

Through Their Eyes and In Their Shoes: Providing Group Awareness During Collaboration Across Virtual Reality and Desktop Platforms

David Saffo
saffo.d@northeastern.edu
Northeastern University
JPMorgan Chase & Co.

Cody Dunne
c.dunne@northeastern.edu
Northeastern University

Andrea Batch
ajulca@umd.edu
University of Maryland

Niklas Elmqvist
elm@umd.edu
University of Maryland




Figure 1: Desktop and VR views in the VRxD application. Screenshots from our virtual reality(VR)/desktop cross-virtuality analytics (XVA) system, VRxD. This system was used as the medium for a collaborative XVA user study examining the role of abstract vs. natural spatial mapping as well as perspective or interaction sharing on collaborative visualization user behavior. The abstract view features a non-immersive desktop (animated figure: ④) and an immersive VR ④ parallel coordinates visualization. Similarly, the natural spatial mapping view features a non-immersive desktop ④ and an immersive VR ④ pitch trajectory visualization. Furthermore, this system implements our four levels of “eyes-and-shoes” group awareness techniques: L1: Landmarks and Analogous Views ④ ④ ④ ④, L2: Information Cues ④ ④, L3: Interaction Sharing ④ ④ ④, and L4: Perspective Sharing ④ ④.



This work is licensed under a [Creative Commons Attribution International 4.0 License](https://creativecommons.org/licenses/by/4.0/).

CHI '23, April 23–28, 2023, Hamburg, Germany
© 2023 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-9421-5/23/04.
<https://doi.org/10.1145/3544548.3581093>

ABSTRACT

Many collaborative data analysis situations benefit from collaborators utilizing different platforms. However, maintaining *group awareness* between team members using diverging devices is difficult, not least because *common* ground diminishes. A person using head-mounted VR cannot physically see a user on a desktop computer even while co-located, and the desktop user cannot easily

relate to the VR user's 3D workspace. To address this, we propose the “eyes-and-shoes” principles for group awareness and abstract them into four levels of techniques. Furthermore, we evaluate these principles with a qualitative user study of 6 participant pairs synchronously collaborating across distributed desktop and VR head-mounted devices. In this study, we vary the group awareness techniques between participants and explore two visualization contexts within participants. The results of this study indicate that the more visual metaphors and views of participants diverge, the greater the level of group awareness is needed. A copy of this paper, the study preregistration, and all supplemental materials required to reproduce the study are available on [OSF \(link\)](#).

CCS CONCEPTS

• **Human-centered computing** → **User interface management systems**; *Visualization systems and tools*; Visualization theory, concepts and paradigms.

KEYWORDS

Asymmetric collaboration, virtual reality, immersive analytics, ubiquitous analytics.

ACM Reference Format:

David Saffo, Andrea Batch, Cody Dunne, and Niklas Elmqvist. 2023. 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)*, April 23–28, 2023, Hamburg, Germany. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3544548.3581093>

1 INTRODUCTION

Group awareness [22], or the ability to see and understand the activities of the overall team, is a key to collaborating effectively [19, 20, 23]. If I know what you are doing, I can better organize my work, and I can also engage with you more directly when I see your work touching mine, thus increasing our *collaborative coupling* [54]. Achieving this kind of awareness is less challenging when inhabiting the same physical space, collaboration medium, and time instance as your partner (i.e. *co-located synchronous collaboration* [3]). However, when collaboration spans multiple devices, geographic distances, and time intervals, maintaining group awareness invariably becomes more complex [19]. For example, past work has studied group awareness during both fully-distributed [57] as well as in mixed-presence (hybrid distributed and co-located) settings [32, 52, 58]. This complexity is exacerbated when collaborators use different types of devices and communication media because of the diminished *common ground* [9] between platforms. For example, imagine coordinating between a smartphone user and a desktop user on the same visualization dashboard; because of the limited screen size, the mobile user will see much less of the dashboard at any point in time—perhaps only a single chart—and thus spatial relations and deixis [26] across the platforms are inconsistent.

This gap in common ground, a central design consideration for collaborative visualization [24], widens as the platforms diverge in their physical form factor, display mode, and interaction model. It is particularly problematic when collaborating across the reality–virtuality continuum (RVC) [15]—such as across desktop and

virtual reality (VR) head-mounted devices. At the same, engaging multiple device platforms in collaboration is often beneficial [27], especially when combining traditional 2D views with immersive 3D views [28, 59]. With the rise of immersive analytics [38], virtual and mixed reality (MR) has seen increasing use as a platform for visualization and visual analytics, including for graph visualization [11], multidimensional data [10], and even economical analysis [4]. VR devices present unique affordances for stereoscopic vision, spatial interaction, and immersive environments, providing key advantages over desktop platforms for visualization workflows such as high-dimensional or natural spatial mapping data analytics. Likewise, desktop platforms have unique advantages for these same workflows—such as text entry, precise interaction, and easier access to external tools. Through cross-platform collaboration, team members can leverage the strengths of their platforms and balance their weaknesses. For example, consider the following vignette inspired by SieVRt (<https://sievr.com/>), a VR system for remote and collaborative radiology:

Imagine two physicians, Alice and Bob, visualizing cerebral arteries captured using medical imaging (CT scans). Abstract 2D networks can quickly identify areas of interest, but 3D natural spatial mapping renders will always need to be cross-referenced and consulted as a part of the complete radiology workflow [42]. A *cross-platform collaboration* workflow would enable Alice to use a desktop computer and Bob to use a Virtual Reality headset to work together. Each platform can now be used to its strengths, with Alice identifying areas of interest using her 2D overview and Bob investigating them in-situ in his immersive 3D view.

Alice uses the abstract representation of the computers and often quickly search for nothing else, and no monitors present and not stereoscopic. VR/AR displays are fundamentally 3D. Furthermore, the immersion to Bob's perspective neglects the broad range of display abnormalities that is difficult for a VR user to expect to glance at collaborators' 3D view of the data and easily with the same outside or inside the 3D as Alice viewing as if they happen to occupy the same physical space. This challenge raises the question: “how do we help collaborative work across different views of the same data?”

In this paper, we explore a design principle for providing group awareness across heterogeneous devices that we call “*through their eyes and in their shoes*” (eyes-and-shoes). Alice shares the data as a 2D image to Bob, and Bob shares the data as a 3D image to Alice. By detailing 3D representation and how collaborators can use it to demand different things. To clarify, Alice hovers over the abnormality and group question with the 2D view of the data. This interaction, video highlights the remote collaborator's view with a glowing effect, intent based on the 2D and 3D views acting as a *shared ground* context. This is the generalized formulation of the eyes-and-shoes model for further use and its application to cross-platform collaboration between desktop computers and Virtual Reality devices. Furthermore, in this paper, we also explore the design space of eyes-and-shoes by partitioning it into four levels of techniques.

Bob and Alice discuss the abnormality (in question) and visually altering their local visualization state to filter or display different images and data layers. Bob begins to refer to data values and color encodings that do not match Alice's current view of the brain vessel dimensions. Alice uses the XVA system's feature, allowing users to toggle synchronization of the local interaction state with collaborators. This feature sets Alice's filters to what Bob has selected, which reveals that Bob is currently viewing the blood flow data layer. This *interaction sharing* allows Alice to understand the context of Bob's reference and achieve common ground without asking

Alice
abnor
other
munc
featur
collab
that s
view
With
under
with c

To assess the validity of the eyes-and-shoes collaboration principle, we implement its basic philosophy in a cross-platform synchronous collaborative visualization tool for head-mounted virtual reality and desktop platforms called VRxD. The impetus for this specific instantiation of the principle is that many immersive analytics applications involve data that would benefit from both 2D and 3D visualization [38]. This suggests a hybrid approach where one collaborator working on a desktop computer and one in a VR immersive analytics environment could be particularly powerful [37]. Using this VRxD platform, we conducted a remote user study involving 6 pairs of collaborators exploring a baseball pitch and pitcher statistics datasets in tandem. We varied the interaction and view models of the collaborators, thus changing the amount of group awareness and common ground between them. Additionally, we explored two distinct visualization scenarios; abstract and natural spatial mapping visualizations. Through this study, we share our observations on collaborative visualization behavior within this cross-virtuality analytics [43] (XVA) system and the context of our “eyes-and-shoes” design principles. These results will help inform the design of future XVA systems and provide a behavioral baseline for future studies of XVA group awareness techniques. Specifically, this paper contributes:

- (1) The **“through their eyes and in their shoes” design principle for providing group awareness across heterogeneous devices** by letting users step into the views of their collaborators to maintain common ground. While these are not new ideas, our conceptualization and suggested four levels of group awareness generalize them.
- (2) **The design and results of a qualitative user study** validating the “eyes-and-shoes” principle and our four levels of group awareness techniques, using our prototype XVA system VRxD. Our results indicate that the more visual metaphors and views of participants diverge, the more group awareness needs to be supported.

This paper and all supplemental materials are freely available on [OSF \(link\)](#).

2 RELATED WORK

Here we describe the related literature grounding our research. In particular, we cover previous contributions in *collaborative visualization* (CV) and *cross-virtuality analytics* (XVA)—as they relate to *immersive analytics* (IA).

2.1 Collaborative Visualization

Collaborative visualization (CV) is a long-standing research thrust pursued by data visualization researchers and engineers alike [55]. Isenberg et al. [30] define CV across time (synchronous or asynchronous) and space (co-located or distributed). The cross-sections of these two dimensions predominantly determine how, when, and where we think about the design and research of CV systems. Heer and Agrawala [24] present seven areas of consideration for CV: division and allocation of work; common ground and awareness; reference and deixis; incentives and engagement; identity, trust, and reputation; group dynamics; and consensus and decision making. They discuss these considerations within the context of distributed,

asynchronous collaboration. However, the broader ideas of these areas are relevant across the design space of CV—leading researchers to explore them across different contexts and devices [1, 28, 35].

The majority of existing CV literature focuses on non-immersive 2D displays. Immersive 3D display technologies, such as augmented reality (AR) and virtual reality (VR), blur the lines of the CV design space presented by Isenberg et al. [30]. More recent work has begun to refer to collaboration across different devices as “symmetry” [48], distinguishing this aspect from space and time. VR users can be simultaneously physically distributed and virtually co-located—requiring research and design across both categories of CV. This increased complexity has made CV one of the grand challenges for IA [13]. Several past publications have explored CV with homogeneous immersive devices [29, 41]. However, there is a growing interest in CV across heterogeneous immersive and non-immersive devices, leading to the designation of supporting cross-platform collaboration as a significant component of this challenge.

2.2 Cross-Platform Visualization

Cross-platform collaborative visualization utilizes several devices across the reality-virtuality continuum (RVC) [40] to conduct visual analytics synchronously or asynchronously with multiple users. Examples of existing literature exploring this area include work utilizing a mixture of mobile and large-scale displays [27], AR and tabletop displays [49], and VR and desktop displays [28]. Fröhler et al. [15] survey and categorize these works as cross-virtuality analytics (XVA)—or more generally cross-virtuality (XV) [2]. The authors define XVA as “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.”

Core to the philosophy of XVA is the opportunity to provide the optimal techniques, encoding, interactions, and view of data using tailored visual metaphors with various devices depending on the task at hand [15, 43]. In particular, combining different devices in XVA has been shown to enable complementary use, where devices mutually scaffold each other’s weaknesses [27]. Fröhler et al. [15] further categorize XVA works into four categories: spatially agnostic (simultaneous use of devices), augmented (displays extended and spatially orientated), networked (collaborative), and transient (switching between realities). The authors conclude by outlining a series of challenges and opportunities, including collaboration across the RVC—highlighting the importance of group awareness and the lack of generalizable frameworks for XVA awareness cues.

2.3 Cross-Platform Collaboration

Our work in this article represents networked XVA across the desktop and VR stages of the RVC. Combining desktop, and immersive views is often beneficial because it enables an ex-situ 2D analytical view to be combined with an in-situ 3D first-person view. For example, ReLive [28] enables a user to view 3D trace data from a user study both in an immersive VR view, as well as switch to a 2D desktop view for an overview and summative analysis. Similarly, Wang et al. [59] seamlessly combine AR and desktop interfaces in a particle visualization application where the interface work together yet provide unique individual capabilities. However, both of the

above are “hybrid” desktop/immersive interfaces intended for a single user.

Cross-platform collaborative visualization systems enable multiple people to work together using different devices. Only a few studies exist in this space. Johnson et al. [31] present an experiment where a tablet participant, given an abstract scene representation, can assist a AR HMD participant during locate tasks using deixis—finding that a list view out-performed a spatial layout. Schroeder et al. [46], in a study involving desktop and CAVE users working together, find that the display and input technology used by collaborators have a significant impact on leadership and tasking. Tong et al. [56] present a comparative study examining symmetric and asymmetric visualization views across desktop and VR collaborators. The results of their study shows the promise of well designed asymmetric visualization views for task productivity and mental demand.

Most relevant to our work is that of Reski et al. [42] on an empirical evaluation of synchronous collaboration between a desktop user and a VR user using asymmetric visualization design. Collaborators were provided visual cues and shared interactions across the desktop and VR interfaces, such as highlighting shared data and providing the VR participant’s location and field of view for the desktop participant. The authors report results and observations that indicate that users benefit from this collaboration model and report a high level of group awareness during the task.

3 THE EYES-AND-SHOES PRINCIPLES

The common idioms “*Walk a mile in their shoes*” and “*Try seeing it from their point of view*” synthesize the human experience of attempting to understand another person’s perspective, experience, and motivation before passing judgment. This action is a crucial step for decision-making, diplomacy, and—most relevant to this paper—collaboration [18]. In collaborative work in general, and collaborative visualization (CV) in particular, this is known as *common ground* [9]: the sum of mutually known beliefs, knowledge, and suppositions among the participants in an exchange.

Common ground can be thought of as consisting of both *static* and *dynamic* components; knowledge that exists prior to a collaborative session versus time-sensitive knowledge that arises during it, respectively. Establishing the static common ground is fundamental to effective collaboration but trivial from an immersive and collaborative technology perspective. Maintaining the dynamic aspect, however, requires providing *group awareness* [21, 22]: the ability to perceive the activities of the entire team of collaborators. Group awareness can be achieved through local or global perspective sharing in co-located or distributed settings with homogeneous devices [20, 21, 23].

In this paper, we explore the idea that as collaborators’ devices and perspectives diverge, the dynamic common ground diminishes when group awareness is reduced or non-existent. Put differently, when another person’s perspective, experience, and motivations diverge from your own, it becomes increasingly difficult to see things “through their eyes and in their shoes.” Thus, providing group awareness on heterogeneous devices is critical for cross-virtuality analytics [15], such as collaborating across desktop and VR platforms. As discussed in Section 2, providing group awareness

for collaborative visualization to facilitate common ground is not a new idea but draws from many existing systems that provide group awareness, such as map overviews [31], remote presence [32, 42], and avatars [34, 46]. Thus, while the concept is not original, we find our “eyes-and-shoes” moniker to be particularly apt: the “eye” part can represent the visual cues, widgets, and views that provide awareness, and the “shoes” part are interactions that the user performs to gain awareness.

Based on this unified model of group awareness for both input and output, we propose four levels of interventions organized from least to most awareness. Below we first describe each level and cite its inspirations drawn from the literature. We also discuss its intrusiveness, which is relevant when factors other than collaboration efficiency are essential (such as presence and immersion in VR). Finally, we exemplify the level within a hypothetical cross-platform visualization system for two physicians, Alice and Bob, collaborating on a brain scan in 2D and 3D, respectively.

• Level 1: Landmarks and Analogous Views. Divergent devices

Imagine two physicians, Alice and Bob, visualizing cerebral arteries captured using medical imaging (CT scans). Abstract 2D networks can quickly identify areas of interest, but 3D natural spatial mapping renders will always need to be cross-referenced and consulted as a part of the complete radiology workflow [42]. A *cross-platform collaboration* workflow would enable Alice to use a desktop computer and Bob to use a Virtual Reality headset to work together. Each platform can now be used to its strengths, with Alice identifying areas of interest using her 2D overview and Bob investigating them in-situ in his immersive 3D view.

Alice uses the abstract representation of the brain scan to quickly search for and identify an abnormality exposed by visual encodings such as color and size. Alice attempts to relay this information to Bob by referencing the brain region of the abnormality using its color and relative size. Bob has an analogous 3D view of the brain scan, with the same labels and colors as Alice acting as *landmarks*. As a result, Bob can direct his attention to the correct area of interest, establishing common ground with Alice.

• Level 2: Information Cues. Visual cues (i.e. cursor position, focus, notifications) shared across views show where collaborators are looking and working. Cues provide common ground for collaborators and are useful for the division and allocation of tasks and when the 2D and 3D views are providing information for XVA systems is less straightforward as there often are no one-to-one mappings of interactions and input modalities across devices. For example, there is no clear way to represent the cursor position from a 2D visualization to a 3D visualization, even when using the same data. Instead, XVA systems need to contextually represent these information cues, enabling brushing and linking behaviors.

Alice and Bob discuss the setting that collaboration, field of view, and filters. These local visualization features to filter or display by providing context and hints to all. Bob begins activities to alter the current view of the brain vessel dimensions. Alice uses the XVA system’s feature, allowing users to toggle synchronization of the local interaction state with collaborators. This feature sets Alice’s filters to what Bob has selected, which reveals that Bob is currently viewing the blood flow data layer. This *interaction sharing* allows Alice to understand the context of Bob’s reference and achieve common ground without asking for clarification directly.

Alice
abnor
other
munc
featur
collab
that s
view
With
under
with c

Alice uses the abstract representation of the brain scan to quickly search for and identify an abnormality exposed by visual encodings such as color and size. Alice attempts to relay this information to Bob by referencing the brain region of the abnormality using its color and relative size. Bob has an analogous 3D view of the brain scan, with the same labels and colors as Alice acting as *landmarks*. As a result, Bob can

Imagine two physicians, Alice and Bob, visualizing cerebral arteries captured using medical imaging (CT scans). Abstract 2D networks can quickly identify areas of interest, but 3D

Imagine two physicians, Alice and Bob, visualizing cerebral arteries captured using medical imaging (CT scans). Abstract 2D networks can quickly identify areas of interest, but 3D natural spatial mapping renders will always need to be cross-referenced and consulted as a part of the complete radiology workflow [42]. A *cross-platform collaboration* workflow would enable Alice to use a desktop computer and Bob to use a Virtual Reality headset to work together. Each platform can now be used to its strengths, with Alice identifying areas of interest using her 2D overview and Bob investigating them in-situ in his immersive 3D view.

visual encodings such as color and size. Alice attempts to relay

Alice uses the abstract representation of the brain scan to quickly search for and identify an abnormality exposed by visual encodings such as color and size. Alice attempts to relay this information to Bob by referencing the brain region of the abnormality using its color and relative size. Bob has an analogous 3D view of the brain scan, with the same labels and colors as Alice acting as *landmarks*. As a result, Bob can direct his attention to the correct area of interest, establishing common ground with Alice.

detailed 3D representation and does not know which Alice

Bob has located the area of interest Alice wants to discuss. However, Bob can see several potential abnormalities with the detailed 3D representation and does not know which Alice is referencing. To clarify, Alice hovers over the abnormality in question within the 2D view of the data. This interaction highlights the selected blood vessels with a glowing effect in both the 2D and 3D views acting as an *information cue*. Thus, both confirm that they look at the same abnormality, further reinforcing common ground.

different images and data layers. Bob begins to refer to data

Bob and Alice discuss the abnormality in question, individually altering their local visualization state to filter or display different images and data layers. Bob begins to refer to data values and color encodings that do not match Alice’s current view of the brain vessel dimensions. Alice uses the XVA system’s feature, allowing users to toggle synchronization of the local interaction state with collaborators. This feature sets Alice’s filters to what Bob has selected, which reveals that Bob is currently viewing the blood flow data layer. This *interaction sharing* allows Alice to understand the context of Bob’s reference and achieve common ground without asking for clarification directly.

ing,” where the display output of collaborators is shared across devices. However, perspective sharing across heterogeneous devices represents an ample design space that warrants further exploration and possibly requires novel techniques. For XVA systems, perspective sharing is the most intrusive method of group awareness, requiring collaborators to look away from their view to watch the display output of their peers (shifting context).

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

Inspiration and the (Tra)jectory of Clinical Research regarding the (ab)normality of **How One Body** has continuously been referencing other potential abnormalities in the same vicinity that look

Alice and Bob are drawing close to a consensus regarding the abnormality. However, Bob has continuously been referencing other potential abnormalities in the same vicinity that look mundane from Alice's view. Alice uses the XVA system's feature that allows users to view the display output of their collaborator's device. This feature opens a window for Alice that shows the feed from Bob's HMD device, providing a clear view of the potential abnormalities Bob has been referencing. With the detailed 3D view and Bob's *perspective sharing*, Alice understands Bob's perspective and can discuss the matter with complete group awareness.

4 STUDY METHODS

We designed and conducted a synchronous remote qualitative exploratory user study to examine participants' collaborative visualization (CV) behavior across desktop and VR devices. Our goal for this study was not to compare this collaboration model to alternatives or prove collaborative visualization's utility. Instead, we aim to observe, categorize, and abstract collaborative behavior within this model and how different scenarios, decisions, and interventions influence this behavior. These research artifacts will allow us to understand how users collaborate when given asymmetric views of the same data, how our design principles work in practice and contribute to a well-informed discussion on designing better future systems. The following subsections will detail the between and within-subject factors, visualization design, experimental task, recruitment and participants, procedure, measures, and analysis.

4.1 Preregistration

We preregistered our study design on OSF before collecting data. This preregistration is available online at [OSF \(link\)](#). Our preregistration reflects our initial study design, which involved two separate experiments, one using baseball data and the other using food recipe data. We have deviated from this design by using baseball data for both visualization contexts and studying this aspect within-subject rather than between-subject. Furthermore, per a reviewer's suggestion, we appended our analysis to code for leadership behavior. Aside from these changes, our preregistration reflects our study design and methods.

4.2 Evaluation Platform: VRxD

To support this user study, we implemented a dedicated evaluation platform supporting collaborative desktop and virtual reality that we call VRxD. We abstracted the system’s application and visualization state information into a JSON object-store (e.g. filters, focus, data, FOV). The system uses separate network layers for data communication and video streaming. Figure 1 shows the system in action for both desktop and VR views.

VRxD is built using the front-end JavaScript framework *Svelte* for reactive declarations and document object model (DOM) manipulation; *Aframe.js* for a system-entity-component abstraction of WebXR [61] and *Three.js*; and *D3.js* [6] for data visualization

functionality. For data communication, we used real-time document synchronization provided by deepstream.io, a centralized client-server API. With this API, the interaction state object-stores from desktop and VR clients were separately shared and synchronized. Additionally, to implement display output sharing from desktop to VR, we used peer.js, a peer-to-peer network library, to share the live video feed over WebRTC.

4.3 Visualization Contexts, Data, and Design

Immersive analytics (IA) visualization design literature commonly focuses on one of two visualization contexts [33]. The first is abstract data visualization, defined by encoding data with or without a natural structure using abstract visual elements (i.e. bar charts, line graphs, and node-link visualization). The second is natural spatial mapping, defined by encoding data with a natural structure exactly or close to its real-world appearance (i.e. volumetric, scientific, geographic, and trajectory visualization).

Of course, 3D visualizations can be viewed from a 2D display and vice versa. However, doing so comes with limitations, such as spending more time manipulating the perspective of a 3D view [4] in 2D or difficulty reading charts precisely with a current 3D display [60]. A key motivation for this work is that collaboration across devices could be vital to leveraging the strengths of different platforms while minimizing their weakness. Furthermore, while planning the study, we made the following assumptions:

- **Abstract data visualizations** should be *easier to interpret and interact with for a desktop user*, and there would be less divergence in the visualization design between the 2D and 3D views.
- **Natural spatial mapping visualizations** should *grant the advantage to the VR user*, but at the expense of more divergence between the 2D and 3D views.

As a result, we believed that **both** visualization contexts could affect the collaborative behavior of VR and desktop users [46]. Examining these contexts allows us to discover if different behaviors or levels of group awareness should be considered when designing XVA systems. We design a VR and a desktop version for each of these two contexts, yielding a total of **four** separate versions.

4.3.1 Data. We selected baseball pitching data for both visualization contexts. Baseball's multidimensional range of data types, availability, and relative familiarity with U.S. participants made it a suitable choice for our purposes. Pitching data can come in various forms, ranging from summary and predictive statics describing a pitcher's historical performance to descriptive statistics and trajectories describing the details of individual pitches. As a result, many data visualization techniques are suitable for this data—and baseball analytics (a.k.a. sabermetrics) heavily utilizes 2D and 3D visualizations. These characteristics provided us with ample design space to implement within and many examples of existing baseball visualizations in which to ground our design—such as the visualizations created by Baseball Savant.¹

For the abstract data context, summary and predictive statics of the top and bottom 10 Major League Baseball (MLB) pitchers by

xba² from the 2021 season was used. For the natural spatial mapping context, descriptive statics and trajectories of a MLB pitcher's individual pitches from a single game were used. This data was collected from public repositories of baseball data provided by Baseball Savant.³ Relative to data with similar features from other domains (e.g. astronomy, medicine, multi-omics), baseball data requires a lower level of expertise to engage with and will be familiar to a larger population of potential participants.

4.3.2 Abstract Visualization Design. The abstract visualization context features a parallel coordinates visualization (Figure 1 left) encoding summary and predictive statistics of MLB pitcher's performance from a single season. Parallel coordinates are well suited for visualizing this type of multivariate data and will allow participants to explore the relationships between statics and overall pitcher performance. Additionally, immersive 3D parallel coordinates have been explored in academic literature [5, 8, 45, 50, 51] and demonstrate several unique advantages, such as encoding additional dimensions without introducing excessive over-plotting.

The **desktop non-immersive view** presents a standard parallel coordinates plot where, for a given dimension, each pitcher's data is encoded with a circle along an axis, and lines connect a pitcher's encoded data across several dimensions (Figure 1 top left). The visualization allows for additional dimensions to be added or removed and for the order of dimensions to be changed. Furthermore, details on demand are provided when hovering over a line, highlighting the selected pitcher across all dimensions and displaying the encoded data in tabular form.

The **immersive VR view** follows a similar design as the desktop view. However, it features a 3D parallel coordinates encoding inspired by the visualization design presented by Butscher et al. [8]. The design was simplified for the VR environment. Data is plotted on scatterplots fixed to 2D planes in the 3D environment for each dimension. Points on the chart are sorted by value along the X and Y axis (Figure 1 bottom left). Respective points on each neighboring scatterplot are connected with lines. Just like the 2D version, dimensions can be added and removed, and the order of dimensions along the Z axis can be changed. In both views, the color scale encoded on the connecting lines can be changed to the relative value between dimensions or the values in the first dimension.

4.3.3 Natural Spatial Mapping Visualization Design. The natural spatial mapping visualization (Figure 1 right) features baseball pitch trajectories with related summary and predictive statistics as abstract visualizations. This design was inspired by the many examples of baseball pitch statistic and trajectory visualization found on websites such as Baseball Savant⁴ and Fan Graphs.⁵

The **non-immersive desktop view** is designed as a dashboard using popular baseball pitch visualizations commonly found on websites reporting baseball data. The central visualizations are the pitch trajectories separated into three plots representing a front, side, and overhead view of the pitches. The trajectory visualizations

²Expected batting average, or xba, is a predictive statistic describing the average hits batters are expected to achieve versus a pitcher.

³https://baseballsavant.mlb.com/statcast_search

⁴<https://baseballsavant.mlb.com/visuals>

⁵<https://www.fangraphs.com/>

¹<https://baseballsavant.mlb.com/visuals>




Figure 2: Participant pairs were provided either a L3 interaction sharing feature “peek” (animated figure: ④), or a L4 perspective sharing feature “watch” ⑤. The above-mentioned features could be used to provide local interaction state synchronization or global perspective sharing—presented with the pitch trajectory and parallel coordinates visualization designs, respectively.

are supported with additional visualizations and tables encoding summary statics, pitch speed and frequency, and pitch movement.

The **immersive VR view** shares all of the same features as the desktop version, except the trajectories are visualized in 3D real-life scale. The supporting plots encoding summary statics, pitch speed and frequency, and pitch movement can be viewed in organized clusters fixed to 2D planes in the 3D environment placed on either side of the user. These clusters could be freely moved by the user.

Both views allow coloring and filtering pitches by pitch type, pitch outcome, and pitch speed. Furthermore, hovering on visuals on any of the charts make them stand out, and clicking on them animates the pitch real-time speed and movement across the trajectory(s).

4.4 Group Awareness Techniques

We modulate two levels of group awareness techniques provided to participants between subjects so that half of the participant pairs would be given **Level 3** interaction sharing techniques (L3), and the other half **Level 4** perspective sharing techniques (L4). This enables us to observe how each technique provides group awareness and how different levels of group awareness might affect behavior. We opted to only model L3 and L4 in our group awareness factor levels because we felt L1 and L2 were the minimum level of interventions needed for collaboration—and are realistically present in most collaborative systems. Here we outline how we implemented each level of group awareness:

- **L1: Landmarks and Analogous Views** was used extensively throughout both visualization designs. Where possible, user interfaces (UI) (Figure 1 ④ ⑤) and visualization metaphors (Figure 1 ⑥) were shared between desktop and VR implementations. This came from supplementary visualizations, control UI, legends, color scales, titles, captions, and documents.
- **L2: Information Cues** were provided to all participants by enabling shared brushing and linking across views. When one

participant hovered or focused on a data point, it would be highlighted in their collaborator’s view.

- **L3: Interaction Sharing** allows each view’s interactions and filters to be synced and displayed on the opposite view. Half of the participant groups will be provided a “peek” (Figure 1 ④) button, which completely synchronizes the interactions of their collaborator while the feature is toggled and returns to the previous state when released. Furthermore, all participants will be provided a “copy” (Figure 1 ⑤ ⑥) button to allow participants to copy the current filter settings of their collaborator. Figure 2 illustrates how this feature can be used to establish local interaction-sharing synchronization.
- **L4: Perspective Sharing** allows participants to share their current perspective with collaborators and will be provided to the other half of participant groups. The non-immersive desktop view can toggle a window that shows the current perspective of the immersive VR view tracked to the VR participant’s field of view and movements. This window can be resized and moved freely within the 2D view. Likewise, the immersive VR user can view the non-immersive view on a window fixed to a 2D plane that can be moved freely in the 3D environment (Figure 1 ⑤). These screen-sharing views are always synchronized and stream collaborators’ actions in real time. Figure 2 illustrates how this feature can provide global perspective sharing.

4.5 Experimental Task

We designed our experimental tasks with two primary goals in mind. First to emulate real-world collaborative visualization tasks, and second to encourage collaboration between participants. We deemed that quantitative measures, such as time and accuracy, would require tightly controlled and low-level tasks, which would not achieve our goals. Tasks designed this way could be too easy for participants to complete individually, and low-level tasks would rarely be a focus for real-world collaborative visualization. Instead, we focused on qualitative measures and designed open-ended tasks requiring a mix of high, mid, and low-level tasks.

To achieve this, we asked participants to complete a report with a provided template for each visualization context that would require them to collaborate on a series of multi-level visualization tasks. Table 1 presents the exact details of each report question and respective task typology [7]. Since accuracy was not a concern, we designed most questions not to have a particular correct answer. Instead, the question would prompt participants to decide on how they want to define a correct answer before completing the task. Furthermore, participants could reuse the answer and work done to help complete later tasks.

For example, during the parallel coordinates experiment, participants are first asked to find clusters within the data, then what statistics correlate with the expected batting average (XBA), and finally, to identify the best and worst pitcher. To solve these tasks, participants must first decide what defines a cluster, what constitutes a correlation, and what would make a pitcher the best or worst—requiring consensus among participants achieved through conversation and debate. Then participants could use the clusters already found to compare with those found for XBA and use this knowledge to find a suitable best and worst pitcher.

CONTEXT	QUESTION	TASK			
PC	Are there any distinct groups; if so, what defines them?	Discover	→	Explore	→ Summarize
	What data points correlate with expected batting average?	Discover	→	Browse	→ Identify
	Who are the top and bottom performing pitcher and why?	Discover	→	Locate	→ Compare
	Summarize and conclude your findings.	Present	→	Lookup	→ Summarize
PT	What is the pitcher's name, age, and throwing arm?	Present	→	Lookup	→ Identify
	What do the stats tell us about this pitcher's abilities?	Discover	→	Browse	→ Summarize
	Detail the characteristics of each pitch type.	Discover	→	Browse	→ Compare
	Summarize and conclude your findings.	Present	→	Lookup	→ Summarize

Table 1: Experimental tasks. Overview of the tasks that participant pairs were asked to perform during the user study. For Context, PC refers to parallel coordinates and PT to pitch trajectories.

The desktop user could toggle a window to view and edit the report, while the VR user could see the report displayed and updated in real-time on a plane in the VR environment. We chose only to allow the desktop participant to edit the report, as they would be able to complete the input task more efficiently than the VR user (who has no access to a keyboard), and as a way to encourage collaboration.

4.6 Recruitment and Participants

Participants were recruited through the online participant platform [Prolific](#). We used separate recruitment posts for desktop and VR participants. Desktop participants were pre-screened for self-reported baseball knowledge, and VR participants were only pre-screened for self-reported ownership of a VR headset, due to a lack of overlap in the participant pools. Perspective participants were asked to complete a consent and information form that presented all pertinent information about the study. Additionally, participants were asked to self-report their confidence in their baseball and baseball statistics knowledge and any vision impairments that might hinder their ability to participate.

Once these forms were completed, participants were redirected to a Doodle poll and asked to provide their availability over the study's designated running duration. The poll matched VR and desktop participants that indicated compatible availability. Qualified stand-in participants were prepared for each session using personal and professional contacts. If one of the participants could not or failed to participate, a stand-in participant was called to replace them. We recruited six pairs of participants, and four stand-in participants were used (3 desktops, 1 VR). The majority of participants self-identified a high level of confidence in their baseball knowledge. Lastly, all VR participants participated using a Meta Quest 1 or 2 HMD. We tested our application on both headsets before experimenting, and we found the experience to be comparable between devices. Furthermore, past research examining remote VR study methodology had reported broadly-consistent results even when participants utilized several different VR headsets [44].

4.7 Procedure

Once a VR and desktop participant had been paired, they were sent confirmation messages through Prolific and separate calendar invites that indicated the date, time, and video meeting room where

the study would be conducted. Approximately 24 hours before the study was scheduled, participants were sent a follow-up message asking them to confirm they would still be able to participate at the scheduled time. If one or more participants could not participate or confirm, the scheduled study was rescheduled, or a stand-in participant was used.

Each study session was scheduled for 90 minutes of active participation, roughly split into 20 minutes for setup and 35 minutes for each visualization context. One researcher was present for the entirety of the study session to instruct further and oversee participants. After joining the Zoom call, participants were instructed to access the study apparatus via their desktop or VR headset browser. Next, they were instructed to share their screen through Zoom so that the researcher could monitor their collaboration and record the study sessions for future analysis. Before each study context, participants were walked through the visual tool and interactions available to them and instructed to use the visualization tool and collaborative features to complete the provided report template (Figure 1) with their collaborator.

Participants were generally free to complete this task as they saw fit. However, the researcher was available to provide further explanations or technical help and keep the study running on time. For example, to ensure each session did not exceed 90 minutes, the researchers would prompt participants to wrap up a task if the time per task was becoming disproportionate. At the end of the study, participants were given a brief open-ended exit interview and asked for details about how they used the collaborative features and any additional comments they might have. Participants who completed the study received \$23 in remuneration either through Prolific or with an electronic gift card.⁶

4.8 Analysis and Measures

Each session's audio and video were recorded and prepared for further analysis—editing the footage to separate the two visualization context portions and preparing transcripts using automated audio transcription software. The videos were then analyzed using deductive and inductive qualitative coding. The videos were first reviewed for general collaborative behavior and coupling inspired

⁶For a session lasting approximately 90 minutes, this corresponds to roughly a \$15 hourly rate, which is the U.S. federal minimum wage from January 30, 2022.

by the methods employed by Tang et al. [54], an approach successfully used to understand collaboration around tabletop immersive analytics in social VR settings [44]. Next, the videos were coded for collaborative activity matching the considerations presented by Heer and Agrawala [24]. Finally, each common ground and awareness code was further coded for which level(s) of group awareness technique (Sec. 3) was used. We further sub-coded instances of tightly-coupled work (**AVLC**, **STSF**) for leadership behavior among participants. We defined leadership behavior as instances of a participants beginning or leading a topic of discussion, initiating collaborative activities, or directing work. Codes were assigned to the start and end timestamps indicating the approximate duration of the observation. Coding was performed by one researcher and reviewed by a second.

5 RESULTS

Here we detail the results of our coding analysis. This includes descriptions of each code, the frequency and percentage of time-coded (Figure 3), the co-occurrence of codes (Figure 4), and the frequency and percentage of leadership (Figure 5).

5.1 Collaborative Coupling and Behavior

We systematically reviewed session videos and iteratively narrowed down coupling styles and group behavior, finding similar coupling behavior to those identified by Tang et al. [54]. The following section will define these behaviors in the context of our study apparatus. Codes are organized from **tightly-coupled**

→ **loosely-coupled** → **not coupled**.

AVLC: *Active viewing, listening, conversing*: One participant is actively leading the discussion while the other participant listens attentively, often following along by watching the actions of the other through interaction or perspective sharing. During this behavior, participants will engage in back-and-forth discussion, often leading to changes in who is leading, what they are working on, and achieving consensus.

STSF: *Same task same focus*: Both participants actively engage with the visualization tool while working around the same task, data points, filters, and views. During this behavior, participants often identify areas of interest they can explore or discuss co-currently.

STDF: *Same task different focus*: Both participants actively engage with the visualization tool while working on the same task but focus on different data points, filters, and views. During this behavior, participants commonly explore different areas to find observations to share with their collaborators.

PVLC: *Passive viewing, listening conversing*: One participant is speaking or attempting to lead a discussion, while the other participant is not listening closely. This behavior occurs when the conversation is off-topic, for another task, or one participant focuses on their work.

DT: *Different tasks*: Participants work on separate tasks, often with different data points, filters, and views. This behavior often occurred when one participant wrapped up a task while the other worked ahead. For example, the VR user works on a new task while the desktop participants edit the report for a previous task.

D: *Disengaged*: One or both participants are not working toward any relevant visualization tasks. This behavior could occur due to technical issues, outside factors, or other distractions that would distract the participants' attention.

Overall, participants spent most of their time engaged in tightly- or loosely-coupled work and relatively little time in uncoupled behavior, with behaviors sustaining for several minutes. "Peek" participants spent more time in **AVLC**, while "watch" participants coded **AVLC** and **STSF** more proportionally. "Watch" participants exhibited more overall codes, indicating more frequent behavioral shifts. For between-subject factors, "watch" participants exhibited comparable behavior across both conditions. However, "peek" participants spent more time in **STDF** with parallel coordinates and more time in **STSF** with pitch trajectories.

5.2 Collaborative Activity

Beyond high-level collaborative coupling behavior codes, we analyzed videos for lower-level collaborative activities and considerations presented by Heer and Agrawala [24]. We observed activities matching four of the seven considerations outlined in their work: division and allocation of work; reference and deixis; common ground and awareness; and consensus and decision making.

DAW: *Division and allocation of work*: Participants decide which tasks to work on, when to switch, or how to best complete tasks. This activity typically occurred at the beginning or end of tasks and signaled participants to start tightly-coupled collaboration styles.

RND: *Reference and deixis*: Participants verbally reference a data point (name, color, value), visualization, label, or other item relating to a task. Participants used RND during and transition to tightly-coupled behavior to gain or direct their collaborator's attention. In addition to direct references, participants would use filler words (i.e. uh, um, so) and requests (i.e. "can you see what I am seeing") to perform RND.

CGA: *Common ground and awareness*: Participants act to establish common ground and achieve group awareness by adjusting their views or focus—interacting with the visualization tool to see what their collaborator is seeing. This activity was commonly performed after instances of RND and utilized different group awareness techniques depending on the context and availability. This activity was used while transitioning to or maintaining tightly-coupled work.

CDM: *Consensus and decision making*: Participants actively discuss and confirm each other's conclusions or answers for a given task. This activity was commonly achieved verbally, requiring subsequent RND and CGA activities, with one participant re-stating an observation or finding while the other confirmed or added to the statement. This activity commonly concluded periods of tightly-coupled work, leading to new DAW activities and loosely-coupled work.

RND was generally the most frequent collaborative activity for within and between-subject factors. However, "watch" pairs performed RND and CGA roughly equal amounts at 158 and 149 occurrences. Additionally, "watch" pairs performed more overall actions than "peek" pairs (P: 293, W: 439). CGA and CDM roughly tie for the most time spent performing each activity (P: 16%/17%, W: 16%/18%). These trends are consistent for the within-subject factors, with only

	Peek Groups						Code Frequency						Watch Groups						
	P1 PC	P2 PC	P3 PC	P1 PT	P2 PT	P3 PT	Peek PC	Peek PT	Peek	Watch	Watch PC	Watch PT	W1 PC	W2 PC	W3 PC	W1 PT	W2 PT	W3 PT	
Active Viewing/Listening	11	8	16	8	7	19	35	34	69	82	46	36	21	14	11	12	8	16	Active Viewing/Listening
Same Task Same Focus	4	4	4	6	3	14	12	23	35	58	32	26	18	4	10	9	6	11	Same Task Same Focus
Same Task Different Focus	7	5	9	2	2	3	21	7	28	58	35	23	18	9	8	5	8	10	Same Task Different Focus
Passive Viewing/Listening	1	5	2	4	1	3	8	8	16	26	14	12	6	6	2	6	4	2	Passive Viewing/Listening
Different task	7	3	2	5	2	3	12	10	22	15	7	8	6	1	0	2	1	5	Different task
Disengaged	0	3	0	1	0	0	3	1	4	4	3	1	1	1	1	0	1	0	Disengaged
Division and Allocation of Work	7	4	6	5	2	9	17	16	33	48	20	28	8	5	7	9	7	12	Division and Allocation of Work
Reference and Deixis	15	15	26	12	6	34	56	52	108	158	91	67	23	34	34	15	20	32	Reference and Deixis
Common Ground and Awareness	14	11	19	10	4	29	44	43	87	149	91	58	39	27	25	17	16	25	Common Ground and Awareness
Consensus and Decision Making	10	4	8	13	6	24	22	43	65	84	38	46	17	9	12	18	11	17	Consensus and Decision Making
Level 1: Landmarks and Analogous Views	4	4	7	3	3	21	15	27	42	84	38	46	20	8	10	16	10	20	Level 1: Landmarks and Analogous Views
Level 2: Information Cues	2	2	2	3	1	11	6	15	21	19	12	7	7	2	3	2	2	3	Level 2: Information Cues
Level 3: Interaction Sharing	10	9	13	6	1	6	32	13	45	12	6	6	0	0	6	1	2	3	Level 3: Interaction Sharing
Level 4: Perspective Sharing	0	0	0	0	0	0	0	0	0	66	55	11	27	19	9	4	3	4	Level 4: Perspective Sharing
Percentage of Coded																			
Active Viewing/Listening	59.6%	27.3%	72.3%	41.6%	48%	65.3%	46.5%	54.2%	49.7%	36.5%	39.6%	32.1%	27.9%	52.2%	41.3%	33.9%	23.1%	38.8%	Active Viewing/Listening
Same task same focus	15.6%	12.1%	9.4%	20%	11.1%	23.4%	12%	20.2%	15.4%	28.3%	25.4%	32.7%	32.1%	8.9%	33.6%	40.7%	23.8%	33.4%	Same task same focus
Same task different focus	24.6%	27%	11.9%	9.3%	30.7%	4.5%	21%	10.5%	16.6%	20.3%	20.1%	20.6%	20.8%	18.3%	21.1%	8.4%	34.1%	19.6%	Same task different focus
Passive Viewing/Listening	2.1%	14.8%	2.9%	13.6%	8.1%	4%	7.4%	8%	7.7%	8.5%	8.1%	9.1%	5.6%	17%	2.3%	14.9%	9.8%	3%	Passive Viewing/Listening
Different task	18%	12.4%	3%	13.8%	4.8%	2.4%	10.4%	6.7%	8.9%	5.1%	5.3%	4.8%	13.2%	0.7%	0%	2%	7.2%	5.2%	Different task
Disengaged	0%	5.6%	0%	1.5%	0%	0%	2.2%	0.5%	1.5%	1.1%	1.5%	0.6%	0.4%	2.7%	1.6%	0%	1.9%	0%	Disengaged
Division and Allocation of Work	5.2%	3.5%	3.2%	3.4%	2.5%	4.9%	3.8%	4%	3.9%	8.1%	8.9%	6.9%	14.5%	2.1%	8.8%	11.7%	3.4%	5.5%	Division and Allocation of Work
Reference and Deixis	5.3%	6.3%	6.1%	5.9%	7.7%	7.1%	6%	6.8%	6.3%	8.6%	9.3%	7.7%	7.1%	13.3%	7.9%	6.1%	8.2%	8.7%	Reference and Deixis
Common Ground and Awareness	23.1%	16.5%	16.2%	22%	4.3%	7.7%	18%	12.1%	15.5%	16.4%	18.9%	12.6%	26.6%	23.2%	5.2%	10.1%	17.7%	10.4%	Common Ground and Awareness
Consensus and Decision Making	15.4%	3.5%	16%	16.5%	38.4%	26.7%	10.9%	25.1%	16.9%	18.5%	17.3%	20.2%	19.2%	19%	13.2%	27.6%	14.7%	18.3%	Consensus and Decision Making
Level 1: Landmarks and Analogous Views	1.4%	2.5%	1.6%	1.1%	3.1%	5%	1.9%	3.3%	2.5%	7.1%	6.7%	7.8%	12.2%	4.7%	2%	10.5%	7.2%	5.7%	Level 1: Landmarks and Analogous Views
Level 2: Information Cues	0.7%	1.4%	0.4%	1.5%	0.8%	2.7%	0.9%	2%	1.3%	1.6%	1.9%	1.2%	3.7%	1%	0.4%	0.9%	1.1%	1.5%	Level 2: Information Cues
Level 3: Interaction Sharing	21.7%	15.1%	14.8%	20.3%	1.2%	2%	16.6%	8.2%	13.1%	0.5%	0.4%	0.6%	0%	0%	1.4%	0.1%	0.8%	0.8%	Level 3: Interaction Sharing
Level 4: Perspective Sharing	0%	0%	0%	0%	0%	0%	0%	0%	0%	9.8%	13.1%	5%	18.4%	17.5%	2.1%	1.9%	9.4%	3.7%	Level 4: Perspective Sharing
Within-Subject Factors						Between-Subject Factors						Within-Subject Factors							

Figure 3: Code frequency and percentage. The frequency of codes (top) and the percentage of time-coded (bottom) for each level of coding and all within and between-subject factors. The left and right grids detail the individual results for each pair organized by their between-subject factor (“peek” left, “watch” right) and within-subject factor (PC = parallel coordinates, PT = pitch trajectories). The center grid represents the aggregate result of groups within and between-subject factors.

a minor difference with more time spent performing CDM during the PT visualization context for both groups.

5.3 Group Awareness Techniques

We further subcode each CGA activity code instance for the level of group awareness technique(s) used. The following section will detail our observation of the use of these techniques.

L1: Landmarks and Analogous Views: Participants utilized L1 when CGA directly proceeded RND, which included a direct reference to a data point or component of a visualization. These instances of CGA were brief, concluding quickly if the reference was clear and easy to find or, otherwise, required higher-level techniques.

L2: Information cues: Participants utilized L2 when they needed help locating a data point or performing an interaction. Participants also used L2 with other techniques during viewing and listening behavior.

L3: Interaction sharing: Participants with the L3 “peek” interaction utilized it in various ways. They would use L3 in short bursts to quickly view what their collaborators were doing when L1 and L2 techniques were insufficient for CGA or during periods of loosely-coupled work to reestablish CGA. Additionally, “peek” was used for longer durations during **AVLC** behavior to follow along as collaborators explored or presented. Desktop and VR participants utilized the peek interaction proportionately. The L3 “copy”

interaction was provided to all participants and used to establish CGA quickly and continue with either **STSF** or **STDF** behaviors.

L4: Perspective sharing: This behavior was coded when it was apparent to the researcher that the participant was viewing the window perspective sharing window. For example, when desktop users move their head or gaze toward the watch window while stopping other interactions, or when VR users place the window in the center of their gaze.

Participants with the L4 “watch” interaction utilized it similarly to participants with the L3 “peek” interaction during AVCL behavior for sustained CGA while viewing or presenting. However, VR participants accounted for most of the uses of L4 “watch” for shorter bursts of CGA. Additionally, after determining that one of their views was better suited to a particular task, these participants also used this technique during **AVLC**. For example, when searching for clusters with 3D parallel coordinates or viewing the spatial pitch trajectories of outliers.

Comparing CGA techniques across between-subject factors, “watch” pairs utilized their L4 “watch” interaction more frequently than their counterparts utilized their L3 “peek” (P: 45, W: 66). However, both groups utilized these techniques at roughly the same rate (P: 13%, W: 10%). Additionally, “watch” pairs relied on L1 techniques far more than L4 techniques (W L1: 84, W L4: 66), while “peek” pairs relied on them proportionally (P L1: 42, P L4: 45).

Examining these techniques across within-subject factors reveals some stark differences. Both groups of participants (“watch” and “peek”) utilized their respective L3 or L4 interaction technique more frequently and for longer during PC (P: 32 and 17%, W: 55 and 13%) compared to PT (P: 13 and 5%, W: 11 and 6%). Instead, both participant groups relied more heavily upon L1 techniques when performing CGA for the PT factor.

5.4 Code Co-Occurrences

We also examined the co-occurrence of collaborative activity and group awareness techniques with the high-level coupling behavior codes. The resultant adjacency matrices can be seen in Figure 4.

As expected for collaborative activities, tightly-coupled collaboration behavior co-occurs significantly more often than loosely- or uncoupled behavior. Furthermore, CGA was the most common activity during tightly-coupled work. Notably, CDM and DAW occurred most frequently during **AVLC** and less so during other behaviors.

For group awareness techniques, L1 occurred most frequently during tightly-coupled work (**AVLC**: 36, **STSF**: 71) and was also common during **STDF**, co-occurring 21 times. Both L3 and L4 techniques occurred most often during **AVLC** (L3: 34, L4: 37) and were common during **STSF** (L3: 18, L4: 22). Finally, L2 techniques were most common during **STSF**, co-occurring 29 times.




Figure 4: Code co-occurrence. The co-occurrence matrix of collaborative activities (top) and group awareness techniques (bottom) with high-level collaborative behavior coupling styles and leadership.

5.5 Leadership

The percentage and total of leadership codes per group and condition can be seen in Figure 5. VR participants trended towards leading more often during the “peek” condition, while desktop participants lead more often during the “watch” condition. We also examined the co-occurrence of leadership during lower-level codes (Figure 4). These results highlight this pattern even further, with




Figure 5: Leadership. The percentage and total distribution of leadership behavior per group and experimental factor (PC = parallel coordinates, PT = pitch trajectories).

desktop participants leading more commonly during L4 interactions and VR participants leading during L3 interactions. Additionally, desktop users led much more frequently during DAW activities, where as other activities co-occurred relatively proportionally.

6 DISCUSSION

Here we discuss the results of our study, share general trends, and explore implications. We then use our findings to derive guidelines for collaborative cross-virtuality analytics systems. Finally, we discuss limitations and future work.

6.1 Behavioral Code Takeaways

Coupling Behavior: Participants spent most of their time in tightly-coupled behavior across all groups and factors. This behavior is a strong indicator that, despite asymmetric views, participants were able to establish enough common ground to facilitate tightly-coupled work. The L3 “peek” interaction led to participants spending more time viewing and listening to collaborators, likely because it controls users’ visualization state while in use. Conversely, L4 “watch” could be utilized in the background, making it easier for collaborators to maintain common ground while working in parallel.

Collaborative Activity: Common ground and awareness (CGA) and reference and deixis (RND) were crucial activities that allowed participants to move between collaborative activities and visualization tasks while remaining or entering tightly-coupled work. This is apparent in the co-occurrence between CGA, RND, and tightly-coupled work (Figure 4). We believe that asymmetric views influenced this behavior and created a demand for more CGA activities.

Group Awareness Techniques: We observed that users could establish group awareness using all four techniques provided. Furthermore, we observed how the different levels of group awareness were used depending on the level of group awareness needed. For example, L1 and L2 techniques were used when interpreting quick references, performing locate tasks, and maintaining previously established group awareness. At higher levels, L3 and L4 were used when transitioning from loosely-coupled work to tightly-coupled work and establishing or re-establishing group awareness.

Leadership: Instances of leadership behavior often accompanied the use of group awareness techniques to facilitate reference, deixis, and common ground. We observed that VR participants were more likely to use L4 techniques while desktop participants were leading. This can be clearly seen in the co-occurrence of desktop leading and L4 codes (Figure 4). Similarly, desktop participants utilized L3 interactions more commonly while VR participants were leading. This suggests that users will use the higher-level group awareness techniques to follow along while being led. Furthermore, it indicates that the technique may play a role in who is more likely to take leadership or follower roles. However, leadership is prone to being influenced by external factors such as expertise, individual differences, and assigned roles. For example, desktop led much more frequently during division and allocation of work, likely due to the input task being assigned to them.

6.2 Experimental Factors

Abstract vs. Natural Spatial Mapping: The impetus for separately investing in abstract and natural spatial mapping visualization contexts was that desktop and VR views have different advantages and, thus, different collaboration behavior. We did observe different collaborative behavior between contexts, most notably in the level of group awareness techniques used. Participants relied heavily on landmarks and analogous views (L1) group awareness techniques during the pitch trajectories (PT) natural spatial mapping context and utilized interaction sharing (L3) and perspective sharing (L4) techniques more frequently during the parallel coordinates (PC) abstract context. Our initial reasoning can explain some of this behavior. Desktop participants had a much easier time performing interactions during the PC context, resulting in the VR participant relying on them for the L3 and L4 techniques.

However, this argument does not fully explain these results. Another factor likely at play here is the differences in visualization design and complexity between PC and PT contexts. The PC visualization design is a single visual metaphor and UI paradigm, while the PT visualization design utilizes several visualizations, UI panels, and tables. The desktop view distributes these elements on a 2D plane, but the VR view has these elements distributed in 3D space. As a result, desktop and VR collaborators do not have a 1:1 mapping for reference and deixis and would have to search for the correct view (L1) more often before using L3 and L4 techniques.

We could potentially elevate these issues by using more extensive information cue (L2) group awareness techniques. Our L2 techniques were primarily semi-shared brushing and linking, i.e., what one collaborator was hovering over was also highlighted for the other. However, we could have taken this further with “pinging” interactions where participants could ping aspects of their views,

highlighting the relevant view for their collaborator. This interaction would help direct attention to the correct location across disparate views.

“Peek” vs. “Watch”: Our goal for examining the “peek” and “watch” interactions between-subjects was not to conclude that one technique is better than the other—and indeed, our results do not suggest this. Instead, we wanted to study how, where, and why participants would use these techniques as a primary collaboration mode. We believe that the results, recordings, and transcripts we collected provide a means to answer those questions.

The perspective-sharing (L4) “watch” interaction was used more frequently than the interaction-sharing (L3) “peek” interaction. However, “watch” was disproportionately used more often by the VR participant. A possible reason for this could be that VR participants have more space to orient the “watch” window within view without blocking other elements, allowing them to glance at it when needed. However, the “watch” window was displayed for desktop participants over the rest of the visualization. Even though the window could be resized and moved around, participants kept it toggled off most of the time.

Observing participants’ behavior, a clear pattern for what each technique was suitable for emerged. Interaction sharing was helpful when a quick look at collaborators’ views was needed, but total group awareness was not. Additionally, participants used interaction sharing when specific interactions were more accessible for their collaborators to perform instead. For example, highlighting a specific pitcher on the parallel coordinates plot was often quicker on the desktop, thanks to mouse precision. On the other hand, perspective sharing was most helpful when one view was more effective for a specific task, for long periods of actively viewing, and for tasks requiring complete group awareness while working in parallel. For example, clusters of pitchers across correlated statistics were easier to see in the parallel coordinates immersive VR view.

6.3 Design Alternatives by Level

While we ultimately opted for the group awareness techniques discussed in Sections 4.4 and 5.3, here we discuss alternatives to these methods.

Level 1: One approach to landmarks, beyond the UI elements and visualization metaphors used in our own work, might include a grid, icons, or glyph-based landmarks [17]. Our participants used L1 techniques mainly for brief instances of CGA; a persistent grid and glyph approach might clutter the view, and it is not clear that toggling would yield the same beneficial results as previously demonstrated.

Level 2: Apart from brushing and linking, other information cues might include “auras” such as those used in ReLive [28] highlighting the attention of the users during the sessions the authors evaluated. Avatars showing the presence of a collaborator were used in Prezi [34]. Such an approach might create either a spatial- or interaction-based mapping between 2D and 3D space and represent the immersive user as a cursor in the desktop user’s view and the desktop user’s cursor as a floating object in the immersive user’s space. This approach could have benefited our users, given that they typically used L2 techniques when they needed help locating a data point or interacting with a UI element. However, brushing and linking are currently more common in data visualization.

Level 3: Beyond “peeking,” other interaction sharing could include relinquishing control, allowing a remote user to take over interaction with the system. However, this is quite intrusive approach. Our participants used L3 techniques in short bursts only when L1 and L2 techniques failed them or they lost the CGA thread. Changing the direction of the control such that one user can hijack their collaborator’s view could be disruptive to the workflow of the user subject to being controlled.

Level 4: Embodying another collaborator’s immersive session—really stepping into their shoes—would be true perspective sharing. However, this approach would not be appropriate for our scenario because one participant only had a desktop environment on hand. Transposing the landscape from 2D to 3D and vice-versa may offer a viable alternative, especially if paired with a “visual differencing” technique (comparable to a view of differences in lines of a text file using a file comparison utility). Given that the participants in our evaluation who used L4 techniques were mainly VR users, this area could be particularly relevant for future immersive implementations.

6.4 Designing Collaborative XVA Systems

We present the following considerations for designing collaborative cross-virtuality analytics (XVA) systems based on our experience in implementing VRxD and conducting this user study.

- (1) **Separate but not disparate.** XVA systems should utilize tailored views that leverage the affordances of each device. However, this should not mean that views do not share any similarities. Landmarks and analogous views group awareness techniques were critical for effective collaborative behavior. They allowed participants to collaborate effectively without constant total group awareness. Furthermore, by reusing visual elements and UI design where possible, the development time of the XVA system was reduced—a vital benefit given their complexity.
- (2) **Follow the leader.** As our leadership study results show, the user with the view best suited for a particular task will often take charge. As such, it will be necessary to determine which views are best for what tasks to design them accordingly. Leadership roles for specific tasks should be encouraged through visualization and interaction design, while other views should be designed with support or viewing in mind.
- (3) **Queue the cues.** We did not thoroughly implement information cue techniques in VRxD and noticed several scenarios where more cues would have been helpful. It is essential to consider how users will utilize information cues for reference and deixis, allowing them to direct their collaborators’ attention without perfect knowledge of their views. Furthermore, designing such information cues will require new techniques for mapping landmarks and analogous views when 1:1 mappings do not exist.
- (4) **Show, don’t tell.** Interaction and perspective group awareness techniques were critical when higher levels of group awareness were required. We should design these interactions around satisfying the following tasks: quickly checking what collaborators are doing, presenting to or assisting collaborators interactively, and viewing collaborators’ activity for sustained periods. These tasks could warrant uniquely designed interaction and perspective-sharing techniques to support them. The techniques

we implemented were general and not optimized to handle these tasks. Providing users with less intrusive methods of perspective sharing should encourage its use more often—leading to more group awareness and better collaboration.

7 LIMITATIONS AND FUTURE WORK

The intention of our visualization design was not to be the most effective or comprehensively implemented version of an XVA system. Instead, it was merely an appropriate platform for participants to complete the experimental task while enabling us to observe collaborative behavior. Obviously, the actual design is a factor that will influence user performance and, as such, could limit the generalizability of our results. Even so, we intend the results from this study to act as a baseline for future domain-specific studies to build upon as a standard group awareness model.

Likewise, individual differences between participants—such as personalities, levels of engagement, and baseball knowledge—can influence the results of such a study. We believe that future studies can better control this factor by thoroughly screening participants’ personalities and domain knowledge to select appropriate pairs.

We forecast several future directions for collaborative cross-virtuality analytics research. There is excellent potential in a design space for information cues, interaction sharing, and perspective-sharing group awareness techniques. While we investigated elements of this design space, our work is not exhaustive. Furthermore, longitudinal studies of collaborative XVA systems are needed to examine how collaborative behavior changes with more experience and evolving tasks. Finally, more domain-specific collaborative XVA applications are needed to understand their utility fully.

8 CONCLUSION

Group awareness is imperative to effective collaboration and collaborative visualization. However, as collaborators’ views diverge, common ground diminishes, making group awareness harder to achieve. This phenomenon is problematic for collaborative cross-virtuality analytics (XVA), as tailoring views and visual metaphors to collaborators’ disparate devices is core to its philosophy. To address this, we presented our “eyes-and-shoes” principles that abstract achieving group awareness into four levels of techniques. We further evaluated these techniques and different visualization contexts, giving us insights into their use and general collaborative XVA behavior. We hope these principles and our documented results and observations can be used to design more effective XVA systems. Furthermore, we think this work can catalyze increased interest in researching and applying this unique modality of collaboration.

ACKNOWLEDGMENTS

This work was partially supported by the U.S. National Science Foundation award IIS-1908605. Any opinions, findings, and conclusions expressed in this material are those of the authors and do not necessarily reflect the views of the funding agency.

This paper was prepared for informational purposes by the teams of researchers from the various institutions identified above, including the Global Technology Applied Research Center of JPMorgan Chase & Co. This paper is not a product of the JPMorgan Chase Institute. Neither JPMorgan Chase & Co. nor any of its affiliates

make any explicit or implied representation or warranty and none of them accept any liability in connection with this paper, including, but limited to, the completeness, accuracy, reliability of information contained herein and the potential legal, compliance, tax or accounting effects thereof. This document is not intended as investment research or investment advice, or a recommendation, offer or solicitation for the purchase or sale of any security, financial instrument, financial product or service, or to be used in any way for evaluating the merits of participating in any transaction.

REFERENCES

- [1] Sriram Karthik Badam and Niklas Elmqvist. 2014. PolyChrome: A Cross-Device Framework for Collaborative Web Visualization. In *Proceedings of the Ninth ACM International Conference on Interactive Tabletops and Surfaces* (Dresden, Germany) (ITS '14). Association for Computing Machinery, New York, NY, USA, 109–118. <https://doi.org/10.1145/2669485.2669518>
- [2] Sriram Karthik Badam, Andreas Mathisen, Roman Rädle, Clemens N. Klokmoose, and Niklas Elmqvist. 2019. Vistrates: A Component Model for Ubiquitous Analytics. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2019), 586–596. <https://doi.org/10.1109/TVCG.2018.2865144>
- [3] Ronald M. Baecker. 1993. *Readings in Groupware and Computer-Supported Cooperative Work*. Morgan Kaufmann Publishers, San Francisco.
- [4] Andrea Batch, Andrew Cunningham, Maxime Cordeil, Niklas Elmqvist, Tim Dwyer, Bruce H. Thomas, and Kim Marriott. 2020. There Is No Spoon: Evaluating Performance, Space Use, and Presence with Expert Domain Users in Immersive Analytics. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (2020), 536–546. <https://doi.org/10.1109/TVCG.2019.2934803>
- [5] Szymon Bobek, Sławomir K. Tadeja, Łukasz Struski, Przemysław Stachura, Timoleon Kipourou, Jacek Tabor, Grzegorz J. Nalepa, and Per Ola Kristensson. 2022. Virtual Reality-Based Parallel Coordinates Plots Enhanced with Explainable AI and Data-Science Analytics for Decision-Making Processes. *Applied Sciences* 12, 1 (2022), 13 pages. <https://doi.org/10.3390/app12010331>
- [6] Michael Bostock, Vadim Ogievetsky, and Jeffrey Heer. 2011. D³: Data-Driven Documents. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (2011), 2301–2309. <https://doi.org/10.1109/TVCG.2011.185>
- [7] Matthew Brehmer and Tamara Munzner. 2013. A Multi-Level Typology of Abstract Visualization Tasks. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2376–2385. <https://doi.org/10.1109/TVCG.2013.124>
- [8] Simon Butscher, Sebastian Hubenschmid, Jens Müller, Johannes Fuchs, and Harald Reiterer. 2018. Clusters, Trends, and Outliers: How Immersive Technologies Can Facilitate the Collaborative Analysis of Multidimensional Data. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) (CHI '18). Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3173574.3173664>
- [9] Herbert H. Clark and Susan E. Brennan. 1991. Grounding in Communication. In *Perspectives on Socially Shared Cognition*, L. B. Resnick, J. M. Levine, and S. D. Teasley (Eds.). American Psychological Association, Washington, DC, USA, 127–149. <https://doi.org/10.1037/10096-006>
- [10] Maxime Cordeil, Andrew Cunningham, Tim Dwyer, Bruce H. Thomas, and Kim Marriott. 2017. ImAxes: Immersive Axes as Embodied Affordances for Interactive Multivariate Data Visualisation. In *Proceedings of the ACM Symposium on User Interface Software and Technology*. ACM, New York, NY, USA, 71–83. <https://doi.org/10.1145/3126594.3126613>
- [11] Maxime Cordeil, Tim Dwyer, Karsten Klein, Bireswar Laha, Kim Marriott, and Bruce H. Thomas. 2017. Immersive Collaborative Analysis of Network Connectivity: CAVE-style or Head-Mounted Display? *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (2017), 441–450. <https://doi.org/10.1109/TVCG.2016.2599107>
- [12] Ciro Donalek, S. G. Djorgovski, Alex Cioc, Anwell Wang, Jerry Zhang, Elizabeth Lawler, Stacy Yeh, Ashish Mahabal, Matthew Graham, Andrew Drake, Scott Davidoff, Jeffrey S. Norris, and Giuseppe Longo. 2014. Immersive and collaborative data visualization using virtual reality platforms. In *Proceedings of the IEEE International Conference on Big Data*. IEEE, Los Alamitos, CA, USA, 609–614. <https://doi.org/10.1109/BigData.2014.7004282>
- [13] Barrett Ens, Benjamin Bach, Maxime Cordeil, Ulrich Engelke, Marcos Serrano, Wesley Willett, Arnaud Prouzeau, Christoph Anthes, Wolfgang Büschel, Cody Dunne, Tim Dwyer, Jens Grubert, Jason H. Haga, Nurit Kirshenbaum, Dylan Kobayashi, Tica Lin, Monsurat Olaosebikan, Fabian Pointecker, David Saffo, Nazmus Saquib, Dieter Schmalstieg, Danielle Albers Szafir, Matt Whitlock, and Yalong Yang. 2021. Grand Challenges in Immersive Analytics. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 459:1–459:17. <https://doi.org/10.1145/3411764.3446866>
- [14] James D. Foley, Andries van Dam, Steven Feiner, and John Hughes. 1990. *Computer Graphics: Principles and Practice* (second ed.). Addison-Wesley, Reading, MA, USA.
- [15] Bernhard Fröhler, Christoph Anthes, Fabian Pointecker, Judith Friedl, Daniel Schwajda, Andreas Riegler, Shailesh Tripathi, Clemens Holzmann, Manuel Brunner, Herbert Jodlbauer, Hans-Christian Jetter, and Christoph Heinzl. 2022. A Survey on Cross-Virtuality Analytics. *Computer Graphics Forum* 41, 1 (2022), 465–494. <https://doi.org/10.1111/cgf.14447>
- [16] Danilo Gasques, Janet G. Johnson, Tommy Sharkey, Yuanyuan Feng, Ru Wang, Zhuoqun Robin Xu, Enrique Zavala, Yifei Zhang, Wanxie Xie, Xinming Zhang, Konrad Davis, Michael Yip, and Nadir Weibel. 2021. ARTEMIS: A Collaborative Mixed-Reality System for Immersive Surgical Telementoring. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 662:1–662:14. <https://doi.org/10.1145/3411764.3445576>
- [17] Sohaib Ghani and Niklas Elmqvist. 2011. Improving Revisitation in Graphs through Static Spatial Features. In *Proceedings of Graphics Interface*. Canadian Human-Computer Communications Society, Waterloo, ON, Canada, 175–182.
- [18] Adam M. Grant and James W. Berry. 2011. The Necessity of Others is The Mother of Invention: Intrinsic and Prosocial Motivations, Perspective Taking, and Creativity. *Academy of Management Journal* 54, 1 (2011), 73–96. <https://doi.org/10.5465/amj.2011.59215085> arXiv:<https://doi.org/10.5465/amj.2011.59215085>
- [19] Carl Gutwin and Saul Greenberg. 1998. Design for Individuals, Design for Groups: Tradeoffs between Power and Workspace Awareness. In *Proceedings of the ACM Conference on Computer-Supported Collaborative Work*. ACM, New York, NY, USA, 207–216. <https://doi.org/10.1145/289444.289495>
- [20] Carl Gutwin and Saul Greenberg. 1998. Effects of Awareness Support on Groupware Usability. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 511–518. <https://doi.org/10.1145/274644.274713>
- [21] Carl Gutwin and Saul Greenberg. 2002. A Descriptive Framework of Workspace Awareness for Real-Time Groupware. *Computer Supported Cooperative Work* 11, 3–4 (2002), 411–446. <https://doi.org/10.1023/A:1021271517844>
- [22] Carl Gutwin, Saul Greenberg, and Mark Roseman. 1996. Workspace Awareness in Real-Time Distributed Groupware: Framework, Widgets, and Evaluation. In *Proceedings of the British HCI Conference*. Springer, London, UK, 281–298. https://doi.org/10.1007/978-1-4471-3588-3_18
- [23] Carl Gutwin, Mark Roseman, and Saul Greenberg. 1996. A Usability Study of Awareness Widgets in a Shared Workspace Groupware System. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work*. ACM, New York, NY, USA, 258–267. <https://doi.org/10.1145/240080.240298>
- [24] Jeffrey Heer and Maneesh Agrawala. 2008. Design considerations for collaborative visual analytics. *Information Visualization* 7, 1 (2008), 49–62. <https://doi.org/10.1057/palgrave.ivs.9500167>
- [25] Jeffrey Heer, Fernanda Viégas, and Martin Wattenberg. 2007. Voyagers and Voyeurs: Supporting Asynchronous Collaborative Information Visualization. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 1029–1038. <https://doi.org/10.1145/1435417.1435439>
- [26] William C. Hill and James D. Hollan. 1991. Deixis and the Future of Visualization Excellence. In *Proceedings of the IEEE Visualization Conference*. IEEE, Los Alamitos, CA, USA, 314–320. <https://doi.org/10.1109/VISUAL.1991.175820>
- [27] Tom Horak, Sriram Karthik Badam, Niklas Elmqvist, and Raimund Dachselt. 2018. When David Meets Goliath: Combining Smartwatches with a Large Vertical Display for Visual Data Exploration. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 1–13. <https://doi.org/10.1145/3173574.3173593>
- [28] Sebastian Hubenschmid, Jonathan Wieland, Daniel Immanuel Fink, Andrea Batch, Johannes Zagermann, Niklas Elmqvist, and Harald Reiterer. 2022. ReLive: Bridging In-Situ and Ex-Situ Visual Analytics for Analyzing Mixed Reality User Studies. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 24:1–24:20. <https://doi.org/10.1145/3491102.3517550>
- [29] Sebastian Hubenschmid, Johannes Zagermann, Simon Butscher, and Harald Reiterer. 2021. STREAM: Exploring the Combination of Spatially-Aware Tablets with Augmented Reality Head-Mounted Displays for Immersive Analytics. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 469:1–469:14. <https://doi.org/10.1145/3411764.3445298>
- [30] Petra Isenberg, Niklas Elmqvist, Jean Scholtz, Daniel Cernea, Kwan-Liu Ma, and Hans Hagen. 2011. Collaborative visualization: definition, challenges, and research agenda. *Information Visualization* 10, 4 (2011), 310–326. <https://doi.org/10.1177/1473871611412817>
- [31] Janet G. Johnson, Danilo Gasques, Tommy Sharkey, Evan Schmitz, and Nadir Weibel. 2021. Do You Really Need to Know Where “That” Is? Enhancing Support for Referencing in Collaborative Mixed Reality Environments. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 514:1–514:14. <https://doi.org/10.1145/3411764.3445246>
- [32] KyungTae Kim, Waqas Javed, Cary Williams, Niklas Elmqvist, and Pourang Irani. 2010. Hugin: A Framework for Awareness and Coordination in Mixed-Presence Collaborative Information Visualization. In *Proceedings of the ACM Conference on Interactive Tabletops and Surfaces*. ACM, New York, NY, USA, 231–240. <https://doi.org/10.1145/1936652.1936694>

- [33] Matthias Kraus, Johannes Fuchs, Björn Sommer, Karsten Klein, Ulrich Engelke, Daniel Keim, and Falk Schreiber. 2022. Immersive Analytics with Abstract 3D Visualizations: A Survey. *Computer Graphics Forum* 41, 1 (2022), 201–229. <https://doi.org/10.1111/cgf.14430>
- [34] László Laufer, Péter Halácsy, and Adam Somlai-Fischer. 2011. Prezi meeting: collaboration in a zoomable canvas based environment. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 749–752. <https://doi.org/10.1145/1979742.1979673>
- [35] Gun A. Lee, Theophilus Teo, Seungwon Kim, and Mark Billinghurst. 2018. A User Study on MR Remote Collaboration Using Live 360 Video. In *Proceedings of the IEEE International Symposium on Mixed and Augmented Reality*. IEEE, Los Alamitos, CA, USA, 153–164. <https://doi.org/10.1109/ISMAR.2018.00051>
- [36] Alexandre Maes, Xavier Martinez, Karen Druart, Benoist Laurent, Sean Guégan, Christophe H. Marchand, Stéphane D. Lemaire, and Marc Baaden. 2018. MinOmics, an Integrative and Immersive Tool for Multi-Omics Analysis. *Journal of Integrative Bioinformatics* 15, 2 (2018), 20180006. <https://doi.org/10.1515/jib-2018-0006>
- [37] Kim Marriott, Jian Chen, Marcel Hlawatsch, Takayuki Itoh, Miguel A. Nacenta, Guido Reina, and Wolfgang Stuerzlinger. 2018. Immersive Analytics: Time to Reconsider the Value of 3D for Information Visualisation. In *Immersive Analytics*. Lecture Notes in Computer Science, Vol. 11190. Springer, Cham, Germany, 25–55. https://doi.org/10.1007/978-3-030-01388-2_2
- [38] Kim Marriott, Falk Schreiber, Tim Dwyer, Karsten Klein, Nathalie Henry Riche, Takayuki Itoh, Wolfgang Stuerzlinger, and Bruce H. Thomas (Eds.). 2018. *Immersive Analytics*. Lecture Notes in Computer Science, Vol. 11190. Springer International Publishing, New York, NY, USA. <https://doi.org/10.1007/978-3-030-01388-2>
- [39] Will McGrath, Brian Bowman, David C. McCallum, Juan David Hincapié-Ramos, Niklas Elmquist, and Pourang Irani. 2012. Branch-explore-merge: facilitating real-time revision control in collaborative visual exploration. In *Proceedings of the ACM Conference on Interactive Tabletops and Surfaces*. ACM, New York, NY, USA, 235–244. <https://doi.org/10.1145/2396636.2396673>
- [40] Paul Milgram, Haruo Takemura, Akira Utsumi, and Fumio Kishino. 1995. Augmented reality: a class of displays on the reality-virtuality continuum. In *Telematics and Telepresence Technologies*, Vol. 2351. SPIE, Bellingham, WA, USA, 282–292. <https://doi.org/10.1117/12.197321>
- [41] Jens Müller, Johannes Zagermann, Jonathan Wieland, Ulrike Pfeil, and Harald Reiterer. 2019. A Qualitative Comparison Between Augmented and Virtual Reality Collaboration with Handheld Devices. In *Proceedings of Mensch und Computer*. ACM, New York, NY, USA, 399–410. <https://doi.org/10.1145/3340764.3340773>
- [42] Nico Reski, Aris Alissandrakis, and Andreas Kerren. 2022. An Empirical Evaluation of Asymmetric Synchronous Collaboration Combining Immersive and Non-Immersive Interfaces Within the Context of Immersive Analytics. *Frontiers in Virtual Reality* 2, Article 743445 (2022), 29 pages. <https://doi.org/10.3389/frvir.2021.743445>
- [43] Andreas Riegler, Christoph Anthes, Hans-Christian Jetter, Christoph Heinzl, Clemens Holzmann, Herbert Jodlbauer, Manuel Brunner, Stefan Auer, Judith Friedl, Bernhard Fröhler, Christina Leitner, Fabian Pointecker, Daniel Schwajda, and Shailesh Tripathi. 2020. Cross-Virtuality Visualization, Interaction and Collaboration. In *Proceedings of the International Workshop on Cross-Reality Interaction*. CEUR Workshop Proceedings, Aachen, Germany, Article 1, 4 pages. <http://ceur-ws.org/Vol-2779/paper1.pdf>
- [44] David Saffo, Sara Di Bartolomeo, Caglar Yildirim, and Cody Dunne. 2021. Remote and Collaborative Virtual Reality Experiments via Social VR Platforms. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 523, 15 pages. <https://doi.org/10.1145/3411764.3445426>
- [45] Hinako Sassa, Maxime Cordeil, Mitsuo Yoshida, and Takayuki Itoh. 2021. Visual Linking of Feature Values in Immersive Graph Visualization Environment. In *Proceedings of the 14th International Symposium on Visual Information Communication and Interaction* (Potsdam, Germany) (VINCI '21). Association for Computing Machinery, New York, NY, USA, Article 13, 6 pages. <https://doi.org/10.1145/3481549.3481555>
- [46] Ralph Schroeder, Anthony Steed, Ann-Sofie Axelsson, Ilona Heldal, Åsa Abelin, Josef Wideström, Alexander Nilsson, and Mel Slater. 2001. Collaborating in networked immersive spaces: as good as being there together? *Computer Graphics* 25, 5 (2001), 781–788. [https://doi.org/10.1016/S0097-8493\(01\)00120-0](https://doi.org/10.1016/S0097-8493(01)00120-0)
- [47] Michail Schwab, David Saffo, Yixuan Zhang, Shash Sinha, Cristina Nita-Rotaru, James Tompkin, Cody Dunne, and Michelle A. Borkin. 2021. VisConnect: Distributed Event Synchronization for Collaborative Visualization. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2021), 347–357. <https://doi.org/10.1109/TVCG.2020.3030366>
- [48] Mickael Sereno, Xiyao Wang, Lonni Besançon, Michael J. McGuffin, and Tobias Isenberg. 2022. Collaborative Work in Augmented Reality: A Survey. *IEEE Transactions on Visualization and Computer Graphics* 28, 6 (2022), 2530–2549. <https://doi.org/10.1109/TVCG.2020.3032761>
- [49] Aaron Stafford, Wayne Piekarski, and Bruce H. Thomas. 2006. Implementation of god-like interaction techniques for supporting collaboration between outdoor AR and indoor tabletop users. In *Proceedings of the IEEE/ACM International Symposium on Mixed and Augmented Reality*. ACM, New York, NY, USA, 165–172. <https://doi.org/10.1109/ISMAR.2006.297809>
- [50] Slawomir Konrad Tadeja, Timoleon Kipouros, and Per Ola Kristensson. 2019. Exploring Parallel Coordinates Plots in Virtual Reality. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 1–6. <https://doi.org/10.1145/3290607.3313068>
- [51] Slawomir K. Tadeja, Timoleon Kipouros, and Per Ola Kristensson. 2020. IPCP: Immersive Parallel Coordinates Plots for Engineering Design Processes. In *AIAA Scitech Forum*. American Institute of Aeronautics and Astronautics, Reston, VA, USA, 20 pages. <https://doi.org/10.2514/6.2020-0324>
- [52] Anthony Tang, Michael Boyle, and Saul Greenberg. 2005. Understanding and Mitigating Display and Presence Disparity in Mixed Presence Groupware. *Journal of Research and Practice in Information Technology* 37, 2 (2005), 71–88.
- [53] Anthony Tang, Carmen Neustaedter, and Saul Greenberg. 2006. VideoArms: embodiments for mixed presence groupware. In *Proceedings of the British HCI Conference*, Vol. 20. Springer, London, UK, 85–102. https://doi.org/10.1007/978-1-84628-664-3_8
- [54] Anthony Tang, Melanie Tory, Barry Po, Petra Neumann, and Sheelagh Carpendale. 2006. Collaborative coupling over tabletop collaboration. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 1181–1190. <https://doi.org/10.1145/1124772.1124950>
- [55] James J. Thomas and Kristin A. Cook (Eds.). 2005. *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. IEEE Computer Society, Los Alamitos, CA, USA.
- [56] Wai Tong, Meng Xia, Kam Kwai Wong, Doug A. Bowman, Ting-Chuen Pong, Huamin Qu, and Yalong Yang. 2023. Towards an Understanding of Distributed Asymmetric Collaborative Visualization on Problem-solving. In *2023 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*.
- [57] Philip Tuddenham and Peter Robinson. 2007. Distributed Tabletops: Supporting Remote and Mixed-Presence Tabletop Collaboration. In *Proceedings of IEEE Tabletop*. IEEE, Los Alamitos, CA, USA, 19–26. <https://doi.org/10.1109/TABLETOP.2007.15>
- [58] Philip Tuddenham and Peter Robinson. 2009. Territorial coordination and workspace awareness in remote tabletop collaboration. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 2139–2148. <https://doi.org/10.1145/1518701.1519026>
- [59] Xiyao Wang, Lonni Besançon, David Rousseau, Mickael Sereno, Mehdi Ammi, and Tobias Isenberg. 2020. Towards an Understanding of Augmented Reality Extensions for Existing 3D Data Analysis Tools. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 1–13. <https://doi.org/10.1145/3313831.3376657>
- [60] Matt Whitlock, Stephen Smart, and Danielle Albers Szafir. 2020. Graphical Perception for Immersive Analytics. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces*. IEEE, Los Alamitos, CA, USA, 616–625. <https://doi.org/10.1109/VR46266.2020.00084>
- [61] World Wide Web Consortium (W3C). 2022. WebXR Device API (working draft). <https://www.w3.org/TR/webxr/>.
- [62] Ji Soo Yi, Youn ah Kang, John Skasko, and Julie A. Jacko. 2007. Toward a Deeper Understanding of the Role of Interaction in Information Visualization. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1224–1231. <https://doi.org/10.1109/TVCG.2007.70515>