these platforms. Our overarching design intent was thus to create two sub-systems that elicit a sense of continuity and cohesion in the user. We based this decision on previous work's findings [19, 36, 54].

Building on ImAxes' "data equals axis" metaphor, we wanted to extend its capabilities, support seamless transitions between virtualities, including cross-virtuality brushing and linking, drill-down, details-on-demand, and filtering. All these operations should follow standard usability guidelines, ideally in embodied form.

To ensure that the users would choose their visualization's virtuality mode (desktop vs. AR) based on the perceived performance, comfort, and properties of the current graph, we aimed to create an environment with a comparable experience across both sub-systems. This way, we tried to discourage users from disregarding a specific mode due to a lack of features and/or performance, empowering them to follow their VR process in whatever way they see fit. However, this does not mean that authoring *every* type of graph has to be equally easy or difficult on both sides. For example, consider that the immersive 3D space is a more natural fit for authoring a 3D graph [7, 17], so this should be easier in 3D. Also, interactions that require detailed manipulation and/or text entry, like creating formulas and annotations, are easier on a desktop using conventional input methods, i.e., keyboard and mouse. [16]. This is partly due to the challenges of mid-air 3D interaction, which can make accurate interactions difficult in AR/VR [3, 16].

Since the majority of today's work spaces still involve real rooms, it is reasonable to bring virtual content, such as data visualizations, into the user's real-world work environment through AR technologies. This is often referred to as offering users an XR/MR experience. In our case, we wanted to offer such an XR experience with matching functionalities on both desktop and in the AR environment, while allowing users to move freely between the two. Thus, it is important that our system works well in both of these environments.

3.2 HybridAxes Implementation

To create HybridAxes we extended the 2017 version of ImAxes [15] from GitHub, used IATK [14], some of DXR's rendering techniques [49], and also added support for UnityXR. We used a Varjo XR-3 as HMD. We decided to support the following graphs in HybridAxes: histograms, 2D/3D scatterplots, SPLOMS, PCPs, and line graphs (time-series), which supports the most common VA scenarios. We added support for dynamic modification of color and size encodings, by adding two interactive drop targets in the left and right half of the substrate of a visualization (seen in Figure 3), respectively. The legend of the graph shows these encodings and also allows the user to modify them. For each axis, we also made sure that the tick marks were readable. Furthermore, we added a history stack to support logging and to enable undo/redo. To align with users' existing mental models and inspired by the visual representation of undo/redo as clockwise/counter-clockwise arrow icons, undo/redo actions are initiated by a counter-clockwise or clockwise swipe on the controller touchpad, respectively.

To improve the framerate, we created a dynamic mesh that holds all the data for a graph, encoded as points, and render it through appropriate shaders [12, 38]. This functionality supports brushing and linking through masking textures, so that most computation happens on the GPU. To accelerate details-on-demand display for data points, we also created a set of shaders to identify the point closest to the controller. While we experimented with bare-hand interaction, pilot participants consistently complained about tracking issues, including unreliable pinch recognition especially when rotating objects in 3D, and also about needing to have the hands up in front of the HMD.

3.3 Transitioning Graphs between Virtualities

To support the transition of graphs between the desktop and the 3D space,we created "panels," which can be mapped to real wall, desk, or monitor surfaces (supported by Varjo's SDK to detect real surfaces) or serve as virtual floating panels (in AR or VR). Such panels can serve different roles as, e.g., a data dimensions shelf (as shown in Figure 7C) or a monitor panel (as shown in Figure 7B). Each of these panels supports transitions between the mentioned virtualities, visualized through appropriate animations upon interaction (e.g the red or squares in Figure 1C.2 and Figure 1C.3 respectively). Users can then pull a 2D graph from the desktop panel to transition it into 3D and vice versa. To assist the user, we added visual and haptic feedback for such interactions, so that users know when a transition will happen.

We only manipulate a graph post-transition when users move a higher-dimensional graph to a lower-dimensional space. For instance, when they transfer a 3D scatterplot to the desktop, the third dimension of this graph will disappear. If available, we then use a "free" graph channel, such as colour or size, to accommodate the third dimension on the desktop. Figure 1 B and C show examples. If a graph has already attributes attached to both size and colour channels, we visually signal to the user that this transition is going to destroy the third dimension. Then, if the user proceeds, we transition as much of the graph into 2D as possible, but leave the "lost" dimension as an axis object in AR.

To improve the user experience, we decrease the perceived latency by anticipating the users' interactions and pre-constructing graphs in the background so that transitions appear fast and seamless. To resize a graph users can grab the substrate area of a graph with both controllers simultaneously, which then resizes the whole graph proportional to the distance between the controllers.

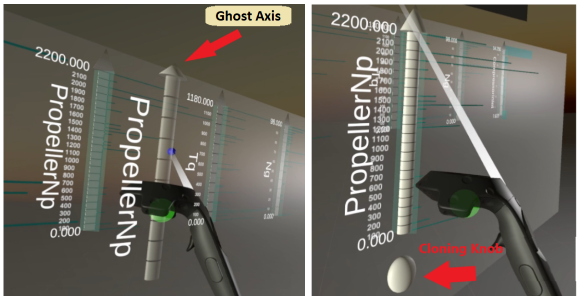

Fig. 2: When pulling an axis from a data dimensions panel, a ghost axis with an animated transition appears (left). Cloning knob that appears below all instantiated graphs (right).

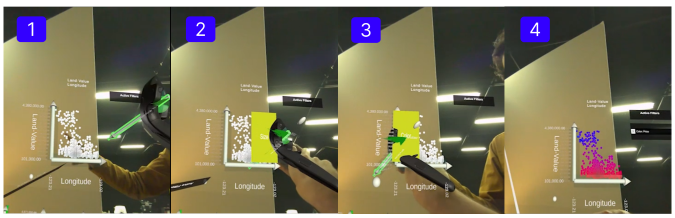

Fig. 3: Size/Color encodings after dropping a dimension. **1)** The user brings the selected axis towards the visualization. **2,3)** When the axis collides with the left/right side of the substrate, a yellow panel labeled "color"/"size" appears, indicating that dropping the axis will assign the dimension to the respective channel. **4)** Changed encoding, with legend on the right side. For a larger version see the appendix.

New Filtering System We designed a new filtering method for IA systems. Adhering to familiar desktop VA conventions and matching previous work [10], we added an interactive filter area at the top of each graph (shown in Figure 4). To reduce space usage, this panel is typically shown in compact form (Figure 4a), which summarizes the filters, but expands into a larger version (Figure 4b) when the user interacts with it. To add new filters, users then only need to drag a data dimension (i.e., a 3D axis object) and drop it into the filter area, which then causes the appropriate type of (editable) filter to be created.

Graph Cloning Cloning enables branching within the workflow, comparing graphs, or creating different versions of the same graph [56]. To enable users to duplicate graphs, we added a new, elliptical cloning "knob" below their origin, Figure 2. To clone a graph (with all filters and encodings), users only need to pull away from this knob.

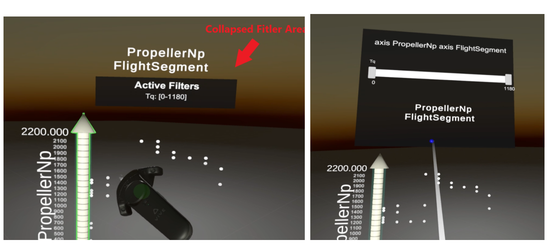

(a) Collapsed filter area. (b) Expanded filter area.

Fig. 4: Collapsed filter area (left) and expanded filter area shown only when pointing at/interacting with it (right).

3.4 Desktop System (CODAP Plugin)

For the desktop system, we chose CODAP [11], an open-source data visualization tool designed for teaching introductory data science. It features a robust API with Javascript/HTML plugins. CODAP uses window panes that can show data in the form of tables and supports drag-and-drop graph creation Figure 5, by dragging data dimensions from a table header and dropping them into a graph window, as well as brushing-and-linking. To prevent the screen from becoming too cluttered, we disabled creation of more than one table and more than one map view. Still, users could create as many graphs as they desired.

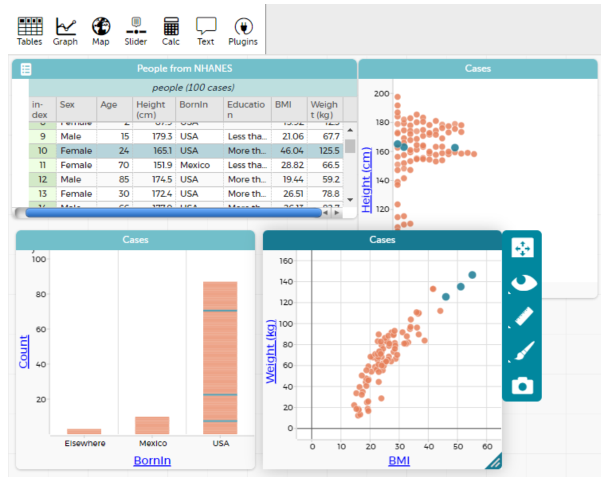

Fig. 5: Window panes inside CODAP.

To support cross-virtuality actions, we created a CODAP plugin that receive messages and then performs the corresponding action and sends messages about actions that happened on the desktop. As desktop window location and graph information is shared with the AR system, this enables HybridAxes to not only transition graphs between the virtualities, but also supports cross-virtuality brushing-and-linking where data points highlighted on the desktop are also highlighted in AR and vice versa. To filter a graph in CODAP, users first brush a series of points and then click on the designated filter button, to then either hide the selected points or the un-selected ones. When performed on a table view, the filter is global and applies to all graphs. For graphs and maps, the action creates a local filter that only applies to that visualization.

We use a node.js relay server to transfer messages, which also enables us to limit the rate of updates for interactions that result in (too) many messages. This relay also logs information.

Our original intent was to use an actual monitor for CODAP. Yet, in pilots we discovered that the Varjo XR-3's video-see-through image quality (particularly in terms of contrast) is just not good enough to enable users to read the small text in CODAP. We illustrate this issue in Figure 6 with an AR snapshot of the real-world monitor.

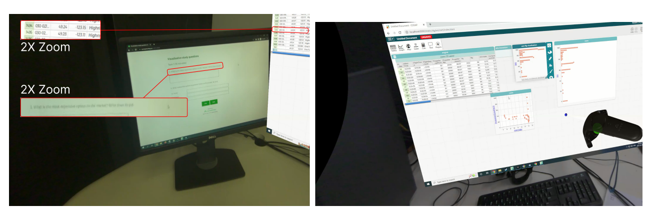

Fig. 6: Readability comparison between real (left) and virtual monitor (right) as seen through the AR headset, when viewed from roughly the same distance. High resolution version in the appendix.

As a remedy, we decided to project a virtual image of the desktop feed onto the video of the real world. Thus, while wearing the HMD, the user sees an image of the desktop overlaid onto the real-world monitor, which avoids taking the users out of their desktop experience. This workaround also allows us to increase the monitor size by 10%, which further addressed the text readability issue. We also place the virtual desktop image as a "monitor panel" a few centimeters in front of the real monitor, to enable users to "grab" and "drop" content from/into the monitor. Another interesting issue identified during pilots was that users tried to use the AR controller to control the cursor on desktop, instead of with the mouse. Yet, the (rotation) jitter of the controller prevented users from controlling the cursor sufficiently accurately. This caused substantial frustration and we thus decided to disable this feature.

3.5 Other Cross-Virtuality Interactions

Graphs transferred between virtualities retain the axes, colour, and size encodings. Due to technical limitations, users currently cannot transfer the filters of an AR graph to the desktop. Instead, the points that were filtered in AR are highlighted after the graph transitions to the desktop. This makes it easier for the user to see the data on the desktop that they filtered in AR. If so desired, the user can then use a single CODAP filtering action to actually filter the data. Further, whenever the user brushes a series of data points in either sub-system, these points are highlighted in all counterpart graphs. Selecting an empty spot within a graph clears the selected points.

4 Evaluation

We aimed to address our main research question around the effects of our hybrid system on users and their analytical process through specific experimental tasks that we discuss in this section.

4.1 User Study

An evaluation of IA systems is affected by the user's cognitive abilities and environmental noise, favouring observational studies [5]. Previous work showed that a mixed-method approach yields valuable insights into user behaviours and system usability [6, 9, 10, 18, 37]. Thus, we conducted a mixed-method study, with observational sessions, questionnaires, and semi-structured interviews. We divided our study participants equally into two groups, where the first one, **D**, used only the **desktop** system to analyze the data, whereas the second group, **H**, used the whole **HybridAxes** system. The dataset, tasks, and time constraints were the same for both groups. Each session took 80-90 min.

4.1.1 Dataset and Study Procedure

We used a multi-dimensional housing dataset to encourage our users to create different types of 2D and 3D graphs with different encodings, e.g., colour and size, opendata.vancouver.ca. In pilots, we used a different dataset to train participants, but they then performed poorly on the main tasks due to unfamiliarity with the data. Thus, we decided to use the same dataset for both the training and the main tasks, also to ensure a reasonable study duration. To combat data familiarity bias, we used only a small subset of the housing data for the training, while using the complete dataset for the main tasks.

Each participant had to perform three tasks. The first two involved directed questions on the data, through which we wanted to see (and measure) how participants use the system to make sense of the data to arrive at a concrete task answer. The third task was exploratory.

In the **first task** participants had to find the most expensive housing option and to name the top three data dimensions that made this option expensive, justifying their answers with the data. This requires exploring most data dimensions and finding ways to assess the effect of each of these on the property price, like a sensitivity analysis. We gave ample time to encourage users to explore the system and its features, and to better their understanding of the data.

In the **second task** participants had to look at the relationship of 7 data dimensions (size, price, age, renovation year, latitude, longitude, and neighbourhood) and describe how they interact with one another. Then, they had to identify the best values for those dimensions for a larger family (5 or more people) and identify which part of the city provides the most options, i.e. they had to consider the following data dimensions for answering the question: latitude, longitude, and neighbourhood. Our aim for this task was to see how users deal with scenarios that involve visualizing four or more dimensions simultaneously. We were also interested to see how they use the system to find a cluster with a given range for these attributes.

To identify emergent interaction patterns and behaviours, we designed the **third task** to be exploratory. There, we asked participants to imagine that they had been recruited by a real estate agency and were asked to first explore and then identify interesting and non-trivial insights within the dataset.

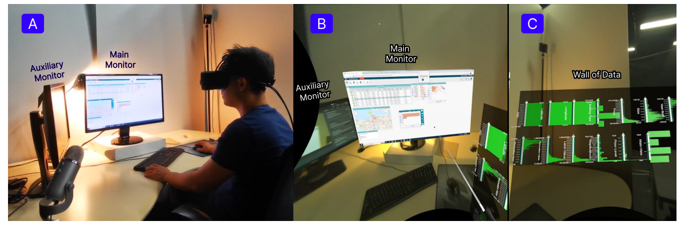

Fig. 7: Depiction of the study setup. A shows the view of the recording camera. B and C are the center and right egocentric views of the participant, respectively, captured within the AR system.

Introduction and Training (15-25 min.): Participants first completed a demographic questionnaire. Then group H was trained on the AR system, with verbal and gestural instructions or assistance, as needed. Both groups D and H were then introduced to the desktop system. After that, participants of group H were also introduced to the cross-virtuality features. For both group D and H sessions, the researcher first introduced a feature, and then asked participants to perform a simple task with it.

Tasks 1 and 2 (15+10 min.): Here, each task description was successively shown on an auxiliary monitor, together with a countdown timer and a text box for the answer, see Figure 7. To ensure participants understood the tasks, we answered any potential questions before they started. Participants were asked to use think-aloud while working on their answers, and could take a break between tasks.

Exploratory Question for Task 3 (20 min.): The question for task 3 was again shown on the monitor. In this task participants were asked to explicitly verbalize each insight they had during this session, which the experimenter noted down.

Post-study Questionnaires (5 min.): Afterwards, we asked participants to fill a questionnaire about their perceived performance with and preferences for the system, and another about their perceived mental and physical fatigue, using the NASA TLX questions.

Interview and Wrap up (15 min.): Finally, we conducted a semi-structured interview with two parts: 1) Follow-up questions to contextualize post-study questionnaire answers and to record perceptions of the system and user performance within it, and 2) followups about a users' behaviours during the study, to confirm experimenter observations or query the reasons behind them.

4.1.2 Participants

We recruited 16 participants (7 female, 9 male) aged between 25 and 38, mostly from the local university as well as some working professionals. All 16 had worked with at least one data visualization tool before (Excel, Tableau, JMP, Observable, D3, Python, or R). As mentioned, we divided these participants into 2 groups of 8. The first group (group D) worked solely on the desktop, while group H used the hybrid system. When asked to rate their data visualization skills and literacy, 7 out of 8 in each group rated themselves 6 or higher on a seven-point scale (in group D only one and two in group H rated themselves 7/7), which indicates high VA familiarity. All participants in group H had prior experience with VR/AR. Three reported using AR/VR more than six times in the last month, while the other five reported using AR/VR at least twice in the last month.

4.1.3 Data Collection and Analysis Procedure

We collected a range of qualitative and quantitative data. First, two experimenters took notes of participant behaviours and their spoken words. Second, we recorded an over-the-shoulder video of each session, see Figure 7. This proved useful for corroborating notes with participant behaviours, especially for group H. Lastly, we captured the desktop and for group H also the AR feed. We voice-recorded the interviews, but made also written notes of answers. Further, we recorded each action and the status of the system when that action was performed. We also logged each head, hand, and gaze movement as well as the coordinates of all the objects in the scene three times every second. To further support attempts at reproduction and verification of our results, we have published an anonymized version of our collected data alongside our system's open source code on Github[1].

Using an open coding method, three coders went over the video and audio recordings and looked for patterns in screen/content organization, switching between sub-systems (AR and Desktop), instances of using AR and desktop together, different strategies in graph creation or insight-generation, and any effects of the hybrid system on the fluidity of the VA process. To complement the findings from our coding, we also analyzed the logged data to quantitatively characterize use patterns.

4.2 Observations and Findings

In this section, we report our main study observations and findings.

4.2.1 General Findings

Findings from Questionnaires: Users rated their perceived speed, accuracy, and difficulty in performing the tasks on a 7-point scale, where lower means faster or easier. To investigate the impact of the participant group on perceived speed, accuracy, and difficulty, we conducted an ANOVA with the participant group as the independent variable. None of these ANOVAs showed a signficant difference, for perceived speed ($F(1, 16) = 0.127$, $p > 0.05$), accuracy ($F(1, 16) = 0.098$, $p > 0.05$), nor difficulty ($F(1, 16) = 0.567$, $p > 0.05$). Thus there are no strong differences in the perceived speed, accuracy, or difficulty between group H and group D.

As reasons that might have slowed them, group H mentioned glitches in both in AR and desktop sub-systems, system unfamiliarity (4/8 participants), and the novelty of the dataset (3/8), while Group D mostly mentioned the novelty of the dataset (4/8) and system unfamiliarity (3/8). In terms of accuracy, group D pointed at the desktop system as the main reason for inaccuracies (6/8). On the other hand, group H mostly reported the unfamiliarity with the dataset and the system as the main factors for their perceived accuracy. Group H reported a slightly higher average difficulty with system (M: 4.38, STD: 0.86) than group D (M: 4, STD: 1). On top of learning a new dataset Group H mostly mentioned the mental load of learning two new sub-systems as the main challenge. The desktop users mostly mentioned dataset unfamiliarity (6/8) and then the system (2/8) as the main difficulty.

Only task 1 had an objectively correct answer. We thus used this task as an indicator of actual user accuracy. This task had four parts, and we graded each answer using a range of 0 to 4. Almost everyone answered the first two questions correctly. Figure 8 shows the aggregated scores, where group H (M: 2.87, STD: 1.125) was higher/better than group D (M: 2.25, STD: 0.88), but not significantly so. There was no significant difference in time to complete task 1 between group D (M: 13.36 minutes, STD: 1.90) and group H (M: 13.43 minutes, STD: 1.64).

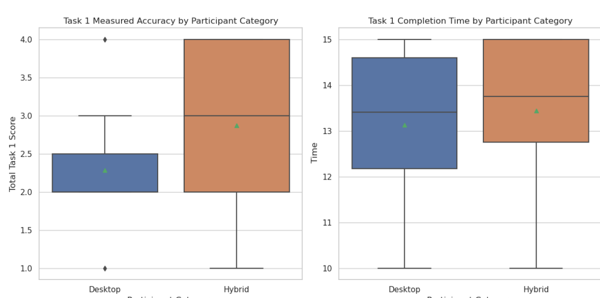

Fig. 8: Accuracy of task 1 across group H and group D.

On a 7-point scale group H users reported a slightly higher degree of tiredness and fatigue (M: 4.88 STD: 0.78) than group D (M: 5, STD: 1.32), but not significantly so, Figure 9. Nor could we find significant effects of fatigue on users' perceived speed, accuracy, or task difficulty. Generally, group D mostly reported cognitive fatigue due to unfamiliarity with the system and the dataset, while hybrid users ranked physical fatigue as more significant (due to the weight and bulkiness of the headset, and eye strain).

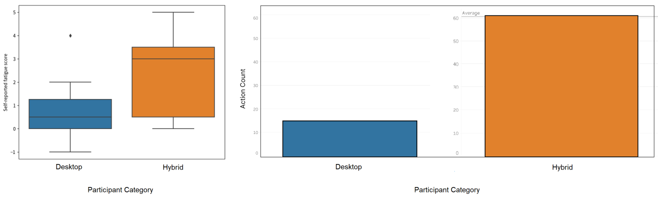

Fig. 9: Left: Boxplot for self-reported fatigue grouped by Participant Category (sub-system), right: Average graph creation count per group (group H on the left and group D on the right).

To better understand our users' behaviours, we analyzed the recorded action data, comparing the following graph actions across desktop and AR: 1) creation, 2) moving and resizing, 3) changing colour and size encodings, 4) brushing and linking, and 5) accessing details-on-demand. In AR, we also analyzed the number of created 2D and 3D graphs and the number of transitions where they moved a graph to or from the desktop. Table 1 shows summary statistics.

	Actions		Created graphs		Moving graphs		Brushing/ selecting items	
	Average	Median	Average	Median	Average	Median	Average	Median
Desktop Group	190	187	16.14	18	36	27	75.85	59
Hybrid Group	319	323	59.01	57.5	138.5	127.5	52.30	38.5

Table 1: Action data for the two participant groups.

Group D performed an average of 190 actions at an average rate of 4.2 actions/min. Group H clearly did more: 319 actions at 7 actions/min. across both sub-systems, shown in Figure 9 and Figure 10, where most happened in AR (235, 5.2 actions/min.), with only 84 actions on the desktop (2.9 actions/min.), see Figure 11. When queried about this behaviour, most group H participants mentioned that it was easy to quickly create and destroy graphs in AR which allowed them to test many different hypotheses. Five of them also mentioned that the game-like nature of the AR sub-system appealed and made the exploration task more "enjoyable".

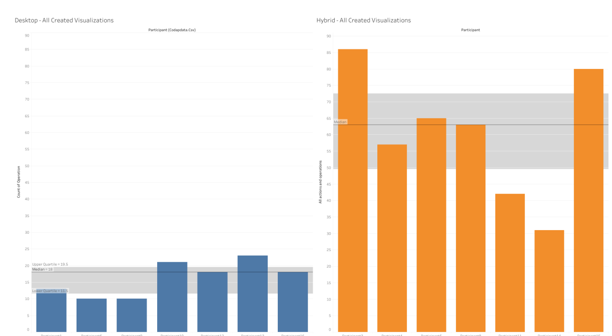

Fig. 10: Comparison of the number of created graphs per participant between group H on the left and group D on the right.

Another interesting finding was that most group H participants created substantially fewer graphs on the desktop compared to group D. Yet, this fact is not well represented in the overall averages shown in Table 2 and Table 1. While it seems that creating 15.87 graphs (group H on the desktop) is close to 16.14 graphs (group D), this is only due to two participants, P3 and P14, who spent much time on the desktop. If we remove them, the average for group H drops to only 6 graph creation actions, which is much lower than group D.

[1]https://github.com/rajabiseraji/NewHybridViz

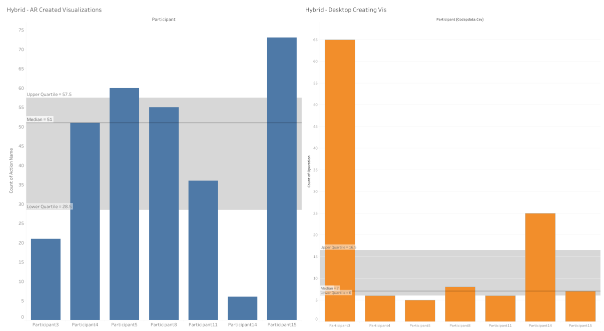

Fig. 11: Comparison of the number of created graph in group H between the AR sub-system (left) and desktop sub-system (right).

	Actions		Created graphs		Moving graphs		Brushing/selecting items	
	Average	Median	Average	Median	Average	Median	Average	Median
Group H on Desktop	84	67	15.87	6.5	24.65	11.5	17.75	10.5
Group H in AR	235.28	266	43.14	51	113.85	116	34.28	28

Table 2: Overview of group H actions on desktop and in AR.

4.2.2 Observed Patterns and Behaviours

We first discuss behavioural patterns observed in both participant groups. Then, we mention unique behaviours and patterns.

Visualization Strategy: Multiple Graphs vs. One with Multiple Dimensions: When addressing each task, participants in both groups seemed to fall roughly into two categories. Those in the first preferred creating two or more graphs and then brushing and linking to find connections (P9, P10, P12). In contrast, the second category created as few graphs as possible and added multiple dimensions onto them with colour and size (P2, P6, P13). P6 specifically mentioned: "*For answering question [2], I want more axes, to check more variables at the same time. Maybe more than just an X and Y.*"

Still, participants in group H tended to create more graphs to answer a question, most of which were created in AR and mentioned the ease of cloning a graph, the tangibility of interacting with AR graphs, and more available visualization space as the main factors. P15 said: *"It made more sense to just create as many graphs as possible and brute force our way into finding the answer.*"

Hybrid Specific: Glancing into the Desktop Sub-system: At some point during the study, all group H participants looked back and forth between their AR graphs and the desktop ones. For a majority (7/8), this was a frequent behaviour, confirmed through their interviews and quantitative system logs. They did this for the following reasons: 1) To find brushed points across the modalities (P5, P7, P8) or 2) to use features that were absent in a sub-system, e.g., 3D graphs on the desktop or maps or tables in AR (P8, P11, P15). P7 explained "*First: working in AR and looking at the map to see what gets highlighted. Second, work on the desktop and see if it is showing something interesting in the AR graphs.*" Our participants also identified that they only exhibited this back-and-forth glancing only because interaction was synchronized between sub-systems.

Hybrid Specific: AR/Desktop Transitions: Seven group H participants transitioned a graph between AR and desktop, i.e., from 3D to 2D, or vice versa, an average of 6.16 times during their session (median: 3.5) and mostly during the second task Figure 12. Compared to the overall number of actions, they used this feature only rarely. When asked, they reported that the process of creating graphs in both AR and desktop was simple and quick enough for them not to need to transition a graph between the two modes. Two of the participants with a strong background in spatial data analysis and visual analytics mentioned that if they were presented with 3D spatial data they would have used this hybrid transition more often.

Only two of the participants from group H (P5, P8) reported the mental load of switching as discouraging. Further, five mentioned that they mostly used cross-virtuality brushing to connect the data across the two sub-systems. Almost all mentioned the "hassle" of switching between controller and mouse discouraged switching. Finally, all of them mentioned that if not for the issue that sometimes prevented filters of a graph from being transitioned between modes, they would have used the AR/Desktop transition more often.

Hybrid Specific: Staying in the Comfort Zone: For group H, all chose one sub-system (AR or desktop) as their main one within their first few minutes, and then spent most of their time within it. Most chose AR, while only P3 and P14 chose desktop. When queried about this behaviour and taking their background into account, we identified two main reasons: 1) participants with a strong VR background preferred AR and 2) participants with a strong data visualization background began using AR out of curiosity and remained there as they got more proficient with the system. P8 stated: "*I used the switching a little bit. If I was more expert with the systems, [...] I would have been more comfortable with some features in one and had our preferences in using them. But after a while, I felt more comfortable in AR and mostly worked there. I found my preferred way, and stayed there!*"

Hybrid Specific: Use of 3D Graphs: Not unexpectedly, participants created fewer 3D graphs than 2D ones in AR. Yet, they created most of these graphs in the second and third tasks. They mentioned that 3D graphs were useful for identifying multi-variable relationships (P8, P11) and to get a more holistic view of the data (P7). However, some also reported that analyzing a 3D graph to be more mentally demanding (P11). Most participants used them as in previous work [38], where they rotated the 3D graphs and used every face of it as a 2D projection (P4, P5, P11, P15). In contrast, P7 and P15 used 3D graphs directly to identify data clusters and trends.

Hybrid Specific: Desktop and AR Usage Patterns: Figure 12 shows the usage patterns of the two sub-systems for group H. P8 and P14, who did not have a strong background in AR/VR, spent their first session mostly on the desktop. Both moved to AR at some point during the second task (at ≈10 min. Figure 12). They stated that for the first task, they felt they had to come up with correct answers within a time limit, and thus they mostly used the interface they were comfortable with. However, for the entire second task and most of the third one, they reported less pressure due to the exploratory nature of the questions and thus could explore the novel interface for looking at the data.

We also observed and later confirmed in the interviews that participants needed to create more graphs for task 2 and thus needed more space (especially those that created multiple graphs, see subsubsection 4.2.2). All group H participants created 3D graphs to be able to see multiple interactions simultaneously. These factors motivate and support the more frequent usage of AR in the second and third tasks. As visible in Figure 12, P3 was an exception, as they spent most of their time on the desktop, as they were "more comfortable with desktop VA tools." They also felt the time pressure, which negatively impacted their experience and limited their use of the AR sub-system.

Based on our observations and interviews, group H mostly interacted with the desktop for one or more reasons: 1) reading data point details in the table (all Ps), 2) using the map (P4, P7, P8, P11, P15), or 3) filtering based on brushed data (P4, P8, P14).

Hybrid Specific: Filtering and Data Querying Strategy: For tasks that required drilling down and multiple filters, our hybrid participants took two different approaches. P5, P7, and P8 applied multiple filters on their AR graphs and drilled down, while P3, P4, P11, and P15 chose a combination of brushing and filtering in both AR and desktop (mostly in AR) to identify their answer. We could not pinpoint any specific factor that might point to a system characteristic, thus this might be due to differences in analysis approaches.

Engagement with AR Sub-system: Many (P4, P5, P7, P8, P11, P15) mentioned the AR system to be "*intuitive*" and "*enjoyable and engaging to work with*". All expressed their interest in further "*playing*" with the system, well after the designated timelimit had expired. When observing these users working in the hybrid system, they at times exhibited playful behaviours and seemed to enjoy the setup. All mentioned that although the system has "*rough edges*", it is also quite "*natural and intuitive*" to work with. At first glance, it might seem that participants' behavior could be linked to the novelty effect, particularly for those with limited AR/VR experience. Yet, even regular AR/VR users (P15,

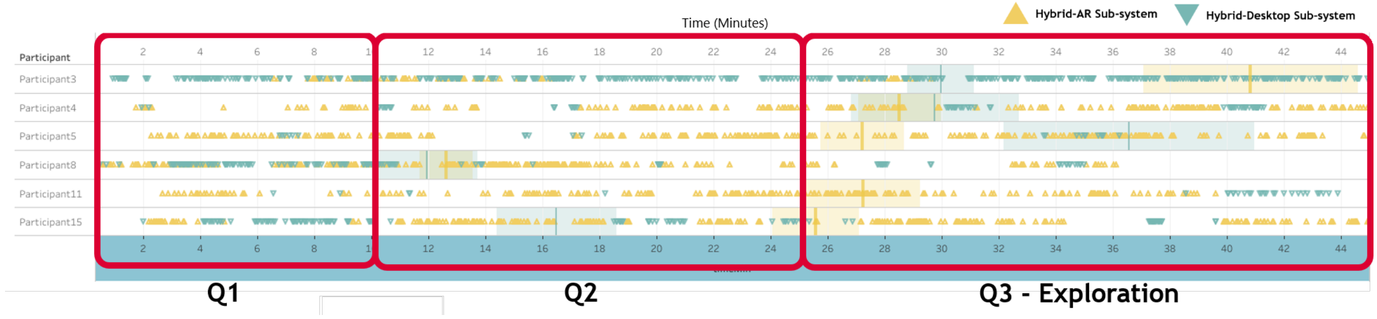

Fig. 12: Timeline of group H users' actions colored by the type of system. Each point in the timeline represents a user action done in the session. Due to a software bug, the timeline data (and only that) for two users was lost and thus cannot be shown here.

P8, P5) expressed similar sentiments. Thus we can state that the user interface and interactions contributed substantially to their sense of enjoyment and playfulness.

Hybrid Specific: Space Organization and Access: We observed and later confirmed through the interviews that all hybrid participants designated a specific 3D space for their AR graphs. Interestingly, most chose areas that put the monitor either just within peripheral vision (P3, P4, P5, P11) or where it was completely out of view (P8, P15). We suspect that they did this to keep themselves focused on the task at hand and to avoid cluttering their view, see previous work [54]. Both P8 and P15 placed their graphs behind their seat, opposite to the monitor, which meant they had to turn the chair completely to access these graphs. They regularly performed actions on their AR graphs and then quickly used the rotating chair to look back at their monitor, check the effect of their work on the desktop, and then back again to AR.

Further, we identified that participants never placed or stacked AR graphs in front of another, where one would obscure another. Instead, they treated the space around them as a giant spherical manifold within arms' reach, and put their graphs (roughly) somewhere on this sphere.

Desktop Specific: Screen Organization: In group D, P2, P6, P12 and P13 constantly re-organized their screen to ensure they had fewer than 3 or 4 active graphs, removing irrelevant ones or re-using graphs by changing axes to get the desired results. P2 and P6 later expressed the need for having more than 2 graph attributes available to them. They felt limited by the 30" monitor. Still, we suspect that the cognitive load of seeing and memorizing multiple graphs might play a role in this behaviour. Unfortunately, we did not collect sufficient data to address this question in more depth.

4.3 Discussion

Overall, our study results reveal interesting insights into the different ways users interact with a hybrid cross-virtuality system that allows them to freely move between different modes of virtuality.

4.3.1 System Usage

When used as complimentary components, desktop and IA systems can be useful in routine (day-to-day) VA tasks. In our study, users picked one mode of virtuality as their dominant one. However, as mentioned, we observed that users used the other sub-system to complement their interaction. Based on users' answers and our observations, the most prevalent reason might be the lack of a feature or view in the dominant sub-system of choice, e.g., the lack of a map in the AR sub-system or the lack of 3D graphs in the desktop sub-system. Still, another popular use case was insight verification, where participants would regularly cross-check an insight that they felt uncertain about by using the other sub-system.

Our findings confirm previous work [54], which identified that desktop users find AR to be a valuable extension to their experience. However, our participants' statements provide a nuanced extensions to Wang et al. [54], who reported that "users appreciate non-synchronized views". Our work suggests that, when users select a dominant sub-system, i.e., AR, and mostly remain in it, it is beneficial for the secondary (non-dominant) sub-system, i.e., desktop in our case, to remain synchronized with the dominant one.

It is currently inadvisable to transition all VA work to AR/VR/XR, also because efficient VA depends on existing software and interaction devices, such as mouse and keyboard. Yet, our study demonstrates that one of the main use cases for hybrid VA systems is to give users the option to use AR/XR for some of their analysis activities, e.g., when they are faced with complex multi-dimensional relationships.

Users of a hybrid system dedicate space (separate from their desktop) for AR operations. As mentioned in the previous section, most of hybrid users dedicated a part of the available space for their analysis tasks. In our study, these areas were mostly in mid-air to the right of the desktop monitor (at an appropriate distance), but not on the table, or behind the participant. Based on the interviews, it seems that this separation helped them create a "mental workspace" (P5, P15). This interesting observation highlights that designers need be aware of such space usage pattern and design the system and users' work spaces accordingly. Further, as users treated the 3D space around them as a big spherical monitor and never stacked graphs, i.e., always made sure that everything was fully visible, XVA designers should design for this imaginary spherical manifold, and only involve stacking/occlusion when needed to transform visualizations, e.g., [36]. We still acknowledge that with more complex datasets it may become impossible to avoid occlusions completely.

Support multiple graph authoring and filtering strategies. We observed that our participants used different strategies to create graphs and filter them. Some of them preferred to create multiple graphs and to use brushing-and-linking. Others preferred to add more channels on fewer graphs. In our study, due to the speed of creation and destruction of graphs in AR and the "tangibility and intuitiveness of AR" (P8, P15, P11), our hybrid users mostly opted for the multiple graphs approach. However, one mentioned that if modifying the colour and size encodings in AR were more "intuitive" (P8), they would have used it potentially more often. Thus, it is important to optimize a system to support both approaches well.

No difference between AR and desktop in perceived performance (speed, accuracy, difficulty). With two unfamiliar systems, users performed more or less the same. This is not unexpected and speaks to the complexity of learning the systems and the task overshadowing any differences (both in interaction and data cognition).

In a Hybrid VA system, AR is the mostly used for exploration and desktop is mostly used for detail-intensive activities. Most of our group H participants expressed that they felt the "*time ticking down*" (P14) or the "*pressure of finding the right answer*" (P4, P5, P7, P8) in the first task. However, in the third task, almost all (7/8) freely explored the data in the AR system, occasionally using the desktop system for verification. This suggests that for exploration tasks AR provides a "*playground-like environment*" (P4), particularly for users who are less familiar or comfortable with AR. For tasks that require drilling down into the data and "*reading exact values*" (P11), most of our participants preferred the desktop. They cited several reasons for this preference, including the more precise nature of interaction through the mouse and keyboard, the availability of tables on the desktop, and their experience with and comfort in working within a desktop environment.

Regarding AR, users also reported positively on being able to see the environment around them, instead of being "locked up" in VR. Yet, since the Varjo's cameras are almost 20 cm in front of the user's eyes, they mentioned the feeling that their hands were not where they should be. Still, as the session went on, they forgot about the offset and adapted to the new view.

4.3.2 Limitations

Virtual Monitor Usage: Due to technical constraints (refer to Figure 3.4), our study employs a video-see-through display and a virtual monitor in our hybrid setup evaluation. Although most users did not notice the virtual nature of the monitor, we recognize our solution as a primary limitation. Future research is encouraged to explore alternatives and strive for a solution with real displays.

Controller Usage over hand tracking: Users noted the "hassle" of switching between mouse/keyboard and controllers, impacting their switching frequency. We acknowledge this as an important limitation and recommend future researchers explore robust hand-tracking methods for smoother switching behaviour.

Sample Size and Study Design: Our number of participants is in line with other VA between-subject research. Still, a different outcome may result from a larger sample and/or a within-subjects design. On the other hand, and as we found several insignificant results, it seems unlikely that one could expect large differences. Most participants were university students with above-average VA experience, which matches reasonably (but not perfectly) with our main target of regular VA users.

Learning Factor and Dataset Familiarity: The observations, interviews, and the quantitative data show that users still became more proficient with the systems as time went on. Clearer results may result from thoroughly familiarizing participants in advance of a study.

Using the same dataset for training and evaluation might have familiarized users with some of the data dimensions, which we acknowledge as a potential confound. We controlled for this by using different data subsets for training and main tasks. Users' post-study questionnaire responses also indicated that they were so busy learning the system during training that they did not get that familiar with the data itself.

Technology Limitations: There was a (limited) number of bugs and tracking glitches that had occasional negative effects on users' performance and their experience. Yet, they were rare enough not to fundamentally affect the outcomes of this study.

Use of Inter-coder Reliability Check: The interviews and video recordings were coded by three researchers, who had collaborated before and were somewhat familiar with each others' coding processes. Thus, they quickly reached a consensus on diverging codes. Yet, it might have been better to use more independent coders and an inter-rater reliability measure, e.g., Cohen's κ.

Better Use of 3D Graphs: Some participants did not use 3D graphs as often as 2D ones as they felt the data did not "fit" 3D graphs well. This could have been addressed by choosing a dataset that relied (even) more on spatial data. That said, we believe that for a fair comparison, one needs a dataset that can also be analyzed in 2D, as otherwise one sub-system would be clearly favoured.

Task time limit: We used time limits to give all participants the same amount of exposure to the system and to ensure the experiment's logistical feasibility. Yet, this created time pressure that might have altered their behaviour, particularly for P3 and P8. We thus acknowledge that if the focus is on observing and analyzing users' behaviours and their exploratory data analysis journey, it is better to use no time limit.

4.3.3 Lessons for Future Hybrid Visualizations Systems

Both sub-systems should match in terms of interaction. If the interaction methods differ between sub-systems, this incurs a mental context-switching effort. Our users also mentioned that due to the differences in filtering, they thus used it more on the system where it was easier (i.e., in AR). Thus, we recommend treating the two sub-systems as parts of a continuous experience and making interaction methods as consistent as possible. User feedback also suggests that adhering to this recommendation even more strictly could potentially enhance user performance in our IA system. Essentially, minimizing the need for switching between vastly different visual and physical interfaces could allows users to focus more on the data/task and less on the interface mechanics. This directly aligns with Nielsen's principle of "consistency and standards" in UX design [42].

Use Imbalance in Sub-system Features to Guide Users to Specific Functionalities. Some functionality makes (generally) less sense in a specific system, e.g., text entry in AR. Yet, a feature's absence in one sub-system can encourage users to use the other one for such functionality, e.g., like the map in our study. Designers of hybrid VA systems can leverage this to guide users accordingly. Overall, the asymmetric nature of hybrid systems could be used to encourage designer-planned behaviours. Yet, we acknowledge that this introduces additional challenges, as this requires designers to comprehensively understand all (future) needs of every user.

Dedicated 3D space for working with AR. As much as our users enjoyed their freedom, they dedicated a part of the 3D space for working with AR, separate from the desktop monitor. Hybrid system designers could design the work space accordingly. For instance, a system might allow users to create visual boundaries to delineate "work spaces." Another approach could boost AR contrast in dedicated spaces to make the virtual content easier to perceive.

5 Conclusion and Future Work

We presented and evaluated HybridAxes, a novel hybrid visual analytics solution with full support for authoring (almost) all data visualizations on the desktop and in AR. Our study thus explores the new design space of cross-virtuality VA systems and we found that desktop and immersive analytics systems serve well as complementary sub-systems for routine VA tasks. Still, we did not identify strong perceived performance differences between the two systems. Interestingly, participants cross-checked their work in the other sub-systems, but transitioned graphs only sometimes, as it was (too) easy to recreate them. Further, we revealed the importance of supporting different strategies for authoring and filtering graphs. Finally, we identified that the designers of hybrid systems should match the features between the sub-systems, unless they want to steer the users to specific usage patterns (if appropriate).

As highlighted in the discussion section, our study outcome matches findings from prior work on spatial user interfaces, e.g., [8, 43, 54]. However, our work extends these findings by identifying the challenges in environments with simultaneous support for interaction with visualizations on both the desktop and in AR as well as user behaviors related to switching between the environments. At first glance, our findings might seem to suggest that integrating an AR environment alongside a desktop merely functions as a glorified secondary large display. Despite the latter being a valid use-case, it is not the main point of this work. Instead, we emphasize the importance of deliberate interaction design for such a multi-modal system with on-demand switching ability, which users chose to cross-checked their analysis outcomes. We also point out how such a system could be used to steer users toward specific functionality. Lastly, our study unveils an emerging need for in-depth exploration of user preferences within such multi-modal systems, particularly in scenarios requiring concurrent use of both subsystems to accomplish specific objectives.

Future improvements to the system include better (hand) tracking and/or more accurate tracking systems. Another aspect for future work is the user study, addressing the unfamiliarity factor with substantial training, to evaluate the hybrid system in a scenario that is more representative of long-term usage.

Acknowledgments

The first author thanks the Canadian NRC for funding.

References

[1] I. J. Akpan and M. Shanker. A comparative evaluation of the effectiveness of virtual reality, 3d visualization and 2d visual interactive simulation: an exploratory meta-analysis. *SIMULATION*, 95(2):145–170, 2019. doi: 10.1177/0037549718757039 1

[2] R. Amar, J. Eagan, and J. Stasko. Low-level components of analytic activity in information visualization. In *IEEE Symposium on Information Visualization, 2005. INFOVIS 2005.*, pp. 111–117, 2005. doi: 10.1109/INFVIS.2005.1532136 1

[3] A. U. Batmaz and W. Stuerzlinger. The effect of rotational jitter on 3d pointing tasks. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI EA '19, p. 1–6. Association for Computing Machinery, New York, NY, USA, 2019. doi: 10.1145/3290607.3312752 3

[4] H. Benko, E. Ishak, and S. Feiner. Collaborative mixed reality visualization of an archaeological excavation. In *Third IEEE and ACM International Symposium on Mixed and Augmented Reality*, pp. 132–140, 2004. doi: 10.1109/ISMAR.2004.23 2

[5] E. Bertini, H. Lam, P. Isenberg, C. Plaisant, and S. Carpendale. Empirical studies in information visualization: Seven scenarios. *IEEE transactions on visualization and computer graphics*, 18, 11 2011. doi: 10.1109/TVCG.2011.279 5

[6] M. Billinghurst, M. Cordeil, A. Bezerianos, and T. Margolis. Collaborative immersive analytics. In K. Marriott, F. Schreiber, T. Dwyer, K. Klein, N. H. Riche, T. Itoh, W. Stuerzlinger, and B. H. Thomas, eds., *Immersive Analytics*, vol. 11190, pp. 221–257. Springer International Publishing, 2018. Series Title: Lecture Notes in Computer Science. doi: 10.1007/978-3-030-01388-2_8 1, 5

[7] S. Blum, G. Cetin, and W. Stuerzlinger. Immersive analytics sensemaking on different platforms. In *Computer Science Research Notes*. Západočeská univerzita, 2019. doi: 10.24132/CSRN.2019.2902.2.9 3

[8] M. Cavallo, M. Dholakia, M. Havlena, K. Ocheltree, and M. Podlaseck. Dataspace: A reconfigurable hybrid reality environment for collaborative information analysis. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 145–153, 2019. ISSN: 2642-5254. doi: 10.1109/VR.2019.8797733 1, 9

[9] M. Cavallo, M. Dolakia, M. Havlena, K. Ocheltree, and M. Podlaseck. Immersive insights: A hybrid analytics system forCollaborative exploratory data analysis. In *25th ACM Symposium on Virtual Reality Software and Technology*, pp. 1–12. ACM, 2019. doi: 10.1145/3359996.3364242 1, 2, 3, 5

[10] S. Chakraborty and W. Stuerzlinger. Vizinteract: Rapid data exploration through multi-touch interaction with multi-dimensional visualizations. In *Human-Computer Interaction INTERACT '21*, vol. LNCS 12934, p. 610–632. Springer, Aug 2021. doi: 10.1007/978-3-030-85613-7_39 2, 3, 4, 5

[11] Concord-Consortium. Common online data analysis platform (CODAP). `https://codap.concord.org`, 2014. 4

[12] M. Cordeil. MaximeCordeil/IATK. Programmers: _:n883 original-date: 2017-02-16T05:25:32Z. 3

[13] M. Cordeil, B. Bach, A. Cunningham, B. Montoya, R. T. Smith, B. H. Thomas, and T. Dwyer. *Embodied Axes: Tangible, Actuated Interaction for 3D Augmented Reality Data Spaces*, p. 1–12. Association for Computing Machinery, New York, NY, USA, 2020. 2

[14] 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. 200–209, 2019. ISSN: 2642-5254. doi: 10.1109/VR.2019.8797978 3

[15] 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. Association for Computing Machinery, 2017. doi: 10.1145/3126594.3126613 2, 3

[16] T. J. Dube and A. S. Arif. Text entry in virtual reality: A comprehensive review of the literature. In M. Kurosu, ed., *Human-Computer Interaction. Recognition and Interaction Technologies*, pp. 419–437. Springer International Publishing, Cham, 2019. 3

[17] T. Dwyer, K. Marriott, T. Isenberg, K. Klein, N. Riche, F. Schreiber, W. Stuerzlinger, and B. H. Thomas. *Immersive Analytics: An Introduction*, pp. 1–23. Springer International Publishing, Cham, 2018. doi: 10.1007/978-3-030-01388-2_1 1, 2, 3

[18] M. El Meseery, Y. Wu, and W. Stuerzlinger. Multiple workspaces in visual analytics. In *Symposium on Big Data Visual Analytics*, BDVA '18, pp. 70–81, Oct 2018. doi: 10.1109/BDVA.2018.8534019 3, 5

[19] B. Ens, B. Bach, M. Cordeil, U. Engelke, M. Serrano, W. Willett, A. Prouzeau, C. Anthes, W. Büschel, C. Dunne, T. Dwyer, J. Grubert, J. H. Haga, N. Kirshenbaum, D. Kobayashi, T. Lin, M. Olaosebikan, F. Pointecker, D. Saffo, N. Saquib, D. Schmalstieg, D. A. Szafir, M. Whitlock, and Y. Yang. Grand challenges in immersive analytics. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI '21. Association for Computing Machinery, New York, NY, USA, 2021. doi: 10.1145/3411764.3446866 1, 2, 3

[20] B. Ens, S. Goodwin, A. Prouzeau, F. Anderson, F. Y. Wang, S. Gratzl, Z. Lucarelli, B. Moyle, J. Smiley, and T. Dwyer. Uplift: A tangible and immersive tabletop system for casual collaborative visual analytics. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1193–1203, 2021. doi: 10.1109/TVCG.2020.3030334 2

[21] J. A. W. Filho, W. Stuerzlinger, and L. Nedel. Evaluating an immersive space-time cube geovisualization for intuitive trajectory data exploration. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):514–524, 2020. doi: 10.1109/TVCG.2019.2934415 2

[22] M. Fisherkeller, J. Friedman, and J. Tukey. Prim-9: An interactive multidimensional data display and analysis system. slac-pub-1408. *Stanford Linear Accelerator Center, Stanford, Calif. FITCH, WM and MARGOLIASH, E.(1967). Construction of phylogenetic trees. Science*, 155:279–284, 1974. 1

[23] A. Fonnet and Y. Prié. Survey of immersive analytics. *IEEE Transactions on Visualization and Computer Graphics*, 27(3):2101–2122, 2021. doi: 10.1109/TVCG.2019.2929033 1

[24] B. Fröhler, C. Anthes, F. Pointecker, J. Friedl, D. Schwajda, A. Riegler, S. Tripathi, C. Holzmann, M. Brunner, H. Jodlbauer, H.-C. Jetter, and C. Heinzl. A survey on cross-virtuality analytics. *Computer Graphics Forum*, 41(1):465–494, 2022. doi: 10.1111/cgf.14447 1, 2

[25] D. Garrido, J. Jacob, and D. C. Silva. Building a prototype for easy to use collaborative immersive analytics. In M. Paszynski, D. Kranzlmüller, V. V. Krzhizhanovskaya, J. J. Dongarra, and P. M. A. Sloot, eds., *Computational Science – ICCS 2021*, Lecture Notes in Computer Science, pp. 628–641. Springer International Publishing, 2021. doi: 10.1007/978-3-030-77961-0_50 1

[26] T. Gesslein, V. Biener, P. Gagel, D. Schneider, P. O. Kristensson, E. Ofek, M. Pahud, and J. Grubert. Pen-based interaction with spreadsheets in mobile virtual reality, 2020. 2, 3

[27] L. Grammel, M. Tory, and M.-A. Storey. How information visualization novices construct visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):943–952, 2010. doi: 10.1109/TVCG.2010.164 1

[28] R. Grasset, J. Looser, and M. Billinghurst. Transitional interface: concept, issues and framework. In *2006 IEEE/ACM International Symposium on Mixed and Augmented Reality*, pp. 231–232, 2006. doi: 10.1109/ISMAR.2006.297819 2

[29] K. Hinckley, K. Yatani, M. Pahud, N. Coddington, J. Rodenhouse, A. Wilson, H. Benko, and B. Buxton. *Pen + Touch = New Tools*, p. 27–36. Association for Computing Machinery, New York, NY, USA, 2010. 2

[30] S. Hubenschmid, J. Wieland, D. I. Fink, A. Batch, J. Zagermann, N. Elmqvist, and H. Reiterer. Relive: Bridging in-situ and ex-situ visual analytics for analyzing mixed reality user studies. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, CHI '22. Association for Computing Machinery, New York, NY, USA, 2022. doi: 10.1145/3491102.3517550 1

[31] H.-C. Jetter, J.-H. Schröder, J. Gugenheimer, M. Billinghurst, C. Anthes, M. Khamis, and T. Feuchtner. Transitional interfaces in mixed and cross-reality: A new frontier? In *Companion Proceedings of the 2021 Conference on Interactive Surfaces and Spaces*, ISS '21, p. 46–49. Association for Computing Machinery, New York, NY, USA, 2021. doi: 10.1145/3447932.3487940 3

[32] J. Jo, S. L'Yi, B. Lee, and J. Seo. *TouchPivot: Blending WIMP & Post-WIMP Interfaces for Data Exploration on Tablet Devices*, p. 2660–2671. Association for Computing Machinery, New York, NY, USA, 2017. 1

[33] R. Kosara, F. Bendix, and H. Hauser. Parallel sets: interactive exploration and visual analysis of categorical data. *IEEE Transactions on Visualization and Computer Graphics*, 12(4):558–568, 2006. doi: 10.1109/TVCG.2006.76 2

[34] R. S. Laramee and R. Kosara. *Challenges and Unsolved Problems*, pp. 231–254. Springer Berlin Heidelberg, Berlin, Heidelberg, 2007. doi: 10.1007/978-3-540-71949-6_5 2

[35] 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. 335–338. Association for Computing Machinery, New York, NY, USA, 2019. doi: 10.1145/3343055.3360746 3

[36] B. Lee, M. Cordeil, A. Prouzeau, B. Jenny, and T. Dwyer. A design space for data visualisation transformations between 2d and 3d in mixed-reality environments. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, CHI '22. Association for Computing Machinery, New York, NY, USA, 2022. doi: 10.1145/3491102.3501859 3, 8

[37] B. Lee, X. Hu, M. Cordeil, A. Prouzeau, B. Jenny, and T. Dwyer. Shared surfaces and spaces: Collaborative data visualisation in a co-located immersive environment. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1171–1181, 2021. doi: 10.1109/TVCG.2020.3030450 3, 5

[38] B. Lee, X. Hu, M. Cordeil, A. Prouzeau, B. Jenny, and T. Dwyer. Shared surfaces and spaces: Collaborative data visualisation in a co-located immersive environment. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1171–1181, 2021. Conference Name: IEEE Transactions on Visualization and Computer Graphics. doi: 10.1109/TVCG.2020.3030450 3, 7

[39] 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, 2014. doi: 10.1109/3DVis.2014.7160093 2

[40] P. Milgram and F. Kishino. A taxonomy of mixed reality visual displays. *IEICE Trans. Information Systems*, vol. E77-D, no. 12:1321–1329, 12 1994. 1

[41] P. Millais, S. L. Jones, and R. Kelly. Exploring data in virtual reality: Comparisons with 2d data visualizations. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI EA '18, p. 1–6. Association for Computing Machinery, New York, NY, USA, 2018. doi: 10.1145/3170427.3188537 1

[42] J. Nielsen. Usability 101: Introduction to usability. *Jakob Nielsen's Alertbox*, 2003. 9

[43] P. Reipschlager, T. Flemisch, and R. Dachselt. Personal augmented reality for information visualization on large interactive displays. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1182–1192, 2021. doi: 10.1109/TVCG.2020.3030460 2, 9

[44] N. Reski, A. Alissandrakis, J. Tyrkkö, and A. Kerren. "oh, that's where you are!" – towards a hybrid asymmetric collaborative immersive analytics system. In *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society*. Association for Computing Machinery, New York, NY, USA, 2020. 3

[45] A. Riegler, C. Anthes, H.-C. Jetter, C. Heinzl, C. Holzmann, H. Jodlbauer, M. Brunner, S. Auer, J. Friedl, B. Fröhler, C. Leitner, F. Pointecker, D. Schwajda, and S. Tripathi. Cross-virtuality visualization, interaction and collaboration. In *XR@ISS*, 11 2020. 2

[46] J. M. Rzeszotarski and A. Kittur. Kinetica: Naturalistic multi-touch data visualization. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '14, p. 897–906. Association for Computing Machinery, New York, NY, USA, 2014. doi: 10.1145/2556288.2557231 1

[47] R. Sadana, T. Major, A. Dove, and J. Stasko. Onset: A visualization technique for large-scale binary set data. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1993–2002, 2014. doi: 10.1109/TVCG.2014.2346249 2

[48] D. Saffo, A. Batch, C. Dunne, and N. Elmqvist. Through their eyes and in their shoes: Providing group awareness during collaboration across virtual reality and desktop platforms. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23. Association for Computing Machinery, New York, NY, USA, 2023. doi: 10.1145/3544548.3581093 3

[49] R. Sicat, J. Li, J. Choi, M. Cordeil, W.-K. Jeong, B. Bach, and H. Pfister. Dxr: A toolkit for building immersive data visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):715–725, jan 2019. doi: 10.1109/TVCG.2018.2865152 3

[50] D. Snowdon, E. F. Churchill, and A. J. Munro. *Collaborative Virtual Environments: Digital Spaces and Places for CSCW: An Introduction*, pp. 3–17. Springer London, London, 2001. doi: 10.1007/978-1-4471-0685-2_1 1

[51] W. Stuerzlinger, T. Dwyer, S. Drucker, C. Görg, C. North, G. Scheuermann, K. Marriott, F. Schreiber, K. Klein, N. H. Riche, T. Itoh, and B. H. Thomas. *Immersive Human-Centered Computational Analytics*, pp. 139–163. Springer International Publishing, Cham, 2018. doi: 10.1007/978-3-030-01388-2_5 1, 2

[52] J. J. Thomas and K. A. Cook. *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. National Visualization and Analytics Ctr, 2005. 1

[53] W. Tong, M. Xia, K. K. Wong, D. A. Bowman, T.-C. Pong, H. Qu, and Y. Yang. Towards an understanding of distributed asymmetric collaborative visualization on problem-solving. In *2023 IEEE Conference Virtual Reality and 3D User Interfaces (VR)*, pp. 387–397, 2023. doi: 10.1109/VR55154.2023.00054 3

[54] X. Wang, L. Besançon, D. Rousseau, M. Sereno, M. Ammi, and T. Isenberg. *Towards an Understanding of Augmented Reality Extensions for Existing 3D Data Analysis Tools*, p. 1–13. Association for Computing Machinery, New York, NY, USA, 2020. 2, 3, 8, 9

[55] K. Wongsuphasawat, Z. Qu, D. Moritz, R. Chang, F. Ouk, A. Anand, J. Mackinlay, B. Howe, and J. Heer. *Voyager 2: Augmenting Visual Analysis with Partial View Specifications*, p. 2648–2659. Association for Computing Machinery, New York, NY, USA, 2017. 2

[56] L. Zaman, C. Neugebauer, W. Stuerzlinger, and R. Woodbury. GEM-NI+: Leveraging difference visualization and multiple displays for supporting multiple complex generative design alternatives. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI EA '18, p. 1–6. Association for Computing Machinery, New York, NY, USA, 2018. doi: 10.1145/3170427.3188593 4

A Codes and Thematic Analysis

In this appendix we list the detailed results of the thematic analysis that was mentioned in this section. This list can be seen in Figure 13.

B Full-Size Charts and Figures

To accommodate length constraints, we sized the included visuals above accordingly, but we acknowledge that some readers may find it challenging to discern numbers and labels in some of the images. To address this concern, we show all such figures in this appendix at full size for optimal readability. The following list shows the correspondence of each figure with their full-size version.

- *Figure* 3 → *Figure* 14
- *Figure* 6 → *Figure* 15
- *Figure* 8 → *Figure* 16
- *Figure* 9 → *Figure* 17
- *Figure* 10 → *Figure* 18
- *Figure* 11 → *Figure* 19

Theme	Sub Theme	Code	Numbers of participants where the code occurred
Switching between sub-systems (Total number of code occurances: 37)	Interaction	Looking between AR and Desktop	5
		Using raycaster from controller for analysis in AR	1
		Cognitive load of switching between controller and mouse	1
		Bringing out a graph from desktop	1
	Natural Interactions	Using the controller on the desktop instead of the mouse	3
	Screen Organization	Resizing graph in desktop Sub-system	8
		Needs a notepad	1
		Overlapping graphs, crowded view	1
		Need to have multiple graphs on screen	1
		Organizing the window charts on desktop	2
	Suggestions for AR	Asking to having two controllers	1
	Why Desktop?	Accuracy was better on Desktop	2
		Reading details is easier on desktop	4
		User thinks speed in desktop is higher	1
	Why not switch?	Moving between AR and desktop has cognitive cost	1
		Not trusting the transition of all vis attributes between AR and Desktop	2
		User could answer all questions using the desktop.	1
		Forgot about the switching feature	2
		Filtering does not transfer	2
	Why switch?	Speed on AR/ accuracy on desktop	2
AR sub-system Interactions (Total number of code occurances: 40)	3D Visualization	Rotating the 3D visualization to read it as multiple 2D graphs	4
		Using 3D for overall patterns and clusters	3
	Environment Organization	Arranging graphs in different sides of the environment	6
		Many graphs, still not cluttered in AR	1
		side by side graphs around the user in AR	1
	Natural Interactions,Why AR?	More intuitive or "Natural"	1
	Problems in AR	Filtering is not Precise in AR	2
		Brushing is not supported in AR Histograms	1
		Handling details on AR is difficult	4
		Interaction with graphs: they stop at a distance	1
		Zooming in and stretching the plots buggy or not possible	1
	Why AR?	Data in AR is more tangible	2
		More space in the visual field: lots of graph at the same time	3
		Speed was better with the AR	2
		Working in AR was intuitive, enjoyable, fun	3
		AR was more comfortable and easier for the user	1
		The Graphs are larger(easier to read)	1
		Easier to find dimensions	1
		Needed multi variant analysis and found AR better for that	4
	Why Desktop?	Not comfortable in AR	1
Data analysis and Visualization manipulation (Total number of code occurances: 12)	2+ dimensions	Having multiple layer of filtering(means having multiple dimension of data)	2
		Adding the third dimension makes the analysis complex	3
	Suggestions for AR	Maybe AI help with identifying clusters	1
		Aggregation and sorting data on AR	2
User adaptability and learning (Total number of code occurances: 6)	Why not switch?	Not switching because being more comfortable in one medium	3
		Time pressure make user use the system they are familiar	3

Fig. 13: The four main themes, their sub categories, and the number of the participants that were observed in each code section.

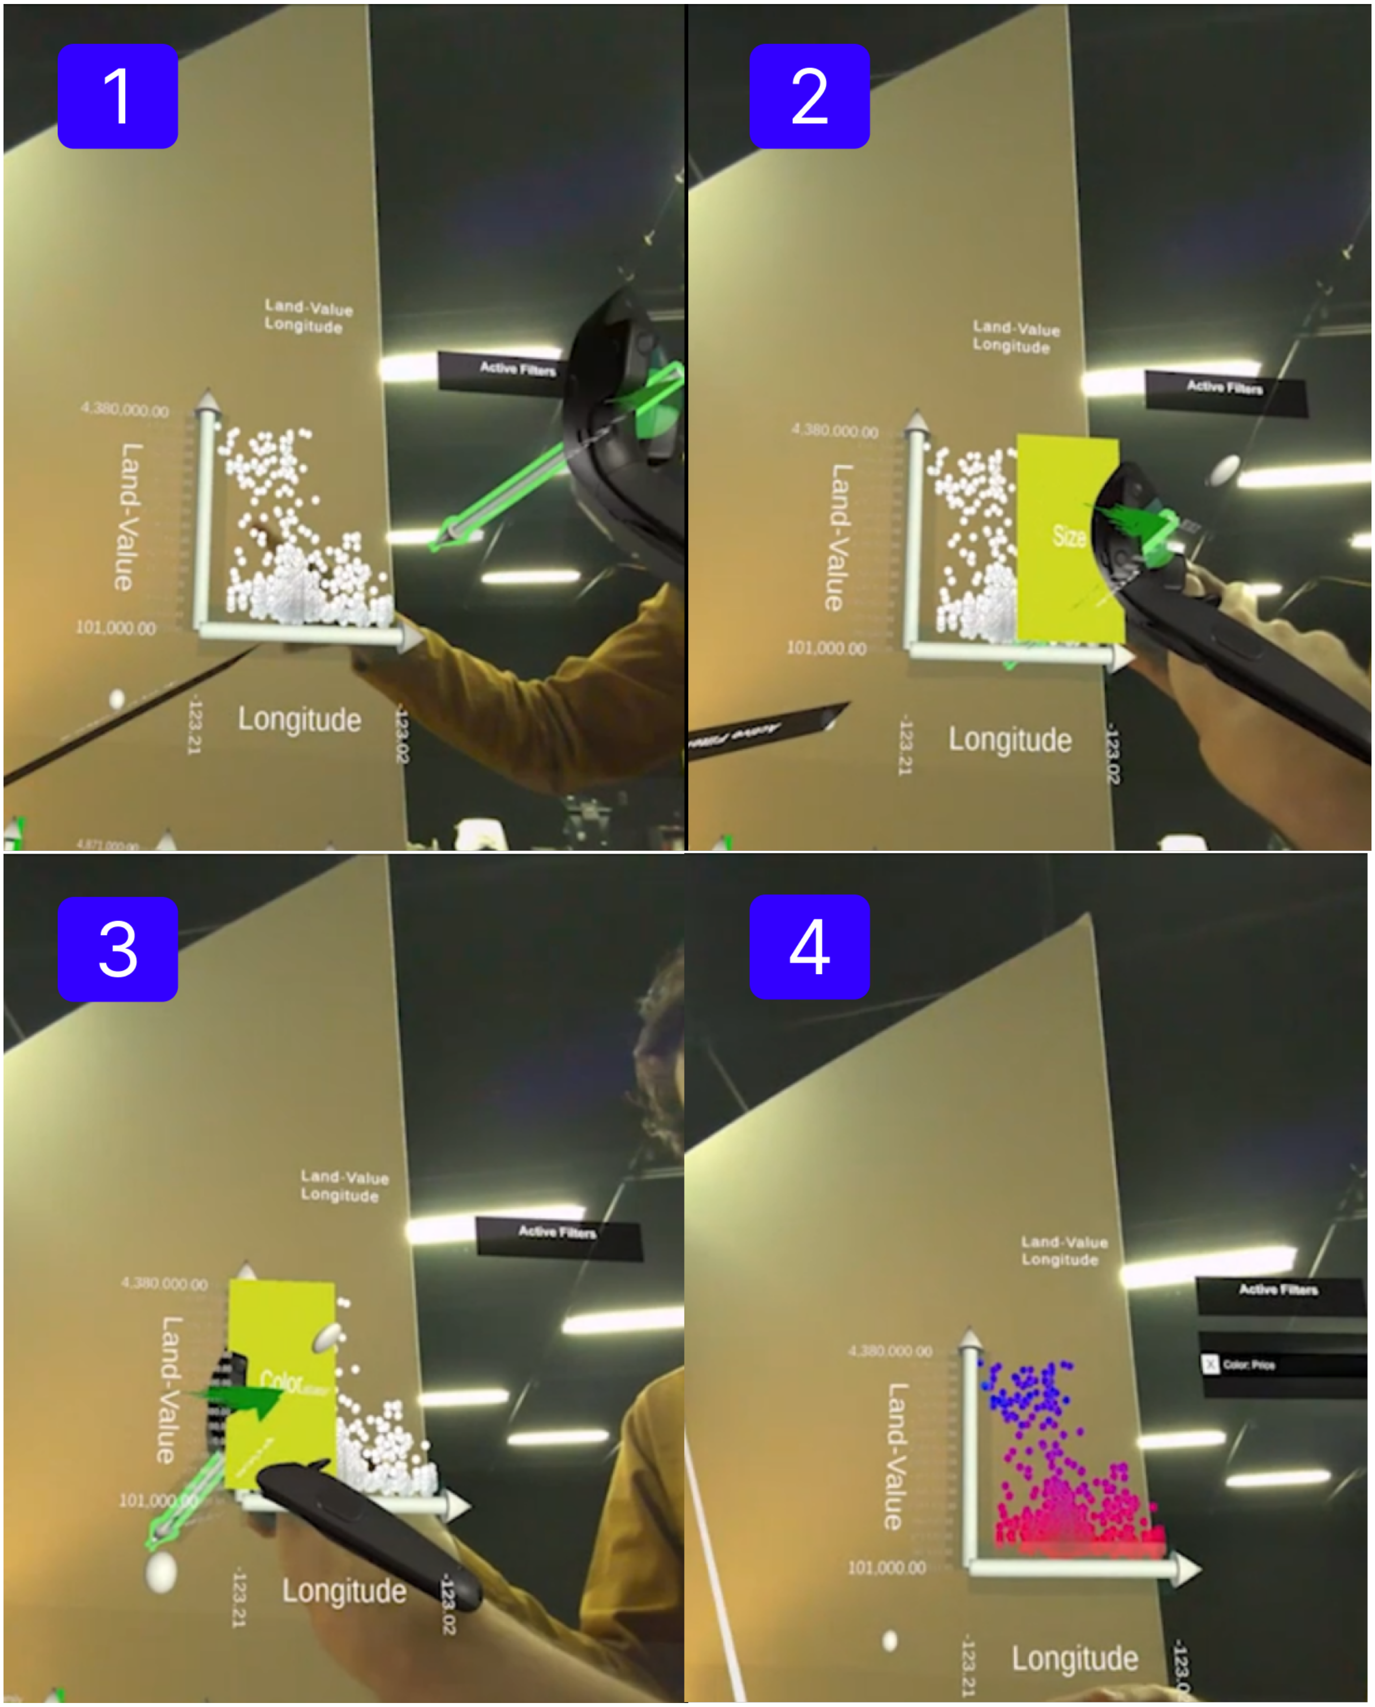

Fig. 14: Enlarged version of Figure 3, showing steps for modifying size/color encodings.

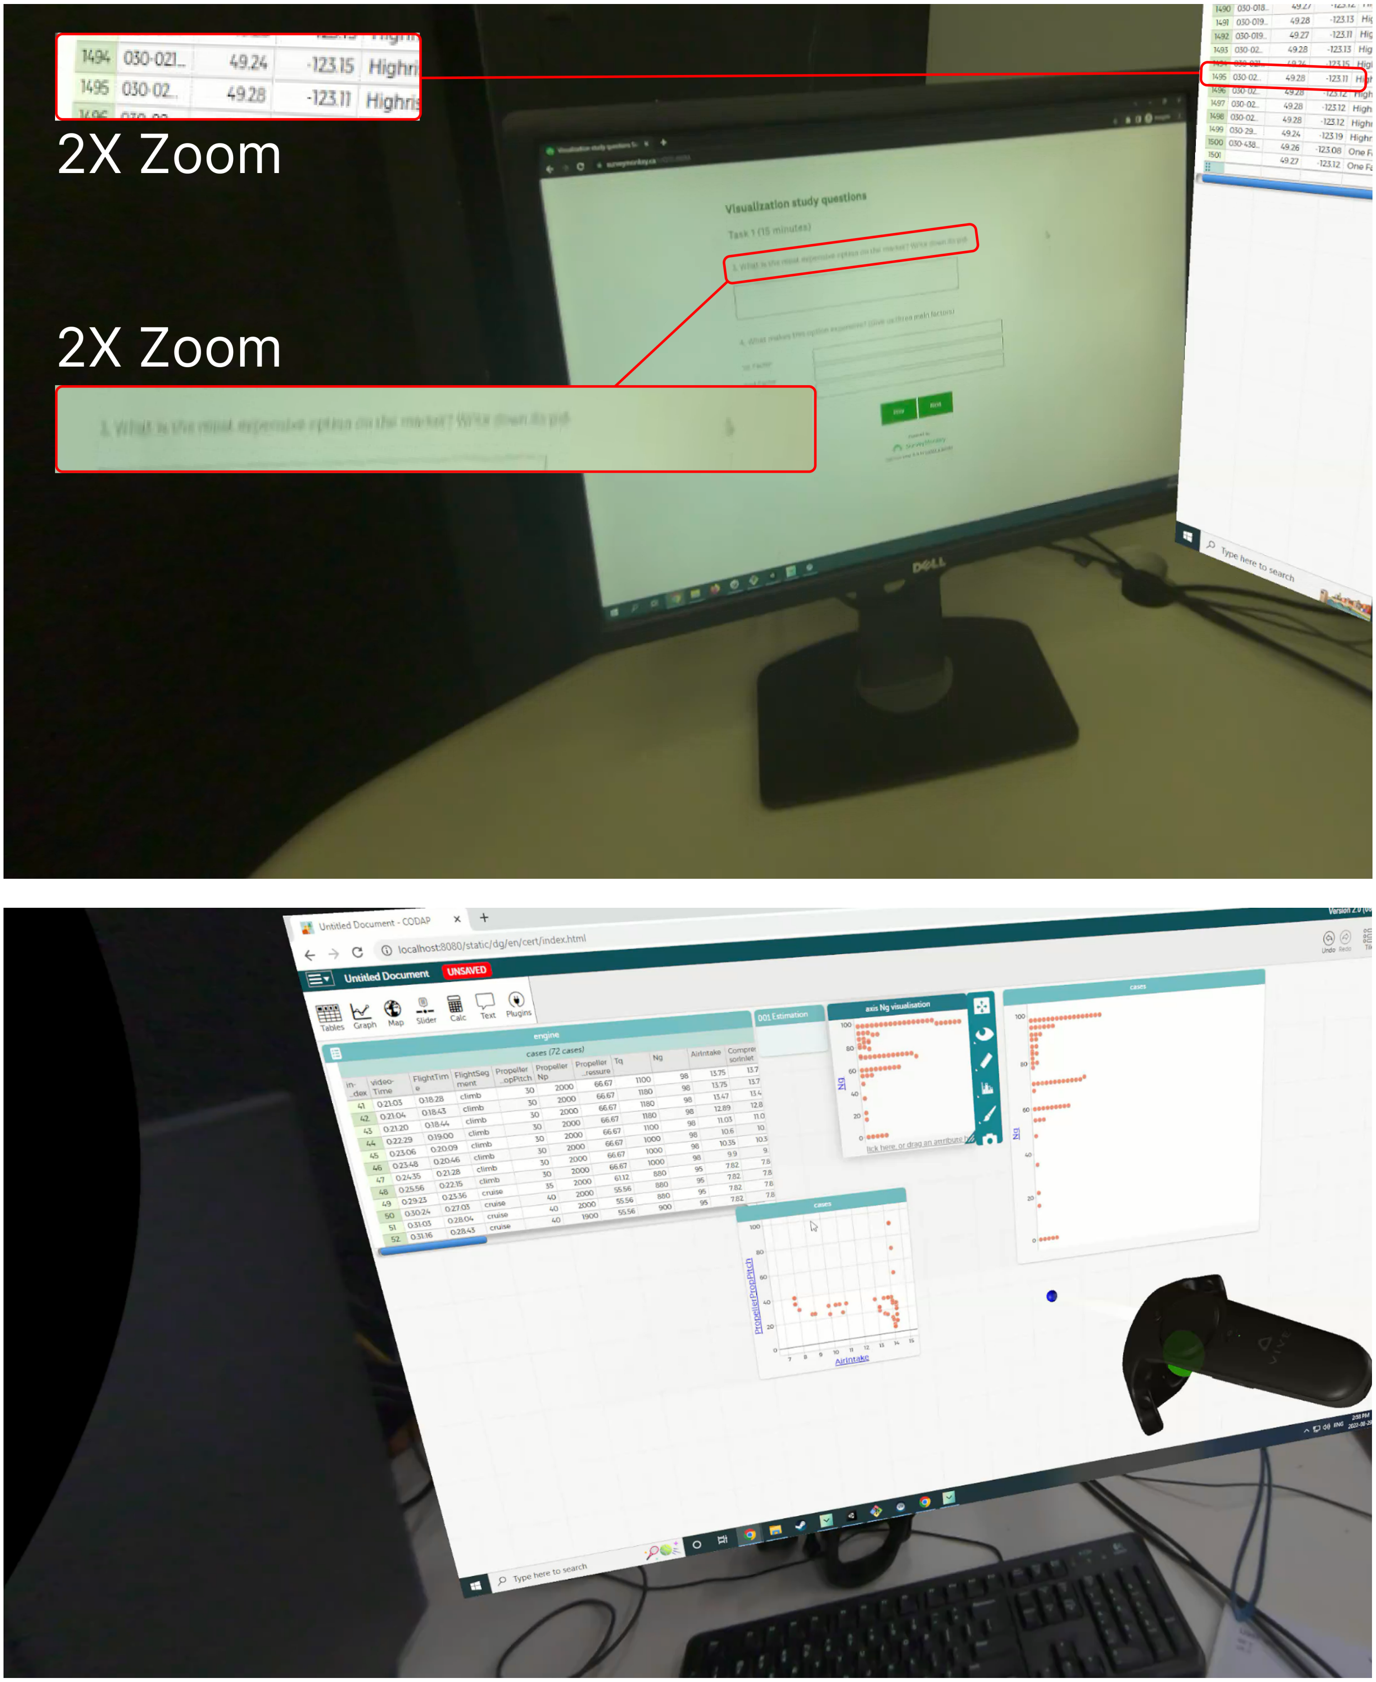

Fig. 15: Enlarged version of Figure 6, showing the readability difference between the real and virtual monitor as seen through the AR headset, when viewed from roughly the same distance.

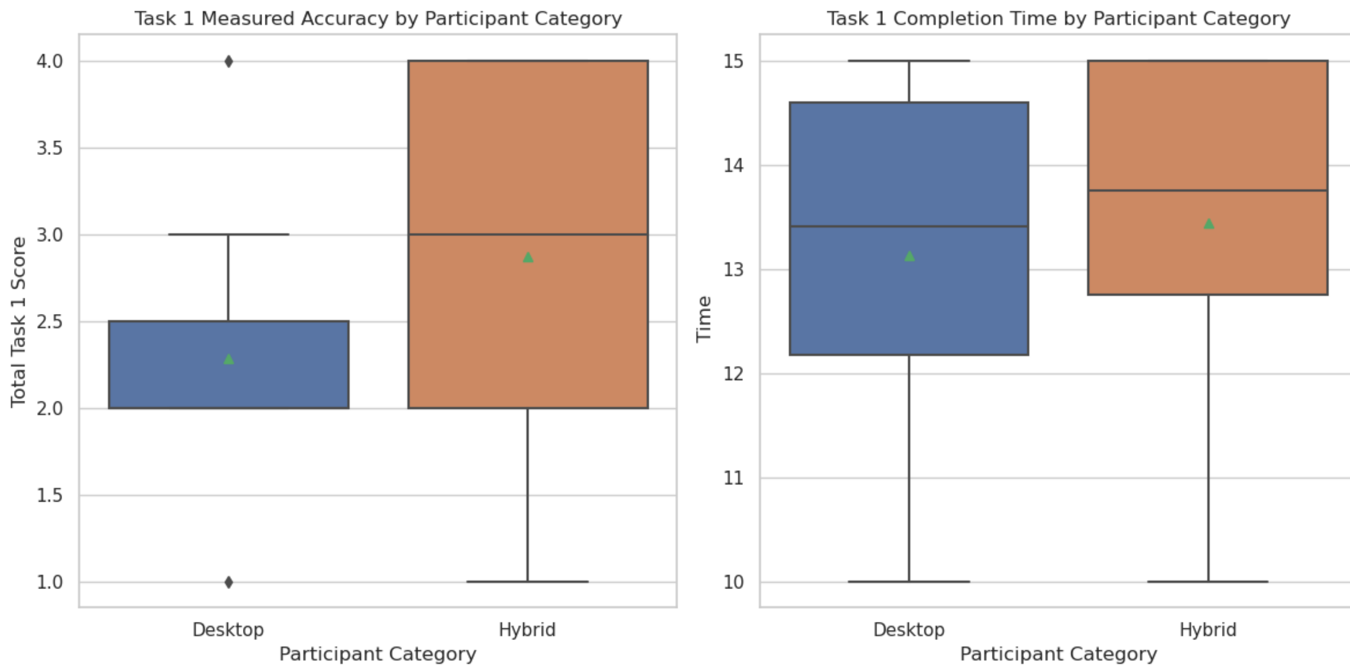

Fig. 16: Enlarged version of Figure 8, showing

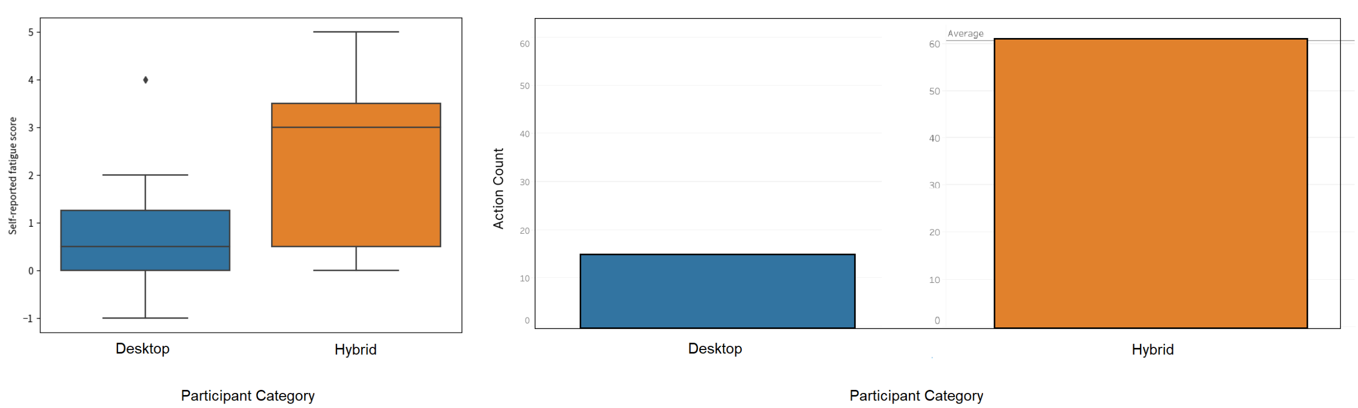

Fig. 17: Enlarged version of Figure 9, showing the accuracy of task 1 across group H and group D.

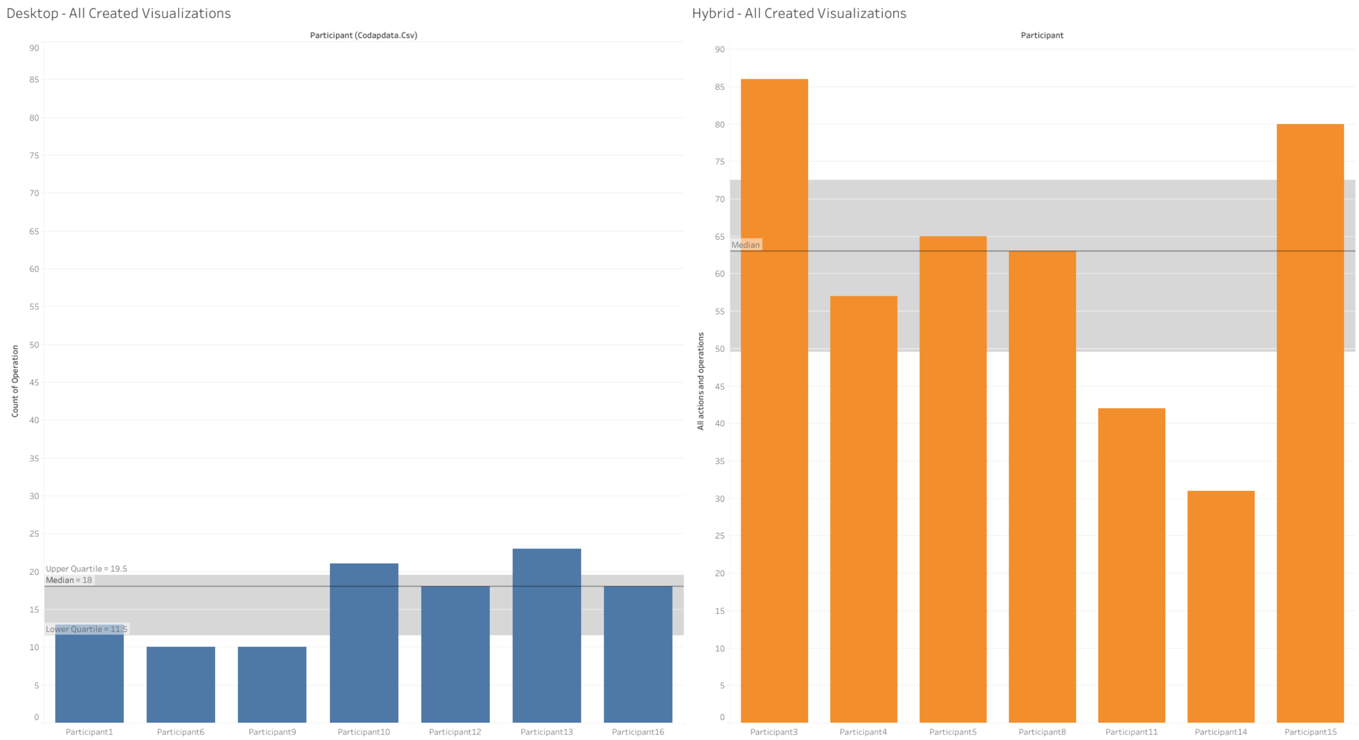

Fig. 18: Enlarged version of Figure 10, showing the comparison of the number of created graphs per participant between group H on the left and group D on the right.

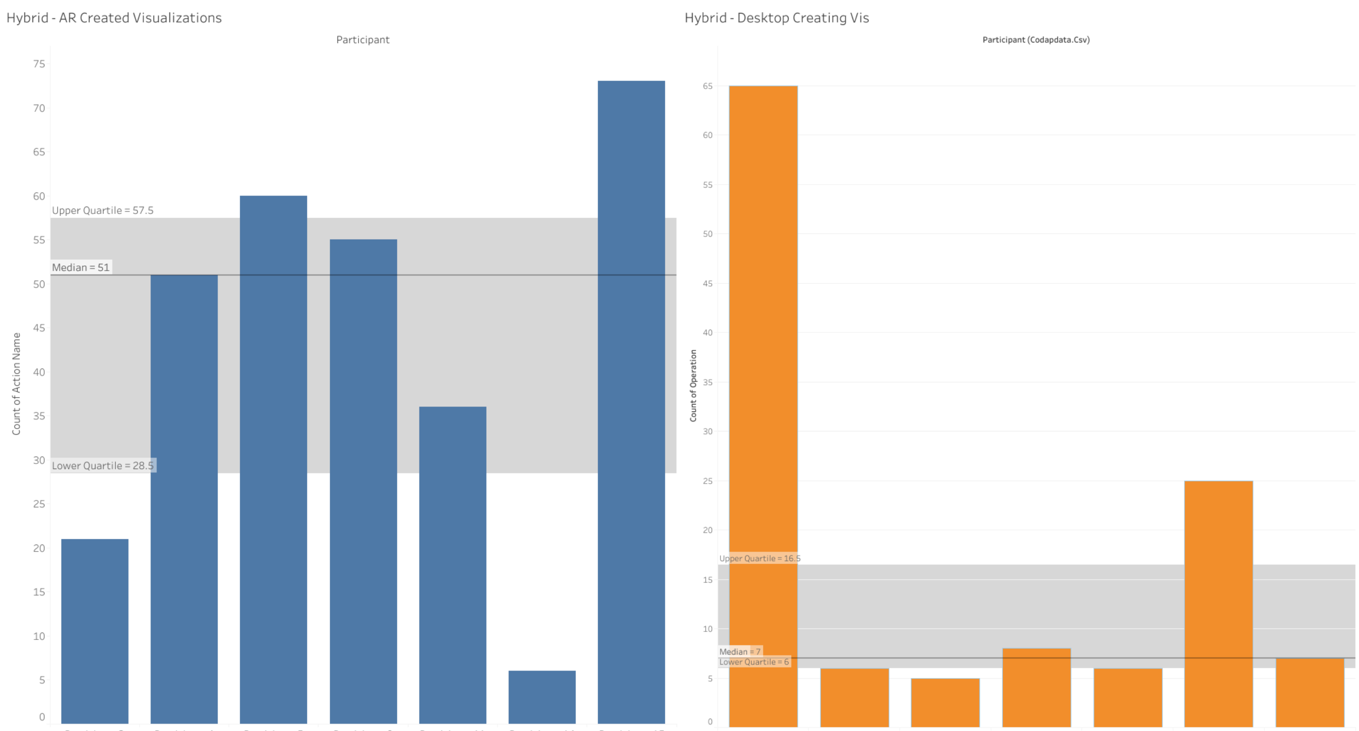

Fig. 19: Enlarged version of Figure 11, showing the comparison of the number of created graph in group H between the AR sub-system (left) and desktop sub-system (right).