

Personal Augmented Reality for Information Visualization on Large Interactive Displays

Patrick Reipschläger*, Tamara Flemisch*, Raimund Dachzelt




Fig. 1. Impressions from our prototype for extending visualizations with Augmented Reality: (a) Two analysts working with our environment, (b) display showing our *AR Brushing and Linking*, *Embedded AR Visualizations*, and *Extended Axis Views*, (c) *Hinged Visualizations* to improve perception of remote content, and (d) *Curved AR Screen* providing an overview of the entire display.

Abstract—In this work we propose the combination of large interactive displays with personal head-mounted Augmented Reality (AR) for information visualization to facilitate data exploration and analysis. Even though large displays provide more display space, they are challenging with regard to perception, effective multi-user support, and managing data density and complexity. To address these issues and illustrate our proposed setup, we contribute an extensive design space comprising first, the spatial alignment of display, visualizations, and objects in AR space. Next, we discuss which parts of a visualization can be augmented. Finally, we analyze how AR can be used to display personal views in order to show additional information and to minimize the mutual disturbance of data analysts. Based on this conceptual foundation, we present a number of exemplary techniques for extending visualizations with AR and discuss their relation to our design space. We further describe how these techniques address typical visualization problems that we have identified during our literature research. To examine our concepts, we introduce a generic AR visualization framework as well as a prototype implementing several example techniques. In order to demonstrate their potential, we further present a use case walkthrough in which we analyze a movie data set. From these experiences, we conclude that the contributed techniques can be useful in exploring and understanding multivariate data. We are convinced that the extension of large displays with AR for information visualization has a great potential for data analysis and sense-making.

Index Terms—Augmented Reality, Information Visualization, InfoVis, Large Displays, Immersive Analytics, Physical Navigation, Multiple Coordinated Views

1 INTRODUCTION

Interactive information visualization on large or wall-sized displays has been shown to have great potential for data analysis [3, 38, 51]. Due to their size, they can visualize large amounts of data or several visualizations at once which are often linked to each other. Such multiple coordinated views are useful to help understanding complex, multivariate data sets by highlighting connections between views through Brushing and Linking [74]. However, working in front of large displays also leads to awareness and perception issues regarding the peripheral areas of the display [4, 14]. This results in users having to step back from the display to gain an overview of the data [4, 51] which in turn means that details within specific visualizations are no longer perceivable. While multiple coordinated visualizations are a valuable technique for analyzing multivariate data sets, they still require extensible context switches. Therefore, it can be beneficial to visualize all available information directly in place. Due to their size, large displays are also suited for collaborative work [42, 43] enabling several users to interact with the display at the same time. However, as users interact

with the display, they block parts of it with their body which can disturb other analysts [45, 64]. Highlights, which result from Brushing and Linking, can also be distracting to others. The same is true for other people's tools that might obstruct important details of a visualization.

Besides large displays, the recent field of *Immersive Analytics* has shown great potential for improving data analysis and mitigating problems commonly associated with three dimensional visualizations on traditional displays [88]. By using immersive technologies, *Immersive Analytics* aims to overcome barriers between people, their data, and the tools they use for analysis and decision making [24]. In this regard, physical navigation [8, 28, 47], spatial immersion [3], and embodied data exploration [50, 88] have proven to help in understanding complex data [24]. *Immersive Analytics* applications also have limitations regarding interaction and the visual quality of visualizations [7, 11, 24]. By combining *Immersive Analytics* and large displays, their individual limitations can be mitigated and their advantages strengthened [7].

Based on these promising two avenues of research, we propose to combine a large, interactive display with head-mounted, personal Augmented Reality (AR). We imagine one or more analysts in front of a large, vertical, interactive display which shows one large or several linked visualizations of multivariate data. Each analyst wears an AR head-mounted display (HMD) which is used to overlay the screen with additional, user-specific information. We are convinced that this combination can mitigate the aforementioned issues for both, large displays and pure immersive analysis techniques. The display serves as a frame of reference for all AR content and facilitates natural pen and touch interaction [85]. Augmented Reality, however, can help to guide analysts' awareness, show remote parts of the screen, and to display additional content directly embedded in the visualization [90]. Furthermore, this

* The first two authors contributed equally to this work.

All authors are with the Interactive Media Lab at the Technische Universität Dresden, Germany. Raimund Dachzelt is with the Centre for Tactile Internet (CeTi) and the Cluster of Excellence Physics of Life, TU Dresden. E-mails: {patrick.reipschlaeger, tamara.flemisch}@tu-dresden.de, dachzelt@acm.org.

Manuscript received 30 Apr. 2020; revised 31 July 2020; accepted 14 Aug. 2020.
Date of publication 14 Oct. 2020; date of current version 15 Jan. 2021.
Digital Object Identifier no. 10.1109/TVCG.2020.3030460

enables personal views which foster undisturbed and unobstructed interaction and allow for individual, personalized selections, annotations, and extensions to the data shown on the display. In this work, we contribute an analysis of the design space of the combination of a large, interactive display, Augmented Reality, and information visualizations. In this design space, we first analyze the spatial alignment of the display, the analyst, and the AR content. Second, we discuss which parts of a visualization can be augmented. Third, we describe how to leverage personal views to show additional information and minimize the mutual disturbance of analysts. To illustrate the potential of our conceptual design space, we contribute several augmentation techniques, which address specific visualization problems. To examine our concepts, we developed a generic, data-driven AR visualization framework as well as a prototype implementing our proposed techniques (see Fig. 1). We demonstrate their potential in a use case scenario in which we analyze a movie data set. Lastly, we discuss technical and perceptual aspects and give an outlook on how the presented concepts can be extended in the future.

2 RELATED WORK AND BACKGROUND

Our work touches upon the research fields of Information Visualization (InfoVis) and Human-Computer Interaction (HCI). In particular, it refers to the use of (1) large, interactive displays and (2) Mixed Reality (MR) for information visualization as well as (3) the general combination of interactive surfaces and Augmented Reality.

2.1 Information Visualization on Large Displays

Large, interactive displays have been increasingly used for data analysis and visualization with positive impact [3, 71]. Their screen size, the vast physical space in front of them [3, 8, 44], and their interaction capabilities make them a powerful foundation for exploring large amounts of data. Possible interaction modalities for information visualization beyond the traditional desktop setup have been proposed [56, 75], with touch interaction being the most common one. Data analysis benefits from touch input due to a more direct interaction with the data [20, 38, 43, 49, 51, 68], although some use cases profit from more distant interaction [49–51]. The increased screen estate of the displays can be used for showing large amounts of data, such as node-link diagrams [50, 68], or to incorporate multiple views of data [51, 55]. Due to their size, large displays also facilitate collaborative data analysis and multiple analysts working in parallel [42, 50, 51, 68, 82]. Even though the benefits of large, interactive surfaces for visualization is evident, they also contribute specific challenges, affordances, and requirements [4] which we further discuss in Sect. 2.4.

2.2 Immersive Analytics

The field of *Immersive Analytics* makes use of various display technologies to facilitate and support embodied data analysis and decision making [24]. Recent research in *Immersive Analytics* focuses especially on how Mixed Reality environments can help to explore and understand data. Cordeil et al. [22] present a system to explore multidimensional data in Virtual Reality (VR) through manipulating and arranging linked axes. Approaches on how to ideally route links in immersive spaces have been analyzed by Prouzeau et al. [69]. Alternatively, Liu et al. [58] adapted the concept of small multiples to an immersive environment. Overall, different MR technologies have their own strengths and weaknesses [7]. In this regard, Augmented Reality shows great potential for data analysis since it allows for embedding data representations into the physical space and therefore facilitates the simultaneous exploration of large amounts of data [11, 90].

Whitlock et al. [88] studied how visualization techniques and different immersive display modalities affect data analysis tasks. They found that analysts benefit from AR and VR when estimating depth and height but also that color perception in video see-through AR is challenging. The impact of different edge types for node-link diagrams in AR was investigated by Büschel et al. [16]. Batch et al. [11] found that analysts are highly interested, present, and engaged when analyzing data in immersive environments. Similar to visualization on large displays, *Immersive Analytics* also presents new affordances and

requirements regarding the perception, design, and interaction of and with visualizations [24, 88].

Interaction in immersive environments is often achieved through mid-air gestures. Although they allow for rich input modalities, they lack haptic feedback [24]. Therefore, recent work has combined AR with multi-touch displays. Wang et al. [86] propose to integrate Augmented Reality visualizations into traditional desktop setups. Their findings suggest that this helps analysts to better understand their data. Work by Chen et al. [21] and *VisAR* by Kim et al. [48] both combine personal AR views with small as well as large displays to generate interactivity for otherwise static visualizations. Sun et al. [80] incorporate visualizations on wall-sized displays and AR HMDs for collaborative visual analysis and personalized views. Likewise, Butcher et al. [17] apply a video see-through AR HMD to augment a large interactive tabletop for analyzing 2D visualizations. AR visualizations are directly situated on the display and analysts interact with proxies on the interactive surface. Furthermore, work that shows display-like surfaces in a pure VR environment [30, 31, 69] may be adapted to a setup consisting of a physical display and head-mounted AR.

2.3 Combining Interactive Surfaces with Mixed Reality

Starting with the early works *Hybrid Displays* [29] and *Augmented Surfaces* [73], the recent development of commercial VR and AR HMDs has led to a renewed interest in combining Mixed Reality with interactive or non-interactive displays. Bodgan et al. [15] and Mendes et al. [61] use a stereoscopic display to investigate the differences between 2D and 3D interaction for modeling and 3D manipulation of objects, but there is no consensus which technique is preferable. Grubert et al. [34] use head-mounted Augmented Reality to extend the screen around mobile devices. *DualCAD* by Milette and McGuffin [62] loosely combines a traditional desktop CAD application with head-mounted AR. In *SymbiosisSketch* by Arora et al. [5] users can draw objects in AR using a tablet. However, the interaction on the tablet is spatially decoupled from the AR surface that is sketched onto. In a similar way, Sereno et al. [78] propose to manipulate an AR 3D visualization by using touch on a tablet. Again, there is no spatial connection between the tablet and the visualization. Cavallo et al.'s *DataSpace* [18] combines multiple large displays with AR and VR visualizations but with no strong spatial integration between content in AR and on the displays (also cf. [19]). Reipschläger and Dachselt [72] combine an interactive surface with an AR HMD for a stereoscopic 3D modeling tool. In this work, there is a strong connection between the display and the AR content and a particular focus is on how the borders of the display can be utilized to place additional views and how to interact with them. They have also coined the term *Augmented Displays* for seamlessly combining interactive displays with head-mounted AR, to which our approach can also be attributed.

2.4 Challenges for Visualization on Large Displays

Based on previous works on visualization for large, interactive displays, we compiled several important challenges, which we describe in detail in the following sections.

[C.1] Perceptual Issues on Large Displays

Large displays have special requirements regarding human perception due to their size. Endert et al.'s [28] list of perceptual issues with large displays consists of the importance of peripheral vision, distortion due to extreme viewing angles and long distances, and the different effects of visual variables on large displays compared to traditional displays. **Change Blindness:** Change blindness is a big problem on large displays due to the screen extending into peripheral vision [28] which makes acquiring an overview of the data [65] or comparisons [4] much more difficult. Motion and animation can be used to create awareness in peripheral areas [10], e.g., by temporary flashing, afterglow effects [13], or ambient light around the display [65]. However, while they increase awareness, they do not improve perception of the data itself and users have to move to the actual position to see what exactly changed. **Distortion due to Viewing Angles:** Extreme viewing angles can distort the shape [89] and orientation [28] of data marks as well as

the perception of aspect ratios and color encodings [4, 14]. In contrast, others found color to be more robust to distortion [28, 92] as well as length [14] and position [89]. Dynamically correcting the user's perspective [63], images with multiple scales [40], lenses [49, 50], and potentially using curved displays [67] have been explored to address the issue of distortion. **Physical Navigation:** Users move naturally in front of large displays [8, 71] to utilize the increased space [3]. This physical navigation [8] usually seems to improve performance compared to virtual navigation [47] and increases spatial memory [3, 44]. Generally, users step back from the display to gain an overview of the data [8, 28] and move closer to access details [14, 28, 40]. **Effects of Visual Encodings:** However, perceptual issues emerge when stepping back from the display as visual encodings will be lost at a distance or are perceived differently [28, 40, 92]. This can be mitigated by visual aggregation through using color or creating light and dark regions [28, 92]. Proxemics [9], i.e., adapting the information displayed according to a user's proximity, is another technique to handle this issue, e.g., by adjusting the level of detail [46, 53] or by using geometrical zooming [50]. Hybrid-image visualization [40] takes another direction by embedding different levels of detail into one visualization which are only visible from a particular distance. Another strategy is offloading detail views on additional devices that can be used from afar [38, 49].

C.2 Managing Density and Complexity

The size of large displays is their most important characteristic but at the same time also represents the challenge how the obtained space is best used. Andrews et al. [4] provide several approaches to using the available space to increase the amount of displayed information. **Appropriate Use of Space:** While just increasing the number of data items is an obvious option [20, 80], adding more data dimensions [82] might result in a more comprehensive visualization. Further details, like labels and values [1, 80] and multiple levels of data [52], can be embedded or incorporated into the visualization. Beside scaling visualizations and using the space to enhance one visualization, the acquired space can be used to show multiple, coordinated visualizations [1, 20, 51, 74, 82]. However, given a big enough data set, even large screens might not suffice to display them [27, 39]. Therefore, it is important to consider alternative ways of visualizing additional data besides increasing the screen space. **Scaling of Visual Encodings:** Managing the density and complexity of data is not a problem exclusive to visualization on smaller screens, but poses an issue on large displays as well. Some visual encodings, such as the number of data marks and their size [3, 92] as well as color itself [28], scale well with size and benefit from the additional space. Others, such as the number of perceivable colors and network connectivity [4], i.e., edge crossings, do not improve with additional space. **Overplotting:** Problems such as overplotting also remain on large displays. For scatter plots, clutter reduction, displacement, zooming, and density plots [76] as well as pixel-based techniques [70] can be used to mitigate them. For dense node-link diagrams, edge bundling [57] and cutting [25] can be used to minimize edge crossings and reduce edge congestion. **Data Complexity:** Other common approaches to reduce the complexity of data are clustering and dimensionality reduction [59, 87] as well as utilizing focus + context techniques, e.g., magic lenses [83] or space-folding [27], overview + detail techniques [39], and semantic zooming [49]. For zoomed information spaces, off-screen visualization techniques [12, 35, 66] and offloading information to auxiliary devices to increase screen estate [38, 49] are also beneficial.

C.3 Support for Multiple Users

Large displays are naturally suitable for multi-user scenarios. Their large size and the thereby defined physical space on the screen and in front of it makes them ideal for co-located work and collaboration [41, 43, 50]. **Obstruction and Distraction:** Large displays also lead to challenges as users will inevitably physically obstruct, disturb, and distract each other when working in groups [43]. Prouzeau et al. [68] suggest to leverage interactions that have a large visual footprint to encourage close collaboration and facilitate the frequent switch [43] from closely to loosely coupled collaboration [81].

Furthermore, users can easily lose awareness for what others are doing while working on a different area of the large display [36, 43, 60, 81]. Gutwin & Greenberg [36] state that to maintain workspace awareness, users need to monitor each other's behavior, for instance by showing past interactions of others [37], incorporating visual connections between information [42], using ambient light [65], or embedding animation and off-screen visualization techniques [35, 66]. **Personal Territories:** Studies have shown that users value a personal territory to interact with while still maintaining shared spaces [77, 81]. On vertical displays, territories are mostly transient and dynamic [43] when moving around and nonetheless similar to those on horizontal displays [6]. Having a personal space can either be achieved by a designated space on the display itself or an additional personal device, such as a smartphone [20, 49] or a smart watch [38].

So far, we discussed issues for visualizing data on large, interactive displays and common strategies for addressing them. These presented strategies might sufficiently resolve a particular challenge. However, we now introduce a novel approach and a promising alternative to address many of these challenges and potentially mitigate shortcomings of already existing techniques.

3 GENERAL CONCEPT AND DESIGN SPACE

Our goal with this work is to address the aforementioned challenges and improve the visualization of multivariate data on large, interactive displays. We assume a scenario consisting of one or multiple data analysts in front of a large, vertical, interactive display that either shows a single large visualization or multiple coordinated views (MCV) of multivariate data. We propose to equip each of these analysts with an AR HMD that extends visualizations or parts of visualizations with personalized information. As a delimitation to other work, we solely focus on setups consisting of a single large display and exclude multi-display environments (cf. [49, 53]). We only consider vertically oriented screens of larger size, e.g., digital whiteboards up to wall-sized displays. Consequently, we assume that the display remains stationary during use and does not change its position. Our concepts can be used in a multi-user environment and we discuss cases with both single and multiple users. However, we do not focus on such techniques fostering communication and collaboration between users.

We contribute a design space that provides a systematic overview defining and exploring the boundaries of our proposed setup. It can be used to inspire and guide the development of new visualization and interaction techniques. It consists of (1) the spatial alignment of the display, the users, and the AR content, (2) how different parts of a visualization can be augmented, and (3) how to provide personal content in multi-user scenarios.

3.1 Spatial Alignment of Display and AR Content

The placement of AR content with respect to its spatial alignment to the display is an essential design aspect for augmenting visualizations. In this section, we present possible placement options and outline their respective implications. Reipschläger and Dachsel [72] have previously defined three levels of spatial proximity that relate to our design space: AR objects situated directly on or in front of the display, objects close or directly at the edge of the display, and objects with no spatial relation to the display. Our work focuses mostly on the first two levels, as we consider a strong connection and a precise spatial alignment of AR and display content to be of essential importance.

We classify the display's screen as a frame of reference for the placement of AR objects, defining a planar space around the screen. Content positioned on this plane represents a natural extension of the screen and thus conveys a strong relation to the display. It can be divided based on the cardinal directions to create nine different principal zones to place objects, which can be further extended orthogonally into the third dimension for placing content in front or behind the display¹

¹While these spatial zones are considered primarily for the placement of AR content in the context of this work, they also describe the general relationship of visualizations to their surrounding environment, regardless of the medium used.




Fig. 2. Planar zones of spatial alignment of AR content with the display, which extend in front and behind the display's screen.




Fig. 3. Overview of the components of a visualization that can be extended with AR content and the symbols used throughout the paper to represent those components.

	Spatial Alignment	Visualization Parts	Personal Views
Embedded AR Visualizations	☰	📊	👤
Hinged Visualizations	☰	📊	👤
Curved AR Screen	☰	all parts	👤
AR Brushing and Linking	☰	📊 🔗	👤
Extended Axis Views	☰	📊	👤
AR Visualization Layers	☰	📊 🔗	👤
Magic AR Lenses	☰	📊 🔍	👤
AR Annotations	☰	📊	👤

Fig. 4. Overview of our design space dimensions and their relation to the techniques presented in Sect. 4.

(see Fig. 2). Although in principle, each of the zones is suitable for placing arbitrary AR content, there are special relationships that should be considered. Of all zones, the central zone ☰ is the most important one because it has the strongest alignment to the display and to content on the display. For visualizations with a rectangular shape, the left zone ☰ and right zone ☰ are particularly suitable for content that relates to the vertical dimensions of the visualization. Consequently, the top zone ☰ and bottom zone ☰ are particularly suitable for content that refers to the horizontal dimensions. An example for this is to extend a data dimension shown in the visualization on the screen to the left and right or top and bottom to provide additional data. Since the four corner zones have the weakest spatial connection to the display, the other zones should be preferred for placing objects. AR content may also occupy more than one zone, e.g., top and bottom ☰, left and right ☰, or even all of them ☰. Furthermore, Content on, behind, or in front of a spatial zone is considered to be within that specific zone. With increasing display size, all spatial zones besides the central zone become less important because of their increasing distance to the center of the screen and therefore also to the likely position of the analysts. Conversely, the smaller the display gets, the more important the peripheral zones become. We are therefore convinced that these spatial zones are also useful for the design of techniques for smaller screens. The current position of an analyst in relation to the display can also influence the spatial positioning of the AR content (cf. Proxemic Interaction [9]). For example, if an analyst is closer to the left edge of the display than to the right, AR content can be displayed preferably in the left zone compared to the right zone. In practice, some zones may be unsuitable for placement of AR content if they conflict with or are constrained by the physical condition of the room. To avoid this, the spatial zones can be adapted to the room geometry, e.g., using the floor of the room for the bottom zone.

3.2 Augmentations for Parts of a Visualization

While the previous considerations could be applied to every type of content, this section concentrates on how visualizations themselves can be augmented. A visualization consists of components, such as axes \mathcal{L} , data marks \mathcal{M} , and labels, that are common to almost every visualization and which can be extended using AR techniques (see Fig. 3 for a comprehensive list). In general, augmenting specific parts of a visualization can address problems specific to large displays, which we describe in Sect. 2.4. Especially data marks benefit vastly from using the space around and in front of them to present additional, potentially multivariate data. We can harness the property that every analyst has their own personal view to accommodate for legends \mathcal{L} , labels, tooltips, personal tools as well as selections \mathcal{S} . Annotations \mathcal{A} of the data can be created by, e.g., using a digital pen to write on the display. Another common issue in visualization is overplotting and occlusion (\mathcal{C}) through data marks, for instance in dense scatter plots, or nodes and links \mathcal{L} in node-link diagrams. Using the third dimension, i.e., central zone ☰, and incorporating Augmented Reality can help to address this issue by providing the increased space of another spatial dimension. Additional UI elements \mathcal{U} for interaction with a visualization are fundamental for the functionality of interactive

visualizations. Therefore, we consider them as part of a visualization which can be augmented as well.

We have previously described how content in AR can be spatially aligned with the display. However, this spatial alignment does not solely apply to a display but to individual visualizations and their components as well. As a consequence, depending on the component, zones can either be very small (e.g., in case of a narrow axis \mathcal{L}), restricted by other components in proximity (e.g., a bar that is next to another bar), or may not be available at all (e.g. a line covers only one dimension). Larger zones are more suitable to be extended with AR content. For instance, whereas the zones at the beginning and at the end of an axis can be used for auxiliary components, like UI elements, such as filters, labels, drop down menus, or legends, the adjacent and central zones of the axis are beneficial for displaying additional information for the visualization. Furthermore, general UI elements or tools \mathcal{T} can be completely offloaded into AR if necessary and beneficial. Besides augmenting parts of the visualization, augmentations can be used to enhance the visualization as a whole \mathcal{L} , e.g., by highlighting recently updated visualizations or by extending the whole visualization into the third dimension. Large displays are especially well suited for showing more than one visualization at a time due to their size [4, 51]. As a consequence, the space around a single visualization may be constrained by other visualizations on the display (cf. Sect. 3.1). To address this, the spatial zones can be folded by 90 degrees so that they are now orthogonal to the visualization. Due to this spatial adaptation, the resulting augmentation space becomes similar to a cuboid.

3.3 Personal Views in Multi-User Scenarios

In addition to showing multiple visualizations, large displays are suitable for being used by multiple people at the same time [4]. This can result in several issues discussed in \mathcal{C} (Sect. 2.4). Combining displays with AR HMDs allows for **public views** as well as **personal views**. By default, all content on the display can be seen by all users and all content on the AR HMDs is only perceptible by individual users. However, this is more of a technological consequence and there is a distinction between content on the AR glasses being private because others simply cannot see it or content being private by design. While it is difficult to hide content on the display for individual users, content in AR can be shared among multiple users to foster collaboration. However, there can be benefits in not sharing AR content with others. **Preventing Obstruction:** The additional views created by individual users might obstruct or confuse others. An example is selections created by analysts which highlight data items in other visualizations through Brushing and Linking. Since these highlights would be visible to everyone regardless of whether they are interesting for others, they should not be shared by default. **Personalized Views:** Personalized views and tools can be configured to an individual analyst's needs. For example, Magic Lenses [83] are personal tools which can apply a variety of filters to a data set to support the current analyst's task. Since these views are specific to every person, they should not be shared with others as they have their own personalized views. **Private Information:** People might consider certain information private and therefore want to keep it for themselves. One example is annotations created by analysts which may




Fig. 5. *Embedded AR Visualizations* directly situated on a visualization on the display.




Fig. 6. *Hinged Visualizations* in AR space to address limited perception of far away objects.

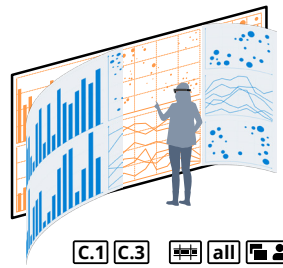


Fig. 7. *Curved AR Screen* to enable an overview of the whole display even when close to the screen.




Fig. 8. *AR Brushing and Linking* with 3D Beziér links indicating connections between visualizations.

contain content that people might feel uncomfortable in immediately sharing with others. These aspects illustrate how AR can be used to enable personal views and in which cases they can be useful. Of course, it is also important for collaborative work, how personal information and views can be shared with other analysts and how they can support social protocols and foster discourse. We discuss this aspect in Sect. 7.

The symbols we have introduced in this section for the different dimensions of the design space will be used in the conceptual illustrations for each technique to reference their relation to the design space. Additionally, Fig. 4 provides an overview of the techniques and their associated parts in our design space.

4 VISUALIZATION TECHNIQUES

In this section, we present techniques that serve as examples to illustrate, explore, and give an overview of our design space. We categorized our techniques by the challenges [C.1] and [C.2] depending on which of them is mainly addressed by a technique (see Sect. 2.4). Each technique also addresses challenge [C.3] as user-specific content is enabled through the use of personal AR HMDs. Due to space constraints, we opted to show the breadth of our proposed concepts by briefly describing a multitude of techniques instead of an in-depth analysis of a few techniques. As a consequence, specific visualizations used to illustrate techniques are also exemplary and the same or similar augmentations can be potentially used for other suitable visualization techniques. Furthermore, we will describe the associated interactions in general, instead of presenting a fine-grained interaction design. Due to the combination of interactive displays and head-mounted AR, we envision the primary use of touch and pen interaction, both for the data and for controlling the individual techniques themselves.

4.1 Perceptual Issues on Large Displays

In this section, four techniques are presented that primarily address perceptual problems on large displays ([C.1]) and make it easier for analysts to better understand connections between the data and maintain an overview of the entire data set.

4.1.1 Embedded AR Visualizations

In contrast to a single visualization, multiple coordinated views are good at representing multivariate data sets. However, studies have shown that this is associated with many context and focus switches and that embedded visualizations may be preferable [71, 90]. We therefore propose to display additional data in AR space located directly on corresponding data objects of an existing visualization (see Fig. 5 and Fig. 14c). These *Embedded AR Visualizations* extend orthogonally from the display. In principle, any type of visualization can be embedded. To ensure a strong spatial connection they can be adapted to the original form of the data marks they extend. For example, a scatter plot can be extended by a cylinder-shaped visualization which encodes additional data dimensions similar to a stacked bar chart. Due to the additional information displayed by the embedded visualization, reading the original visualization may become more difficult. Therefore, it would be sensible to limit *Embedded AR Visualizations* to data marks selected by analysts. They can either be toggled individually with a touch tap

or be shown on-demand for the duration they are touched. *Embedded AR Visualizations* address all three challenges, however primarily [C.1] by avoiding focus switches. They also support [C.2] through their ability to incorporate multiple data levels and additional data into a single data mark. Since they are user-specific, they can be configured by each analyst individually regarding what additional data is represented and how it is visualized ([C.3]). They use the central spatial zone in front of a visualization. We recommend to utilize the potential of the third dimension added by AR instead of simply showing another 2D visualization on a data mark. As this technique is focused mainly on the extension of data marks, size needs to be considered. The smaller a data mark is and the denser they are grouped together, the less space is available for showing embedded visualizations.

4.1.2 Hinged Visualizations

Perception issues through distortion caused by extreme viewing angles are a big problem on large displays (see [C.1]). Therefore, we propose to support awareness for remote parts of the display by rotating these parts in AR space towards the analyst similar to a hinged door (see Fig. 1c and 6). The larger the distance of an analyst towards the *Hinged Visualizations*, the more it is tilted until it is aligned orthogonally to the display and faces the analyst. Accordingly, as analysts move closer, the angle gradually decreases until the AR visualization has vanished to give way to the original visualization on the display. Angles also decrease and the *Hinged Visualizations* ultimately disappear the further the analyst steps back from the display to perceive the visualizations as a whole. *Hinged Visualizations* are automatically generated based on the current selection of an analyst to allow for an overview and a rough comparison of connected values between remote visualizations. However, analysts can also explicitly toggle *Hinged Visualizations* for specific visualizations regardless of their current selection. A potential disadvantage of *Hinged Visualizations* is that they can occlude other parts of the screen. However, we consider this a minor issue since the occluded parts are difficult to detect by the current position of an analyst anyway. Each analyst has individual *Hinged Visualizations* based on their current data selection and position ([C.3]). *Hinged Visualizations* take up the left, center, and right spatial zones and refer to the visualization as a whole. They are aimed at use cases where there is a large number of connected visualizations on the display. It is only of limited assistance when few visualizations are present, since large parts of the screen have to be tilted.

4.1.3 Curved AR Screen

Similar to *Hinged Visualizations*, this technique addresses the perceptual problem of parts of the display being in the analyst's peripheral vision ([C.1]). However, instead of individual visualizations, this technique aims to provide an overview of and awareness for the entire screen. Curved displays improve the perception of peripheral content, as shown by Shupp et al. [79]. Therefore, we propose to simulate such a curved display in AR. The screen, with the exception of an area of configurable size directly in front of the individual analyst, is transferred into AR space onto two curvatures, which extend to the left and right of an analyst (see Fig. 7 and 1d). The curvatures are influenced by an analyst's position. The further away they are from a certain display




Fig. 9. Expanding the axis of visualizations to create *Extended Axis Views* and to show additional information, such as aggregations.




Fig. 10. Showing *AR Visualization Layers* in front of the display to offload data or to show additional data for an existing visualization.




Fig. 11. *Magic AR Lenses* as an example for personalized tools transferred into AR space.




Fig. 12. Personal *AR Annotations* created by pen input in green for one analyst and blue for the other.

edge, the stronger the respective side's curvature becomes. Accordingly, the curvature flattens and even vanishes completely as an analyst approaches the corresponding display edge. The distance of an analyst to the display also reduces the curvature as they gradually move from a detailed view to an overview position. *Curved AR Screen* supports multiple users by providing a personalized view based on each user's position and allows to see content that might otherwise be occluded by other users (C.3). *Curved AR Screen* affects the whole central area of the display but may also extend into the left and right areas adjacent to the display. However, it combines well with other techniques, because content created by other techniques can be transferred onto the curvature as well. *Curved AR Screen* affects all parts of a visualization, in fact the whole display. As the whole screen is curved, the size and number of visualizations on the display is not relevant and the technique works well even for a single large visualization.

4.1.4 AR Brushing and Linking

When an analyst selects one or more items, connected items in other visualizations are often highlighted by Brushing and Linking. However, they can be hard to perceive due to impaired color perception in the peripheral field of view [4] and an analyst's ability to track changes over long distances [28] (see C.1). To address this problem, three-dimensional AR links, like Bézier curves, can be displayed whenever an analyst selects data marks (see Fig. 1a,b, Fig. 8 and Fig. 14d). This provides a clear visual link that extends from the original selection to the linked data and is therefore more easily perceptible than highlights. These links can encode additional values, for example using the color and height [91] or the shape of the curve [16]. Analysts can also filter the links, e.g., to display only links of certain attributes, thus further enhancing the analysis process and contributing to C.2 Managing Density and Complexity. With the number of selections or their individual size growing, the amount of links can lead to visual clutter. To avoid this, we suggest displaying links on demand, e.g., showing them while an analyst touches a data mark, or displaying them only for a short amount of time to establish the necessary awareness. Alternatively, the links can be clustered or bundled [25, 57, 69] but this makes it difficult to follow individual links. In general, AR links can be used to augment data marks, links, and selections. They reside in the central zone of the component they augment and the central zone of the display itself because they can cross several visualizations. Links are created individually for each person based on their current selection and therefore do not disturb others (C.3). As their main purpose is to show connections to other visualizations, AR links work best in a multiple coordinated view environment. However, they might be used in single visualizations as well to show additional relations between individual data marks. Additionally, AR links can indicate completely different types of connections, such as the dynamic information flow between visualizations.

4.2 Managing Density and Complexity

We now present four exemplary techniques that are primarily aimed at making dense and complex data easier to experience and explore (C.2).

4.2.1 Extended Axis Views

The huge amount of data presented on a large display can be overwhelming and make it challenging to get an overview or identify patterns at first (see C.2). Aggregations are an effective way to address this problem [92]. They can be shown for specific axes, i.e., data dimensions, of a visualization, for instance by providing a histogram of the distribution of values along an axis. In order to use the screen space exclusively for the core visualizations, we suggest moving the *Extended Axis Views* into AR space (see Fig. 1a,b, Fig. 9, and Fig. 14a). *Extended Axis Views* are situated directly on the axis of a visualization and expand it to the corresponding adjacent spatial zone, e.g., the bottom spatial zone for the bottom horizontal axis. This creates a strong spatial alignment between the axis and the extended view and thus allows to easily see which data is aggregated. In case of wall-sized displays, the aggregation can be visualized directly on the floor in front of the wall and therefore integrates nicely with the room geometry. When a multitude of adjacent visualizations results in constrained space, the *Extended Axis Views* can be folded as described in Sect. 3.2.

Besides aggregations, this technique can also be used to show additional, personalized (C.3), axis-related data for a visualization. For example, a 2D scatter plot can be extended to a scatter plot matrix with three dimensions. The additional dimension is extended directly at the axis of the original visualization so that a cuboid is created (cf. ScatterDice [26]). This results in a strong spatial relation between the individual parts of the matrix. The cuboid could be rotated with a touch gesture, so that one of the parts formerly in AR is now shown on the display in better detail and can be interacted with using touch input.

4.2.2 AR Visualization Layers

Dense visualizations may result in overplotting [76] and an overall unreadable visualization (see C.2). The increased space provided by large displays can reduce the problem of overlapping data marks but leaves remaining issues, such as edge crossings in node-link diagrams [4]. To tackle these issues, we propose to transfer elements of the original visualization into AR onto multiple layers or show additional data on AR visualization layers parallel to the screen (see Fig. 10). On the one hand, the layers can be used to group data marks semantically, like filtering by specific attributes. In case of node-link diagrams, overlapping edges also benefit from AR by resolving the ambiguity among the edges. Edge crossings are further addressed by moving a single node, a group of nodes, or semantic clusters onto layers in front of the display. Thereby, the connectivity of the transferred nodes to the remaining graph is evident since the connecting edges are shown between the layers. On the other hand, additional visualizations of the same type can be placed in layers in front of or behind the display in AR allowing for superimposed comparison [33] and thus addressing C.2. For example, if a bar chart visualizes the number of movies released in each year in a specific genre, the development for other genres can be displayed as *AR Visualization Layers* for comparison. By scrolling through the layers using a designated touch gesture, analysts can move a certain layer onto the display to interact with it and see it in more detail. As a potential problem, the readability of the original visualization is affected by the superimposed AR layers. We propose two solutions to overcome this problem: First, the *AR Visualization Layers* are only displayed

on-demand to provide a quick summary of similarities among the data. Second, the transparency of the AR visualizations can be reduced in favor of improved readability of the display. Both variants can be combined, i.e., showing the AR visualizations in high transparency and switching to full opacity on-demand. The layers themselves are either directly on, in front of, or behind the display or a specific visualization, taking up the central spatial zone. *AR Visualization Layers* affect data marks, links, and selections, and may even augment the whole visualization. While the additional AR visualizations are only visible to individual analysts, scrolling through the *AR Visualization Layers* may obstruct other analysts that intend to work with previous layers (C3).

4.2.3 Magic AR Lenses

Complex data sets benefit from using specialized tools to facilitate analysis and exploration. These tools can also be transferred to AR space. In the following, this is illustrated by the example of magic lenses, which have been successfully applied to the field of data visualization in the past [83]. Magic lenses can apply arbitrary filters to the underlying data set, e.g., to reduce its complexity (C2), which makes them a useful, personalizable tool for data analysis. However, these tools take up a lot of space on the visualization and can thus disturb other analysts in multi-user scenarios during exploration. They may also extend the current data by private information that analysts do not want to share. By transferring a lens into AR space, it becomes invisible to other analysts (see Fig. 11). Not only are they no longer hindered by the lens, but the lens becomes a personal, user-specific tool that can be freely configured according to an analyst's needs (C3). Controlling the lens can still be achieved via the touch interface of the display. If the spatial alignment between the display and the AR device is synchronized, the analyst can interact with the AR object as if it were still on the screen. In addition, the filters assigned to a lens can now use the third dimension added by AR which enables an entirely new class of filters, such as resolving the overlapping of nodes (similar to *AR Visualization Layers*). *Magic AR Lenses* mostly affect data marks but may, depending on the applied filters, affect other parts of the visualization as well. The lenses themselves reside directly on and are aligned with the display or individual visualization but filters may expand the lenses into the space in front or even behind the screen. Magic Lenses are suitable for single as well as multiple visualizations on the display, although an individual lens is usually confined to only a single visualization, e.g., a map.

4.2.4 AR Annotations

While analyzing data, it is often helpful for analysts to create notes, sketches, and drawings directly on the visualizations to better capture insights [84, 85]. However, the notes of one analyst can be irrelevant and even disturbing for others and influence their own analysis process, which is an essential part of (C3). At the same time, analysts may want to keep their notes private and not immediately share them with others. Therefore, we suggest moving the annotations to AR space. Analysts create their annotations writing on the display, for example using a digital pen. However, these annotations are created solely in AR space (see Fig. 12 and Fig. 14b). Provided the quality of the AR content is sufficiently good, there is no difference between the presentation of annotations on the display from the perspective of an individual analyst. Furthermore, analysts should be enabled to share their annotations with others using additional touch interaction techniques on the display to foster collaboration.

4.3 Combining Techniques

Many of the previously presented techniques can be adapted or even combined according to the requirements of a specific use case. Techniques that affect different parts of a visualization can be used together without conflict. The same applies to techniques that affect different spatial zones. In addition, techniques can be adapted to be compatible with each other to create a mutually beneficial combination. For example, *Magic AR Lenses* can show *Embedded AR Visualizations* for all data marks in the lens area allowing to select whole regions instead of individual data marks. *Magic AR Lenses* can also be combined

with *AR Visualization Layers* to show layers only for a specific area rather than for the whole visualization, which can be advantageous for large visualizations. *AR Brushing and Linking* can be adapted to not only link visualizations on the display but also visualizations residing in AR like *Hinged Visualizations*, *Curved AR Screen*, and *Extended Axis Views*. Furthermore, principles like *Hinged Visualizations* can be applied to techniques such as *Extended Axis Views* and even *Magic AR Lenses*. Therefore, those visualizations would face the analyst as well and ensure better readability over far distances.

5 AR VISUALIZATION FRAMEWORK AND PROTOTYPE

In order to execute and examine our techniques, we first developed a visualization framework that enables traditional visualizations on a display as well as their free placement in the 3D AR space and a combination of the two. Furthermore, it provides extensive capabilities for manipulation and interaction with the visualizations. This was used as a foundation for the prototype combining a display and AR HMDs and implementing the techniques presented in the previous section.

Technical Setup As a large display we used a Microsoft *Surface Hub* (84", 3840x2160 pixel, 1.872x1.053m, touch and pen enabled) and our own tiled display wall (7680x3240 pixel, 4.86x2.06m, touch, pen, and tangible enabled). For Augmented Reality, we provide a MS *HoloLens 1* head-mounted display for each user. We used the *Unity* 3D engine and C# to implement all parts of the prototype and framework. The application on the display and those on the HMDs are stand-alone and initially independent. Since the impression of a single, seamless application is of great importance for the user experience, we developed a custom client server solution using TCP and Open Sound Control (OSC) to synchronize them especially regarding the touch interaction of the users.

u2vis: A Universal Unity Visualization Framework Although many frameworks for information visualization exist, they are mostly limited to traditional 2D visualizations. We investigated whether existing frameworks for *Immersive Analytics* are suitable for our goals. IATK² looked promising but we found it difficult to integrate the interaction capabilities, especially regarding touch, that were essential for our prototype. DXR³ only supports *Unity 2017* and seems to not be maintained anymore. Therefore, we implemented our own, data-driven framework for interactive information visualization called *u2vis*⁴. It provides a joint data basis for all visualizations, separates data and its presentation, and allows for the manipulation of visualizations by touch and other interaction modalities. Our framework currently supports 2D and 3D bar charts, scatter plots, and line charts, as well as parallel coordinates and pie charts. While Fig. 13 only depicts the 2D projection shown on the display, all visualizations can be freely positioned in 3D as well. Our framework is written with an explicit focus on easy extensibility, making it simple and straightforward to add new visualization types. All visualizations are composed by combining several components in the *Unity* editor by drag and drop and configured with the help of a graphical interface. The loose coupling between these components makes it possible to use only needed parts and to easily replace components with alternatives tailored to the current use case.

Prototype Based on this foundation, we implemented the techniques presented in Sect. 4 as a prototype (see Fig. 1). Visualizations can be placed in any arrangement on the display and each user is able to access their desired techniques independent of other users. Since the individual techniques build upon the common foundation of our visualization framework, combining them with each other is easy. To ensure the spatial alignment of the AR content and the screen, the AR content has a shared anchor on the lower left corner of the display. All objects in AR are positioned relative to the anchor. For interaction, we focused on touch input to interact with the visualizations and our techniques and use the digital pen to create personal annotations. When a user touches the display, the corresponding action, e.g., creating a

²<https://github.com/MaximeCordeil/IATK>

³<https://sites.google.com/view/dxr-vis>

⁴Sources available at <https://github.com/im1ldresden/u2vis>

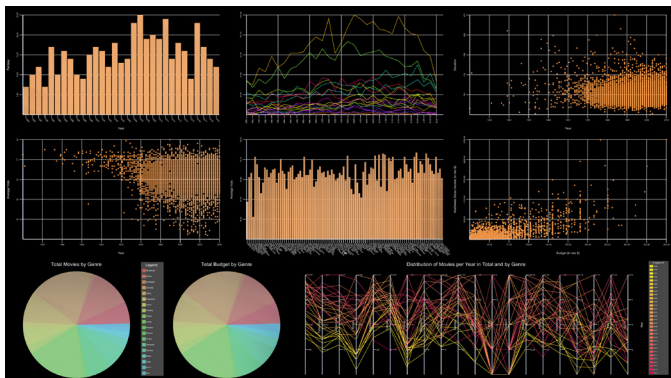


Fig. 13. Screenshot of the *Unity* application running on the display, showcasing the capabilities of our visualization framework. Several visualizations of a movie data set are shown that serve as the foundation for our use case.

selection, is computed on the display and the result, e.g., highlighting a specific data object, is transmitted to the HMDs. Within our framework, each visualization consists of a display part, which can function on its own, and the HMD part, that reacts to events from the display's counterpart.

6 USE CASE: EXPLORING A MOVIE DATA SET

We examine our concepts as well as our implemented techniques in the form of a cognitive walkthrough based on well-established analysis tasks [2] and a specific use case scenario [54], that we implemented in our prototype. We visualize a movie data set⁵ consisting of multiple coordinated views (see Fig. 13). In the following, we describe how Sue and Matt solve analysis and exploration tasks utilizing our techniques.

Sue and Matt decide to investigate the typical duration of movies and how it changed over the years. They approach a scatter plot showing the release date from 1920 to 2019 on the x axis and the duration of a movie on the y axis. To better identify clusters, both choose our *Extended Axis Views* technique to show a histogram at the y axis (see Fig. 14a). Because the histogram is displayed in AR, no other visualizations need to be replaced, moved, or overlaid. With the help of the histogram Matt detects that the typical movie duration is between 90 and 108 minutes. The scatter plot itself reveals that movies nowadays tend to have similar durations but with fewer outliers than in the 20th century. To accentuate this development, Matt sketches with his digital pen to create personal *AR Annotations* (see Fig. 14b).

The two analysts now focus on separate tasks, so they decide to start working independently. Sue is further interested in the development of movies over time, especially in fantasy movies, whereas Matt wants to gain knowledge about specific movies. Therefore, Sue moves towards a bar chart showing the number of fantasy movies over time. A hold gesture on a data mark, i.e., a specific year, reveals another bar chart orthogonal to the display, using the *Embedded AR Visualizations* technique, which shows the number of movies in all genres for that particular year (see Fig. 14c). While the most fantasy movies (25) were produced in 2007, Sue finds out that this is still vastly below the number of produced dramas in that year (148). Through *Embedded AR Visualizations*, Sue can access this information on demand directly within the context of an existing visualization, avoiding focus switches and saving valuable screen space.

Meanwhile, Matt looks at the scatter plot showing budget and worldwide gross income and selects two movies with similar budget but a huge difference in income, which are *Star Wars: Episode VII - The Force Awaken* and *The Chronicles of Narnia: Prince Caspian*. With our *AR Brushing and Linking* technique, the selected movies are highlighted in all other relevant visualizations with 3D links in AR (see Fig. 14d). This, together with *Hinged Visualizations* (or alternatively the *Curved AR Screen*) guides his awareness to the relevant visualizations. Thus,

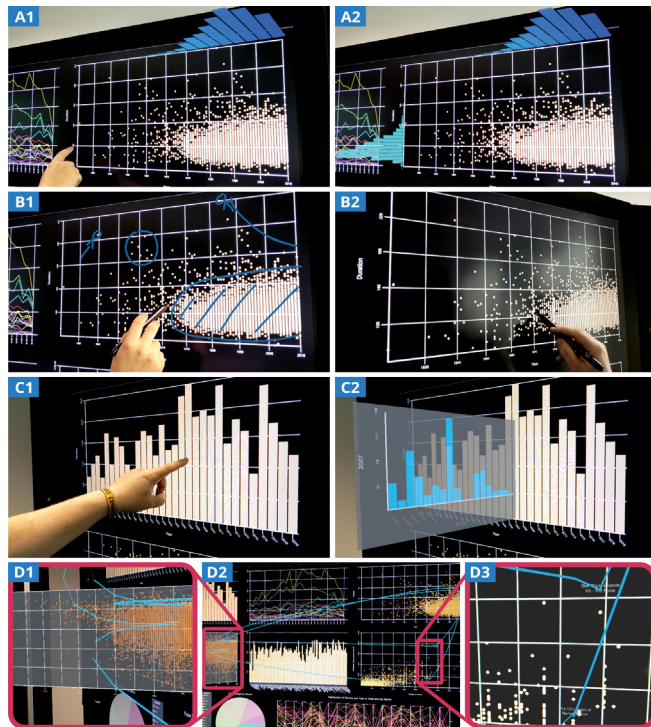


Fig. 14. Illustrating our use case: (a1) Selecting *Extended Axis Views* that (a2) show an aggregation of movie duration. (b1) One analyst annotating a scatter plot, (b2) without the sketches being visible to others. (c1) Selecting the year 2007 to show (c2) *Embedded AR Visualizations* of the amount of movies released in each genre for this year. (d2) *AR Brushing and Linking* between visualizations, (d3) a selection of two movies in a scatter plot visualizing budget and worldwide gross income, (d1) *Hinged Visualizations* showing the average votes for those movies.

he does not have to move to the other side of the display to obtain an overview and an approximate position of the selected data. Retrieving exact values, however, requires him to approach the visualization. One of the *Hinged Visualizations* shows a scatter plot of release date in comparison to average user vote. Based on the selection and the accompanying links, it is easy to spot the desired movies in this chart and to derive that *Star Wars* has an average rating of 7.9 and *Narnia* of 6.5. Since only Matt sees the selection, the links, and the curved or hinged parts, Sue can keep exploring the data without visual clutter or disturbance.

This cognitive walkthrough shows how our proposed setup supports common analysis tasks and highlights the potential benefits compared to traditional setups, such as saving valuable screen space, integrating additional information directly in place, and not obstructing other analysts. Based on this use case and our own experiences, we are confident that our approach has high potential in assisting data analysis involving multiple users. However, this claim must be evaluated in a user study with specific analysis tasks that need to be solved using our setup in comparison to conventional multiple coordinated views systems (cf. [51]).

7 DISCUSSION

To summarize our experiences based on the prototype and the cognitive walkthrough from the use case, we discuss the perception of AR content situated on the display as well as the technical limitations of our setup and state-of-the-art AR in general. Furthermore, we consider collaborative work and possible future applications.

Perception of AR Content and Display An obvious question when extending or overlaying a display with AR content is how well each medium can be perceived and if they influence each other. In the case of AR content, this largely depends on the technology. The *HoloLens 1*

⁵available on Kaggle (<https://www.kaggle.com/stefano1eone992/imdb-extensive-dataset>)

we used creates a relatively opaque image with little bleed-through of the background. However, resolution and field of view are limited, so displaying very small or detailed content is problematic. This is especially true for very small text. We recommend to show such content only on the display as its high resolution predestines it for these fine details. Furthermore, it should be noted that the perception of colors in AR is inferior to conventional displays [88]. Therefore, visual attributes, such as shape or size, should be preferably used since they are well perceptible in AR. Otherwise, displaying AR content superimposed on the display is usually unproblematic with respect to perception.

Regarding the display itself, a possible issue is that AR content makes it difficult to see the visualizations *on* the display. AR content should therefore not be placed arbitrarily but always in relation to objects on the display, so that they support and extend them. Nonetheless, switching between AR and display content may still lead to focus and attention issues. Thus, we recommend to carefully examine which content to transfer to AR and whether its benefits outweigh its costs. However, AR can also reduce focus switches, for instance through *Embedded AR Visualizations*. Furthermore, we suggest the use of on-demand techniques to reduce the number of objects displayed simultaneously. In general, Augmented Reality technology still has lots of room for improvement and we expect it to further advance over the coming years. We believe that the visual quality of AR displays might become similar to the quality of high-resolution displays.

One question that arises in this context is: Why is an interactive display still needed at all if the quality of an AR HMD is sufficiently good? First, the display has a haptic affordance not readily available in AR and analysts can interact with data using natural touch and pen input. Second, the display is visible to everyone without requiring additional technology and can be used as a shared space for data exploration. Third, the display offers a physical point of reference which can be used for the orientation and alignment of AR content.

Combining Environments for Data Analysis Using different interactive environments has great potential for data analysis which was demonstrated by previous work on large display environments [3, 51], physical visualizations [23], and immersive environments [24]. We believe that combining different environments can leverage their advantages and mitigate their disadvantages at the same time. In this case, interaction constitutes a special challenge due to switches between interaction paradigms. Similar to Butcher et al. [17], our concepts address the missing haptic feedback of mid-air interaction commonly used in AR by using touch interaction to control AR. However, possibly having to reach through AR objects may nevertheless seem strange to users due to obscuring their hands. Furthermore, the optimal balance between content on the display and content situated in AR is an interesting question that calls for further study. In this context, we have proposed an initial approach through our concepts. Nevertheless, we agree with Bach et al. [7] on not advocating for a complete substitution of existing environments and techniques. We consider our approach to be an interesting addition to other established strategies to address the known issues (see Sect. 2.4) of large display environments and to make them an even more powerful and sophisticated tool for data analysts, especially for multi-user scenarios.

Alignment of AR Content in Relation to the Display One of our assumptions in this work is that the alignment of AR content with the display has a positive effect on analysts' performance. However, it is unclear how this impacts the user's experience and it should be further examined in comparison to alternative placement options. As mentioned in Sect. 3.1, a precise spatial alignment of the AR content with the content shown on display is essential for many of the presented techniques. This places great demands on the tracking technology used to ensure that, for instance, AR objects bound to specific data marks in a dense visualization are not mistakenly associated with adjacent data marks. With the *HoloLens 1* we used for our prototype, limited accuracy and drift are major issues that get worse with increasing display sizes. Feature-based tracking, as used by the *HoloLens*, reaches its limits when users are extremely close to a large or wall-sized display as only few features remain that can be tracked. Furthermore, the

alignment of objects is not stable regarding rotation and translation. As a consequence, objects constantly shift their position as analysts move around in front of the display. We expect these problems to improve and eventually be solved with subsequent technologies.

Collaboration Between Multiple Analysts By using AR HDMs, our setup currently has a technology-caused intrinsic personal view for all analysts. A system that would be focused on fostering communication and collaboration requires dedicated techniques for sharing content between analysts. However, simply sharing all content between everyone wastes much of the inherent potential of the system, as personalized views can vary greatly from person to person or analysts may want to keep their annotations private. The resulting interesting design question is which content should be visible to everyone or only to certain analysts and which content should remain private. This, however, is beyond the scope of this work and left for future investigation. In addition, we refer to the large body of work in the field of computer-supported cooperative work (CSCW) for general techniques to support collaboration.

Connecting Multiple Visualizations with AR While AR can be used to better illustrate relationships between several visualizations (see *AR Brushing and Linking*), there is potential in using AR to visualize the connection in more detail. AR could be used to create temporary, user-specific views derived from the connection between existing visualizations. For example, an AR view can highlight the difference between two visualizations and aid in comparison tasks [32, 33]. We think there is merit in investigating the types of connections that can be visualized and how to best represent them using AR.

Further Potential Utilization of AR In the context of this work, we have so far talked about the direct extension of the display with AR content. However, the use of AR content that is independent of the display but is instead based on the spatial position of the user can be useful as well. One possible approach is to move away from simply extending content on the display and to focus on the AR visualizations. In this case, the display is still valuable for naturally interacting with the AR content. For example, a 3D AR visualization, such as a height map, could be positioned on the display. Specialized tools, such as slicing planes, can then be manipulated using touch interaction to facilitate a convenient way for the exploration of 3D visualizations. Although our chosen setup has a particular set of affordances and requirements, our techniques can serve as an inspiration for further adaption. For example, they can be extended to horizontal displays and even mobile devices as well. They are not limited to information visualization but can be applied to data and scientific visualization as well.

8 CONCLUSION

We presented the combination of a large interactive display with head-mounted, personal Augmented Reality for information visualization. Our contribution consists of a design space, techniques for extending visualizations with AR, and the data-driven AR visualization framework *u2vis* as well as the accompanying prototype we developed. We further investigated our approach with a cognitive walkthrough of a use case scenario, which indicates that our techniques can help to reduce perceptual issues in large displays for data exploration and to facilitate the analysis of dense data sets. However, further and more extensive evaluation is required to validate these indications. We are therefore convinced that the extension of large displays with AR has great potential for information and data analysis. We hope that our work inspires the exploration of this new exciting class of visualization applications, which combine interactive displays with Augmented Reality.

ACKNOWLEDGMENTS

We want to thank Mats Ole Ellenberg, Severin Engert, and Remke Dirk Albrecht for their support in creating this publication. This work was funded by DFG grant 389792660 as part of TRR 248 (see <https://perspicuous-computing.science>) and the DFG as part of Germany's Excellence Strategy Clusters of Excellence EXC-2068 390729961 *Physics of Life* and EXC 2050/1 390696704 *Centre for Tactile Internet with Human-in-the-Loop (CeTI)* of TU Dresden.

REFERENCES

- [1] M. Agarwal, A. Srinivasan, and J. Stasko. VisWall: Visual data exploration using direct combination on large touch displays. In *VIS '19*, pp. 26–30, 2019. doi: 10.1109/VISUAL.2019.8933673
- [2] R. Amar, J. Eagan, and J. Stasko. Low-level components of analytic activity in information visualization. In *INFOVIS '05*, pp. 111–117, Oct 2005. doi: 10.1109/INFVIS.2005.1532136
- [3] C. Andrews, A. Endert, and C. North. Space to think: Large high-resolution displays for sensemaking. In *CHI '10*, p. 55–64. ACM, 2010. doi: 10.1145/1753326.1753336
- [4] C. Andrews, A. Endert, B. Yost, and C. North. Information visualization on large, high-resolution displays: Issues, challenges, and opportunities. *InfoVis*, 10(4):341–355, 2011. doi: 10.1177/1473871611415997
- [5] R. Arora, R. Habib Kazi, T. Grossman, G. Fitzmaurice, and K. Singh. SymbiosisSketch: Combining 2D & 3D sketching for designing detailed 3D objects in situ. In *CHI '18*, pp. 185:1–185:15. ACM, 2018. doi: 10.1145/3173574.3173759
- [6] A. Azad, J. Ruiz, D. Vogel, M. Hancock, and E. Lank. Territoriality and behaviour on and around large vertical publicly-shared displays. In *DIS '12*, p. 468–477. ACM, 2012. doi: 10.1145/2317956.2318025
- [7] B. Bach, R. Sicat, J. Beyer, M. Cordeil, and H. Pfister. The hologram in my hand: How effective is interactive exploration of 3D visualizations in immersive tangible augmented reality? *IEEE TVCG '18*, 24:457–467, 2018. doi: 10.1109/TVCG.2017.2745941
- [8] R. Ball, C. North, C. North, and D. A. Bowman. Move to improve: Promoting physical navigation to increase user performance with large displays. In *CHI '07*. ACM, 2007. doi: 10.1145/1240624.1240656
- [9] T. Ballendat, N. Marquardt, and S. Greenberg. Proxemic interaction: Designing for a proximity and orientation-aware environment. In *ITS '10*, pp. 121–130. ACM, 2010. doi: 10.1145/1936652.1936676
- [10] L. Bartram, C. Ware, and T. Calvert. Moticons: Detection, distraction and task. *International Journal of Human-Computer Studies*, 58:515–545, 2003. doi: 10.1016/S1071-5819(03)00021-1
- [11] A. Batch, A. Cunningham, M. Cordeil, N. Elmquist, T. Dwyer, B. H. Thomas, and K. Marriott. There is no spoon: Evaluating performance, space use, and presence with expert domain users in immersive analytics. *IEEE TVCG '20*, 26:536–546, 2020. doi: 10.1109/TVCG.2019.2934803
- [12] P. Baudisch and R. Rosenholtz. Halo: A technique for visualizing off-screen objects. In *CHI '03*, pp. 481–488. ACM, 2003. doi: 10.1145/642611.642695
- [13] P. Baudisch, D. Tan, M. Collomb, D. Robbins, K. Hinckley, M. Agrawala, S. Zhao, and G. Ramos. Phosphor: Explaining transitions in the user interface using afterglow effects. In *UIST '16*, p. 169–178. ACM, 2006. doi: 10.1145/1166253.1166280
- [14] A. Bezerianos and P. Isenberg. Perception of visual variables on tiled wall-sized displays for information visualization applications. *IEEE TVCG '12*, 18:2516–2525, 2012. doi: 10.1109/TVCG.2012.251
- [15] N. Bogdan, T. Grossman, and G. Fitzmaurice. HybridSpace: Integrating 3D freehand input and stereo viewing into traditional desktop applications. In *3DUI '14*, pp. 51–58, 2014. doi: 10.1109/3DUI.2014.6798842
- [16] W. Büschel, S. Vogt, and R. Dachsel. Augmented reality graph visualizations: Investigation of visual styles in 3D node-link diagrams. *CG&A '19*, pp. 29–40, 2019. doi: 10.1109/MCG.2019.2897927
- [17] S. Butscher, S. Hubenschmid, J. Müller, J. Fuchs, and H. Reiterer. Clusters, trends, and outliers: How immersive technologies can facilitate the collaborative analysis of multidimensional data. In *CHI '18*, pp. 90:1–90:12. ACM, 2018. doi: 10.1145/3173574.3173664
- [18] M. Cavallo, M. Dholakia, M. Havlena, K. Ocheltree, and M. Podlasek. Dataspace: A reconfigurable hybrid reality environment for collaborative information analysis. In *IEEE VR '19*, pp. 145–153, March 2019. doi: 10.1109/VR.2019.8797733
- [19] M. Cavallo, M. Dholakia, M. Havlena, K. Ocheltree, and M. Podlasek. Immersive insights: A hybrid analytics system for collaborative exploratory data analysis. In *VRST '19*. ACM, 2019. doi: 10.1145/3359996.3364242
- [20] M. Chegini, K. Andrews, T. Schreck, and A. Sourin. Multiple linked-view exploration on large displays facilitated by a secondary handheld device. In *IWAIT '19*. SPIE, 2019. doi: 10.1117/12.2521648
- [21] Z. Chen, W. Tong, Q. Wang, B. Bach, and H. Qu. Augmenting static visualizations with PapARVis designer. In *CHI '20*, p. 1–12. Association for Computing Machinery, 2020. doi: 10.1145/3313831.3376436
- [22] M. Cordeil, A. Cunningham, T. Dwyer, B. H. Thomas, and K. Marriott. ImAxes: Immersive axes as embodied affordances for interactive multi-variate data visualisation. In *UIST '17*, pp. 71–83. ACM, 2017. doi: 10.1145/3126594.3126613
- [23] P. Dragicevic, Y. Jansen, and A. Vande Moere. Data physicalization. In *Springer Handbook of Human Computer Interaction*. Springer, 2020.
- [24] T. Dwyer, K. Marriott, T. Isenberg, K. Klein, N. Riche, F. Schreiber, W. Stuerzlinger, and B. H. Thomas. *Immersive Analytics: An Introduction*, pp. 1–23. Springer, 2018. doi: 10.1007/978-3-030-01388-2_1
- [25] D. Edge, J. Larson, M. Mobius, and C. White. Trimming the hairball: Edge cutting strategies for making dense graphs usable. In *Big Data '18*, pp. 3951–3958, 2018. doi: 10.1109/BigData.2018.8622521
- [26] N. Elmquist, P. Dragicevic, and J. D. Fekete. Rolling the dice: Multidimensional visual exploration using scatterplot matrix navigation. *IEEE TVCG '08*, 14:1539–1148, 2008. doi: 10.1109/TVCG.2008.153
- [27] N. Elmquist, N. Henry, Y. Riche, and J.-D. Fekete. Melange: Space folding for multi-focus interaction. In *CHI '08*, p. 1333–1342. ACM, 2008. doi: 10.1145/1357054.1357263
- [28] A. Endert, C. Andrews, Y. H. Lee, and C. North. Visual encodings that support physical navigation on large displays. In *GI '11*, p. 103–110. Canadian Human-Computer Communications Society, 2011.
- [29] S. Feiner and A. Shamash. Hybrid user interfaces: Breeding virtually bigger interfaces for physically smaller computers. In *UIST '91*, p. 9–17. ACM, 1991. doi: 10.1145/120782.120783
- [30] J. A. W. Filho, C. M. D. S. Freitas, and L. Nedel. Comfortable immersive analytics with the VirtualDesk metaphor: Case studies and perspectives. *IEEE CG&A*, pp. 1–1, 2019. doi: 10.1109/MCG.2019.2898856
- [31] J. A. W. Filho, W. Stuerzlinger, and L. Nedel. Evaluating an immersive space-time cube geovisualization for intuitive trajectory data exploration. *IEEE TVCG '20*, 26:514–524, 2020. doi: 10.1109/TVCG.2019.2934415
- [32] M. Gleicher. Considerations for visualizing comparison. *IEEE TVCG '18*, 24:413–423, 2018. doi: 10.1109/TVCG.2017.2744199
- [33] M. Gleicher, D. Albers, R. Walker, I. Jusufi, C. D. Hansen, and J. C. Roberts. Visual comparison for information visualization. *InfoVis*, 10:289–309, 2011. doi: 10.1177/1473871611416549
- [34] J. Grubert, M. Heinisch, A. Quigley, and D. Schmalstieg. MultiFi: Multi fidelity interaction with displays on and around the body. In *CHI '15*, p. 3933–3942. ACM, 2015. doi: 10.1145/2702123.2702331
- [35] S. Gustafson, P. Baudisch, C. Gutwin, and P. Irani. Wedge: Clutter-free visualization of off-screen locations. In *CHI '08*, pp. 787–796. ACM, 2008. doi: 10.1145/1357054.1357179
- [36] C. Gutwin and S. Greenberg. A descriptive framework of workspace awareness for real-time groupware. *CSCW '02*, 11:411–446, 2002. doi: 10.1023/A:1021271517844
- [37] A. H. Hajizadeh, M. Tory, and R. Leung. Supporting awareness through collaborative brushing and linking of tabular data. *IEEE TVCG '13*, 19:2189–2197, 2013. doi: 10.1109/TVCG.2013.197
- [38] T. Horak, S. K. Badam, N. Elmquist, and R. Dachsel. When David meets Goliath: Combining smartwatches with a large vertical display for visual data exploration. In *CHI '18*, pp. 19:1–19:13. ACM, 2018. doi: 10.1145/3173574.3173593
- [39] A. Ion, Y.-L. B. Chang, M. Haller, M. Hancock, and S. D. Scott. Canyon: Providing location awareness of multiple moving objects in a detail view on large displays. In *CHI '13*, p. 3149–3158. ACM, 2013. doi: 10.1145/2470654.2466431
- [40] P. Isenberg, P. Dragicevic, W. Willett, A. Bezerianos, and J. D. Fekete. Hybrid-image visualization for large viewing environments. *IEEE TVCG '13*, 19:2346–2355, 2013. doi: 10.1109/TVCG.2013.163
- [41] P. Isenberg, N. Elmquist, J. Scholtz, D. Cernea, Kwan-Liu Ma, and H. Hagen. Collaborative visualization: Definition, challenges, and research agenda. *InfoVis*, 10:310–326, 2011. doi: 10.1177/1473871611412817
- [42] P. Isenberg, D. Fisher, S. A. Paul, M. R. Morris, K. Inkpen, and M. Czerwinski. Co-located collaborative visual analytics around a tabletop display. *IEEE TVCG '12*, 18, 2012. doi: 10.1109/TVCG.2011.287
- [43] M. R. Jakobsen and K. Hornbæk. Up close and personal: Collaborative work on a high-resolution multitouch wall display. *ACM Transactions on Computer-Human Interaction*, 21, 2014. doi: 10.1145/2576099
- [44] M. R. Jakobsen and K. Hornbæk. Is moving improving? Some effects of locomotion in wall-display interaction. In *CHI '15*, p. 4169–4178. ACM, 2015. doi: 10.1145/2702123.2702312
- [45] M. R. Jakobsen and K. Hornbæk. Negotiating for space? Collaborative work using a wall display with mouse and touch input. In *CHI '16*, p. 2050–2061. ACM, 2016. doi: 10.1145/2858036.2858158
- [46] M. R. Jakobsen, Y. Sahlemariam Haile, S. Knudsen, and K. Hornbæk. Information visualization and proxemics: Design opportunities and empir-

- ical findings. *IEEE TVCG '13*, 2013. doi: 10.1109/TVCG.2013.166
- [47] Y. Jansen, J. Schjerlund, and K. Hornbæk. Effects of locomotion and visual overview on spatial memory when interacting with wall displays. In *CHI '19*, pp. 291:1–291:12. ACM, 2019. doi: 10.1145/3290605.3300521
- [48] T. Kim, B. Saket, A. Endert, and B. MacIntyre. VisAR: Bringing interactivity to static data visualizations through augmented reality. In *IEEE VIS '17 Workshop on Immersive Analytics*, 2017.
- [49] U. Kister, K. Klamka, C. Tominski, and R. Dachsel. GraSp: Combining spatially-aware mobile devices and a display wall for graph visualization and interaction. *Computer Graphics Forum*, 36:503–514, 2017. doi: 10.1111/cgf.13206
- [50] U. Kister, P. Reipschläger, F. Matulic, and R. Dachsel. BodyLenses: Embodied magic lenses and personal territories for wall displays. In *ITS '15*, pp. 117–126. ACM, 2015. doi: 10.1145/2817721.2817726
- [51] R. Langner, U. Kister, and R. Dachsel. Multiple coordinated views at large displays for multiple users: Empirical findings on user behavior, movements, and distances. *IEEE TVCG '19*, 25:608–618, 2019. doi: 10.1109/TVCG.2018.2865235
- [52] R. Langner, U. Kister, C. Engel, T. Horak, and R. Dachsel. Towards a fluid interaction concept bringing star plots to interactive displays. In *ITS '15*, pp. 349–354. ACM, 2015. doi: 10.1145/2817721.2823505
- [53] D. Ledo, S. Greenberg, N. Marquardt, and S. Boring. Proxemic-aware controls: Designing remote controls for ubiquitous computing ecologies. In *MobileHCI '15*. ACM, 2015. doi: 10.1145/2785830.2785871
- [54] D. Ledo, S. Houben, J. Vermeulen, N. Marquardt, L. Oehlberg, and S. Greenberg. Evaluation strategies for HCI toolkit research. In *CHI '18*, pp. 36:1–36:17. ACM, 2018. doi: 10.1145/3173574.3173610
- [55] A. Lee, D. Archambault, and M. Nacenta. Dynamic network plaid: A tool for the analysis of dynamic networks. In *CHI '19*, pp. 130:1–130:14. ACM, 2019. doi: 10.1145/3290605.3300360
- [56] B. Lee, P. Isenberg, N. H. Riche, and S. Carpendale. Beyond mouse and keyboard: Expanding design considerations for information visualization interactions. *IEEE TVCG '12*, 18:2689–2698, 2012. doi: 10.1109/TVCG.2012.204
- [57] A. Lhuillier, C. Hurter, and A. Telea. State of the art in edge and trail bundling techniques. *Computer Graphics Forum*, 36:619–645, 2017. doi: 10.1111/cgf.13213
- [58] J. Liu, A. Prouzeau, B. Ens, and T. Dwyer. Design and evaluation of interactive small multiples data visualisation in immersive spaces. In *IEEE VR '20*, pp. 1–8, 2020. doi: 10.1109/VR46266.2020.00-23
- [59] S. Liu, D. Maljovec, B. Wang, P. Bremer, and V. Pascucci. Visualizing high-dimensional data: Advances in the past decade. *IEEE TVCG '17*, 23:1249–1268, 2017. doi: 10.1109/TVCG.2016.2640960
- [60] M. Mateescu, C. Pimmer, C. Zahn, D. Klinkhammer, and H. Reiterer. Collaboration on large interactive displays: A systematic review. *Human-Computer Interaction*, 2019. doi: 10.1080/07370024.2019.1697697
- [61] D. Mendes, F. Fonseca, B. Araújo, A. Ferreira, and J. Jorge. Mid-air interactions above stereoscopic interactive tables. In *3DUI '14*, pp. 3–10, 2014. doi: 10.1109/3DUI.2014.6798833
- [62] A. Millette and M. J. McGuffin. DualCAD: Integrating augmented reality with a desktop GUI and smartphone interaction. In *ISMAR '16*, pp. 21–26, 2016. doi: 10.1109/ISMAR-Adjunct.2016.0030
- [63] M. A. Nacenta, S. Sakurai, T. Yamaguchi, Y. Miki, Y. Itoh, Y. Kitamura, S. Subramanian, and C. Gutwin. E-Conic: A perspective-aware interface for multi-display environments. In *UIST '07*, p. 279–288. ACM, 2007. doi: 10.1145/1294211.1294260
- [64] P. Peltonen, E. Kurvinen, A. Salovaara, G. Jacucci, T. Ilmonen, J. Evans, A. Oulasvirta, and P. Saarikko. It's mine, don't touch! Interactions at a large multi-touch display in a city centre. In *CHI '08*, p. 1285–1294. ACM, 2008. doi: 10.1145/1357054.1357255
- [65] F. Perteneder, E.-M. B. Grossauer, J. Leong, W. Stuerzlinger, and M. Haller. Glowworms and fireflies: Ambient light on large interactive surfaces. In *CHI '16*, p. 5849–5861. ACM, 2016. doi: 10.1145/2858036.2858524
- [66] J. Petford, I. Carson, M. A. Nacenta, and C. Gutwin. A comparison of notification techniques for out-of-view objects in full-coverage displays. In *CHI '19*. ACM, 2019. doi: 10.1145/3290605.3300288
- [67] A. Prouzeau, A. Bezerianos, and O. Chapuis. Visual immersion in the context of wall displays. In *ISS '16*, p. 33–39. ACM, 2016. doi: 10.1145/3009939.3009945
- [68] A. Prouzeau, A. Bezerianos, and O. Chapuis. Evaluating multi-user selection for exploring graph topology on wall-displays. *IEEE TVCG '17*, 23:1936–1951, 2017. doi: 10.1109/TVCG.2016.2592906
- [69] A. Prouzeau, A. Lhuillier, B. Ens, D. Weiskopf, and T. Dwyer. Visual link routing in immersive visualisations. In *ISS '19*, p. 241–253. ACM, 2019. doi: 10.1145/3343055.3359709
- [70] R. G. Raidou, M. E. Gröller, and M. Eisemann. Relaxing dense scatter plots with pixel-based mappings. *IEEE TVCG '19*, 25:2205–2216, 2019. doi: 10.1109/TVCG.2019.2903956
- [71] F. Rajabiyazdi, J. Walny, C. Mah, J. Brosz, and S. Carpendale. Understanding researchers' use of a large, high-resolution display across disciplines. In *ITS '15*, pp. 107–116. ACM, 2015. doi: 10.1145/2817721.2817735
- [72] P. Reipschläger and R. Dachsel. DesignAR: Immersive 3D-modeling combining augmented reality with interactive displays. In *ISS '19*, pp. 29–42. ACM, 2019. doi: 10.1145/3343055.3359718
- [73] J. Rekimoto and M. Saitoh. Augmented surfaces: A spatially continuous work space for hybrid computing environments. In *CHI '99*, pp. 378–385. ACM, 1999. doi: 10.1145/302979.303113
- [74] J. C. Roberts. State of the art: Coordinated multiple views in exploratory visualization. In *CMV '07*, pp. 61–71, 2007. doi: 10.1109/CMV.2007.20
- [75] J. C. Roberts, P. D. Ritsos, S. K. Badam, D. Brodbeck, J. Kennedy, and N. Elmqvist. Visualization beyond the desktop—the next big thing. *CG&A '14*, 34:26–34, 2014. doi: 10.1109/MCG.2014.82
- [76] A. Sarikaya and M. Gleicher. Scatterplots: Tasks, data, and designs. *IEEE TVCG '18*, 24:402–412, 2018. doi: 10.1109/TVCG.2017.2744184
- [77] S. D. Scott, M. S. T. Carpendale, and K. M. Inkpen. Territoriality in collaborative tabletop workspaces. In *CSCW '04*, pp. 294–303. ACM, 2004. doi: 10.1145/1031607.1031655
- [78] M. Sereno, L. Besançon, and T. Isenberg. Supporting volumetric data visualization and analysis by combining augmented reality visuals with multi-touch input. In *EuroVis '19 - Posters*, 2019. doi: 10.2312/eurovis.2019.1136
- [79] L. Shupp, C. A. PhD, M. D.-K. PhD, B. Yost, and C. North. Shaping the display of the future: The effects of display size and curvature on user performance and insights. *Human-Computer Interaction*, 24:230–272, 2009. doi: 10.1080/07370020902739429
- [80] T. Sun, Y. Ye, I. Fujishiro, and K. Ma. Collaborative visual analysis with multi-level information sharing using a wall-size display and see-through HMDs. In *PacificVis '19*, pp. 11–20, 2019. doi: 10.1109/PacificVis.2019.00010
- [81] A. Tang, M. Tory, B. Po, P. Neumann, and S. Carpendale. Collaborative coupling over tabletop displays. In *CHI '06*, pp. 1181–1190. ACM, 2006. doi: 10.1145/1124772.1124950
- [82] M. Tobiasz, P. Isenberg, and S. Carpendale. Lark: Coordinating co-located collaboration with information visualization. *IEEE TVCG '09*, 15:1065–1072, 2009. doi: 10.1109/TVCG.2009.162
- [83] C. Tominski, S. Gladisch, U. Kister, R. Dachsel, and H. Schumann. Interactive lenses for visualization: An extended survey. *Computer Graphics Forum*, 36:173–200, 9 2017. doi: 10.1111/cgf.12871
- [84] J. Walny, S. Huron, C. Perin, T. Wun, R. Pusch, and S. Carpendale. Active reading of visualizations. *IEEE TVCG '18*, 24(1):770–780, 2018.
- [85] J. Walny, B. Lee, P. Johns, N. Henry Riche, and S. Carpendale. Understanding pen and touch interaction for data exploration on interactive whiteboards. *IEEE TVCG '12*, 18, 2012. doi: 10.1109/TVCG.2012.275
- [86] X. Wang, L. Besançon, D. Rousseau, M. Sereno, M. Ammi, and T. Isenberg. Towards an understanding of augmented reality extensions for existing 3D data analysis tools. In *CHI '20*, CHI '20, p. 1–13. ACM, 2020. doi: 10.1145/3313831.3376657
- [87] J. Wenskovich, I. Crandell, N. Ramakrishnan, L. House, S. Leman, and C. North. Towards a systematic combination of dimension reduction and clustering in visual analytics. *IEEE TVCG '18*, 24:131–141, 2018. doi: 10.1109/TVCG.2017.2745258
- [88] M. Whitlock, S. Smart, and D. A. Szafir. Graphical perception for immersive analytics. In *IEEE VR '20*, pp. 616–625, 2020. doi: 10.1109/VR46266.2020.00084
- [89] D. Wigdor, C. Shen, C. Forlines, and R. Balakrishnan. Perception of elementary graphical elements in tabletop and multi-surface environments. In *CHI '07*, p. 473–482. ACM, 2007. doi: 10.1145/1240624.1240701
- [90] W. Willett, Y. Jansen, and P. Dragicevic. Embedded data representations. *IEEE TVCG '17*, 23:461–470, 2017. doi: 10.1109/TVCG.2016.2598608
- [91] Y. Yang, T. Dwyer, B. Jenny, K. Marriott, M. Cordeil, and H. Chen. Origin-destination flow maps in immersive environments. *IEEE TVCG '19*, 25:693–703, 2019. doi: 10.1109/TVCG.2018.2865192
- [92] B. Yost, Y. Haciahmetoglu, and C. North. Beyond visual acuity: The perceptual scalability of information visualizations for large displays. In *CHI '07*, p. 101–110. ACM, 2007. doi: 10.1145/1240624.1240639