

MARVIS: Combining Mobile Devices and Augmented Reality for Visual Data Analysis

Ricardo Langner
 Technische Universität
 Dresden
 Dresden, Germany
 langner@acm.org

Marc Satkowski
 Technische Universität
 Dresden
 Dresden, Germany
 msatkowski@acm.org

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

Raimund Dachzelt*[†]
 Technische Universität
 Dresden
 Dresden, Germany
 dachzelt@acm.org




Figure 1: Data visualization using mobile devices and Augmented Reality head-mounted displays: (a) Envisioned usage scenario; (b) 2D scatterplot extended with superimposed 3D trajectories/paths; (c) 3D wall visualization in AR aligned with the mobile device; (d) Use of AR for seamless display extension around a geographic map; (e) Combining visualizations with an AR view between the devices.

ABSTRACT

We present MARVIS, a conceptual framework that combines mobile devices and head-mounted Augmented Reality (AR) for visual data analysis. We propose novel concepts and techniques addressing visualization-specific challenges. By showing additional 2D and 3D information around and above displays, we extend their limited screen space. AR views between displays as well as linking and brushing are also supported, making relationships between separated visualizations plausible. We introduce the design process and rationale for our techniques. To validate MARVIS' concepts and show their versatility and widespread applicability, we describe six implemented example use cases. Finally, we discuss insights from

expert hands-on reviews. As a result, we contribute to a better understanding of how the combination of one or more mobile devices with AR can benefit visual data analysis. By exploring this new type of visualization environment, we hope to provide a foundation and inspiration for future mobile data visualizations.

KEYWORDS

data visualization, mobile devices, head-mounted augmented reality, cross-device interaction, data analysis, mobile data visualization, augmented displays, immersive analytics

ACM Reference Format:

Ricardo Langner, Marc Satkowski, Wolfgang Büschel, and Raimund Dachzelt. 2021. MARVIS: Combining Mobile Devices and Augmented Reality for Visual Data Analysis. In *CHI Conference on Human Factors in Computing Systems (CHI '21)*, May 8–13, 2021, Yokohama, Japan. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3411764.3445593>

1 INTRODUCTION

Exploration and interpretation of data is part of many daily activities and well supported by visualizations. Due to technical and research advances, this is no longer limited to traditional desktop setups [33, 66, 91]. As part of this, *Mobile Data Visualization* (MobileVis, e.g., [24, 65]) seems to be particularly significant because

*Also with, Cluster of Excellence Physics of Life, Technische Universität Dresden.

[†]Also with, Centre for Tactile Internet with Human-in-the-Loop (CeTI), Technische Universität Dresden.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

CHI '21, May 8–13, 2021, Yokohama, Japan

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8096-6/21/05.

<https://doi.org/10.1145/3411764.3445593>

nowadays, many people use their phones or tablets as their primary computer. Many MobileVis examples aim at casual use (cf. Personal Information Visualization [49]) and are available to a wide audience [93, 97]. Even though rarely explored, the usage of single or multiple mobile devices also for professional visual data analysis is promising, as the work by Plank et al. [85] or Langner et al. [64] demonstrate. Thereby, visualizations on mobile devices can be literally picked up and explored, and multiple devices can be physically moved and arranged, which seems especially useful in multi-user settings.

Additionally, *Augmented Reality* (AR) represents an exciting alternative to common desktop workplaces. Though not arrived in the mass market yet, Augmented Reality head-mounted displays (AR HMDs), such as the Microsoft HoloLens¹, show the interest and intent of major technology companies to develop AR products. They already provide a high visual quality and allow for a seamless integration into our real world. This makes them also interesting for visualization researchers. Using AR for data analysis gained much research interest in the last years, which in a broader sense belongs to the research area of Immersive Analytics [31].

In this paper, we explore the fundamentals and investigate the usefulness of combining mobile devices and head-mounted AR for data visualization. We are convinced that bringing these different technologies together is promising since their combination could compensate for the weaknesses of each component. Smartphones or tablets are omnipresent, have traditional, fixed-sized but high-resolution displays, are tangible and thus provide haptic feedback, and allow for precise interaction via touch or pen input. In contrast, AR HMDs provide highly flexible and virtually unlimited screen space, enhanced 3D perception through stereoscopy, and a truly personal display without a fixed form-factor. While the underlying idea for this work—extending mobile devices with 2D or 3D information shown in AR—is similar to recent approaches that augment desktop or large display workplaces [87, 89, 115] (also called Augmented Displays [87]), the key difference is our focus on mobile devices.

We envision scenarios in which users—each wearing ordinary-looking glasses with AR functionality—can sit with their mobile device(s) to perform visual data analysis tasks. Phones or tablets show visualization views to be interacted with. Devices can be held in the hand, moved around on a table, or physically arranged with each other easily. AR, on the other hand, supports and extends mobile device views by providing additional 2D and 3D information *around, above, and even between* them.

Specifically, our main contribution is a proposed conceptual framework, called MARVIS (Mobile Devices and Augmented Reality for Data Visualization). This framework comprises exemplary visualization approaches and general principles that take advantage of using AR to extend mobile 2D displays. We show how such a setting can be used to, for example, provide additional context information while navigating information spaces on mobile devices, integrate 2D and 3D stereoscopic visualizations, support cross-device linking and brushing, or make relationships between separated views intelligible.

MARVIS is divided into four core parts and contributions: First, we systematically describe the *design process and design rationale* (section 3), which is based on seven expert interviews. Second, we present novel *visualization and interaction concepts* (section 4) that focus on using either a single mobile device or a combination of two or more mobile devices. Third, we validated our concepts by implementing a *prototype application consisting of six example use cases* (section 5) that adapt different visualization techniques and data sets, thereby demonstrating the versatility and widespread applicability of the MARVIS approach. Finally, we report on feedback collected from seven experts in *hands-on experience sessions* and reflect on *lessons learned* (section 6). With our exploratory and constructive work (cf. constructive research [83]) we aim at providing inspiration and an informed basis for the development of future visualization applications in such novel augmented visualization environments.

2 RELATED WORK

The work we present builds on existing research at the intersection of the fields of information visualization (InfoVis) and human-computer interaction (HCI). In the following, we review prior work in the relevant areas of mobile devices for InfoVis, AR for InfoVis, as well as the combination of commodity displays with AR.

2.1 Mobile Devices for InfoVis

Simple visualization techniques have already been in focus for devices such as PDAs [7, 15, 58] where the limited performance and input was a major constraint. Today's bandwidth of mobile devices ranges from smartphones and tablets to wearables such as smartwatches and fitness trackers, which all are considered in the area of MobileVis [24, 65]. While a basic approach is to transfer and adapt existing visualization techniques (e.g., [109]), a majority of the work on InfoVis beyond the desktop [33, 66, 91] builds on the use of modern mobile devices. Since the most prominent input modality of those devices is touch, much work focused on deploying multi-touch interactions to different visualization techniques [6, 30, 94, 96]. One key difference is that instead of using traditional user interface widgets, such as menus, functionality can be accessed by directly touching data items, axes, legends, or other elements of a visualization. Furthermore, touch is often combined with pen input. For example, Jo et al. [57] developed pen and touch gestures for a data exploration application on tablets and found this type of interaction beneficial for novices. Spatial input represents another interesting form of input, because it can outperform classic interactions for zooming and panning [92, 103].

Responsive data visualization [1, 2] aims at developing strategies for adapting layout, information density, and interaction to specific devices. This is especially crucial considering the limited display space of mobile devices. Studies on efficiency and performance of visualizations on mobile devices present a vital foundation for these adaptations [9, 12, 13]. Due to their connectivity, mobile devices can be easily combined with other displays. Approaches that make use of multiple mobile devices for data visualization are most relevant for our work. While Woźniak et al. [118] designed specifically for two mobile devices, Langner et al. [63, 64] considered multiple devices used in combination for visual data analysis. However, a

¹Microsoft HoloLens, <https://www.microsoft.com/en-us/hololens/>

study by Plank et al. [85] revealed that users seem to be “hesitant to use multiple tablets in parallel” and often start with just one device. Zagermann et al. [122] showed that, on the other hand, extensive training can lead to high device utilization.

Overall, this prior work shows the potential of mobile devices for data visualization. There still are open challenges, such as designing effective interactions for small screens and complex visualizations, further developing visual techniques to address the limited screen space, and integrating mobile devices into existing device ecologies [26] in a useful way.

2.2 Augmented Reality for InfoVis

With new immersive display technologies, such as Virtual and Augmented Reality head-mounted displays, there has been an increased interest in making these technologies available for data analysts and to study their usefulness in InfoVis. For example, Bach et al. [5] examined the effectiveness of AR HMDs compared to a desktop setting and tablet-based AR for typical tasks on 3D scatterplots. Recently, Reipschläger et al. [89] designed several techniques for augmenting and enhancing visualizations shown on a large vertical display with AR. Similar, Mahmood et al. [70] suggest to enhance data analysis with vertical displays by utilizing the surrounding space for showing coordinated AR visualizations. Others have studied perceptual advantages and challenges for immersive visualizations (e.g., [69, 117]) or looked into the design of individual AR visualizations (e.g., [19]). More specifically, there has been research on mobile Augmented Reality for InfoVis. This includes the work by Büschel et al. [18], investigating interaction for AR visualizations with spatially-tracked tablets, and work done by Drochters et al. [29], in which the authors used tablets placed on a table to anchor nodes of a 3D node-link diagram that could be explored with mobile AR.

Currently, these trends have culminated in the formation of the field of Immersive Analytics (IA) [23, 36, 74]. IA combines the use of novel, immersive display technologies with spatial or embodied interaction [16] to facilitate data analysis. Skarbez et al. [102] define IA “as the science of analytical reasoning facilitated by immersive human-computer interfaces.” A possible advantage of IA is the use of the whole body for spatial interaction in an immersive environment [38, 59].

This is highly relevant to our research: The spatial configuration of mobile devices and the relations between them are central to our concepts and lend themselves to scenarios where users can frequently change their viewing angles and positions, possibly benefiting from IA concepts. IA also suggests itself for collaborative settings [8] and can be used to show user-specific information via HMDs [106]. As such, we believe that it fits our use case of augmenting multiple mobile devices.

The potential benefits of using physical surfaces as references for immersive 3D visualizations have been studied by Filho et al. [35], using the example of a space-time cube visualization for trajectory data, rendered above a 2D map in a virtual desk environment. Similarly, research on using mobile devices [17] and custom tangible controllers [28] for the exploration of immersive visualizations suggests that combining mobile devices and AR visualizations is advantageous.

2.3 Integration of AR and other Displays

Early concepts of virtually extending workspaces around conventional displays date back to the early nineties [34, 90]. Since then there was a lot of interest around this idea. Therefore, several use cases have been suggested, such as drawing/sketching [3], CAD or 3D modeling [76, 87], and most importantly data visualization [21, 22, 50, 53, 70, 89, 99, 106, 115]. This idea has also been applied to different types of displays, including desktops [76, 87, 115], tabletops [20], and larger vertical displays [53, 70, 89, 106]. Cavallo et al. [21, 22] even investigate how to integrate AR within a multi-display environment.

Interestingly, there is also work on augmenting mobile devices. Grubert et al.’s MultiFi [42] combines AR with displays on and around the body. They found that, for instance, combining a smartwatch and AR HMD can improve task performance but at the costs of higher workload and reduced usability. Sereno et al. [99] envision that tablets can be used for view manipulation and annotation during the exploration of a volumetric AR visualization. Wu et al. [119] envisioned gestural interactions to pull digital content from a tablet or other display into AR. Normand and McGuffin’s VESAD [80] envisioned to use the co-planar space around a smartphone for, e.g., extended, alternative, or more detailed views of phone content. Very close to the idea behind MARVIS is Hubenschmid et al.’s ARts [50], in which they describe a system for visualizing an AR 3D parallel coordinates plot above multiple tablets. However, the tablets are only used as a points of reference to specify the location of plot axes, thus the displays do not show visualizations.

Overall, it appears that the integration of AR and other displays has great potential and research on this topic is very active. However, we think that besides Hubenschmid et al.’s initial work [50], the combination of AR and multiple mobile devices is largely underrepresented.

3 MARVIS: PROCESS & DESIGN

The motive of this section is to explain how MARVIS’ concepts (section 4) evolved. We describe our development process (see 3.1), report on seven expert interviews (see 3.2), and outline our design decisions (see 3.3).

3.1 Development Process

Our iterative development process can be described as a mixture of ideation, design thinking, and task-oriented design. The process consisted of the following steps: **(1-Ideas)** We started by drafting initial ideas based on our experience with the development of interactive visualization tools. Additionally, we invited and interviewed seven visualization and HCI researchers regarding their visions and ideas. **(2-Scope)** We then discussed results of the first step and decided on fundamental design choices, such as AR having a supportive role, to define the scope of our framework. **(3-Concepts)** We developed a set of novel concepts for the use of mobile devices and AR HMDs for data visualization (section 4). These concepts use and adapt existing visualization mechanisms, such as overview+detail, small multiples, and linking and brushing. They further support different visualization tasks and activities [14, 121], such as identifying, relating, aggregating, and comparing. **(4-Prototype)** We implemented

	Grouped Comments	Interviewed Experts 1-7
Use AR to show additional information based on what's on mobiles.	■ ■ ■ ■ ■ ■ ■ ■	■ ■ ■ ■ ■ ■ ■ ■
Use AR to enable multi-user functionality (e.g., personal views).	■ ■ ■ ■ ■ ■ ■ ■	■ ■ ■ ■ ■ ■ ■ ■
Focus on mobile devices, because they are established already.	■ ■ ■ ■ ■ ■ ■ ■	■ ■ ■ ■ ■ ■ ■ ■
Align AR information flat on one layer (e.g., at table or mobile).	■ ■ ■ ■ ■ ■ ■ ■	■ ■ ■ ■ ■ ■ ■ ■
Use small devices such as smartphones as controllers and for menus.	■ ■ ■ ■ ■ ■ ■ ■	■ ■ ■ ■ ■ ■ ■ ■
Use AR to show relations/connections between views/devices.	■ ■ ■ ■ ■ ■ ■ ■	■ ■ ■ ■ ■ ■ ■ ■
Use spatial device movement for interaction (e.g., slicing, peephole).	■ ■ ■ ■ ■ ■ ■ ■	■ ■ ■ ■ ■ ■ ■ ■
Take advantage of stereoscopy and use AR to show 3D Visualizations.	■ ■ ■ ■ ■ ■ ■ ■	■ ■ ■ ■ ■ ■ ■ ■

Table 1: Comments made during the expert interviews in the design phase. This table lists 8 of overall 16 grouped comments with a minimum of 4 mentions, sorted by frequency.

six exemplary and interactive use cases (section 5) to aid the discussion and revision of our concepts. Each use case adapts a real-world data set and demonstrates a particular concept, such as seamless coupling of a 2D and 3D visualization. **(5-Feedback)** Finally, we invited the same seven experts again and conducted expert reviews using the developed prototype.

3.2 Expert Interviews in Design Phase

Besides developing our own initial ideas, we invited seven other researchers from our institute for individual expert interviews. All experts (four PhD-level, three PhD student) were HCI and visualization researchers and have developed interactive information systems themselves in the past. They all had good knowledge in using data visualizations: on a 5-point scale (no, low, moderate, advanced, expert-level), four rated their experience regarding the use of visualizations as moderate, two as advanced and one as expert-level.² The seven semi-structured interviews lasted approx. 45 minutes and were conducted by one investigator. The goal was to collect specific ideas regarding the use of mobile devices and AR HMDs for data visualization, to identify promising general principles, and to discuss challenges. As impulses for discussion, we provided a list of typical visualization tasks [14, 121] and four application examples as image cards showing existing visualization techniques (i.e., geospatial data visualization, temporal data visualization, network visualization, and multiple coordinated views). Interviews comprised elements of brainstorming and paper prototyping. Besides pen and paper, we printed silhouettes of different mobile devices, which were used by the experts to quickly sketch and explain ideas. We videotaped each session for documentation and took notes for later analysis. Two of the authors discussed, labeled, and grouped these notes. The general themes that emerged from a total of 92 unique notes are: the role of mobile devices and AR, which information to show, how and where to display information in AR, handling of mobile devices, and the interaction with information. The most frequent comments are listed in Table 1. We finally considered these comments for the decision on design choices, design of our concepts, and implementation of the prototype.

3.3 Design Choices & Rationales

The scope and characteristics of MARVIS are defined by several design decisions. We made those based on our own experience,

findings of prior work, envisioned usage and technical aspects, comments from the seven expert interviews, and visualization tasks that could be supported.

Mobile only. We follow an approach that is solely based on mobile devices and AR. We deliberately consider no other stationary displays. Our goal is to explore a device combination that still ensures mobility, allows to be used in various locations, and provides a form of independence.

Seated usage. We focus on situations where users are seated. We already prescribed this for the aforementioned interviews because sitting allows for more concentrated and longer work, which is particularly relevant for professional and analytical settings. When sitting at a table, we assume interactions mainly take place at a table level and one or more devices are moved on or right above a table. Otherwise, when sitting somewhere without a work surface, it is more likely that only a single mobile device is held in hand or resting on the leg or knee. Even though a standing posture could be an alternative, we think that sitting is more comfortable, thus more suited for complex or detailed data analysis.

Focus on mobile device. We suggest to use the mobile device as the primary device. Data analysis is possible with mobile devices only but combining them with AR HMDs will provide advantages. The main reason is because, contrary to AR HMDs, mobile devices are already widely used in everyday life, they are ubiquitous. There are two alternatives to this strategy. First, mobile devices, AR HMDs, and corresponding software application could exist in parallel. This, however, would leave the choice of when to use which tool completely to the user. Moreover, some functions are more practical on mobile devices (e.g., presenting visualizations with many fine details), others for AR (e.g., looking at and exploring a 3D visualization). Second, a focus on AR could make mobile devices expendable. Reasons against this are a currently lower display quality of AR HMDs, social acceptance of gestural interaction [77], and precise input capabilities of mobile devices. We also argue that working with a physical artifact, like a tablet, is a more explicit form of interaction: Mobile devices can be touched and picked up (tangibility), which is immediately visible for the people around and thus preferable over pure virtual interaction.

AR supports and provides context. AR is not used as a substitute but rather as an addition and support, which addresses expert comments such as “display additional data attributes in AR” or “use AR to show connections and relationships”. This is also similar to Wang et al.’s argumentation [115] and leads to the following two principles: (i) Primary output (visualizations) on mobile devices. AR HMDs provide contextual information and additional guidance around, above, and between devices. (ii) To prevent overloading the user’s field of view, we suggest, similar to Grubert et al. [42], AR content to be modest and mostly less detailed. In cases where AR shows a fully and detailed visualization, it either is not the only or main visualization or it directly benefits from being shown in 3D (stereoscopy). While this choice also reflects capabilities of current AR HMDs, it can be reviewed on further significant technical developments.

Touch over mid-air interaction. We propose to realize most interaction with mobile devices. By using touch or pen input or device movement (cf. micro-mobility [73]), we make use of fine motor

²For further details see supplemental material, also available on our project website <https://imld.de/marvis/>

MARVIS: Mobile Devices & Augmented Reality for Data Analysis

CHI '21, May 8–13, 2021, Yokohama, Japan




Figure 2: Use of AR for adapting different *Overview+Detail* and *Focus+Context* techniques: (a) A typical map overview; (b) Marginal histograms around the mobile device; (c) 3D visualization of a Matrix Cube; (d) Navigation support by an off-screen coordinate origin; (e) Zoomed in bar chart with fisheye-style continuation; (f) Mobile device as a detailed lens into a larger map. Note that information shown in AR is colored in purple, that of mobile devices in blue.

skills and sensitivity of our own hands. Moreover, touch capabilities of current mobile devices better suit precision requirements for engaging with often visually detailed visualizations. However, direct interaction with AR content is clearly needed as well. Modern AR HMDs allow for mid-air gestures, such as air-tap or poke. Alternatively, around-device interaction [62] could be useful, because AR content is likely to be displayed in the immediate periphery of a mobile device.

Spatial awareness of devices. We assume positions and orientations of mobile devices and AR HMDs are known and usable for visualization applications. While such a tracking is of course technically challenging, several examples show the technical feasibility with external [64, 86] or internal [41, 48, 55, 75] tracking approaches. MARVIS’ concepts (section 4) would be possible without spatial awareness but in all cases with considerable extra efforts for users. For example, AR content could then only be positioned manually and arrangements of two or more mobile devices would have to be specified by additional touch gestures [63, 78, 82].

4 MARVIS: CONCEPTS

Based on the design choices described above, we propose and systematically describe several new *visualization and interaction concepts*. We present exemplary techniques that are grounded in common visualization principles and tasks, including overview+detail, identification and connection of data items, selection through linking and brushing, data filtering, and visual comparison. By using the combination of mobile devices and AR as an enabling technology, we also adapt and enhance existing visualization techniques. The section is organized into two categories: a single mobile device with AR (see 4.1) and two or more mobile devices with AR (see 4.2).

4.1 Single Mobile Device with AR

The following concepts seek to highlight the benefit of using a single mobile device with AR for working with overview+detail, focus+context, alternative visualizations, visualization user interface components, and superimposed 3D visualizations. Generally, we apply two core ideas for an AR extension of a device: using the space **around** and the 3D space **above**.

4.1.1 Overview+Detail. Overview+detail is characterized by “the simultaneous display of both an overview and detailed view [...]

each in a distinct presentation space” [27]. For mobile devices, however, it is especially difficult to find and use precious screen space for both views. Additionally, a small display makes it harder to keep an overview, e.g., during navigation. We therefore propose to supplement a detailed visualization on a mobile device with useful **virtual overviews shown around the device in AR**. A possible AR overview is a mini map for a geographic map with boundaries indicating the detail view’s location. Figure 2a also shows that in contrast to common practice, this AR overview does not have to be a small world-in-miniature component compared to the mobile device. Another example are marginal histograms for scatterplots or metric views for node-link diagrams [84]. Shown alongside a mobile device in AR, those histograms could support data exploration and discovery by visualizing number of link crossings or link coverage (Figure 2b). An example for more complex data is the 3D visualization of a Matrix Cube [4] shown in AR (Figure 2c). The 3D overview and the highlight of the cube layer currently displayed on the mobile device support understanding and navigation.

These examples illustrate the general idea of showing a distinct overview visualization in AR to help users to explore data, relate views, and make better use of a mobile device’s screen. This approach can be used for many other existing overview+detail techniques (e.g., border visualizations with node proxies for UML diagramming [37], TimeSpan’s overview area [68], or the HEDA component [67]). Moreover, due to stereoscopy AR HMDs can help in cases in which it would be useful to actually display a three-dimensional representation of data for overview (e.g., 3D height maps or 3D displays of medical image data).

4.1.2 Focus+Context & Seamless Visualization Extension. In contrast to overview+detail, the principle behind focus+context is to integrate both the focus and context visualization “into a single display where all parts are concurrently visible” [27]. However, for mobile devices the extent of a focus visualization is critical. It either consumes too much display space and leaves no room for a context visualization, or its scale causes issues of readability and visibility.

Therefore, we suggest to use AR for a **direct and seamless display extension**: The mobile device shows the focus region and AR provides the context visualization. This way, for example, a geographical map on a mobile device can be seamlessly continued in the surroundings of the device. The same applies to off-screen coordinate origins or axes of a scatterplot (Figure 2d). Figure 2e

CHI '21, May 8–13, 2021, Yokohama, Japan

Langner et al.




Figure 3: Use of AR for Alternative Visualization Views, Separated Visualization User Interface Components, and Superimposed 3D Visualizations: (a) SPLOM shows alternative scatterplots configurations; (b) Distributed views of a dashboard; (c) Tilted AR views; (d) Off-loaded legend and menus; (e) Continuous 3D track above a map; (f) 3D wall visualization [112] aligned to a map.

shows a variant of a fisheye-style bar chart [39] with parts of the distorted context area (i.e., thin bars) displayed left and right of the mobile device in AR. As mentioned during the expert interviews, another interesting opportunity for maps or other visualizations of zoomable information spaces is to adapt a *Tangible Views* [104] or *Peephole Display* [120] approach. Visualizations are positioned and fixed in relation to real world objects, such as a table [18]. A handheld mobile device can be used as a window into the data space by moving and rotating it within and around that data (Figure 2f and 6a). In order to support navigation and exploration, the situated information space can be made visible in AR.

Technically, our focus+context approach is not a direct integration of focus and context into a single view. However, we think that AR is very capable of imitating an adequate visual continuity. The perceived outcome depends on the technical abilities of the AR HMD but also on the involved visualization techniques. Techniques, such as bar charts, showing discrete visual elements might tolerate a certain discontinuity, whereas node-link diagrams or maps are more sensitive to visual gaps between AR and a mobile device. However, due to the limited display space, a seamless display extension is relevant for many mobile visualization applications.

4.1.3 Alternative Visualization Views. Many visualization systems allow to choose between alternative views to support and guide data exploration. These alternatives typically result from different visual representations and encodings for the data (cf. level 2 systems [81]) as well as the partitioning of information spaces. However, presenting several alternative views on a mobile device requires a thumbnail-like presentation. Since this influences recognizability, readability, and interaction, Sadana and Stasko [95] investigated multiple coordinated views for tablets using a maximum of three views.

Similar to our overview+detail concept and the *Small Multiples, Large Singles* approach [114], we propose to **show other available views or alternative configurations of a current visualization in AR**: The large single view is displayed on the mobile device, alternatives are arranged around the device. For example, a scatterplot matrix (SPLOM) or trellis plot representing alternative views with different configurations can be shown in AR, while the mobile device displays one specific cell or view (Figure 3a and 5a-c). In case of a dashboard application, primary visualizations can be displayed on the mobile device directly, other available views are

shown around the device instead (Figure 3b). Due to AR’s flexibility, alternative visualizations can also be shown in full size, i.e., adapted to the size of a mobile device. This could support readability and comparison between AR and a mobile device. Figure 3c shows that AR views can also be tilted towards the user or along the border of a device, achieving settings similar to Marquardt et al.’s *SurfaceConstellations* [72].

4.1.4 Separated Visualization User Interface Components. Similar to our concepts for overview+detail and alternative visualization views, we suggest to **offload distinct user interface components of visualizations to AR**. Figure 3d illustrates the idea. This concept is particularly relevant because many visualization systems come with space-consuming components, such as menus, toolbars, controls, and preference panes, but also legends or color scales. Although specific to desktop-sized screens and 3D modeling use cases, DesignAR [87] nicely showcases how such user interface components can be attached around a conventional display using AR. In a follow-up work Reipschläger et al. [89] also describe how components of a visualization, such as axes, links, selections, or legends, can be extended with AR content. However, to specifically address limited screen space of mobile devices, we suggest to first offload separate components to AR and thus release a part of the display area.

4.1.5 Superimposed 3D Visualizations. Until now we have described how AR can augment the space around a mobile device for 2D visualizations. Another essential idea is to use the space above a mobile device to visualize further data by making use of the third dimension. This way AR provides additional, embedded data. At the same time, a number of existing techniques also involve some form of 3D graphics. However, due to missing 3D display capabilities, such 3D visualizations or 3D views are always reduced to a single two-dimensional projection when shown on a mobile device. In addition, effective interactions for 3D visualizations represent a particularly important and ongoing challenge [18].

In this context, we suggest to use an AR HMD to superimpose and couple a 2D visualization on a mobile device with **three-dimensional, hologram-like visualizations in AR**. A simple use case for that are 3D icons or 3D glyphs [108, 113] (Figure 7a). Moreover, techniques using the space-time cube metaphor (e.g., [61]) as an underlying model would also benefit from our setting. The mobile device represents the base of the space-time cube, for example a geographical map, AR is used to visualize the temporal aspects of



Figure 4: Use of AR for Relations Between Visualizations, Combination of Visualizations, and Multi-User Support: (a) Linking and brushing supported by curved AR connections; (b) Ribbons between devices indicate the relative proportions; (c) Icon meta-visualizations reveal view relations; (d) AR bar chart summarizes calculated differences between views; (e) Merging two views in AR; (f) Personal and shared areas for collaborative activities.

data. For instance, AR can show continuous 3D tracks [61] for which time is mapped to the height above the mobile device (Figure 3e and 7b). Figure 3f and 6c illustrate how a 3D wall visualization [112] is aligned to a map displayed on a mobile device. Besides such InfoVis techniques, 3D scientific visualizations could also benefit from this concept. Continuing the example of the map again, AR enables to display wind, weather, and ocean conditions in 3D (above a reference map shown on the mobile device).

We acknowledge the ongoing controversial discourse on the use of 3D for data visualization (e.g., [11, 100, 107]), which is why we argue to limit the use of 3D to when it either is necessary (e.g., due to the data) or has clear benefits for the users. We think 3D visualizations can be made useful by leveraging the combination of mobile devices and AR because AR HMDs provide an enhanced illusion of depth through stereoscopy. They also allow inspecting data from different angles simply by head or device movement, or by literally picking up and rotating a superimposed 3D visualization.

4.2 Two and More Mobile Devices with AR

In the previous section, we described how AR extensions can be useful for a single mobile device. In the following, we explore concepts that illustrate how AR can support data visualization with two or more devices. This is relevant for both single users and multi-user scenarios. We sometimes carry more than one mobile device with us already today; and when two or more people meet, it is likely that several mobile devices will be brought along [64, 85] and used together for collaborative visualization [51]. Moreover, in the future, when tablets become even more lightweight and even less expensive, plenty of such devices might be provided and ready to be used at places such as offices or meeting rooms (borrowing from Weiser’s idea of “pads” [116]). The section comprises concepts which show how relations between visualizations can be extended with AR, feature the arrangement and combination of visualizations, and provide AR support for multiple users. In addition to using the space **around** and **above**, here we also exploit the space **between** multiple devices.

4.2.1 Relation Between Visualizations. In order to make multiple mobile devices useful for data exploration, a user should be allowed to connect and distribute visualizations across those devices [64, 85]. Visualizations can relate to each other in many ways, for instance: In an overview+detail setting the overview does not only show and continuously update the position and size of the detail view,

but allows to control it ‘remotely’. Also, a selection of data items in one view could be used as input for a filter operation applied to a second view. Another relation is the identity of data items across multiple visualizations which often is visible when using linking and brushing. However, the awareness and understanding of these relations can be challenging due to the spatial separation and number of views.

We therefore propose to *provide visual feedback in AR about view relations in the space between and above mobile devices*. As a simple example, Figure 4a shows that linking and brushing can be supported by displaying connections between data items in AR (similar to AR brushing and linking [89], integrated views [54], or spatially-aware visual links [25]). Furthermore, connections can also be drawn in AR between an overview visualization and all its detail views. Connections between views can be represented in many different ways. A minimalistic option for simple cases is to show lines between devices, as directional or non-directional edges. In more refined variants, such lines can be replaced by paths or 3D arcs addressing overlapping and intersections of connections. In addition, the visual links encode further information about the relation itself (Figure 5d). This can be useful to support the understanding of, e.g., how much of an information space is covered in a detail view. Besides bringing existing techniques such as Tobiasz et al.’s icon meta-visualizations [110] to AR (Figure 4c), another approach is to adapt principles of flow maps, sankey diagrams, or parallel sets by displaying connection ribbons of different widths between devices (Figure 4b).

In general, the use of AR provides a large virtual canvas allowing to visualize relations without competing for the limited display space of a mobile device. We think that making relations accessible in this extended way could help users to ascertain, better understand, and even control the state of those connections. The relevance of this increases when more views are involved, more mobile devices are used, and visualizations are spatially more separated.

4.2.2 Combination of Visualizations. So far, the relation between views focused on interaction aspects such as overview+detail or linking and brushing but it can go beyond that and be more complex. We therefore suggest using AR to *support the combination of visualizations*. From a high-level perspective, combinations refer to methods allowing to derive or produce new visualizations based on components of existing visualizations. For example, users might be interested in a visualization showing where two data

subsets overlap or differ. Figure 4d illustrates how a bar chart in AR could summarize differences between the two visualizations. This refers to *data-level combinations* involving logical operators and various forms of aggregation (e.g., conjunction, disjunction, maxima, minima, averages, sums, counts). Alternatively, when data is represented in different visualizations, users might want to better connect this data by bringing them into the same visual space. For example, the difference of two visualizations could be shown in AR by ‘merging’ (Figure 4e) which has some relation to explicit encodings [40] for visual comparison. An AR visualization could also be moved and positioned above another visualization shown on a mobile device to archive the same result. Such *view-level combinations* directly relate to superimposed, overloaded, or nested views, which represent different forms of composite visualizations views [54].

Similar to scatterplot matrices (SPLOM), Figure 7c illustrates the crossing of data dimensions by combining devices via a special device arrangement. This directly builds on Langner et al.’s VisTiles [64] and is motivated by people’s behavior in placing two objects of interest close to each other [52, 111]. Despite diverse possibilities of such combinations, the advantage of AR is that the additional views can again be displayed in any size and orientation, in contrast to the rather unlikely use of a third or even fourth mobile device.

Moreover, AR can improve the discoverability of combined visualizations, helping to introduce that function to users. Using a device’s location and its inertial sensors, the system can try to detect when users intent to combine visualizations. Visual feedback on possible combinations can then be provided. On one hand, it can show possible target positions (“user, you can move your device here or there”). On the other hand, a lightweight preview of what would be displayed could be shown as well (“user, this very combination would be possible here”). That way users can be assisted and provided with a proper feed forward mechanism which supports “awareness of device presence and connectivity” [71].

4.2.3 Multi-User Support. As already mentioned, meetings with two, three, or more people allow for collaborative use of multiple devices. Since we think AR HMDs have great potential for providing visual guidance, similar to visual feedback on possible visualization combinations, we suggest to make use of AR to **support multi-user situations and collaborative activities**. For example, we know that there are different types of territories in collaborative tabletop workspaces [98], which can be visualized in AR (Figure 4f). AR can also enable individual and user-specific information displays during collaboration. When devices are rotated differently, the alignment of text labels or menus can be corrected in AR. Likewise, visualizations of distant devices can be shown in AR in a larger and tilted way. To support phases of parallel work and avoid distraction, linking and brushing can be adapted in a way that users see the result of their interaction in AR only, that is highlighting of selected items is not displayed on mobile devices. We think that there are many of such opportunities for multi-user scenarios, but fully exploring such specialized collaboration techniques would go beyond the scope of this work.

5 MARVIS: PROTOTYPE REALIZATION

In the previous section, we presented new visualization and interaction concepts using the augmentation of the space **around**, **above**, and **between** mobile devices. The goals for this section are to demonstrate the technical feasibility and further highlight the utility of MARVIS. We implemented a prototype application³ consisting of six example use cases. Each example demonstrates parts of our concepts, but also extends and enhances already existing visualization techniques. This systematic approach exemplifies how our proposed concepts can be realized, considering different data and visualization techniques and different user interactions. While the prototype already serves the validation of our concepts, we additionally presented it to seven experts for a hands-on experience (section 6). We first describe our technical implementation and then detail each example use case.

5.1 Technical Implementation & Setup

All six example use cases are implemented in one interactive prototype application. Even though they technically can run in parallel (showcased in Figure 1a), each example loads different data and thus represents another scenario. The application is realized as a client-server model comprising three major components: an AR client, a mobile device client, and an application server.

The *AR client* runs on a Microsoft HoloLens 2, which has a diagonal field of view of 52° and, according to official documentation⁴, provides a resolution of ca. 2.500 light points per radian. Due to built-in hand tracking, mid-air gestures, such as tap, poke, and grasp, allow direct interaction with holograms. The application and all its visualizations are implemented with the Unity game engine. Although interaction with AR content is handled in the AR client, the application server receives notifications about events, such as selections.

The *mobile device client* is a web application written in JavaScript. For the user interface, we use different libraries, e.g., D3.js [10] for visualizations, Materialize for controls such as buttons or sliders, and Hammer.js for touch gesture recognition⁵. We used two Samsung Galaxy Tab S3 tablets (approx. 9.7 inch, 2.048 × 1.536 px, 430 g), but our setup supports any device with a modern web browser to connect and run the application. As with the AR client, interactions with visualizations are handled locally but are also forwarded to the server as events. Upon establishing a connection with the server, the client displays a menu allowing to select and start one of the example use cases.

The *application server* manages the global application state, all connections to client devices, loading of data sets (stored in CSV files), data requests, and spatial positions of devices. It is programmed using Node.js and Express⁶ and uses JSON-RPC on top of WebSockets for client-server communication. We use the motion capturing system OptiTrack⁷ for tracking locations of all client devices. For that, both tablets and the HoloLens are equipped with IR-reflective markers, which are captured by five infrared cameras

³A supplementary video is available on our project website: <https://imld.de/marvis/>

⁴HoloLens 2 hardware <https://docs.microsoft.com/en-us/hololens/hololens2-hardware>

⁵D3.js <https://d3js.org>, Materialize <https://materializecss.com>, Hammer.js <https://hammerjs.github.io>

⁶Node.js <https://nodejs.org/>, Express <https://expressjs.com/>

⁷OptiTrack <https://optitrack.com>




Figure 5: Showing additional visualizations around and between devices: (a) The mobile device is placed inside the SPLOM showing one cell, while AR provides an iconic representation of all other cells; (b) The same SPLOM with a different position of the mobile device; (c) The mobile device shows one cell of the SPLOM with four neighboring tablet-sized cells visible in AR; (d) AR link between two devices; (e) The link shows additional information about the device connection (number of selected and filtered nodes).

mounted to the ceiling. The server receives tracking data from OptiTrack, transforms positions into the coordinate system of the HoloLens, and sends position updates of mobile devices to the AR client.

5.2 Implemented Example Use Cases

5.2.1 Scatterplot Matrix Navigation. Scatterplot matrices (SPLOMs) are widely used and represent a convenient and relatively easy-to-understand multi-dimensional visualization technique. Their main characteristic is that, resulting from all possible mappings of data dimensions to x - y axes, scatterplots are ordered into respective rows and columns. With regard to overview+detail, a SPLOM is often shown next to a larger scatterplot, allowing for a structured exploration of multivariate data. On a mobile device, however, it is hard to show both the SPLOM and the detailed scatterplot at the same time while still ensuring readability and interactivity.

We implemented four different SPLOM designs, each with its own advantages and disadvantages, to explore how AR HMDs can provide support for navigating within a SPLOM. As data we use properties of cars⁸ with a total of eight dimensions. It applies to all designs that the mobile device shows one full-size cell of the SPLOM (beginning at the top left). The general idea for AR is to provide alternative views (see 4.1.3), while the mobile device and the shown cell are part of the SPLOM. The first design shows the entire 8×8 SPLOM in AR, with each cell equally sized as the mobile device view. The second design uses thumbnails instead of full-size cells but still shows the entire SPLOM (Figure 5a). A third design reduces the thumbnail SPLOM to showing only the row and column specified by the mobile device position. The fourth design uses full-size cells again but shows only a single neighbor at each side of the device (Figure 5c). For interaction, users can move orthogonally within the SPLOM (Figure 5a-b) by performing a swipe gesture parallel to one of the scatterplot axes. Touching and encircling data items on the mobile device also allows for linking and brushing, with affected plots shown in AR. Displaying these neighbors as thumbnails of the actual scatterplots (Figure 5b) makes it easier to show an entire SPLOM using reasonable space. The thumbnail size

helps to reduce visual clutter of the user's view. Another display opportunity is a tilted preview attached to a device edge.

Overall, this use case shows how AR can help to work with multiple visualizations. SPLOM navigation also showcases how to compensate for limited mobile display space, thus enabling a structured exploration process that is rather known from visualization solutions for larger, desktop-sized screens. The most relevant concepts for this example are *Alternative Visualization Views* (see 4.1.3) and, as the thumbnail designs also present alternative views in an overview-like manner, partly the *Overview+Detail* concept (see 4.1.1). We think this SPLOM example provides an inspiration to investigate whether other techniques can be adapted in a similar way (e.g., ScatterDice [32], matrix cubes [4], small multiples, dashboards).

5.2.2 Node-Link Diagram and Attribute Visualization. Node-link diagrams represent one of the most common visual representations of network data which can also be associated with additional attributes. In this example, we use Twitter data⁹ of 320 accounts with seven attributes (e.g., follows, followed by, number of tweets). Although various approaches and techniques exist for visualizing such multivariate networks [79], a simple strategy is to use separated views for the topology and attributes, such as node-link layouts plus a parallel coordinates plot [43] or matrix views [60, 105]. In this use case, we show a topology and an attribute view on two different mobile devices. To explore the topology, a user navigates the node-link diagram via drag and pinch-to-zoom. Selected nodes will also be highlighted in the attribute visualization, which is connected through linking and brushing. This linking can be extended to further consider the actual viewport of the node-link diagram to, for example, filter data items in the attribute visualization if these items are, due to zooming and panning, not visible in the node-link diagram anymore (cf. filter-by-viewport [64]). To further enrich the process and to support the user with information about the existence of a cross-device connectivity, additional AR visualizations, like a link (Figure 5d), can be shown. It allows visualizing the effective direction of an connection (in our prototype from the node-link diagram to the attribute visualization) and additional

⁸Source cars data: <https://github.com/RodolfoViana/exploratory-data-analysis-dataset-cars>

⁹Source Twitter data: <https://johnguerra.co/viz/influentials/eurovis2018/>

CHI '21, May 8–13, 2021, Yokohama, Japan

Langner et al.




Figure 6: Extending map visualizations: (a) AR is used to continue the map using less details (street-layer and data items); (b) More abstract visual style (highways, bodies of water, and data items), also illustrating details-on-demand on both the mobile device and in AR; (c) Adaptation of *Great Wall of Space-Time* [112], the wall in AR details unemployment data per selected country over time (bottom to top).

meta information, like the amount of nodes and edges currently visible or selected.

This use case mainly demonstrates a specific possibility of the *Relation Between Visualizations* concept (see 4.2.1). It shows a connection between devices and aims at making the system easier to understand by indicating the relation of a partial view to the overall data.

5.2.3 Map Navigation. Data often has a spatial context and can be visualized in combination with, for example, a geographical map. To show this, we visualize a subset of a real-world victim-based crime data set¹⁰. Each neighborhood of the city of Baltimore is represented by a circle on top of a map (Figure 6a). The size of a circle is mapped to the sum of individual crimes committed in this neighborhood. A single tap selects a circle, while a double tap opens details-on-demand as a tool tip (Figure 6b) which contains additional information about different crime types in that neighborhood. Touching and performing common pinch-to-zoom is used for navigation. Since the small display of a mobile device can make it hard to preserve an overview of the overall information space while zooming and panning a map, parts of the map are displayed around the mobile device in AR (see 4.1.2). For this AR map, we implemented two different abstracted visual styles. Both show data items and details-on-demand in AR but one uses a unicolored street layer (Figure 6a). The other is even more diminished and shows highways and water only but using the typical colors yellow and blue (Figure 6b).

Overall, this use case focuses on the concept *Focus+Context & Seamless Visualization Extension* (see 4.1.2). It illustrates how a display extension for a mobile device with an AR HMD can assist map navigation. We also showcase the use of different levels of detail.

5.2.4 Non-Planar Slices on a Map. Another common case is the visualization of spatio-temporal data. MARVIS' combination of a mobile device and AR HMD is a great opportunity to overlay 3D information, for example, as stacked layers, on a map. As mentioned before (see 4.1.5), the technique *Great Wall of Space-Time* [112] could be adapted directly. For that, we use 21 years of long-term

unemployment data of European countries¹¹. Touch interaction on the mobile device allows users to zoom and pan a map of European countries and select individual countries. These selected countries then represent a path through Europe, which is used to specify a “non-planar slice through space-time” [112]. Finally, the development of long-term unemployment in selected countries over time is visualized in form of an AR 3D wall extruded along the previously defined path (Figure 6c). Furthermore, Tominski and Schulz [112] provided separated controls for adjusting the height and lifting of the wall. Although not realized, we envision this to be realized in a direct fashion by touching and moving AR buttons at the top or bottom of the wall via a mid-air touch or grab gesture.

At the core of this use case is the concept *Superimposed 3D Visualizations* (see 4.1.5). The example aims at showing that transferring 3D graphics shown on top of a map to AR can make the use of 3D visualizations more accessible. The off-loaded color scale (Figure 6c) also demonstrates the idea of *Separated Visualization User Interface Components* (see 4.1.4). Besides interaction opportunities for the map, the mobile device helps to naturally inspect the 3D wall visualization from different angles by either moving the head or picking up and rotating the mobile device. This shows that the two display technologies can nicely complement each other.

5.2.5 Scatterplot with 3D Glyphs and Trajectories. Besides SPLOM, single scatterplots can be used for the visualization of multivariate data. For this purpose, it is typical to not only use the position but other visual channels, such as color, shape, or size, to encode data. Although such a scatterplot can of course be displayed on a mobile device, MARVIS allows different data dimensions to be displayed in AR in the space above the mobile device. To illustrate this, we adapt parts of the Gapminder interface¹², from which we also compiled a data set of approximately 26.000 numeric values. The idea is that the mobile device shows a bubble chart, in which each bubble represents one of 19 countries of the Group of Twenty (G20, excl. European Union). Indicators of economy and politics, such as population, life expectancy, income per person, or CO₂ emissions, are used as data dimensions and can be encoded by x-y position and size of a bubble. While color encodes the affiliation to a continent, each chart shows data for a specific year, spanning 200

¹⁰Source crime data: <https://data.baltimorecity.gov/>, BPD Part 1 Victim Based Crime Data by Baltimore Police Department, last updated 08/25/2017, licensed under CC BY 3.0

¹¹Source unemployment data: <https://www.ilo.org/ilostat> through www.gapminder.org

¹²Gapminder Tool, <https://www.gapminder.org/tools/>



Figure 7: Showing 3D visualizations above and between devices: (a) Scatterplot extended with 3D glyphs; (b) Another scatterplot superimposed with 3D trajectories based on the position (x, y, z) over time; (c) AR heatmap and stacked bars resulting from a combination of visualizations and devices; (d) Since rows and columns are selected, individual cells and stacked bars at the intersection are highlighted.

years from 1810 to 2010. Touching and dragging a slider below the chart triggers animated changes of position and size for each country. In addition to this, our prototype uses AR to add either 3D glyphs or trajectories.

3D Glyphs: On the basis of size or diameter, we extend bubbles into the 3rd dimension by showing 3D glyphs standing on top of them (Figure 7a). Similar to Thakur and Hanson’s 3D visualization [108], each glyph encodes the change of size of its corresponding bubble over all years (from 1810 at bottom to 2010 at top). This improves the overview as it shows how one of the data dimensions develops over time. 3D glyphs are realized using vertically symmetric billboard line graphs (2D graphics sprites always facing the viewer). Once the user selects bubbles, 3D line graphs appear as detail-on-demand. A visual mark in the line graph (height) indicates the current year. When exploring the data dimension time by scrolling the time slider, this visual marker moves up or down accordingly and glyphs follow their corresponding bubbles.

3D Trajectories: As an alternative, a 3D bubble chart can also encode another data attribute using the position (x, y, z). AR can then show the third axis (z), the bubbles and their paths in 3D above the mobile device (Figure 7b). Perpendicular lines from the 3D bubbles connect their 2D counterparts and help users orient. Selections of bubbles on the mobile device also select corresponding 3D bubbles. Upon selection, 3D trajectories appear and provide an overview of how the three data attributes develop over time.

This use case mainly illustrates how AR can be used to encode further data dimensions, improve overall overview, and help with comparison of data items. The most relevant concept for this example is *Superimposed 3D Visualizations* (see 4.1.5). Even though specifically shown here for scatterplots, the same approach could be applied to other visualization techniques easily.

5.2.6 Combined Bar Charts with Heatmap and 3D Stacked Bars. Bar charts represent a simple but widely used technique. Despite their relatively low complexity, using multiple bar charts also allows users to develop an overview of data. Building on multivariate data about university students¹³, their enrolled courses, gender, and country of origin, we show how MARVIS can be used to combine bar charts and thus provide additional, detailed information. The

starting point is a set of two mobile devices: one shows the number of students enrolled per semester (20 winter terms), the other shows the number of students per course of study (16 courses). To inspect further details, the user creates a corner-to-corner device arrangement, activating a heatmap shown between the two devices in AR (Figure 7c). The heatmap results from the combination of the bar charts and thus each cell shows the number of students for the crossing of a specific course of study and semester. This enables users to understand how each item of a bar chart is constructed, without just performing this combination mentally (reducing mental workload). A further selection of one of the bars via tap on the mobile device will reveal even more information for the corresponding heatmap row or column in AR. 3D stacked bars (Figure 7c) display demographic details regarding gender and country of origin (i.e., local or international). In order to focus on individual stacked bars, both a row and a column needs to be selected. As a result, only the stacked bar at the intersection of this selection is highlighted (Figure 7d).

This use case mainly exemplifies the concept *Combination of Visualizations* (see 4.2.2). AR helps to convey further data without altering visualizations on mobile devices, whereas touch interaction is used to stepwise specify where to enrich information.

6 DISCUSSION: EXPERT REVIEWS & LESSONS LEARNED

After having introduced the MARVIS concepts and our implemented prototype in detail, the goal of this discussion section is to report on feedback and reflect on lessons learned.

As another part of our iterative development process and in addition to the initial expert interviews (see 3.2), we invited the same seven experts again (HCI and visualization researchers) to evaluate our concepts and early prototype. We therefore conducted seven individual expert reviews as hands-on sessions allowing them to test the implemented use cases themselves. We deliberately used simple data for the use cases ensuring that specific domain knowledge of the data was not necessary. Each session lasted approx. 50 minutes (eight minutes per use case), was conducted by one investigator, and, after a general introduction, proceeded as follows: The investigator started a use case, explained its data, intended use, and functionality; In a think-aloud style, experts then verbally stated their opinions

¹³Source student data: <https://www-genesis.destatis.de/genesis//online?operation=table&code=21311-0003>

and reported advantages, disadvantages, and problems. While we again videotaped sessions, we also took notes which resulted in much valuable qualitative feedback already. After digitizing a total of 190 individual notes, two of the authors collaboratively labeled and grouped these notes using affinity diagramming. Themes revealed by this are: successful design, design issues and alternatives, missing functionality, AR information display, interaction, device tracking and registration, and user behavior and perception.

Overall, the experts were very positive and interested in seeing how initial ideas evolved into an interactive prototype. They stated, e.g., *“convincing, AR complements and adds value to the exploration process”* and *“interesting idea, the tablet displays a detailed view, AR offers pointers for next steps.”* The SPLOM use case (see 5.2.1) helped to discuss the extent of AR views (*“thumbnails are good for overview, full-size cells are too large”*) and interaction (*“selection in AR via poke gesture is missing”*). The node-link diagram and attribute visualization example (see 5.2.2) showed that relations between views and devices can be designed in various ways (*“the connection should be curved and above devices”*, *“use colored borders to reveal device connections”*). The map navigation use case (see 5.2.3) allowed to reflect on the role of display technologies and the level of detail for AR (*“great, tablet appears like a detail cutout”*, *“the more diminished design is to be preferred, since priority is on data, i.e. tablet”*). Comments about the example of a 3D wall on a map (see 5.2.4) focused on interaction again (*“use mid-air pinch gesture for zooming”*, *“lift the tablet for slicing through the wall”*). The 3D glyphs and trajectories use case (see 5.2.5) emphasized the tangibility of a tablet (*“great to be able to pick up and rotate a 3D visualization”*) but also the importance of alignment (*“hard to recognize which 2D and 3D paths belong together”*). Similarly, the combined bar charts example (see 5.2.6) showed that alignment is crucial (*“gap between views is too large”*, *“show guidelines from bars to columns and rows”*), but also that AR now seemed more central (*“too bad the tablet is only used for selection”*).

Interaction in AR. One of our design choices was to realize most interaction with visualizations through mobile devices. In addition to our awareness that direct interaction with AR content is clearly needed, the use of our prototype emphasizes this need. This was also confirmed by six of seven experts. Essentially, they expected to be able to select or manipulate the AR visualization using mid-air gestures, for example, to *“select a SPLOM cell via air-tap or poke gesture for detailed inspection on a mobile device”*. Another interesting comment refers to interaction with AR content by using the table as an input surface because of its haptics (three experts). One expert even stressed that the actual location and availability as well as the size of a table might have a major influence on the preferred way of interacting with AR content.

As a result, we plan to integrate basic AR interaction into our prototype. Moreover, from the feedback we can conclude that AR interaction should be considered equal to touch interaction.

Mobile Device Interaction. In relation to the design choice spatial awareness of devices and in context of 3D visualizations attached to mobile devices (see 4.1.5), four experts reported advantages of being able to pick up the visualization, easily rotate it, and thereby inspect it from different viewing angles. While the touch interaction simply worked, experts confirmed statements from the first interviews and

suggested to support device movements for selecting or picking a view (two experts), navigating a visualization (three experts), or specifying a parameter such as size of a corresponding AR visualization (one expert). At the same time, one expert noted that such spatial input would be in conflict with the idea *“to pick up a device simply because of comfort reasons or to enjoy a 3D visualization.”*

This illustrates the need to further explore how the inherent mobility of devices can be preserved when device movement is directly used for interaction. In any case, it should be considered that the freedom of movement is rather limited in seated scenarios (design choice seated usage). Therefore, it remains open whether it would perform the same way as it does when standing up [103].

Layout and Alignment in AR. The way information is positioned in AR can rely on physical properties [87], such as the position, orientation, or size of a mobile device as well as the number and location of users. Besides detached and freely positioned information displayed in AR, our own experience and positive expert feedback affirms the primary strategy of positioning AR content relatively to a mobile device, which stems from our design choices to focus on the mobile device and that AR supports and provides context. In this context, 5 experts stated their preference for arranging AR content on a planar layer specified by the device position and orientation. Four experts mentioned that some content could also stand on this plane or being tilted towards the user (e.g., legends, views). Additionally, AR content can also be positioned relatively to the table, i.e., fixed on the table surface (four experts). It also seems to be useful to switch between such points of reference to attach some views to the device and some to the table (three experts).

Overall, we can state that more applications and practical realizations for specific domains are needed. Since there are numerous possible designs of ‘where to place AR views?’ and ‘how to orientate and arrange them?’, this will allow us to better understand consequences of design decisions case by case and then to derive underlying principles.

AR Regions and Content. In relation to layout and alignment, another question is ‘What content should be displayed in AR, and where?’ While this can depend on the actual visualization technique and task, two experts provided interesting insights in this regard. The region on top of a mobile device (cf. planar zones of spatial alignment of AR content [89]) was easily visible in addition to the actual display (without extra head movement). Therefore, they suggest to use this space for overview visualizations. Then, the area to the left and right of a mobile device could be used for alternative views, because occasional left-right head movements appear reasonable for that. Finally, the two experts agreed that the potentially thin region below a device (bottom region, basically between the user and display) should not be used for essential visualizations, but for other interface components such as menus or controls.

We learned that it might be useful to connect different areas around a device to certain purposes—at least when AR is used to display distinct views or user interface components. In principle we have similar strategies for desktop applications, such as a PDF reader: a menu bar is at the top, bookmarks, table of content, comments at the sides, and sometime a status bar at the bottom.

Extent and Level of Detail of AR Information. The map navigation use case (see 5.2.3) shows that we decided to limit the AR visualization to immediate device proximity. Two experts disagreed on this, one saying the extent was just right, the other telling it was too closely. A third expert proposed to “give users control on the extent of an AR view around a mobile device”, allowing to relate it to a meaningful parameter such as maximum travel distance. In our SPLOM use case we used different design variants to also address the extent of AR content, including the two extremes of showing only a single neighboring cell up to showing a full-size 8×8 SPLOM that consumes more space than a typical desk. However, one expert noted that maybe it is not about the extent itself but the question whether it is intended to let users physically walk or move through the data space.

With regard to the level of detail, we outlined in our design choices that AR supports and provides context. We argue for a careful trade-off between what is necessary and possible—moderate and less prominent versus fully detailed. The use of an abstracted visual style for displaying a geographic map in AR was positively mentioned by four experts. Due to its display resolution, two experts raised issues regarding the readability of details or text labels within an AR view. One of them stated that, for example, tool tips or details-on-demand “should have another visual representation adequate for current AR HMDs.”

As a result, we suggest to further explore the spectrum and transition between small and large AR visualizations coupled to a mobile device, especially because both extents appear to be relevant. It is also unclear yet, how other specific visualizations could be seamlessly extended around a device in AR—how to design an abstracted parallel coordinates plots, node-link diagram, or treemap visualization.

Technical Challenges. Similar to the advancements of mobile devices, we think AR HMDs also have the potential to be used in many future everyday situations. However, some technical aspects are crucial to MARVIS concepts. A major issue in relation to *display quality of AR HMDs* is the field of view (FoV). Developing and using our prototype showed that AR visualizations attached to a hand-held mobile device are hard to realize without a large FoV. Additionally, visualizations are often visually fine-grained and detailed, which is why AR HMDs need sufficiently high resolution in order to prevent jagged and pixelated output. Even though neither we nor our participants recognized this as an actual problem, environmental aspects or differences between display technologies might also be relevant. For example, lighting, table pattern, or color accuracy might affect concepts such as the seamless visualization extension (see 4.1.2).

Another technical challenge is the accurate localization of mobile devices and their registration in relation to an AR HMD. Four experts explicitly stated that “this alignment is very important”. Moreover, our current setup with its outside-in tracking is not sufficient for a deployment in non-laboratory environments. Unfortunately, even modern devices, such as the latest Microsoft HoloLens 2, do not allow developers to track physical objects such as mobile devices. Besides a possible AR HMD inside-out tracking similar to how Virtual Reality headsets track their controllers, bluetooth-based approaches [55, 56] could be used to provide basic positioning

information of mobile devices. This seems to be sufficient when it is not about high precision, such as to just display something in AR alongside a mobile device rather than visually connecting two tiny objects shown on mobile devices.

Applications Beyond MARVIS. The idea behind our work can also be generalized as a *combination of commodity 2D displays with AR*—a concept that Reipschläger and Dachsel [87] refer to as Augmented Displays. On a user interface level, it is about offloading user interface components, such as menus or windows, to AR and providing additional screen space to show more or different information to users. In order to reflect on the connection to related work, it is useful to take a closer look at device types, display sizes, and application domains. MARVIS is focused on mobile devices. The concepts of this work, therefore, deal with hand-held displays that can be freely moved and arranged. This differs to most prior work using either body-referenced interaction spaces [42] or stationary displays, such as desktop-like workplaces [87, 115], tabletops [20], or wall-sized displays [89, 106]. At the same time, a shared element is to prioritize the direct extension of a display. This extension is also what distinguishes MARVIS from other approaches which primarily use a mobile device as a tangible interaction proxy [17, 50]. Even though our work presents techniques that are particularly relevant for data visualization, recent concepts for 3D modeling and architecture [88] or sketching [3] show the applicability to other use cases.

Nevertheless, we think there is also a great potential for using AR HMDs in cross-device and multi-display environments (MDE), i.e., where several and possibly heterogeneous devices are used together. As explained in our concepts for two or more mobile devices (see 4.2), AR HMDs can provide assistance for connectivity and particularly for collaboration. Future research can further explore how and in which ways the augmentation of devices and displays can help to, for example, identify information across devices, illustrate relationships between windows or devices, or in relation to workspace awareness [44, 45, 47, 101] and awareness support [46], indicate the source and effect of individual interactions.

7 CONCLUSION

With MARVIS, we contributed a conceptual framework for a new class of mobile visualization environments that combine mobile devices and head-mounted Augmented Reality for the purpose of visual data analysis. In an iterative design process involving several experts we developed the basic concepts for combining one or multiple mobile displays with 2D and 3D augmentations around, above, and even between them. Specifically, we designed a series of interactive visualization concepts for this setup that address typical visualization challenges and illustrate potential benefits.

We validated them by implementing six example use cases, which were tested and reviewed in seven hands-on sessions with experts. Based on their valuable feedback and our own explorations, we were able to derive lessons learned and reflect on research challenges. They will also inform future work in this area with regard to, for example, technical improvements of the prototype and more comprehensible application examples to be evaluated in formal studies. Our concepts and prototype clearly demonstrate the interesting potential of the MARVIS framework and the rich variety

of possible applications. We hope to have laid the foundations for conducting more research on augmented mobile visualizations and that we could inspire future data visualizations.

ACKNOWLEDGMENTS

We thank Christian Tominski for his early advice on MARVIS' concepts and structure of the work, Tamara Flemisch and Konstantin Klamka for supporting the preparation of the paper, and the anonymous reviewers for their constructive feedback and suggestions to improve the paper.

This work was supported in part by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC-2068 - 390729961 - Cluster of Excellence "Physics of Life" of TU Dresden, by DFG under grant number 319919706/RTG2323 "Conductive Design of Cyber-Physical Production Systems", by DFG grant 389792660 as part of TRR 248 (see <https://perspicuous-computing.science>), and by DFG as part of Germany's Excellence Strategy EXC 2050/1 - Project ID 390696704 - Cluster of Excellence "Centre for Tactile Internet with Human-in-the-Loop" (CeTI) of Technische Universität Dresden.

REFERENCES

- [1] Keith Andrews. 2018. Responsive Visualisation. *Proceedings of MobileVis 2018 Workshop at CHI 2018* (April 2018).
- [2] Keith Andrews and Aleš Smrdel. 2017. Responsive Data Visualisation. In *EuroVis 2017 - Posters*, Anna Puig Puig and Tobias Isenberg (Eds.). The Eurographics Association. <https://doi.org/10.2312/eurp.20171182>
- [3] Rahul Arora, Rubaiat Habib Kazi, Tovi Grossman, George Fitzmaurice, and Karan Singh. 2018. *SymbiosisSketch: Combining 2D & 3D Sketching for Designing Detailed 3D Objects in Situ*. Association for Computing Machinery, New York, NY, USA, 1–15. <https://doi.org/10.1145/3173574.3173759>
- [4] Benjamin Bach, Emmanuel Pietriga, and Jean-Daniel Fekete. 2014. Visualizing Dynamic Networks with Matrix Cubes. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Toronto, Ontario, Canada) (CHI '14). Association for Computing Machinery, New York, NY, USA, 877–886. <https://doi.org/10.1145/2556288.2557010>
- [5] Benjamin Bach, Ronell Sciat, Johanna Beyer, Maxime Cordeil, and Hanspeter Pfister. 2018. The Hologram in My Hand: How Effective is Interactive Exploration of 3D Visualizations in Immersive Tangible Augmented Reality? *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (Jan 2018), 457–467. <https://doi.org/10.1109/TVCG.2017.2745941>
- [6] Dominikus Baur, Bongshin Lee, and Sheelagh Carpendale. 2012. TouchWave: Kinetic Multi-touch Manipulation for Hierarchical Stacked Graphs. In *Proceedings of the 2012 ACM International Conference on Interactive Tabletops and Surfaces* (Cambridge, Massachusetts, USA) (ITS '12). ACM, New York, NY, USA, 255–264. <https://doi.org/10.1145/2396636.2396675>
- [7] Benjamin B. Bederson, Aaron Clamage, Mary P. Czerwinski, and George G. Robertson. 2004. DateLens: A Fisheye Calendar Interface for PDAs. *ACM Transactions on Computer-Human Interaction* 11, 1 (March 2004), 90–119. <https://doi.org/10.1145/972648.972652>
- [8] Mark Billinghurst, Maxime Cordeil, Anastasia Bezerianos, and Todd Margolis. 2018. Collaborative Immersive Analytics. In *Immersive Analytics*, Kim Marriott, Falk Schreiber, Tim Dwyer, Karsten Klein, Nathalie Henry Riche, Takayuki Itoh, Wolfgang Stuerzlinger, and Bruce H. Thomas (Eds.). Springer International Publishing, Cham, 221–257. https://doi.org/10.1007/978-3-030-01388-2_8
- [9] Tanja Blaschek, Lonni Besançon, Anastasia Bezerianos, Bongshin Lee, and Petra Isenberg. 2019. Glanceable Visualization: Studies of Data Comparison Performance on Smartwatches. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (Jan 2019), 630–640. <https://doi.org/10.1109/TVCG.2018.2865142>
- [10] Michael Bostock, Vadim Ogievetsky, and Jeffrey Heer. 2011. D³ Data-Driven Documents. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (Dec 2011), 2301–2309. <https://doi.org/10.1109/TVCG.2011.185>
- [11] Richard Brath. 2014. 3D InfoVis is here to stay: Deal with it. In *2014 IEEE VIS International Workshop on 3DVis (3DVis)*, 25–31. <https://doi.org/10.1109/3DVis.2014.7160096>
- [12] Matthew Brehmer, Bongshin Lee, Petra Isenberg, and Eun Kyoung Choe. 2019. Visualizing Ranges over Time on Mobile Phones: A Task-Based Crowdsourced Evaluation. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (Jan 2019), 619–629. <https://doi.org/10.1109/TVCG.2018.2865234>
- [13] Matthew Brehmer, Bongshin Lee, Petra Isenberg, and Eun Kyoung Choe. 2020. A Comparative Evaluation of Animation and Small Multiples for Trend Visualization on Mobile Phones. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (Jan 2020), 364–374. <https://doi.org/10.1109/TVCG.2019.2934397>
- [14] Matthew Brehmer and Tamara Munzner. 2013. A Multi-Level Typology of Abstract Visualization Tasks. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (Dec 2013), 2376–2385. <https://doi.org/10.1109/TVCG.2013.124>
- [15] Thorsten Buring and Harald Reiterer. 2005. ZuiScat: Querying and Visualizing Information Spaces on Personal Digital Assistants. In *Proceedings of the 7th International Conference on Human Computer Interaction with Mobile Devices & Services* (Salzburg, Austria) (MobileHCI '05). ACM, New York, NY, USA, 129–136. <https://doi.org/10.1145/1085777.1085799>
- [16] Wolfgang Büschel, Jian Chen, Raimund Dachsel, Steven Drucker, Tim Dwyer, Carsten Görg, Tobias Isenberg, Andreas Kerren, Chris North, and Wolfgang Stuerzlinger. 2018. Interaction for Immersive Analytics. In *Immersive Analytics*, Kim Marriott, Falk Schreiber, Tim Dwyer, Karsten Klein, Nathalie Henry Riche, Takayuki Itoh, Wolfgang Stuerzlinger, and Bruce H. Thomas (Eds.). Lecture Notes in Computer Science, Vol. 11190. Springer International Publishing, Cham, 95–138. https://doi.org/10.1007/978-3-030-01388-2_4
- [17] Wolfgang Büschel, Annett Mitschick, and Raimund Dachsel. 2019. Investigating Smartphone-based Pan and Zoom in 3D Data Spaces in Augmented Reality. In *Proceedings of the 21st International Conference on Human-Computer Interaction with Mobile Devices and Services* (Taipei, Taiwan) (MobileHCI '19). ACM, 13 pages. <https://doi.org/10.1145/3338286.3340113>
- [18] Wolfgang Büschel, Patrick Reipschläger, Ricardo Langner, and Raimund Dachsel. 2017. Investigating the Use of Spatial Interaction for 3D Data Visualization on Mobile Devices. In *Proceedings of the 2017 ACM International Conference on Interactive Surfaces and Spaces* (Brighton, United Kingdom) (ISS '17). ACM, New York, NY, USA, 62–71. <https://doi.org/10.1145/3132272.3134125>
- [19] Wolfgang Büschel, Stefan Vogt, and Raimund Dachsel. 2019. Augmented Reality Graph Visualizations. *IEEE Computer Graphics and Applications* 39, 3 (May 2019), 29–40. <https://doi.org/10.1109/MCG.2019.2897927>
- [20] 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, Article 90, 12 pages. <https://doi.org/10.1145/3173574.3173664>
- [21] Marco Cavallo, Mishal Dholakia, Matous Havlena, Kenneth Ocheltree, and Mark Podlasek. 2019. Dataspace: A Reconfigurable Hybrid Reality Environment for Collaborative Information Analysis. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, 145–153. <https://doi.org/10.1109/VR.2019.8797733>
- [22] Marco Cavallo, Mishal Dolakia, Matous Havlena, Kenneth Ocheltree, and Mark Podlasek. 2019. Immersive Insights: A Hybrid Analytics System For Collaborative Exploratory Data Analysis. In *25th ACM Symposium on Virtual Reality Software and Technology* (Parramatta, NSW, Australia) (VRST '19). Association for Computing Machinery, New York, NY, USA, Article 9, 12 pages. <https://doi.org/10.1145/3359996.3364242>
- [23] Tom Chandler, Maxime Cordeil, Tobias Czauderna, Tim Dwyer, Jaroslaw Glowacki, Catagay Goncu, Matthias Klapperstueck, Karsten Klein, Falk Schreiber, and Elliot Wilson. 2015. Immersive Analytics. In *2015 big data visual analytics (BDVA 2015) : Hobart, Australia, 22-25 September 2015*. IEEE, Piscataway, NJ, 73–80. <https://doi.org/10.1109/BDVA.2015.7314296>
- [24] Luca Chittaro. 2006. Visualizing information on mobile devices. *Computer* 39, 3 (March 2006), 40–45. <https://doi.org/10.1109/MC.2006.109>
- [25] Haeyong Chung and Chris North. 2017. SAVIL: cross-display visual links for sensemaking in display ecologies. *Personal and Ubiquitous Computing* 22, 2 (dec 2017), 409–431. <https://doi.org/10.1007/s00779-017-1091-4>
- [26] Haeyong Chung, Chris North, Joshi Sarang, and Jian Chen. 2015. Four considerations for supporting visual analysis in display ecologies. In *2015 IEEE Conference on Visual Analytics Science and Technology (VAST)*. IEEE Computer Society, 33–40. <https://doi.org/10.1109/VAST.2015.7347628>
- [27] Andy Cockburn, Amy Karlson, and Benjamin B. Bederson. 2009. A Review of Overview+Detail, Zooming, and Focus+Context Interfaces. *ACM Computing Surveys (CSUR)* 41, 1, Article 2 (Jan. 2009), 31 pages. <https://doi.org/10.1145/1456650.1456652>
- [28] Maxime Cordeil, Benjamin Bach, Yongchao Li, Elliott Wilson, and Tim Dwyer. 2017. Design space for spatio-data coordination: Tangible interaction devices for immersive information visualisation. In *2017 IEEE Pacific Visualization Symposium (PacificVis)*, 46–50.
- [29] Daniel Drechtort and Christian Geiger. 2015. Collaborative Magic Lens Graph Exploration. In *SIGGRAPH Asia 2015 Mobile Graphics and Interactive Applications* (Kobe, Japan) (SA '15). Association for Computing Machinery, New York, NY, USA, Article 25, 3 pages. <https://doi.org/10.1145/2818427.2818465>
- [30] Steven M. Drucker, Danyel Fisher, Ramik Sadana, Jessica Herron, and m.c. schraefel. 2013. TouchViz: A Case Study Comparing Two Interfaces for Data Analytics on Tablets. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Paris, France) (CHI '13). ACM, New York, NY, USA,

- 2301–2310. <https://doi.org/10.1145/2470654.2481318>
- [31] Tim Dwyer, Kim Marriott, Tobias Isenberg, Karsten Klein, Nathalie Riche, Falk Schreiber, Wolfgang Stuerzlinger, and Bruce H. Thomas. 2018. *Immersive Analytics: An Introduction*. Springer International Publishing, Cham, 1–23. https://doi.org/10.1007/978-3-030-01388-2_1
- [32] Niklas Elmqvist, Pierre Dragicevic, and Jean-Daniel Fekete. 2008. Rolling the Dice: Multidimensional Visual Exploration using Scatterplot Matrix Navigation. *IEEE Transactions on Visualization and Computer Graphics* 14, 6 (Nov 2008), 1539–1148. <https://doi.org/10.1109/TVCG.2008.153>
- [33] Niklas Elmqvist, Andrew Vande Moere, Hans-Christian Jetter, Daniel Cernea, Harald Reiterer, and T.-J. Jankun-Kelly. 2011. Fluid Interaction for Information Visualization. *Information Visualization* 10, 4 (2011), 327–340. <https://doi.org/10.1177/1473871611413180>
- [34] Steven Feiner and Ari Shamash. 1991. Hybrid User Interfaces: Breeding Virtually Bigger Interfaces for Physically Smaller Computers. In *Proceedings of the 4th Annual ACM Symposium on User Interface Software and Technology* (Hilton Head, South Carolina, USA) (*UIST '91*). Association for Computing Machinery, New York, NY, USA, 9–17. <https://doi.org/10.1145/120782.120783>
- [35] Jorge A. Wagner Filho, Wolfgang Stuerzlinger, and Luciana Nedel. 2020. Evaluating an Immersive Space-Time Cube Geovisualization for Intuitive Trajectory Data Exploration. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (2020), 514–524.
- [36] Adrien Fonnet and Yannick Prié. 2019. Survey of Immersive Analytics. *IEEE Transactions on Visualization and Computer Graphics* (2019), 1–1.
- [37] Mathias Frisch and Raimund Dachsel. 2010. Off-screen Visualization Techniques for Class Diagrams. In *Proceedings of the 5th International Symposium on Software Visualization* (Salt Lake City, Utah, USA) (*SOFTVIS '10*). ACM, New York, NY, USA, 163–172. <https://doi.org/10.1145/1879211.1879236>
- [38] Bruno Fruchard, Arnaud Prouzeau, Olivier Chapuis, and Eric Lecolinet. 2019. Leveraging Body Interactions to Support Immersive Analytics. In *CHI Workshop on Immersive Analytics*. Glasgow, United Kingdom, 10 pages. <https://hal.archives-ouvertes.fr/hal-02095993>
- [39] Peter S. Games and Alark Joshi. 2014. Visualization of off-screen data on tablets using context-providing bar graphs and scatter plots. In *Visualization and Data Analysis 2014*, Pak Chung Wong, David L. Kao, Ming C. Hao, and Chaomei Chen (Eds.), Vol. 9017. International Society for Optics and Photonics, SPIE, 132 – 146. <https://doi.org/10.1117/12.2038456>
- [40] Michael Gleicher, Danielle Albers, Rick Walker, Ilir Jusufi, Charles D. Hansen, and Jonathan C. Roberts. 2011. Visual comparison for information visualization. *Information Visualization* 10, 4 (10 2011), 289–309. <https://doi.org/10.1177/1473871611416549> arXiv:<http://dx.doi.org/10.1177/1473871611416549>
- [41] Jens Emil Grønbaek and Kenton O'Hara. 2016. Built-In Device Orientation Sensors for Ad-Hoc Pairing and Spatial Awareness. Workshop Paper at ISS 2016 Cross-Surface Workshop.
- [42] Jens Grubert, Matthias Heinisch, Aaron Quigley, and Dieter Schmalstieg. 2015. MultiFi: Multi Fidelity Interaction with Displays On and Around the Body. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) (*CHI '15*). Association for Computing Machinery, New York, NY, USA, 3933–3942. <https://doi.org/10.1145/2702123.2702331>
- [43] Diansheng Guo. 2009. Flow Mapping and Multivariate Visualization of Large Spatial Interaction Data. *IEEE Transactions on Visualization and Computer Graphics* 15, 6 (Nov 2009), 1041–1048. <https://doi.org/10.1109/TVCG.2009.143>
- [44] Carl Gutwin and Saul Greenberg. 1996. Workspace Awareness for Groupware. In *Conference Companion on Human Factors in Computing Systems* (Vancouver, British Columbia, Canada) (*CHI '96*). ACM, New York, NY, USA, 208–209. <https://doi.org/10.1145/257089.257284>
- [45] Carl Gutwin and Saul Greenberg. 1998. Design for Individuals, Design for Groups: Tradeoffs Between Power and Workspace Awareness. In *Proceedings of the 1998 ACM Conference on Computer Supported Cooperative Work* (Seattle, Washington, USA) (*CSCW '98*). ACM, New York, NY, USA, 207–216. <https://doi.org/10.1145/289444.289495>
- [46] Carl Gutwin and Saul Greenberg. 1998. Effects of Awareness Support on Groupware Usability. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Los Angeles, California, USA) (*CHI '98*). ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 511–518. <https://doi.org/10.1145/274644.274713>
- [47] Carl Gutwin, Gwen Stark, and Saul Greenberg. 1995. Support for Workspace Awareness in Educational Groupware. In *The First International Conference on Computer Support for Collaborative Learning* (Indiana Univ., Bloomington, Indiana, USA) (*CSCL '95*). L. Erlbaum Associates Inc., Hillsdale, NJ, USA, 147–156. <https://doi.org/10.3115/222020.222126>
- [48] Ken Hinckley. 2003. Synchronous Gestures for Multiple Persons and Computers. In *Proceedings of the 16th Annual ACM Symposium on User Interface Software and Technology* (Vancouver, Canada) (*UIST '03*). ACM, New York, NY, USA, 149–158. <https://doi.org/10.1145/964696.964713>
- [49] Dandan Huang, Melanie Tory, Bon Adriel Aseniero, Lyn Bartram, Scott Bateman, Sheelagh Carpendale, Anthony Tang, and Robert Woodbury. 2015. Personal Visualization and Personal Visual Analytics. *IEEE Transactions on Visualization and Computer Graphics* 21, 3 (March 2015), 420–433. <https://doi.org/10.1109/TVCG.2014.2359887>
- [50] Sebastian Hubenschmid, Johannes Zagermann, Simon Butscher, and Harald Reiterer. 2018. Employing Tangible Visualisations in Augmented Reality with Mobile Devices. In *MultimodalVis '18 Workshop at 2018 International Conference on Advanced Visual Interfaces* (AVI '18). <http://nbn-resolving.de/urn:nbn:de:bsz:352-2-1100enfo4fofm8>
- [51] 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> arXiv:<http://dx.doi.org/10.1177/1473871611412817>
- [52] Petra Isenberg, Anthony Tang, and Sheelagh Carpendale. 2008. An Exploratory Study of Visual Information Analysis. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Florence, Italy) (*CHI '08*). ACM, New York, NY, USA, 1217–1226. <https://doi.org/10.1145/1357054.1357245>
- [53] Raphaël James, Anastasia Bezerianos, Olivier Chapuis, Maxime Cordeil, Tim Dwyer, and Arnaud Prouzeau. 2020. Personal-Context navigation: combining AR and shared displays in Network Path-following. In *Proceedings of Graphics Interface 2020* (University of Toronto) (*GI 2020*). Canadian Human-Computer Communications Society, 267 – 278. <https://doi.org/10.20380/GI2020.27>
- [54] Waqas Javed and Niklas Elmqvist. 2012. Exploring the design space of composite visualization. In *2012 IEEE Pacific Visualization Symposium* (PacificVis), IEEE, 1–8. <https://doi.org/10.1109/PacificVis.2012.6183556>
- [55] Haojian Jin, Christian Holz, and Kasper Hornbæk. 2015. Tracko: Ad-hoc Mobile 3D Tracking Using Bluetooth Low Energy and Inaudible Signals for Cross-Device Interaction. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology* (Daegu, Kyungpook, Republic of Korea) (*UIST '15*). ACM, New York, NY, USA, 147–156. <https://doi.org/10.1145/2807442.2807475>
- [56] Haojian Jin, Cheng Xu, and Kent Lyons. 2015. Corona: Positioning Adjacent Device with Asymmetric Bluetooth Low Energy RSSI Distributions. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology* (Charlotte, NC, USA) (*UIST '15*). ACM, New York, NY, USA, 175–179. <https://doi.org/10.1145/2807442.2807485>
- [57] Jaemin Jo, Sehi L'Yi, Bongshin Lee, and Jinwook Seo. 2017. TouchPivot: Blending WIMP & Post-WIMP Interfaces for Data Exploration on Tablet Devices. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (*CHI '17*). ACM, New York, NY, USA, 2660–2671. <https://doi.org/10.1145/3025453.3025752>
- [58] Thomas Kirste and Uwe Rauschenbach. 1996. A presentation model for mobile information visualization. *Computers & Graphics* 20, 5 (1996), 669 – 681. [https://doi.org/10.1016/S0097-8493\(96\)00041-6](https://doi.org/10.1016/S0097-8493(96)00041-6)
- [59] Ulrike Kister, Patrick Reipschläger, Fabrice Matulic, and Raimund Dachsel. 2015. BodyLenses: Embodied Magic Lenses and Personal Territories for Wall Displays. In *Proceedings of the 2015 International Conference on Interactive Tabletops & Surfaces* (Madeira, Portugal) (*ITS '15*). Association for Computing Machinery, New York, NY, USA, 117–126. <https://doi.org/10.1145/2817721.2817726>
- [60] Sungahn Ko, Shehzad Afzal, Simon Walton, Yang Yang, Junghoon Chae, Abish Malik, Yun Jang, Min Chen, and David Ebert. 2014. Analyzing high-dimensional multivariate network links with integrated anomaly detection, highlighting and exploration. In *2014 IEEE Conference on Visual Analytics Science and Technology* (VAST). 83–92. <https://doi.org/10.1109/VAST.2014.7042484>
- [61] Menno Jan Kraak. 2003. The Space-Time Cube Revisited From a Geovisualization Perspective. In *Proceedings of the 21st International Cartographic Conference* (ICC 2003). International Cartographic Association, 1988–1996.
- [62] Sven Kratz and Michael Rohs. 2009. HoverFlow: Expanding the Design Space of around-Device Interaction. In *Proceedings of the 11th International Conference on Human-Computer Interaction with Mobile Devices and Services* (Bonn, Germany) (*MobileHCI '09*). Association for Computing Machinery, New York, NY, USA, Article 4, 8 pages. <https://doi.org/10.1145/1613858.1613864>
- [63] Ricardo Langner, Tom Horak, and Raimund Dachsel. 2018. Demonstrating Vistiles: Visual Data Exploration Using Mobile Devices. In *Proceedings of the 2018 International Conference on Advanced Visual Interfaces* (Castiglione della Pescaia, Grosseto, Italy) (*AVI '18*). ACM, New York, NY, USA, Article 69, 3 pages. <https://doi.org/10.1145/3206505.3206583>
- [64] Ricardo Langner, Tom Horak, and Raimund Dachsel. 2018. VisTiles: Coordinating and Combining Co-located Mobile Devices for Visual Data Exploration. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (Jan 2018), 626–636. <https://doi.org/10.1109/TVCG.2017.2744019>
- [65] Bongshin Lee, Matthew Brehmer, Petra Isenberg, Eun Kyoung Choe, Ricardo Langner, and Raimund Dachsel. 2018. Data Visualization on Mobile Devices. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) (*CHI EA '18*). ACM, New York, NY, USA, Article W07, 8 pages. <https://doi.org/10.1145/3170427.3170631>
- [66] Bongshin Lee, Petra Isenberg, Nathalie Henry Riche, and Sheelagh Carpendale. 2012. Beyond Mouse and Keyboard: Expanding Design Considerations for Information Visualization Interactions. *IEEE Transactions on Visualization and Computer Graphics* 18, 12 (2012), 2689–2698. <https://doi.org/10.1109/TVCG.2012.2359887>

- 2012.204
- [67] Mona Hosseinkhani Loorak, Charles Perin, Christopher Collins, and Sheelagh Carpendale. 2017. Exploring the Possibilities of Embedding Heterogeneous Data Attributes in Familiar Visualizations. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (Jan 2017), 581–590. <https://doi.org/10.1109/TVCG.2016.2598586>
- [68] Mona Hosseinkhani Loorak, Charles Perin, Noreen Kamal, Michael Hill, and Sheelagh Carpendale. 2016. TimeSpan: Using Visualization to Explore Temporal Multi-dimensional Data of Stroke Patients. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (Jan 2016), 409–418. <https://doi.org/10.1109/TVCG.2015.2467325>
- [69] Martin Luboschik, Philip Berger, and Oliver Staadt. 2016. On Spatial Perception Issues In Augmented Reality Based Immersive Analytics. In *Proceedings of the 2016 ACM Companion on Interactive Surfaces and Spaces* (Niagara Falls, Ontario, Canada) (ISS '16 Companion). Association for Computing Machinery, New York, NY, USA, 47–53. <https://doi.org/10.1145/3009939.3009947>
- [70] Tahir Mahmood, Erik Butler, Nicholas Davis, Jian Huang, and Aidong Lu. 2018. Building Multiple Coordinated Spaces for Effective Immersive Analytics through Distributed Cognition. In *2018 International Symposium on Big Data Visual and Immersive Analytics (BDVA)*. 1–11. <https://doi.org/10.1109/BDVA.2018.8533893>
- [71] Nicolai Marquardt, Till Ballendat, Sebastian Boring, Saul Greenberg, and Ken Hinckley. 2012. Gradual Engagement: Facilitating Information Exchange Between Digital Devices As a Function of Proximity. In *Proceedings of the 2012 ACM International Conference on Interactive Tabletops and Surfaces* (Cambridge, Massachusetts, USA) (ITS '12). ACM, New York, NY, USA, 31–40. <https://doi.org/10.1145/2396636.2396642>
- [72] Nicolai Marquardt, Frederik Brudy, Can Liu, Ben Bengler, and Christian Holz. 2018. SurfaceConstellations: A Modular Hardware Platform for Ad-Hoc Reconfigurable Cross-Device Workspaces. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) (CHI '18). ACM, New York, NY, USA, Article 354, 14 pages. <https://doi.org/10.1145/3173574.3173928>
- [73] Nicolai Marquardt, Ken Hinckley, and Saul Greenberg. 2012. Cross-device interaction via micro-mobility and f-formations. In *Proceedings of the 25th annual ACM symposium on User interface software and technology* (Cambridge, Massachusetts, USA) (UIST '12). ACM, New York, NY, USA, 13–22. <https://doi.org/10.1145/2380116.2380121>
- [74] Kim Marriott, Falk Schreiber, Tim Dwyer, Karsten Klein, Nathalie Henry Riche, Takayuki Itoh, Wolfgang Stuerzlinger, and Bruce H Thomas. 2018. *Immersive Analytics*. Vol. 11190. Springer.
- [75] David Merrill, Jeevan Kalanithi, and Pattie Maes. 2007. Siftables: Towards Sensor Network User Interfaces. In *Proceedings of the 1st International Conference on Tangible and Embedded Interaction* (Baton Rouge, Louisiana) (TEI '07). ACM, New York, NY, USA, 75–78. <https://doi.org/10.1145/1226969.1226984>
- [76] Alexandre Millette and Michael J. McGuffin. 2016. DualCAD: Integrating Augmented Reality with a Desktop GUI and Smartphone Interaction. In *2016 IEEE International Symposium on Mixed and Augmented Reality (ISMAR-Adjunct)*. 21–26. <https://doi.org/10.1109/ISMAR-Adjunct.2016.0030>
- [77] Calkin S. Montero, Jason Alexander, Mark T. Marshall, and Sriram Subramanian. 2010. Would You Do That?: Understanding Social Acceptance of Gestural Interfaces. In *Proceedings of the 12th International Conference on Human Computer Interaction with Mobile Devices and Services* (Lisbon, Portugal) (MobileHCI '10). ACM, New York, NY, USA, 275–278. <https://doi.org/10.1145/1851600.1851647>
- [78] Heidi Selmer Nielsen, Marius Pallisgaard Olsen, Mikael B. Skov, and Jesper Kjeldskov. 2014. JuxtaPinch: An Application for Photo Sharing. In *Proceedings of the 16th International Conference on Human-computer Interaction with Mobile Devices & Services* (Toronto, ON, Canada) (MobileHCI '14). ACM, New York, NY, USA, 417–420. <https://doi.org/10.1145/2628363.2633569>
- [79] Carolina Nobre, Miriah Meyer, Marc Streit, and Alexander Lex. 2019. The State of the Art in Visualizing Multivariate Networks. *Computer Graphics Forum* 38, 3 (2019), 807–832. <https://doi.org/10.1111/cgf.13728> arXiv:<https://doi.org/10.1111/cgf.13728>
- [80] Erwan Normand and Michael J. McGuffin. 2018. Enlarging a Smartphone with AR to Create a Handheld VESAD (Virtually Extended Screen-Aligned Display). In *2018 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. 123–133. <https://doi.org/10.1109/ISMAR.2018.00043>
- [81] Chris North and Ben Shneiderman. 2000. Snap-together Visualization: A User Interface for Coordinating Visualizations via Relational Schemata. In *Proceedings of the Working Conference on Advanced Visual Interfaces* (Palermo, Italy) (AVI '00). ACM, New York, NY, USA, 128–135. <https://doi.org/10.1145/345513.345282>
- [82] Takashi Ohta and Jun Tanaka. 2012. Pinch: An Interface That Relates Applications on Multiple Touch-Screen by “Pinching” Gesture. In *Advances in Computer Entertainment*, Anton Nijholt, Teresa Romão, and Dennis Reidsma (Eds.). Lecture Notes in Computer Science, Vol. 7624. Springer Berlin Heidelberg, 320–335. https://doi.org/10.1007/978-3-642-34292-9_23
- [83] Antti Oulasvirta and Kasper Hornbæk. 2016. HCI Research as Problem-Solving. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (CHI '16). Association for Computing Machinery, New York, NY, USA, 4956–4967. <https://doi.org/10.1145/2858036.2858283>
- [84] Alexandros Panagiotidis, Michael Burch, Oliver Deussen, Daniel Weiskopf, and Thomas Ertl. 2014. Graph Exploration by Multiple Linked Metric Views. In *2014 18th International Conference on Information Visualisation*. 19–26. <https://doi.org/10.1109/IV.2014.51>
- [85] Thomas Plank, Hans-Christian Jetter, Roman Rädle, Clemens N. Klokmoose, Thomas Luger, and Harald Reiterer. 2017. Is Two Enough?! Studying Benefits, Barriers, and Biases of Multi-Tablet Use for Collaborative Visualization. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '17). ACM, New York, NY, USA, 4548–4560. <https://doi.org/10.1145/3025453.3025537>
- [86] Roman Rädle, Hans-Christian Jetter, Nicolai Marquardt, Harald Reiterer, and Yvonne Rogers. 2014. HuddleLamp: Spatially-Aware Mobile Displays for Ad-hoc Around-the-Table Collaboration. In *Proceedings of the Ninth ACM International Conference on Interactive Tabletops and Surfaces* (Dresden, Germany) (ITS '14). ACM, New York, NY, USA, 45–54. <https://doi.org/10.1145/2669485.2669500>
- [87] Patrick Reipschläger and Raimund Dachsel. 2019. DesignAR: Immersive 3D-Modeling Combining Augmented Reality with Interactive Displays. In *Proceedings of the 2019 ACM International Conference on Interactive Surfaces and Spaces* (Daejeon, Republic of Korea) (ISS '19). ACM, New York, NY, USA, 29–41. <https://doi.org/10.1145/3343055.3359718>
- [88] Patrick Reipschläger, Severin Engert, and Raimund Dachsel. 2020. Augmented Displays: Seamlessly Extending Interactive Surfaces With Head-Mounted Augmented Reality. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI EA '20). Association for Computing Machinery, New York, NY, USA, 1–4. <https://doi.org/10.1145/3334480.3383138>
- [89] Patrick Reipschläger, Tamara Flemisch, and Raimund Dachsel. 2020. Personal Augmented Reality for Information Visualization on Large Interactive Displays. *IEEE Transactions on Visualization and Computer Graphics* (2020), 1–1. <https://doi.org/10.1109/TVCG.2020.3030460>
- [90] Jun Rekimoto and Masanori Saitoh. 1999. Augmented Surfaces: A Spatially Continuous Work Space for Hybrid Computing Environments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Pittsburgh, Pennsylvania, USA) (CHI '99). Association for Computing Machinery, New York, NY, USA, 378–385. <https://doi.org/10.1145/302979.303113>
- [91] Jonathan C. Roberts, Panagiotis D. Ritsos, Sriram Karthik Badam, Dominique Brodbeck, Jessie Kennedy, and Niklas Elmqvist. 2014. Visualization beyond the Desktop—the Next Big Thing. *IEEE Computer Graphics and Applications* 34, 6 (Nov 2014), 26–34. <https://doi.org/10.1109/MCG.2014.82>
- [92] Michael Rohs, Johannes Schöning, Martin Raubal, Georg Essl, and Antonio Krüger. 2007. Map Navigation with Mobile Devices: Virtual versus Physical Movement with and without Visual Context. In *Proceedings of the 9th International Conference on Multimodal Interfaces* (Nagoya, Aichi, Japan) (ICMI '07). Association for Computing Machinery, New York, NY, USA, 146–153. <https://doi.org/10.1145/1322192.1322219>
- [93] Irene Ros. 2014. MobileVis: Examples of Data Visualization Usage on Mobile Devices. <http://mobilevis.is/>. Last checked 2019-12-18.
- [94] Ramik Sadana and John Stasko. 2014. Designing and Implementing an Interactive Scatterplot Visualization for a Tablet Computer. In *Proceedings of the 2014 International Working Conference on Advanced Visual Interfaces* (Como, Italy) (AVI '14). ACM, New York, NY, USA, 265–272. <https://doi.org/10.1145/2598153.2598163>
- [95] Ramik Sadana and John Stasko. 2016. Designing Multiple Coordinated Visualizations for Tablets. *Computer Graphics Forum* 35, 3 (2016), 261–270. <https://doi.org/10.1111/cgf.12902>
- [96] Ramik Sadana and John Stasko. 2016. Expanding Selection for Information Visualization Systems on Tablet Devices. In *Proceedings of the 2016 ACM on Interactive Surfaces and Spaces* (Niagara Falls, Ontario, Canada) (ISS '16). ACM, New York, NY, USA, 149–158. <https://doi.org/10.1145/2992154.2992157>
- [97] Sebastian Sadowski. 2018. Mobile Infovis and Dataviz Pattern: Best of Data and Information Visualisations for Mobile Devices. <https://mobileinfovis.com/>. Last checked 2019-12-18.
- [98] Stacey D. Scott, M. Sheelagh T. Carpendale, and Kori M. Inkpen. 2004. Territoriality in Collaborative Tabletop Workspaces. In *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work* (Chicago, Illinois, USA) (CSCW '04). ACM, New York, NY, USA, 294–303. <https://doi.org/10.1145/1031607.1031655>
- [99] Mickael Sereno, Lonni Besançon, and Tobias Isenberger. 2019. Supporting Volumetric Data Visualization and Analysis by Combining Augmented Reality Visuals with Multi-Touch Input. *EuroVis Posters*, 21–23 pages. <https://doi.org/10.2312/eurp.20191136> Poster.
- [100] Ben Shneiderman. 2003. Why not make interfaces better than 3D reality? *IEEE Computer Graphics and Applications* 23, 6 (Nov 2003), 12–15. <https://doi.org/10.1109/MCG.2003.1242376>
- [101] Anton Sigitov. 2016. Effects of Workspace Awareness and Territoriality in Environments with Large, Shared Displays. In *Proceedings of the 2016 ACM Companion on Interactive Surfaces and Spaces* (Niagara Falls, Ontario, Canada) (ISS Companion '16). ACM, New York, NY, USA, 1–6. <https://doi.org/10.1145/2858036.2858283>

- 3009939.3009940
- [102] Richard Skarbez, Nicholas F. Polys, J. Todd Ogle, Chris North, and Doug A. Bowman. 2019. Immersive Analytics: Theory and Research Agenda. *Frontiers in Robotics and AI* 6 (2019). <https://doi.org/10.3389/frobt.2019.00082>
 - [103] Martin Spindler, Martin Schuessler, Marcel Martsch, and Raimund Dachselt. 2014. Pinch-drag-flick vs. Spatial Input: Rethinking Zoom & Pan on Mobile Displays. In *Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems* (Toronto, Ontario, Canada) (CHI '14). ACM, New York, NY, USA, 1113–1122. <https://doi.org/10.1145/2556288.2557028>
 - [104] Martin Spindler, Christian Tominski, Heidrun Schumann, and Raimund Dachselt. 2010. Tangible Views for Information Visualization. In *ACM International Conference on Interactive Tabletops and Surfaces* (Saarbrücken, Germany) (ITS '10). ACM, New York, NY, USA, 157–166. <https://doi.org/10.1145/1936652.1936684>
 - [105] John Stasko, Carsten Görg, and Zhicheng Liu. 2008. Jigsaw: Supporting Investigative Analysis through Interactive Visualization. *Information Visualization* 7, 2 (2008), 118–132. <https://doi.org/10.1057/palgrave.ivs.9500180> arXiv:<https://doi.org/10.1057/palgrave.ivs.9500180>
 - [106] Tianchen Sun, Yucong Ye, Issei Fujishiro, and Kwan-Liu Ma. 2019. Collaborative Visual Analysis with Multi-level Information Sharing Using a Wall-Size Display and See-Through HMDs. In *2019 IEEE Pacific Visualization Symposium (PacificVis)*, 11–20. <https://doi.org/10.1109/PacificVis.2019.00010>
 - [107] Danielle Albers Szafr. 2018. The Good, the Bad, and the Biased: Five Ways Visualizations Can Mislead (and How to Fix Them). *Interactions* 25, 4 (June 2018), 26–33. <https://doi.org/10.1145/3231772>
 - [108] Sidharth Thakur and Andrew J. Hanson. 2010. A 3D Visualization of Multiple Time Series on Maps. In *Proceedings of the 14th International Conference Information Visualisation (IV'10)*. IEEE Computer Society, Los Alamitos, CA, USA, 336–343. <https://doi.org/10.1109/IV.2010.54>
 - [109] Jesse Thomason and Jingtao Wang. 2012. Exploring Multi-Dimensional Data on Mobile Devices with Single Hand Motion and Orientation Gestures. In *Proceedings of the 14th International Conference on Human-Computer Interaction with Mobile Devices and Services Companion* (San Francisco, California, USA) (MobileHCI '12). Association for Computing Machinery, New York, NY, USA, 173–176. <https://doi.org/10.1145/2371664.2371702>
 - [110] Matthew Tobiasz, Petra Isenberg, and Sheelagh Carpendale. 2009. Lark: Coordinating Co-located Collaboration with Information Visualization. *Visualization and Computer Graphics, IEEE Transactions on* 15, 6 (Nov 2009), 1065–1072. <https://doi.org/10.1109/TVCG.2009.162>
 - [111] Christian Tominski, Camilla Forsell, and Jimmy Johansson. 2012. Interaction Support for Visual Comparison Inspired by Natural Behavior. *Visualization and Computer Graphics, IEEE Transactions on* 18, 12 (Dec 2012), 2719–2728. <https://doi.org/10.1109/TVCG.2012.237>
 - [112] Christian Tominski and Hans-Jörg Schulz. 2012. The Great Wall of Space-Time. In *Proceedings of the Workshop on Vision, Modeling & Visualization (VMV)*, Michael Goesele, Thorsten Grosch, Holger Theisel, Klaus Toennies, and Bernhard Preim (Eds.). The Eurographics Association, 199–206. <https://doi.org/10.2312/PE/VMV/VMV12/199-206>
 - [113] Christian Tominski, Petra Schulze-Wollgast, and Heidrun Schumann. 2005. 3D Information Visualization for Time Dependent Data on Maps. In *Proceedings of the Ninth International Conference on Information Visualisation (IV'05)*. IEEE Computer Society, Los Alamitos, CA, USA, 175–181. <https://doi.org/10.1109/IV.2005.3>
 - [114] Stef van den Elzen and Jarke J. van Wijk. 2013. Small Multiples, Large Singles: A New Approach for Visual Data Exploration. *Computer Graphics Forum* 32, 3pt2 (2013), 191–200. <https://doi.org/10.1111/cgf.12106> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.12106>
 - [115] Xiyao Wang, David Rousseau, Lonni Besançon, Mickael Sereno, Mehdi Ammi, and Tobias Isenberg. 2020. Towards an Understanding of Augmented Reality Extensions for Existing 3D Data Analysis Tools. In *CHI 2020 - ACM Conference on Human Factors in Computing Systems*. Honolulu, United States. <https://doi.org/10.1145/3313831.3376657>
 - [116] Mark Weiser. 1991. The Computer for the 21st Century. *Scientific American* 265, 3 (1991), 94–105.
 - [117] Matt Whitlock, Stephen Smart, and Danielle Albers Szafr. 2020. Graphical Perception for Immersive Analytics. In *Proceedings of the IEEE Conference on Virtual Reality (VR '20)*. IEEE.
 - [118] Paweł Woźniak, Lars Lischke, Benjamin Schmidt, Shengdong Zhao, and Morten Fjeld. 2014. Thaddeus: A Dual Device Interaction Space for Exploring Information Visualisation. In *Proceedings of the 8th Nordic Conference on Human-Computer Interaction: Fun, Fast, Foundational* (Helsinki, Finland) (NordCHI '14). ACM, New York, NY, USA, 41–50. <https://doi.org/10.1145/2639189.2639237>
 - [119] Shengzhi Wu, Daragh Byrne, and Molly Wright Steenson. 2020. "Megereality": Leveraging Physical Affordances for Multi-Device Gestural Interaction in Augmented Reality. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI EA '20). Association for Computing Machinery, New York, NY, USA, 1–4. <https://doi.org/10.1145/3334480.3383170>
 - [120] Ka-Ping Yee. 2003. Peephole Displays: Pen Interaction on Spatially Aware Handheld Computers. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Ft. Lauderdale, Florida, USA) (CHI '03). ACM, New York, NY, USA, 1–8. <https://doi.org/10.1145/642611.642613>
 - [121] Ji Soo Yi, Youn ah Kang, J.T. Stasko, and J.A. Jacko. 2007. Toward a Deeper Understanding of the Role of Interaction in Information Visualization. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (Nov 2007), 1224–1231. <https://doi.org/10.1109/TVCG.2007.70515>
 - [122] Johannes Zagermann, Ulrike Pfeil, Philipp von Bauer, Daniel Fink, and Harald Reiterer. 2020. "It's in My Other Hand!" – Studying the Interplay of Interaction Techniques and Multi-Tablet Activities. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3313831.3376540>