

4. Interaction for Immersive Analytics

Wolfgang Büschel¹, Jian Chen², Raimund Dachzelt¹, Steven Drucker³, Tim Dwyer⁴, Carsten Görg⁵, Tobias Isenberg⁶, Andreas Kerren⁷, Chris North⁸, and Wolfgang Stuerzlinger⁹

¹ Technische Universität Dresden, Germany
[bueschel,dachselt]@acm.org

² The Ohio State University, USA
chen.8028@osu.edu

³ Microsoft Research, USA
sdrucker@microsoft.com

⁴ Monash University, Australia
tim.dwyer@monash.edu

⁵ University of Colorado, USA
carsten.goerg@ucdenver.edu

⁶ Inria & Université Paris-Saclay, France
tobias.isenberg@inria.fr

⁷ Linnaeus University, Växjö, Sweden
kerren@acm.org

⁸ Virginia Tech, USA
north@vt.edu

⁹ Simon Fraser University, Canada
w.s@sfu.ca

Abstract. In this chapter, we briefly review the development of natural user interfaces and discuss their role in providing human-computer interaction that is *immersive* in various ways. Then we examine some opportunities for how these technologies might be used to better support data analysis tasks. Specifically, we review and suggest some interaction design guidelines for immersive analytics. We also review some hardware setups for data visualization that are already archetypal. Finally, we look at some emerging system designs that suggest future directions.

Keywords: natural user interfaces, embodied interaction, post-WIMP interfaces, visual analytics, data visualization

4.1. Introduction

Just as being able to interact with your environment is critical to your sense of ‘being in’ the environment, interaction is critical in immersive analytics systems. Donald Norman introduced the notion that the primary challenges in interactive system design were to reduce the *Gulf of Execution* (the barriers preventing people from completing actions when working with a computer interface) and also the *Gulf of Evaluation* (the aspects of the interface that make it difficult

for people to understand the output of the system) [63,118]. It is arguable that the primary focus of visualization systems historically has been on minimizing the Gulf of Evaluation, that is, making visualizations as understandable as possible and that the Gulf of Execution in this domain is less well explored. However, supporting the fundamental tasks of visual analytics (e. g., exploring the visualization, filtering data, and adjusting parameters) with appropriate, intuitive interaction techniques is necessary for the development of efficient and pleasurable visualization systems.

We would further argue that reducing the Gulf of Execution (for example, through direct manipulation of elements of an information display [63]) and bringing the input and output spaces closer (or even uniting them) strongly contributes to a sense of immersion. Before embarking on an investigation of this premise with any degree of specificity, however, it is worth noting that display and interaction technologies are developing very rapidly, making it difficult to assess the practicability of interaction mechanisms, or indeed to predict what interaction with computers will look like in a few years. Thus, we take a moment to reflect on the history of human-computer interaction, and how the interaction devices and techniques that have been developed have been adopted by visual data analytics systems currently in use.

We can identify three broad and separate generations of interactive computing, characterized by distinct interface styles, each optimized to the available hardware capacity during the same period:

Online computing – beginning in the 1960s;

Graphical user interfaces – spearheaded by research in the 1960s but achieving popular use in the 1980s;

Natural user interfaces – again there is much earlier research, including research for 3D user interfaces, but this movement has garnered widespread attention since the 2000s, with the mainstream adoption of touch computing.

At the dawn of computing, communicating with computers was achieved through very low-bandwidth modalities, such as commands punched on paper tape. In the early 60s, computers “came online” with video, keyboard and (occasionally) light-pen interaction allowing operators to edit programs and documents directly in the systems’ memory [113]. Most fundamental elements of interaction, as they are currently used in visual data analysis software, were demonstrated at this time. Examples include mouse control and the window metaphor (introduced by Engelbart and others in 1968¹) and direct interaction with graphical objects (Sutherland’s seminal SketchPad in 1963 [143]). The first commercially successful graphical user interfaces – as we know them today – appeared in the early 80s. Since then, their WIMP (Windows/Icons/Menus/Pointer) interaction paradigm has become mainstream. It has proved successful in the workplace and more specifically in visual analytics software. In fact, most current visualization and analytics applications still rely on classical desktop interfaces.

¹ At “The Mother of All Demos” <https://www.youtube.com/watch?v=yJDv-zdhzMY>

Meanwhile, innovation on alternative interface technologies has been ongoing. These post-WIMP interfaces do not use traditional menus, toolbars, and windows, but rely on, e. g., touch input, gestures, speech, or tangible interaction. Specifically, the *Natural User Interface* (NUI) movement has sought to teach computers to respond to peoples' gestures on multi-touch screens or in midair following very literal physical metaphors for "natural" interaction (see also Section 4.4.). While gesture control has been investigated for a long time, since at least the 1980s [29], the NUI movement has in recent years gained momentum with the massive popularity of touch-screen phones and tablets [10]. NUI also becomes more important in computing environments where a mouse or keyboard is impractical, such as extremely portable devices or in Augmented or Virtual Reality [98]. In addition, for the last few years, there have been efforts to open up (information) visualization to new devices making use of natural interaction techniques [100]. It clearly follows that such natural interaction techniques are promising for immersive analytics.

Incorporating NUIs into specific applications, however, is not trivial [117]. In general, the goal we aim for with any user interfaces is to minimize the potentially cumbersome mechanics of input and the cognitive distance between users' intent and the execution of that intent as provided by the system, i. e., reducing the Gulf of Execution. Despite already receiving strong emphasis in virtual environment research, interaction in immersive analytics brings in new challenges: How interaction augments insight discovery in large, dense, spatial and non-spatial data is still an active area of research (e. g., [47, 83, 138]). While speculation has begun regarding how this translates to immersive environments (e. g., [128]) there remains a great deal of fundamental research to be done to better understand the challenges of "immersing people in their data."

In Section 4.2. we look at some of the exciting developments in immersive human-computer interaction technologies, both systems that are already commercially available, and ones that are only now emerging in research labs; we list the benefits of immersion that we see for various data analytics applications; and, we motivate some requirements and challenges to be met to create effective immersive interaction devices and techniques.

In Section 4.3. we break down the key tasks that must be supported by any interactive visual analytics system. Then, in Section 4.4., we explore the novel interaction modalities in emerging immersive interaction techniques and discuss their viability in terms of these tasks. Then, in Section 4.5. we look in more detail at how these visual analytics tasks map to immersive interactions and we consider design guidelines for interaction for immersive analytics.

In Section 4.6. we describe in more detail some setups for immersive data visualization that have become 'archetypal' in the sense of being widely investigated, at least in research labs. Then, in Section 4.7. we look at new setups and use-cases for immersive analytics that may indicate future directions. Section 4.8. concludes with some research opportunities.

4.2. Opportunities and Requirements

Our vision for immersive analytics is to augment human cognition when working with data using immersive technologies. Since at least the work of Sutherland [144], the assumption has been that the “ultimate display” technology can take advantage of a broader range of senses and capabilities than traditional WIMP interfaces. In the early days of formal information visualization research, pioneers such as Card, Robertson, and Mackinlay painted a long-term vision for an ultimate “Information Visualizer” that was exploiting 3D graphics and made use of the *information workspace* metaphor [30]. While they discussed interaction at a high-level, basically acknowledging that the system needed to be fluid and performant in response to user actions in order to support workflow and user cognition, they gave little indication of what the mechanics (and the devices) for interaction might look like. Arguably, at the time (1991), while 3D graphics rendering technology was making rapid progress, there was little appreciation generally of the challenges that lay ahead in creating more natural and effective ways to interact with these immersive spaces.

We have come a long way since the early 1990s and there exists a plethora of promising technology for interacting with the surfaces and spaces in our environments rather than interaction with computers via a fixed workstation (see, e. g., [100]). Thus, before exploring specific immersive interaction techniques for data analytics, we look in Subsection 4.2.1. at some general motivating examples of the state of the art in immersive interaction techniques. In light of these examples, we consider in Subsection 4.2.2. some of the potential benefits of immersive interaction and then in Subsection 4.2.3. some desirable interface qualities for immersive analytics, hopefully being more specific than Card *et al.* were able to be in their vision from a quarter of a century ago.

4.2.1. Motivating Examples of Immersive Interaction

A lot of the current excitement around immersion and Virtual Reality (VR) is centered around the new generation of VR headsets, their relatively low cost, and their potential for entertainment in the home. Among the rush of new software for these devices, there are some interesting examples of natural interaction design. For example, *Tilt Brush*² is a VR paint app for the HTC Vive from Google. Tilt Brush features an intuitive “palette” interface. The operator holds the virtual palette in their left hand and uses their right to select different brushes, colors, and other functionality. Different faces of the palette have different suites of tools.

Still in the consumer space, we now have examples of a commodity, functional Augmented Reality (AR) headset systems. The most prominent of them is the Microsoft *HoloLens* (see Figure 1). The HoloLens represents a significant advance on previous systems in being a fully self-contained (no tether) wearable computing device with sufficient computing power to render stereo 3D graphics and to provide accurate optical spatial and gesture tracking through on-board

² <https://www.tiltbrush.com/>



Fig.1: Microsoft HoloLens is among a new generation of consumer-level Augmented Reality devices that feature interaction through gesture and voice recognition. *Courtesy Microsoft - used by permission.*

cameras and image processing. Other systems include Epson's *Moverio* and the *Meta* AR headsets.

Moving to more experimental interface devices, researchers are beginning to create fully dynamic physical display and interaction devices. MIT Media Lab's *InForm* system [50] is one compelling example, featuring a table mounted array of vertically actuated pins (see Figure 2). Potentiometers used for actuation can also detect input forces, and additional optical tracking of users' gestures is achieved via a Microsoft Kinect sensor. The creators give a number of mixed-reality and embodied interaction use-cases for such a system, including actuation of physical objects and devices placed on the surface, haptic input, and also data display and interaction. Input and output space being so closely intertwined, a sense of directness and engagement is created.

Another glimpse of possible future interaction technologies is offered by Google's experimental *Project Soli* (Figure 3). A radar sensor is used to demonstrate the accurate identification of fine, in-air hand and finger gestures not limited by the usual constraints of optical systems. The Soli radar sensor has been miniaturized to a single chip that can easily be incorporated into many different devices. Such devices can potentially increase the vocabulary and the precision of control available for immersive analytics tasks as described below.




Fig. 2: MIT Media Lab’s InForm system is a fully dynamic tangible display and interaction device with many applications. *Courtesy Daniel Leithinger, Creative Commons License.*




Fig. 3: Project Soli from Google is developing an experimental hardware device that uses radar to identify precise gestures. *Courtesy Ivan Poupyrev - used with permission.*

4.2.2. What Are Some Benefits of Immersion?

Talking about interaction for immersive analytics, one should not only ask how interaction can support visual analytics but also how *immersion* may improve interaction. In the following we look into some of these potential benefits.

Immersion → Presence → Engagement. Since at least the time of Slater and Wilbur’s [135] work, the VR literature has made a distinction between *immersion* and *presence* (See Chapter 1). Presence is the user’s sense of “being in” the virtual environment. Immersion is used to describe the technological affordances of the system that determine the degree of presence experienced by the user (*technological immersion*, see also the extended discussion on immersion and engagement in the introduction to this book). It is then assumed that greater immersion (or alternatively, “immersiveness”) leads to a stronger sense of presence in the user, which leads to them being more engaged in the virtual experience [42].

In computer game design (psychological) “immersion” is used more freely as a synonym for “engagement,” for example, as discussed by McMahan [110].

Further, it is understood to apply not only in games that are *spatially immersive* (by providing the user a first-person 3D view of the world), but in any game that is highly engaging. Games researchers differentiate additional dimensions of immersion independent of the spatial setting. For example, sensory-motoric immersion as identified by Bjork and Holopainen [16] is a function of the degree of ease with which the user can direct or interact with the game such that their motor skills are rewarded by immediate and appropriate feedback.

It is such questions of immersion related specifically to interaction, and ultimately how they affect people's presence and engagement with data analysis tasks, that are the concern of this chapter.

Embodiment. In his 2001 treatise [45], Paul Dourish gives a phenomenological motivation for exploring human-computer interaction paradigms that are *embodied*. Embodiment may be through physical affordances encoded in the design of the interface or the realization of the computing interface as a part of the built or social environment. While this is largely a philosophical argument rather than psychological, it is a compelling thesis that, since we are creatures evolved to thrive in a physical world, we are most engaged and effective when acting physically. Closely related to this are themes such as *proprioception* and *spatial memory* that describe how we perceive and understand our bodies and our environment. We will revisit this in Section 4.4. when we look into *Reality Based Interaction* and *Physical Navigation*.

Mobility. Until relatively recently, immersive and natural user interfaces were characterized by large, heavy lab equipment such as tethered AR and VR headsets, projected large walls (such as CAVE), or tabletops supporting multitouch interaction. The breakout technology for multitouch turned out not to be tables or walls, but small portable devices like smartphones and tablets. Most recently, untethered (in fact, fully self-contained) AR headsets like the HoloLens offer the promise of a new type of portable immersive computing. Long imagined scenarios for Augmented Reality are finally achievable with commercially available technology. An example is offering mechanics X-ray vision by overlaying a CAD model directly over the machine they are repairing [112], wherever they are, even outdoors or down a mine. In the future, AR headsets might also allow immersive information retrieval for serendipitous, context-aware search interfaces [25] and, in general, the possibility of *in-situ* data analytics in all kinds of new situations, as described in Chapter 7 on situated analytics.

With Augmented Reality and high-quality hand and environment tracking, every surface can become an interactive display. For example, where currently operations centers require workstations with large and expensive screens and interaction is typically done via conventional keyboard and mouse, in future all that may be required is for an operator to wear a headset. They can then surround themselves with arbitrarily large display “walls” and, for example in the case of emergencies, work from anywhere. Thus, the advantage in this example scenario is not primarily immersion but simply cost and convenience.

Collaborative visualization. One of the most exciting aspects of new display and interaction technologies is that they can potentially make it easier for people

to work together. For example, head and body tracking offers remote collaborators the possibility of much greater presence in a shared environment. In the context of collaborative data exploration tasks, immersive display and interaction technologies offer users new ways to share and interact with information displays or to have their own private views and work areas. On the other hand, social interactions that follow from collaboration and make use of our social awareness and skills [73] can also increase user engagement and the feeling of presence in comparison to isolated desktop interfaces. However, many challenges remain for researchers in realizing this vision for collaborative data analytics, as described in Chapter 8 on collaboration.

Hands-free opportunities. Through voice or gaze control [115], for example, medical professionals can explore patient historical data, imagery or other records to better inform on the spot decisions, even during surgery. Similar hands-free interaction can offer machine operators or engineers additional diagnostic information via an HMD.

The ‘ultimate display.’ Sutherland’s original vision imagined a distant future where: “The ultimate display would, of course, be a room within which the computer can control the existence of matter” [144]. While we have made advances in (for example) 3D printing, such technology remains out of reach. However, film-makers and animators (e.g., [19]) are beginning to imagine the potential of VR and AR that is indistinguishable from reality - which is now not such a remote future possibility.

4.2.3. Desirable Interface Qualities

There are many guidelines available for the design of user interfaces, probably the best known of which are those defined by Norman in his book “The Design of Everyday Things” [118]. While these are valuable design principles in general, the challenge is to reconsider them in the light of immersive and natural interaction. The focus of this section cannot be to state a complete list of requirements but only to suggest interface qualities that system designers may wish to consider.

For instance, *affordances* and *signifiers*, often misunderstood when talking about (virtual) user interfaces [116], become even more important when the user is interacting with an AR system that combines physical and virtual artifacts, and where the affordances of a real object may not always match the interface’s signifiers. Similarly, in a system using VR headsets, *constraints* have to be designed with both the visualization itself in mind as well as the physical objects that naturally constrain the user’s movement.

Elmqvist *et al.* [47] define fluidity in the context of visualization as an “elusive and intangible concept characterized by smooth, seamless, and powerful interaction; responsive, interactive and rapidly updated graphics; and careful, conscientious, and comprehensive user experience.” They hypothesize that interactions that exhibit these types of characteristics will help users to stay in the flow, which is closely related to achieving a deeper sense of immersion. Thus, *Fluid Interaction* is particularly interesting in designing interaction for immersive analytics.

Elmqvist *et al.* present several design guidelines for creating visualization tools that support fluid interactions.

Use smooth animated transitions between states. Used appropriately, animated transitions can help the user maintaining a mental model of the system by avoiding abrupt switches between different modes. In the context of immersive systems, animations have to be handled with care. Forced viewport transitions, e.g., can easily lead to cybersickness.

Provide immediate visual feedback during interaction. Visual feedback should be provided in real-time for every interaction, including low-level interactions such as pressing a key or moving the mouse. This is especially true for continuous, direct input used in immersive analytics, such as head and body movement, gestures, or tracked 3D input devices.

Minimize indirection in the interface. In line with the concept of natural interaction, direct manipulation should be used whenever possible so that interaction operations become part of the visual representation instead of being separated out in control panels.

Integrate user interface components into the visual representation. If direct manipulation is not possible, try to integrate traditional user interface elements seamlessly. In immersive systems this also means to integrate 2D UIs in 3D environments. This poses challenges as to where to put them and how to cope with perspective distortion, occlusion, and text readability.

Reward interaction. Initially, the novelty of immersive systems may motivate users and thus “reward” them. However, this will quickly wear off with continued use. Interaction embodies a dialog between the user and system. Visual indications about how and when interactions can be performed can keep the user stimulated over the course of the exploration and act as rewarding effects.

Ensure that interaction never ‘ends.’ Continued exploration should always be possible and the user should never reach a dead end. At least regarding navigation, this is ensured in any immersive analytics systems in which the user freely moves through some virtual information space, e.g., in VR or AR environments.

Reinforce a clear conceptual model. Operations should be reversible to allow the user to return to a previous state and thus keep a clear idea of a system’s state.

Avoid explicit mode changes. Mode changes may break the user’s flow and should be avoided. If possible, all operations should be integrated and accessible in the same mode.

In addition, there are also *technical requirements* to immersive systems. These include a decent, real-time responsive tracking of the environment and, if applicable, good multi-modal interaction – voice, touch, gesture, eye-gaze, or head tracking. One area in which interaction technologies have seen major improvements that directly affect their immersion is in terms of *latency*, that is, the delay between the user initiating an action and receiving feedback from the system.

Other challenges include scalable interaction, issue of attention management and awareness, as well as unifying multiple device classes while coordinating interactions across devices.

4.3. Basic Tasks for Analytics

This section describes fundamental tasks for interactive visualizations and refers to existing task taxonomies in the literature. These tasks need to be supported by an immersive analytics user interface, and we will use them to anchor the discussion on designing interactions for immersive analytics in Section 4.5..

There exist several taxonomies that capture the individual activities users perform when using visualizations for exploring and understanding any kind of datasets. Most of these taxonomies cover a related set of concepts: Amar and Stasko [2], for instance, describe a set of low-level components gathered by observing students interacting with visualizations, while Heer and Shneiderman [55] break down that interaction into view specification and manipulation. In this chapter, we use the latter taxonomy as a basis for exploring the design of interactions for immersive analytics (see Section 4.5.) due to its focus on interactive dynamics in visual analytics. We review that taxonomy in the following paragraphs and refer the reader to the literature to learn more about alternative interaction taxonomies that range from more general (e. g., [98,154]) to more specific taxonomies (e. g., with a focus on multivariate network analysis [122]).

Heer and Shneiderman [55] distinguish between three high-level categories: (1) *data and view specification* that essentially corresponds to the data transformation and visual mapping steps in the Information Visualization Reference Model proposed by Card *et al.* [30], (2) *view manipulation* that is reflected by the view transformation step in the Information Visualization Reference Model, and finally (3) *process and provenance* for supporting analysis processes.

The category *data and view specification* can be further subdivided into the following lower-level tasks.

Encode/Visualize. Here, the analyst chooses a suitable visualization or visual representation for the data records, i. e., this is the actual creation of the visualization.

Filter. A visualization becomes more scalable by filtering out data records based on specific criteria. A popular interaction technique is dynamic queries, for example.

Sort. This task corresponds to the already mentioned basic interactions mentioned before. Sorting data items is good for many purposes, such as item lists ordered according to measurements (centrality measures, degree-of-interest, etc.).

Derive. The task of deriving information from the primary input data is more related to visual analytics than to pure visualization. For example, the analyst might want to integrate results from automatic computations (aggregations, clusterings, simulations, etc.) into the visual display.

Kerren and Schreiber [85] extended this category by two additional tasks that both address typical interactions needed in visual analytics:

Reconfigure. Changing or reconfiguring the graphical display for specific data records during the analysis process is a standard task performed by analysts.

Adjust. When deriving secondary data from the primary input data, the analyst usually has to modify parameter settings of automated analyses.

The second category discussed by Heer and Shneiderman is *view manipulation* as summarized below.

Select. Selection is a fundamental interaction concept that is often used prior to a filter operation.

Navigate/Explore. One of the core tasks that reflects very well the explorative character of Information Visualization. Navigation and exploration approaches often follow the information-seeking mantra *overview first, zoom and filter, then details on demand* [134] and encompass methods such as focus & context, zooming & panning, semantic zooming, etc.

Coordinate/Connect. In visual analytics systems, analysts are often enabled to relate views or even individual graphical entities to each other. Brushing and linking can then be used to build a bridge between connected views/items.

Organize. More complex visual analytics tools consist of many different components, views, or tabs. These components have to be arranged and grouped on the screen in such a way that the resulting arrangement supports the analysis process efficiently.

The last category on *process and provenance* goes beyond the most traditional information visualization taxonomies as it addresses typical issues that are more related to supporting the analysis process in general, and not tasks specifically related to interactive visualization. We briefly describe the tasks in this category here following that of Ragan *et al.* [124], but cover them in more detail in Chapter 5 where we discuss them in the context of *visual analytics*.

Record. Visualization tools should keep a *record* of the analysts' operations throughout the analysis process. Then, various interactions should be provided to allow them to step back through this record in order to understand the provenance of individual findings. At the very least such stepping is enabled by *undo/redo* operations, but more sophisticated analysis of the provenance, e. g., visualization of operations against a timeline, may also be beneficial.

Annotate. Adding labels, highlights or other annotations to visualizations allows analysts to keep track of significant insights in the context in which they were found. This is useful for explaining the results of an analysis to others, but also, again, to support the analyst's own working memory of the provenance of insights.

Share. Visualization can be a great aid to collaborative analysis, enabling communication about data as well as visual inspection. Thus, tools should support the sharing of various views together with the provenance history and annotations described above in order to make such collaboration as easy

as possible. Collaborative visual analytics in the context of immersive systems is described in great detail in Chapter 8.

Guide. A particularly important use-case of shared visualization is guided explanation of insights within a dataset. To support this scenario, tools should enable generation of visualizations that are static yet self-explanatory (i. e., like the kind of infographics found in news-media reports), animated, or stepped interactions for guided presentation. Enabling such *storytelling* scenarios as effectively as possible is becoming a research area in its own right and is discussed in greater detail in Chapter 6.

We will instantiate a subset of these tasks for the concrete case of immersive analytics in Section 4.5.: select, filter, sort, navigate, reconfigure, and annotate. The next section introduces basic novel interaction modalities, their opportunities and limitations, and their contribution to engagement and immersion.

4.4. Natural User Interfaces (NUI) and Post WIMP Interaction

Interaction techniques for immersive analytics have to be powerful enough to support users in all the tasks mentioned in the previous section, while at the same time not distracting them from their work or breaking immersion. Techniques based on the WIMP metaphor or most classic 3D input devices (see, e. g., [98]) are usually not suitable for interaction in immersive environments. Several of their limitations were listed by van Dam [43]: WIMP interfaces are increasingly complex, can lead to screen clutter, and are inherently indirect, especially for 3D interaction. Maybe even more important for immersive analytics, they are primarily designed for use on desktop computers and do not make use of our rich senses and input capabilities. Novel post-WIMP interfaces are filling this gap. They are typically defined as containing at least one non-classical form of input (e. g., touch, gestures, speech, or tangibles), mostly process parallel input streams, and often support multiple users. They are closely connected to the concept of *Natural User Interfaces* (NUI) introduced earlier in Section 4.1.. Dachsel & Preim define NUIs as “those interfaces where people interact with by means of intuitive and mostly direct actions grounded in real-world, everyday human behavior. Natural does not mean innate but learned and familiar actions, which appear appropriate to the user in the moment of interaction.” [121, p. 472] (translated by Dachsel & Preim). While there is certainly some overlap between the terms post-WIMP and NUI, post-WIMP is a term defined by the forms of input [43] but NUI is about creating an interface that “makes your user act and feel like a natural.” [152, p. 14] It should also be noted that what makes an interface “natural” is debatable, and the term *Natural User Interface* itself has also been criticized, e. g., by Norman [117].

An overview of NUIs as well as guidelines on how to design them are given by Wigdor & Wixon [152] and Fikkert *et al.* [49]. One framework of post-WIMP interaction that is closely related to NUIs is the concept of *Reality Based*

Interaction [73], which was proposed by Jacob *et al.* in 2008. They identify four main themes from the real world that inform and inspire the design of post-WIMP interfaces:

- *Naïve Physics*, informal, common sense knowledge about the physical world, including aspects such as inertia, gravity, etc.
- *Body Awareness & Skills*, the user’s ability to control and be aware of their own movement and the position of their limbs (proprioception).
- *Environment Awareness & Skills*, e.g., the ability to select or grasp objects, and estimate distances or sizes of objects.
- *Social Awareness & Skills*, knowledge about social protocols and the ability to collaborate with other people.

In the specific context of Information Visualization, post-WIMP interaction was analyzed both by Lee *et al.* [100] and Jansen & Dragicevic [78]. Lee *et al.* examined design considerations for “natural” information visualization interactions regarding the individual, the technology being used, social interactions between users, and the relation of technology and the user. Based on those considerations, they also identified opportunities and challenges for NUIs in information visualization. Jansen & Dragicevic adapted and extended the well-known information visualization pipeline for visualizations that go beyond the desktop. They also proposed a visual notation for interactive visualization systems that they applied to different case studies of post-WIMP visualization setups.

In the following, we describe different relevant interaction modalities in more detail. While we focus on NUI, it should be noted that there is also a large body of more “classical” 3D interaction research, including the fields of VR and AR. Here, more traditional 3D input devices such as ART’s Flystick, VR controllers, or even Wii Remotes are being used. These input devices, especially when bundled and natively supported by commercial VR headsets, provide stable and tested means of input and are often well-suited for immersive 3D interaction. Besides the general overview of 3D user interfaces, including 3D input devices, by LaViola *et al.* [98], an overview of (non-immersive) 3D interaction is given by Jankowski and Hachet [76]. A recent survey of Augmented Reality, including a chapter on interaction, was written by Billinghurst *et al.* [15]. In Section 4.4.2., we discuss typical limitations of 3D interaction.

4.4.1. Interaction Modalities

Let us first introduce and discuss the major non-classical interaction methodologies that are commonly used in immersive visualization. Most of them are inspired by how we typically interact with everyday objects, hence the label “natural interaction techniques” sometimes attributed to them. In our case, however, we are interested mostly in how well they support immersion in the visualization task and how well they maintain (and do not interrupt) the flow.

Touch-based (tactile) interaction. Along with the surge of touch-sensing display technologies and hardware of the past two decades came a plethora of

work on its application to interactive data exploration. Because a two-dimensional touch surface lends itself intuitively to the manipulation of 2D representations, a large number of touch interaction paradigms were explored for visualization sub-fields that represent data typically in the two-dimensional plane [100]. However, interaction techniques for inherently three-dimensional datasets have also been explored [67, 84], either for monoscopically (e. g., [52, 105, 155–157]) or stereoscopically displayed data (e. g., [38, 72, 103]). The benefits of haptic feedback are typically cited as a benefit supporting the use of touch interaction as opposed to traditional PC-based environments. In particular, somesthetic feedback [129] gives people the feeling that they actually are manipulating items and, thus, increases the immersion into the interaction with the data. A survey of touch interaction techniques, even if just for the application of data visualization, would be beyond the scope of this text, so we refer to the recent survey by Isenberg & Isenberg [66].

Sketching and pen Interaction. Pen interaction with computers actually predates mouse interaction. Widely regarded as a precursor to hyperlinking (and eventually, therefore, the web), Vannevar Bush’s 1945 concept for an electromechanical document exploration machine (the *Memex*) featured a stylus for annotations [27]. Sutherland’s Sketchpad system [143] was an early demonstration of interactive computer support for technical drawing built around a light-pen system. However, the practicality and ergonomics of holding a tethered pen up to a vertical screen for extended periods of time meant that pen computing took a long time to find favor. Chen *et al.* studied the use of a hand-held tablet as a remote control for viewpoint control and search tasks in an immersive CAVE environment [36]. They found that the tablet’s mobility helps reduce context switching cost due to the form factor of the tablet input. In more recent years, the development of tabletop and smart board computing, smaller tablet devices, and the gradual advancement in machine recognition of hand-writing and drawings has led to renewed interest in pen-based interaction (e. g., [51, 59]).

Much more recently, information visualization researchers have begun examining how people communicate about data and complex systems by sketching on whiteboards. In a study of whiteboard sketches found throughout a large organization, Walny *et al.* found many instances of whiteboards being used to manually create quite sophisticated information visualizations, such as charts and diagrams [150]. Clearly, whiteboards are ubiquitous in workplaces and classrooms, but besides being cheap and convenient there are a number of other aspects that make them ideal in supporting people as they communicate complex ideas, for example, people seem very comfortable organizing their thoughts spatially on a whiteboard in a very freeform way. The observation that people routinely think and communicate about data visually by sketching has led to automated systems for integrating computer-generated charts into people’s hand-drawn sketches.

Tangible interaction and data physicalization. As an alternative to touch interaction, tangible user interfaces, first introduced in the 1990s [70, 133], have also been explored for the control of data exploration. Early examples focused, in particular, on tangible controls for traditional interfaces (e. g., [58, 71])

and on geographical visualization (e. g., [69,119]). For the exploration of spatial data, in particular data defined in 3D space, several approaches demonstrated that tangible interaction can be useful by maintaining the spatial mapping between the exploration tool and a physical representation of the data (e. g., [13, 53, 94, 95]). Some recent work has even demonstrated the use of shape(-changing) displays (e. g., [102]). Beyond spatial data, however, a number of authors has also demonstrated the usefulness of tangible interaction with abstract data representations (e. g., [37, 77, 93, 137]). A closely related form of interaction is the exploration of physical visualizations [79]. In fact, it can be argued that there is a continuum from completely virtual/digital, via tangible interaction, to physical interaction [77].

The advantages of tangible interaction over, in particular, touch interaction lie in the physical nature of the manipulated tokens, facilitating a richer expression of interaction intents than touch interfaces. On the other hand, most tangible interfaces are highly specific, inferior to the flexibility offered by fully digital touch interfaces. The previously mentioned shape displays [50, 102] as well as recent micro-robotic technology (e. g., [99]), however, promise to make tangible or physical data displays more flexible.

Gestural interaction. Gestures are a form of non-verbal communication, often accompanying or sometimes even replacing speech. We use them, for example, to point at objects, to illustrate concepts, or to support cadence. Many taxonomies for gestures exist, such as by McNeill [111] or Karam & schraefel [82]. With gestures being part of our daily lives, gestural interaction systems can feel natural and intuitive. As such, gestures have been proposed for HCI at least since the 1980s [17]. An advantage of such interfaces is that they allow manipulation from a distance without physical contact or the use of any hand-held input device. Hence, they lend themselves for settings such as large display installations [109], public displays [151], or medical applications demanding sterility [80].

An early example in the field of data visualization is the *Information Cube* by Rekimoto & Green [126], a cuboid visualization of hierarchical data. Users wearing datagloves can rotate the visualization to inspect it from all sides and select nodes using hand gestures. Kirmizibayrak *et al.* [87] used gestures to control a visualization of medical volume data. Their studies show that gestures can outperform mouse interaction in rotation tasks. Chen *et al.* designed a set of numerical input techniques to model architectural modeling process [34], ranged from virtual sliders [35] to virtual keypad and gestural input [33]. Hybrid interfaces that combine 2D multi-touch and 3D gestural interaction have been examined by Benko *et al.* [12]. New and affordable depth-sensing cameras such as the *Kinect* or the *Leap Motion* allow gestural interaction to be included in consumer products and, at the same time, also provide researchers with the means to develop low-cost prototypes for this interaction modality. Despite this, several challenges remain. Gestural interaction is often perceived as unergonomic and can easily lead to fatigue (casually called the gorilla-arm effect [57], see also Subsection 4.4.2.). Similar to touch gestures, free-hand gestures also lack discoverability and usually have to be learned and trained for effective use. Finally,

one has to acknowledge that gesture recognition accuracy affects user performance and is also negatively affected by the strategies that users adopt to deal with recognition errors [5].

While we often associate it with hand gestures, the term “gesture” can be applied to many other forms of input. For example, foot gestures [106], touch gestures (see above), and even gestures made with physical props, pens, 3D controllers, or other devices (part of tangible interaction, see above) can be found. However, the use of the term “gestural interaction”—in particular in the context of data visualization—has also been criticized [68]: both the task specification and the parameter manipulation during exploration can and are specified using gestural interaction, but both have different constraints. In particular, for immersive interaction, the specification of the task should happen in an instance in order to not interrupt a user’s thought processes, and the transition from task specification to parameter control needs to be fluid. An example for this are the *MultiLenses* by Kister *et al.* [89]. There, graph lenses are instantiated with a gesture and the main parameter can be immediately manipulated by continuing this gesture in a fluid way.

Gaze interaction. Eye-tracking systems are often used for psychological studies or to evaluate user interfaces; an overview of eye tracking specifically for the *evaluation of visual analytics* is given in [97]. However, already in the 1980s, the first *gaze interaction* interfaces were presented. One such system is Bolt’s *Gaze-Orchestrated Dynamic Windows* [18]. In this system, several video streams were presented on a single large display. By looking at a particular video, the user could start or stop the playback.

The main advantage of gaze interaction is that it does not rely on the user’s hands. As such, gaze interaction is regularly used in systems for the physically disabled (e. g., [64]). However, as Bolt’s system shows, gaze can also be used as a natural interaction technique to indicate a person’s attention or simply keep their hands free for other tasks.

Early eye trackers were cumbersome and required the user to be stationary. However, advances in mobile eye tracking hardware now make it possible to, e. g., integrate them into commercial VR headsets³, enabling their use in immersive analytics applications. Measuring (and reacting to) the user’s attention is a key aspect of *Attentive User Interfaces* [148], an interface strategy that tries to address information overload. Also, visualizing a user’s awareness can be helpful in collaborative systems. For example, coarse eye tracking is used by Dostal *et al.* [44] to measure awareness in a collaborative visualization system using large, wall-sized displays.

When gaze interaction is used directly as an input method, it is often combined with other modalities. While it is possible to design interfaces that solely or mainly use gaze (e. g., Adams *et al.* [1] presented techniques to explore large 2D information spaces using gaze), both precision and efficiency can be hampered by gaze interaction. Rapid, small eye movements (saccades) limit the precision of the eye gaze. The duality of using eyes as an input channel versus their normal task as

³ See, e. g., <https://www.tobiipro.com/product-listing/vr-integration/>

the most important human sense is another limitation of gaze interaction, leading to problems such as the so-called *Midas Touch*: Do users of a gaze interaction interface look at an object or UI element to trigger an action after a certain dwell-time or just to visually inspect it? Additional modalities, e. g., touch, are used to circumvent these limitations. We call such interfaces *gaze-supported* [139]. Stellmach *et al.* [140] presented a system for the gaze-supported exploration of large image databases. Employing a large remote display, gaze was used to control a fisheye lens and fine target selection was done using a multitouch device. A combination of gaze and foot input was used by Klamka *et al.* [91] to explore zoomable information spaces, in this case Google Earth. Still, gaze-based pointing methods, including gaze-supported input with the (space) button for confirmation, can currently not match the performance achievable with a mouse for pointing tasks [107] on smaller screens.

An overview of eye tracking technology and applications including many more examples of gaze interfaces can be found in Majaranta & Bulling’s survey [108]. Cernea and Kerren provide a survey article on technologies for emotion-enhanced interaction that also discusses basics of eye tracking technologies [32].

Physical navigation and locomotion. A key advantage of immersive analytics is the ability to display a large amount of information, potentially across multiple levels of scale. This provides the opportunity to navigate this information efficiently through physical or virtual means. Information foraging theory [120] models the cost of information access, and suggests office and computational metaphors in which frequently accessed information is kept nearby and can be accessed through quick operations, at the expense of less frequently used information stored at greater “distances”. The concept of “information scent” models the visual hints that guide the user in finding and accessing targeted data in an information space.

Physical navigation exploits human embodiment to navigate the information space directly via body movement, such as head rotation to naturally control the view frustum. In contrast, virtual navigation uses some form of indirect interactive control to manipulate the viewpoint, such as a joystick. Beyond this binary categorization, a spectrum of fidelity of interaction ranges from high-fidelity that mimic real-world interactions to low-fidelity that do not attempt to mimic real-world behavior [98]. Jakobsen *et al.* [74] explored interaction techniques based on the spatial relationship of a user and large visualizations, including the adaptation of the visualization based on the user’s position.

In immersive analytics, physical navigation spans two primary levels of scale. Micro-level physical navigation consists primarily of efficient rotational movements of the eyes, head, and body, to rotate the view frustum in the space, and can also include limited translational movements, such as leaning. These forms of physical navigation are very efficient and have been shown to be advantageous over virtual navigation, in terms of both access time as well as improving cognitive understanding of spatial relationships in the data [9]. This is typically implemented using head-mounted displays combined with head tracking, or large high-resolution fixed displays with or without head tracking.

Macro-level physical navigation deals with translational movement for locomotion in the space. Physical navigation for locomotion, such as walking, is typically limited due to the range of the tracked interaction space or the physical size of the display device. Exceptions include Augmented Reality methods based on GPS or computer vision in which users navigate the real world. Hence, longer range locomotion in immersive environments often requires some form of simulated physical navigation, such as walking in place or treadmills, or virtual navigation methods such as teleportation or joystick control [98].

4.4.2. Challenges and Fundamental Limitations of 3D Interaction

In many cases, immersive analytics applications may involve 3D data sets or 3D input and output technology. While the idea of interacting directly in 3D is attractive, many currently available solutions are often not practical and/or inefficient. In this subsection, we list a couple of common issues around 3D interaction that affect interaction performance in non-trivial ways.

One commonly held thought posits that human or technology limits could be easily addressed through navigation or better interaction methods. Yet, this is rarely true in a general sense. For example, interaction-at-a-distance, i. e., ray-based pointing (e. g., [81]), is fundamentally limited by the fact that pointing accuracy is negatively affected by limitations of orientation tracking or human limits on hand stability. Then, any inaccuracies in orientation result in errors that increase with distance. While techniques such as snapping the ray to the nearest target [4] can compensate for this in some instances, such solutions do not scale to dense object arrangements or work for large distances. Moreover, the cost for undoing incorrect selections through snapping is also often ignored. Disambiguation interfaces that temporarily magnify content to facilitate selection achieve this through an increase in interaction cost. Another way to compensate for pointing inaccuracies is to navigate to a position where the problem becomes simpler, e. g., by moving closer to the desired objects. This is true, but this solution adds the cost of navigation to the interaction, which again increases time. The user may then also have to spend additional time to navigate back to the original position. For interfaces that mix interaction and navigation modes, there is also the insidious issue that the need to switch between interaction and navigation increases the potential for mode errors.

Spatial abilities, especially for 3D, are learned skills, which include spatial perception (important for 3D navigation), mental rotation (important for 3D rotation), spatial visualization (important for planning 3D operations), and spatial working memory (important for 3D creation). Such abilities correlate well with STEM (Science, Technology, Engineering, and Mathematics) performance [149], which is only high in a limited subset of the general population. As the spatial abilities of the general population are below the levels observed in the STEM subpopulation, this fundamentally affects the potential generalizability of any interface that requires good spatial interaction skills.

Another fundamental issue is that one cannot interact with occluded content with any degree of precision, as the feedback loop that is central to human

actions (and HCI) is disrupted. This is commonly visible in the real world, where only highly-trained users can manipulate things they cannot see directly with precision. Again, one can compensate for this through interaction or navigation techniques, but only at additional cost.

The human visual system can only focus on a single plane at any given time. In practice this means that one can either focus on the finger/hand or on a display, but not both simultaneously (unless they are close in depth), e. g., [23]. This fundamentally limits interaction accuracy that is achievable away from a display surface, including common stereo display and Augmented Reality systems. A related issue is the vergence-accommodation conflict, e. g., [60], which is the inability of commonly used stereo display systems, including head-mounted displays for Virtual Reality, to allow correct eye-vergence and focal distance for the perceived objects. This conflict affects depth perception, which in turn impacts interaction performance negatively. An easily observable symptom is that users start to “search” for the correct depth of a target in stereo display systems or that the distributions of hit points for 3D target shows an elongated profile in the depth dimension [24], which affects pointing performance [22]. Finally, any form of human stereo vision deficiency also affects interaction performance with systems that rely on accurate depth perception through stereo. Some researchers estimate that more than 30% of the population may have some form of stereo deficiency [56]. Moreover, as people age, average stereo acuity decreases, too [160]. This puts some limits on stereo-based systems being a general form of user interface mechanism.

Throughput is a good measure to characterize interaction performance, as it takes both speed and accuracy in pointing tasks in to account. Thus, throughput enables comparisons across different user strategies, such as slow and accurate vs. fast and sloppy. In terms of throughput, a mouse or touchscreen is significantly better than any currently available 3D interaction device (see, e. g., [20,21,23,146]). While interaction speed in 3D is usually not much affected, interaction accuracy is typically much worse. The list of reasons for this include the lack of human hand stability in mid-air interaction, 3D tracking limitations, depth perception limitations, or any combination thereof. Thus, it is a particular challenge to achieve interaction that feels natural yet is accurate enough.

Interaction in 3D typically involves more than two degrees of freedom (DOF). One good option to address this issue is to reduce the number of degrees of freedom through the use of constraints [132,141,142]

Another issue affecting interaction with 3D is the “gorilla-arm” syndrome, i. e., the fact that holding one’s hands in the air for extended periods of time leads to fatigue and all its associated side-effects [57,75].

Finally, interaction through a single cursor, pen-tip or even multiple touch point(s) on a touch screen is much less rich compared to the interaction capabilities posed by human hands. One illustrative example is achieving the task of tying one’s shoe-laces with any computer-based system.

4.5. Designing Interactions for Immersive Analytics

After describing the typical tasks associated with visual analytics in Section 4.3. and discussing different interaction modalities and their limitations in Section 4.4., we can now combine them and map techniques to tasks. Guidelines such as presented in Section 4.2.3., e.g., Elmqvist’s fluid interaction concepts [47], can be used to inform the design of these interactions. This section outlines a first attempt at matching tasks to interaction techniques. Where appropriate, we explain how fluid interaction can be supported by these mappings. It should be noted that the following should only be seen as initial proposals, showing the opportunities of rich, fluid interaction for immersive analytics.

Selection: Selection is the basis for determining on what elements many subsequent analytic tasks operate. These tasks can include selection of data for more information, reconfiguring of elements, or sorting along a dimension. In immersive environments, techniques such as mouse interaction are not appropriate. Gestural interaction (e.g., pointing to and circling an element) can be used, though mid-air gesturing can be fatiguing. Great care should be taken to minimize strain on the user, e.g., by designing comfortable gestures [57] or providing alternative forms of input. For example, voice along with Natural Language Processing (NLP) can be used to indicate attributes and ranges of interest (e.g., ‘Select all elements with negative profit’) but is itself less suitable for selecting or deselecting individual elements. Eye-gaze, on the other hand, is well suited to indicate the user’s point of regard and can be used to choose an object of interest, though *Midas Touch* usually means that other modalities would need to be engaged to choose the appropriate interaction (e.g., a button press on a VR controller or a voice command). Multi-selection techniques such as lasso selection can in principle be used with each modality that provides a form of pointer, e.g., touch, pen, mid-air pointing, or to some extent gaze. These techniques can provide *immediate, continuous feedback* [104]. It may not be enough to just separate out selections from the subsequent activity since some interactions may involve combining the act of selecting with that subsequent activity to design more fluid interactions [46]. Some work has also been done in combining multiple devices (touch displays, tablets) with immersive displays for more fine-tuned control.

Filter: If items to be filtered are already known, filtering is simply the triggering of an operation and can be accomplished through a semantic gesture (e.g., a wave of the hand), or voice. Filtering often is either used to filter out selected objects, or filter out all non-selected objects. The act of specifying the query may be combined with the interaction itself (e.g., from voice, we might filter out all values less than the average or focus only on outliers). On the other hand, in some cases, filtering needs to be dynamic. For example, a user would like to remove outliers to remove visual clutter in a complex dataset but does not yet have a clear understanding of what constitutes outliers in the specific data. In such use cases, *continuous* changes of the filter threshold and *smooth transitions* between filtered and unfiltered views may be required to keep

the *flow of interaction*. Many natural interaction techniques, like hand or body movements, support continuous input with, however, varying degrees of precision.

Sort: Typically, a sort action is specified by an attribute upon which the data is to be sorted. This attribute may already be selected (see above), or selected during the sort. Some work involved a gesture on a surface or in mid-air, by starting the gesture on a particular axis [46]. Voice can be used as well. Sorting typically involves specifying whether the order is ascending or descending by value, or in some other order (e.g., alphabetical). Using *smooth transitions* to the new order can help users to track both individual items and the general arrangement. While these animations can also easily be triggered after discrete input events (such as a voice command or a button press), continuous input techniques allow for immediate, continuous feedback and may help users to abort erroneous sorting operations.

Navigation: There is a wide history of research on navigating through virtual environments (e.g., [98]) and many of these techniques apply here. Physical navigation, i.e., the user literally moving to change the view on the data, has been examined, e.g., for wall-sized display setups [9, 90]. Walking closer to data in a space can allow precise inspection of details while walking back allows for an overview. Additionally, Rädle *et al.* [123] studied the effect of physical navigation in zoomable user interfaces on large display walls and found not only benefits for navigation performance but also evidence of improved spatial memory. Limited space can be problematic for physical navigation in large datasets. Especially in VR, techniques such as teleporting the user to a new position can help to increase the apparent size of the interaction volume. These techniques, however, are potentially immersion breaking and can negatively affect the user's *mental model*. Gestural interactions can be used such as to shrink or zoom the overall world, or rotate it to a more appropriate view. Maps can be used in conjunction with a visualization to assist in navigating a large visualization.

Proxies that represent the data and make it more tangible can be used if orienting the data is of particular use. For example, Beaudouin-Lafon *et al.* [11] used a prop of a human brain to rotate MRI data on a large display wall, similar to those used by Hinckley *et al.* [58]. Similarly, hand-held devices that serve as peepholes into the virtual scene (e.g., [26, 137]) can be moved around to navigate the visualization.

Reconfiguration: Reconfiguration often includes assigning attributes to visual variables, e.g., to the axes of a visualization. While voice interaction can in principle be used for this task, this requires the correct recognition of spoken, previously unknown data attribute names. A good example of gestural reconfiguration is the ImAxes [39] project. In this HMD-based VR system, users can rearrange data axes with the VR controllers to compose different visualizations, e.g., scatterplots or parallel coordinates. This demonstrates how a close coupling between the user interface and the visual representation can be beneficial: Instead of *explicitly* switching between modes, users

can *directly manipulate* the type of visualization and get *immediate system feedback*.

Labeling and annotating: Annotating data, assigning labels or highlighting parts of a visualization are important actions supporting the discussion and analysis of the data. Text input can be supported by voice recognition. Alternatively and depending on the setup, touch keyboards on personal mobile devices or displayed on a shared interactive display can be used as well. Pen input has also been proposed to annotate visualizations directly, e. g., using tangible lenses with Anoto technology [137]. An overview and a taxonomy of annotations in (outdoor) Augmented Reality has been presented by Wither *et al.* [153].

Given the mappings above, we can identify some fundamental interaction types that can be combined to create higher order interactions. We need to support selection of data and visualization elements. Similarly, we need to support discrete choices from lists. Whenever these items are graphically represented in the system, some form of *pointing* can be used to indicate/choose them, with options such as 3D input devices, touch, gaze, or mid-air pointing. Confirmation can then either use the same modality or a different form of interaction in multi-modal interfaces, e. g., gaze selection with voice confirmation. On the other hand, discrete items (including mode switches) can also be selected by more *symbolic, discrete input methods* such as voice commands, symbolic mid-air gestures, etc. Examples such as multi-selection or camera manipulation show the need for *continuous interaction*. Clearly, voice commands do not lend themselves for this type of interaction. Instead, touch, tracked 3D input devices, some forms of gestures, and physical navigation are good examples for continuous control.

4.6. Archetypal Setups

Systems for immersive analytics can take many and varied forms, including, for example, differences in the number, size, and types of displays used as well as input devices. In this section, some typical setups (i. e., combinations of output and input technology) that showcase this variety are described along with example research applications and upcoming trends. It is worth noting that many of these systems have, to date, seen more active use in Scientific Visualization applications than in the visualization of abstract data, i. e., Information Visualization.

4.6.1. Large Screen Collaborative Space

Large screens, especially wall-sized vertical displays and to lesser degree tabletops, provide the space needed to show complex datasets in their entirety. Their size also gives room for several people to interact with them at once. Therefore, setups of one or more large displays have been used extensively for collaborative visual analytics tasks.

This archetype has several challenges both regarding interaction and perception. One issue is that metaphors for interaction close to the display, e. g.,



Fig. 4: Distribution of a visual analytics system across different display geometries [92].

using touch, often cannot be applied directly to interaction from afar using, e. g., gestures. Another interaction challenge of such systems is to differentiate between the input of individual users. Different methods for touch user identification have been proposed, such as optical tracking in front of display walls [159], capacitive fingerprinting [54], or active IR beacons [130]. Awareness, both in regard to the display content and the collaborators' actions, is more of a perceptual challenge: While standing close to the display, other users and parts of the display can be blocked or simply outside the user's field of view. Stepping back to get an overview, on the other hand, prevents users from direct interaction with the display and may limit inspection of small details in the data as well.

In the last decade, the costs of display panels have fallen sharply. In 2007, LCD panels larger than 40" were not available commercially. Plasma panels larger than 50" cost in excess of USD 5,000 and weighed as much as a person (e. g., [131]). In 2017, full HD 50" LCD panels can be bought for only a few hundred dollars and weigh only a few kilograms. Ultra High Definition (UHD) panels with $3,840 \times 2,160$ resolution are only a little more expensive. Further, multiple software solutions now exist to provide *cluster rendering*, the capability to split the responsibility of rendering a continuous 2D desktop or 3D environment across a number of networked machines, each driving its own set of display panels.⁴ Thus, the cost and difficulty of creating tiled display walls with arrays of such panels have fallen dramatically and with this increased convenience, researchers

⁴ Popular cluster rendering software includes OmegaLib [48], Unity [147], and Sage2 [127].

have begun to investigate the opportunities for visual analytics using such display walls.

In an experimental study evaluating tiled vertical displays of various sizes for visual analytics tasks, Ball and North [8] found that small targets could be acquired and compared significantly more quickly on larger walls. A later, more qualitative study by Andrews *et al.* [3], found that analysts change their approach to managing documents and information when moving from a standard desktop display to a larger, tiled display. Basically, they use a more spatial approach to document organization, structuring the display space to externalize their analysis process.

In addition to the physical area, the other important property of display walls (compared to, for example, large projected displays) is their resolution. Ni *et al.*'s work is among the first to empirically demonstrate the benefits of high-resolution display for information search with text in spatial data [114]. With clustered rendering, resolution is theoretically unlimited. In practice, the cluster architecture has to provide a scalable way to distribute the data to the nodes and, from each node, to render only the part of the view needed for that node. Using this method researchers have been able to investigate new ideas to take advantage of massive resolution spread across a large wall. For example, Isenberg *et al.* [65] create static visualizations that allow inspection of very fine-grained detail when standing close to the wall, but on stepping further back, gross details (labels for larger regions and so forth that were blurred and therefore “hidden in plain sight” when standing close) become apparent (see Figure 5).

These results suggest that high-resolution, large-area display walls can change the interaction paradigm for data analytics, especially with respect to the task of navigation as described in Section 4.3.. That is, to some degree, the user(s) of large display walls are able to navigate by changing their physical position with respect to the data displayed, rather than changing the viewpoint through an interaction device. Arguably, the most “natural” user interface is no interface at all.

4.6.2. Personal Displays + Large Screen Collaborative Space

Setups consisting of one or more large displays can be combined with individual, personal displays. Typically, these personal displays come in the form of handheld smartphones or tablets. For example, several scenarios (and infrastructure to support those scenarios) are explored by Klapperstuck *et al.* [92] for multi-site collaboration, using a shared large area display wall together with handheld and tabletop devices for interaction. However, arguably any additional display (e.g., AR glasses or smartwatches) used by one person in addition to the main display(s) fits this category. An example is the system presented by Butscher *et al.* [28], which combines a large multitouch table with several HMDs for the collaborative analysis of multidimensional data.

This combination can help to address some of the challenges of large display immersive analytics systems: The use of personal devices and their own interaction capabilities provides an additional, independent view for the user and also allows

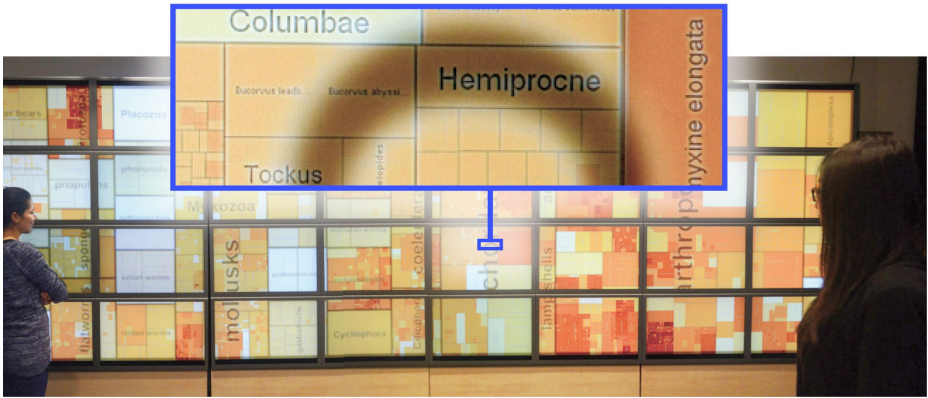


Fig. 5: Use of Hybrid-Image techniques to create combined overview and detail views in static large-wall visualization [65]. *Used by permission. Copyright Petra Isenberg.*

her to interact from a greater distance. For example, with a tablet used as a handheld lens into the data, a user could explore a visualization even if another user stands directly in front of the main display, blocking the view [88]. A user could also utilize the mobility of the tablet to reduce the context switching cost with multiple screens [36]. Additionally, as described above, these personal devices can be used to facilitate tracking and identification of the users, e. g., with RFID technology or using optical tracking such as in Google’s *Project Tango* (<https://get.google.com/tango/>). An overview of current technologies for mobile device localization is given by Horak *et al.* [61].

Still, challenges remain and new issues emerge. One problem of multi-display environments in general and for the combination with handhelds specifically are gaze shifts, i. e., frequent changes of attention between multiple displays. They can interrupt the workflow and affect task performance [36,125]. Minimizing gaze shifts should, therefore, be one of the design goals during the development of such systems. To this end, one can also consider setups that only consist of personal devices (Figure 6). Without the use of shared displays, screen space is limited and the common context that would otherwise be given by the display needs to be defined differently, e. g., by virtually placing the scene on a physical surface [26].

4.6.3. CAVE

The “cave automatic virtual environment” or CAVE (for cleverness) was originally demonstrated in 1992 by Cruz-Neira *et al.* [41] as a VR system using five rear-projected surfaces covering three walls, floor, and ceiling. The CAVE concept has since spawned many variants, including the YURT [158] which features 360° wall, domed-ceiling, and floor projection; and CAVE2 [31], which replaces projection with LCD panels for higher resolution, but eschews ceiling and floor projection.




Fig. 6: Tracked tablet used as a peephole into a 3D information space that is located on a table [26].

Interaction in CAVE-style environments is usually done through optically tracked devices such as hand-held controllers and the primary user's tracked head position. Modern variants such as the CAVE2 and the YURT are large enough to accommodate a reasonable group of people. They would, therefore, be ideal for collaborative use, but a technological limitation is that usually only one user has a correct, undistorted stereoscopic perspective. However, there are experimental methods for multiplexing projected displays that could be used to create true multiuser caves, e. g., [96].

Historically, CAVEs have been used to visualize complex 3D models from domains such as mechanical engineering, architecture, or urban planning as well as scientific visualization applications. They provide a large field-of-view and their size often makes them suitable to display content such as cars, machinery or architectural models at their actual size, or to show small molecular structures at a larger scale for knowledge discovery. On the other hand, a CAVE is usually a major investment of both money and floor space. This limits their use to larger companies or research labs.

Thus, there remains room for improvement in technology for CAVE-style VR setups, but such technology now has to compete with self-contained head-mounted displays that offer comparable experiences at a fraction of the cost.

4.6.4. Head-Mounted Displays (AR/VR)

While head-mounted displays (HMDs) are nearly as old as Virtual Reality research itself [145], technological advances and resulting consumer systems such as the HTC Vive, Oculus Rift, or Microsoft’s HoloLens now make VR and AR headsets widely available (and affordable) for immersive analytics. The systems are usually easy to set up and some HMDs even allow untethered use in non-instrumented environments. In a recent study comparing collaborative network visualization tasks in CAVE2 and HMD conditions, Cordeil *et al.* found participants were able to work more quickly and communicate just as well in the HMD setup [40]. However, the great advantage of the latter is cost.

Compared to using, e. g., mobile devices, such HMDs are usually more immersive. Virtual Reality headsets, however, can easily cause a feeling of disembodiment, negatively influencing presence and limiting collaboration. In any case, an HMD is inherently a single user device. Thus, multiuser setups need to provide an HMD for each user and scalability of HMD based setups is still a challenge – although Facebook and Samsung have demonstrated massively shared HMD VR environments for social computing [136]. Also compare the limitations explained in Subsection 4.4.2., many of which particularly apply to HMDs. Light-field displays [62] are a potential solution to the vergence-accommodation conflict described in Subsection 4.4.2..

4.6.5. Other Setups and Combinations

Besides the archetypes already described, other form factors exist that may also play a role in future immersive analytics systems. For example, wearables such as smart watches or even computers integrated into clothing could allow for personal tool palettes and clipboards, provide notifications, or act as additional input devices. Combinations of different form factors, such as using a display wall in conjunction with mobile devices which has been discussed above, generally allow a “best of both worlds” approach, addressing issues present with each device class. Today’s AR headsets, for example, have only very limited input capabilities that can be extended with additional personal controllers.

The wide range of the archetypes discussed in this section shows that a universal solution does not exist. For example, in-the-wild analytics (see Chapter 7) could hardly be supported by stationary CAVE setups. Instead, advantages and limitations of the different form factors need to be carefully weighed depending on the requirements of the system and are key part of the immersive analytics design space (see Chapter 9).

4.7. Example Systems

This section identifies recent systems that are representative of an emerging class of immersive data visualization systems that feature natural interaction techniques, as explored in this chapter, as a central design philosophy.

(Section 4.7.2.) and the pen-based SketchStory system where the task was data-storytelling (Lee *et al.* [101], see Chapter 6), the modality of the interaction device lends itself to a very “task-focused” interaction paradigm for data exploration; focused in the sense that the user is able to get on with the activity at hand without being distracted by the arcana of the interface, such as hunting through menus.

Data analysis with ImAxes is demonstrated for one type of data (multivariate/tabular). However, the general idea of reusable, reconfigurable and composable, embodied data visualization elements should be applicable to many other types of data.




Fig. 8: A user working with the ImAxes system. Here, the user has arranged four axes to form parallel coordinates.

4.7.2. BodyLenses

Magic lenses, introduced by Bier *et al.* [14], are a focus and context technique that is often used in visualization systems. They provide localized (and in collaborative settings: personalized) alternative views into the data. Besides their use in traditional, desktop-based systems, tangible lenses on tabletops have also been proposed, for example by Kim & Elmqvist [86]. Spindler *et al.* [137] even presented *spatially tracked* tangible magic lenses, which are used in the space above a tabletop.



Fig. 9: A graph visualization using the BodyLenses system. The lens provides a personal embodied territory. The user can configure her lens with a touch menu.

Another novel concept for magic lenses, specifically for the use in front of interactive wall-sized displays, are *BodyLenses* by Kister *et al.* [90]. *BodyLenses* are flexible, personal work territories that can provide diverse functions and tools. They are body-controlled magic lenses appearing on the display wall in front of the users. “Body-controlled” encompasses three different forms of interaction: (a) body movement relative to the display, (b) gestures, mainly arms and hands, but also any other body part, and (c) direct interaction on the display, using touch, pen or tangible input. *BodyLenses* move with the users, supporting implicit navigation within an information space. Additionally, gestures or direct input on the wall can be used to explicitly manipulate them. Furthermore, they serve as personal territories and support mutual awareness of co-located users. Supporting a continuous flow of interaction, with appropriate interaction techniques depending on the distance to the display wall, *BodyLenses* aim to make data exploration an immersive and engaging experience.

The authors explored the design space by examining design aspects such as appearance, function, interaction, and multi-user contexts. They present different shapes of lenses including classic, geometries such as circles and rectangles; body-centric lenses like shadows; and content-aware, data-driven shapes.

The shape of the lens can either be continuously changed depending on the user’s movement and posture (e. g., for shadow-like lenses) or can explicitly be controlled by the user. In addition to changing the lens’ position, the user’s distance to the display can also control other parameters. For example, users can move through time in a time-series visualization or control the properties of the lens function, e. g., a zoom level, abstraction, or displacement factor. Distance-based interaction can also address the notion of *proxemics*, a concept describing the spatial relationships governing social interactions. Thus, the distance between




Fig. 10: BodyLenses allows for multiple lenses at the same time. Here, two users inspect a time series of a biological data set.

a user and the wall can indicate the degree of engagement and can influence the lens accordingly, e. g., fading it out after some threshold.

Wall-sized displays are suitable for multiple users, therefore, BodyLenses support several lenses at a time. Overlapping lenses (and their effects) can be combined and separated to create common embodied territories sharing the same properties and elements (see also [7]). They also continuously convey which part of the data a user is currently investigating, providing mutual awareness.

Kister *et al.* implemented several example applications that address visualization use cases. The first is a graph explorer and includes domain-specific lens functions specific, e. g., local edge, bring neighbors, and fish-eye lenses. This application can be seen in Figure 9, which also shows the concept of using personal menus around the lens that allow changing lens parameters.

Figure 10 shows the second application. It allows the exploration of time series of, e. g., microscopy image sets. The distance to the wall is mapped to time, consequently, the users can easily “step” through time by moving in front of the wall.

ImAxes gives users “presence” in the data, and gives the data an embodied “presence” in the users’ physical space. BodyLenses, on the other hand, show that presence is not limited to fully immersive virtual experiences. The emergent quality of these systems, as well as the minimal interference they place between the user and the data analysis task, are compelling evidence that immersive

environments can create truly new ways for users to experience as well as explore data.

4.8. Conclusion

Modern analysts are well-versed in the specific devices, as well as interface widgets such as the WIMP desktop metaphor used in current visual analytics systems. These hardware and interface components, however, are often inappropriate for the non-traditional immersive analytics environments and applications under development today. New interaction techniques and metaphors must be designed and we must also provide guidance in choosing these new designs based on empirical evidence.

Visualization and virtual reality researchers have been successful in identifying user tasks and some interaction metaphors, are beginning to evaluate the usability of 3D interaction techniques for immersive analytics applications, and are trying to improve the usability of techniques for analytics tasks. However, the usability of 3D interfaces in real-world applications is still not at a desirable level. Scientists perhaps often perceive these technologies as good for demonstrations but without sufficient benefit to be used daily for insight discoveries. It is still an open question under which circumstances immersive analytics interfaces should be 2D or 3D (see Chapter 2) and how these interfaces need to be designed to be intuitive and engaging. Therefore, it is vital to focus on the use and to understand the design and evaluation of interface and interaction techniques to show how immersive analytics can actually increase efficiency, facilitate team collaboration, and reduce cost.

This chapter serves as an early step in this direction. We examined opportunities for the use of natural user interfaces to support immersive data analysis. To this end, we discussed typical analysis tasks and how natural user interfaces can be used to support these tasks. It is clear that no single input method is ‘perfect’ and suitable for all tasks. Instead, trade-offs between them have to be explored to choose techniques based on the specific requirements regarding, e. g., precision or physical demand. We also reviewed different, archetypal hardware setups for data visualization, which suggest that we will see a multitude of system designs in future immersive analytics systems. Finally, we looked at two example system designs that suggest such possible future directions.

One particular challenge is to integrate techniques designed for a single display type, a single task type, or a single user group into a seamless cross-display, cross-task, and cross-use multi-sensory environment. In such a system, realism is perhaps not the key but the magic interactivity integrated with visuals to understand what are the best mappings between these system factors to facilitate scientists’ and users’ decision making and insight discovery process.

A next step toward quantifying the benefits of interaction could be the creation of a taxonomy that separates all variables related to factors such as head tracking, immersion, display sizes, users, and tasks, to classify which characteristics are truly beneficial. There are multiple efforts to advance these frontiers, as discussed

in this chapter. A recent study by Bach *et al.* [6] evaluated what they considered to be the state-of-the-art for mixed-reality data visualization. However, their findings, while generally favorable for the mixed-reality condition, were inconclusive. They describe a number of limitations to their study: the fact that they were testing only one immersive interaction design from a huge space of possibilities; rapidly evolving capabilities of the current technology; and, the limited experience of their participants with immersive environments. In a way, this work neatly summarizes the current state of interaction for immersive analytics. There is considerable potential but precisely what form interaction with immersive analytics will ultimately take is uncertain.

In this chapter, we have surveyed the various technologies that currently enable immersive experiences and how these can be used for data visualization. While there have been immersive VR interfaces in the past, with current technological developments, access to these technologies becomes easier. For more and more domain experts, immersive analytics is within reach. Even everyday users will be able to afford future systems, enabling immersive, personal analytics. As the technologies improve and the costs come down, there will surely be a convergence or at least clearer winners in terms of which technologies are adopted in both the professional and the consumer space. Similarly, particular interaction techniques will emerge as standard as people's familiarity with these environments grows. The interaction techniques that we propose now and the studies that we perform to evaluate and refine them have the opportunity to influence these emerging standards, and hence, significantly impact and shape the future of data analytics in immersive environments.

Acknowledgements

Büschel acknowledges funding by the German Federal Ministry of Education and Research, grant no. 03ZZ0514C and Dwyer acknowledges support by the Australian Research Council Discovery Scheme, project DP180100755.

References

1. Adams, N., Witkowski, M., Spence, R.: The inspection of very large images by eye-gaze control. In: Proceedings of the Working Conference on Advanced Visual Interfaces (AVI). pp. 111–118. ACM, New York (2008) doi: 10.1145/1385569.1385589
2. Amar, R., Eagan, J., Stasko, J.: Low-level components of analytic activity in information visualization. In: Proceedings of the IEEE Symposium on Information Visualization (InfoVis). pp. 111–117. IEEE Computer Society, Los Alamitos (2005) doi: 10.1109/INFVIS.2005.1532136
3. Andrews, C., Endert, A., North, C.: Space to think: Large high-resolution displays for sensemaking. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI). pp. 55–64. ACM, New York (2010) doi: 10.1145/1753326.1753336

4. Argelaguet, F., Andujar, C.: A survey of 3D object selection techniques for virtual environments. *Computers & Graphics* 37(3), 121–136 (2013) doi: 10.1016/j.cag.2012.12.003
5. Arif, A.S., Stuerzlinger, W.: User adaptation to a faulty unistroke-based text entry technique by switching to an alternative gesture set. In: *Proceedings of Graphics Interface (GI)*. pp. 183–192. Canadian Information Processing Society, Toronto (2014) doi: 10.20380/GI2014.24
6. Bach, B., Sicat, R., Beyer, J., Cordeil, M., Pfister, H.: 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), 457–467 (2018) doi: 10.1109/TVCG.2017.2745941
7. Badam, S.K., Amini, F., Elmqvist, N., Irani, P.: Supporting visual exploration for multiple users in large display environments. In: *Proceedings of the IEEE Conference on Visual Analytics Science and Technology (VAST)*. pp. 1–10. IEEE Computer Society, Los Alamitos (2016) doi: 10.1109/VAST.2016.7883506
8. Ball, R., North, C.: Effects of tiled high-resolution display on basic visualization and navigation tasks. In: *Extended Abstracts on Human Factors in Computing Systems (CHI EA)*. pp. 1196–1199. ACM, New York (2005) doi: 10.1145/1056808.1056875
9. Ball, R., North, C., Bowman, D.A.: Move to improve: Promoting physical navigation to increase user performance with large displays. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*. pp. 191–200. ACM, New York (2007) doi: 10.1145/1240624.1240656
10. Ballmer, S.: CES 2010: A transforming trend—the natural user interface. *The Huffington Post* (2010), http://www.huffingtonpost.com/steve-ballmer/ces-2010-a-transformingt_b_416598.html/
11. Beaudouin-Lafon, M., Huot, S., Nancel, M., Mackay, W., Pietriga, E., Primet, R., Wagner, J., Chapuis, O., Pillias, C., Eagan, J., Gjerlufsen, T., Klokmoose, C.: Multisurface interaction in the WILD room. *Computer* 45(4), 48–56 (Apr 2012) doi: 10.1109/MC.2012.110
12. Benko, H., Ishak, E.W.: Cross-dimensional gestural interaction techniques for hybrid immersive environments. In: *Proceedings of the IEEE Conference on Virtual Reality (VR)*. pp. 209–216, 327. IEEE Computer Society, Los Alamitos (2005) doi: 10.1109/VR.2005.1492776
13. Besançon, L., Issartel, P., Ammi, M., Isenberg, T.: Hybrid tactile/tangible interaction for 3D data exploration. *IEEE Transactions on Visualization and Computer Graphics* 23(1), 881–890 (Jan 2017) doi: 10.1109/TVCG.2016.2599217
14. Bier, E.A., Stone, M.C., Pier, K., Fishkin, K., Baudel, T., Conway, M., Buxton, W., DeRose, T.: Toolglass and magic lenses: The see-through interface. In: *Conference Companion on Human Factors in Computing Systems*. pp. 445–446. ACM, New York (1994) doi: 10.1145/259963.260447
15. Billinghurst, M., Clark, A., Lee, G.: A survey of augmented reality. *Foundations and Trends® in Human–Computer Interaction* 8(2–3), 73–272 (2015) doi: 10.1561/11000000049
16. Bjork, S., Holopainen, J.: *Patterns in Game Design* (Game Development Series). Charles River Media, Inc., Rockland, MA, USA (2004)
17. Bolt, R.A.: “Put-that-there”: Voice and gesture at the graphics interface. *ACM SIGGRAPH Computer Graphics* 14(3), 262–270 (Jul 1980) doi: 10.1145/965105.807503
18. Bolt, R.A.: Gaze-orchestrated dynamic windows. *ACM SIGGRAPH Computer Graphics* 15(3), 109–119 (1981) doi: 10.1145/965161.806796

19. Branit, B.: World Builder. Online video (2009), <https://vimeo.com/3365942>
20. Brown, M.A., Stuerzlinger, W.: Exploring the throughput potential of in-air pointing. In: Proceedings of the International Conference on Human-Computer Interaction (HCI). pp. 13–24. Springer, Berlin/Heidelberg (2016) doi: 10.1007/978-3-319-39516-6_2
21. Brown, M.A., Stuerzlinger, W., Mendonça Filho, E.J.: The performance of un-instrumented in-air pointing. In: Proceedings of Graphics Interface (GI). pp. 59–66. Canadian Information Processing Society, Toronto (2014) doi: 10.20380/GI2014.08
22. Bruder, G., Steinicke, F., Stuerzlinger, W.: Effects of visual conflicts on 3D selection task performance in stereoscopic display environments. In: Proceedings of the IEEE Symposium on 3D User Interfaces (3DUI). pp. 115–118. IEEE Computer Society, Los Alamitos (2013) doi: 10.1109/3DUI.2013.6550207
23. Bruder, G., Steinicke, F., Stuerzlinger, W.: To touch or not to touch? Comparing 2D touch and 3D mid-air interaction on stereoscopic tabletop surfaces. In: Proceedings of the 1st Symposium on Spatial User Interaction (SUI). pp. 9–16. ACM, New York (2013) doi: 10.1145/2491367.2491369
24. Bruder, G., Steinicke, F., Stuerzlinger, W.: Touching the void revisited: Analyses of touch behavior on and above tabletop surfaces. In: Proceedings of Human-Computer Interaction (INTERACT). pp. 278–296. Springer, Berlin/Heidelberg (2013) doi: 10.1007/978-3-642-40483-2_19
25. Büschel, W., Mitschick, A., Dachsel, R.: Here and now: Reality-based information retrieval. In: Proceedings of the 2018 Conference on Human Information Interaction & Retrieval. pp. 171–180. CHIIR '18, ACM, New York, NY, USA (2018) doi: 10.1145/3176349.3176384
26. Büschel, W., Reipschläger, P., Langner, R., Dachsel, R.: 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. pp. 62–71. ISS '17, ACM, New York, NY, USA (2017) doi: 10.1145/3132272.3134125
27. Bush, V.: As we may think. *The Atlantic Monthly* 176(1), 101–108 (July 1945), <https://www.theatlantic.com/magazine/archive/1945/07/as-we-may-think/303881/>
28. Butscher, S., Hubenschmid, S., Müller, J., Fuchs, J., Reiterer, H.: 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. pp. 90:1–90:12. CHI '18, ACM, New York, NY, USA (2018) doi: 10.1145/3173574.3173664
29. Buxton, B.: Multi-touch systems that I have known and loved. Tech. rep., Microsoft Research (2007), <http://www.billbuxton.com/multitouchOverview.html>
30. Card, S.K., Robertson, G.G., Mackinlay, J.D.: The information visualizer, an information workspace. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI). pp. 181–186. ACM, New York (1991) doi: 10.1145/108844.108874
31. EVL CAVE2 homepage, <https://www.evl.uic.edu/entry.php?id=2016>
32. Cernea, D., Kerren, A.: A survey of technologies on the rise for emotion-enhanced interaction. *Journal of Visual Languages and Computing* 31, Part A, 70–86 (Dec 2015) doi: 10.1016/j.jvlc.2015.10.001
33. Chen, J., Bowman, D.A.: Effectiveness of cloning techniques for architectural virtual environments. In: IEEE Virtual Reality Conference. pp. 103–110. IEEE (2006) doi: 10.1109/VR.2006.57

34. Chen, J., Bowman, D.A.: Domain-specific design of 3D interaction techniques: An approach for designing useful virtual environment applications. *Presence: Teleoperators and Virtual Environments* 18(5), 370–386 (2009) doi: 10.1162/pres.18.5.370
35. Chen, J., Bowman, D.A., Lucas, J.F., Wingrave, C.A.: Interfaces for Cloning in Immersive Virtual Environments. In: *Eurographics Symposium on Virtual Environments*. The Eurographics Association (2004) doi: 10.2312/EGVE/EGVE04/091-098
36. Chen, J., Narayan, M.A., Manuel, Pérez-Quñones, A.: The use of hand-held devices for search tasks in virtual environments. *The IEEE Symposium on 3D User Interfaces* pp. 15–18 (2005)
37. Claes, S., Moere, A.V.: The role of tangible interaction in exploring information on public visualization displays. In: *Proceedings of the International Symposium on Pervasive Displays (PerDis)*. pp. 201–207. ACM, New York (2015) doi: 10.1145/2757710.2757733
38. Coffey, D., Malbraaten, N., Le, T., Borazjani, I., Sotiropoulos, F., Erdman, A.G., Keefe, D.F.: Interactive Slice WIM: Navigating and interrogating volume datasets using a multi-surface, multi-touch VR interface. *IEEE Transactions on Visualization and Computer Graphics* 18(10), 1614–1626 (2012) doi: 10.1109/TVCG.2011.283
39. Cordeil, M., Cunningham, A., Dwyer, T., Thomas, B.H., Marriott, K.: ImAxes: Immersive axes as embodied affordances for interactive multivariate data visualization. In: *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*. pp. 71–83. UIST '17, ACM, New York, NY, USA (2017) doi: 10.1145/3126594.3126613
40. Cordeil, M., Dwyer, T., Klein, K., Laha, B., Marriott, K., Thomas, B.H.: Immersive collaborative analysis of network connectivity: CAVE-style or head-mounted display? *IEEE Transactions on Visualization and Computer Graphics* 23(1), 441–450 (2017) doi: 10.1109/TVCG.2016.2599107
41. Cruz-Neira, C., Sandin, D.J., DeFanti, T.A., Kenyon, R.V., Hart, J.C.: The CAVE: audio visual experience automatic virtual environment. *Communications of the ACM* 35(6), 64–72 (Jun 1992) doi: 10.1145/129888.129892
42. Cummings, J.J., Bailenson, J.N.: How immersive is enough? A meta-analysis of the effect of immersive technology on user presence. *Media Psychology* 19(2), 272–309 (2016) doi: 10.1080/15213269.2015.1015740
43. van Dam, A.: Post-WIMP user interfaces. *Communications of the ACM* 40(2), 63–67 (Feb 1997) doi: 10.1145/253671.253708
44. Dostal, J., Hinrichs, U., Kristensson, P.O., Quigley, A.: Spidereyes: Designing attention- and proximity-aware collaborative interfaces for wall-sized displays. In: *Proceedings of the International Conference on Intelligent User Interfaces (IUI)*. pp. 143–152. ACM, New York (2014) doi: 10.1145/2557500.2557541
45. Dourish, P.: *Where the Action Is: The Foundations of Embodied Interaction*. MIT Press (2001)
46. Drucker, S.M., Fisher, D., Sadana, R., Herron, J., schraefel, m.c.: TouchViz: A case study comparing two interfaces for data analytics on tablets. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*. pp. 2301–2310. ACM, New York (2013) doi: 10.1145/2470654.2481318
47. Elmqvist, N., Vande Moere, A., Jetter, H.C., Cernea, D., Reiterer, H., Jankun-Kelly, T.J.: Fluid interaction for information visualization. *Information Visualization* 10(4), 327–340 (Oct 2011) doi: 10.1177/1473871611413180

48. Febretti, A., Nishimoto, A., Mateevitsi, V., Renambot, L., Johnson, A., Leigh, J.: Omegalib: A multi-view application framework for hybrid reality display environments. In: Proceedings of the IEEE Conference on Virtual Reality (VR). pp. 9–14. IEEE Computer Society, Los Alamitos (2014) doi: 10.1109/VR.2014.6802043
49. Fikkert, W., D'Ambros, M., Bierz, T., Jankun-Kelly, T.: Interacting with visualizations. In: Kerren, A., Ebert, A., Meyer, J. (eds.) Human-Centered Visualization Environments, LNCS, vol. 4417, chap. 3, pp. 77–162. Springer, Berlin/Heidelberg (2007) doi: 10.1007/978-3-540-71949-6_3
50. Follmer, S., Leithinger, D., Olwal, A., Hogge, A., Ishii, H.: inFORM: Dynamic physical affordances and constraints through shape and object actuation. In: Proceedings of the Annual ACM Symposium on User Interface Software and Technology (UIST). pp. 417–426. ACM, New York (2013) doi: 10.1145/2501988.2502032
51. Frisch, M., Heydekorn, J., Dachzelt, R.: Diagram editing on interactive displays using multi-touch and pen gestures. In: Proceedings of the International Conference on Diagrammatic Representation and Inference (Diagrams). pp. 182–196. Springer, Berlin/Heidelberg (2010) doi: 10.1007/978-3-642-14600-8_18
52. Fu, C.W., Goh, W.B., Ng, J.A.: Multi-touch techniques for exploring large-scale 3D astrophysical simulations. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI). pp. 2213–2222. ACM, New York (2010) doi: 10.1145/1753326.1753661
53. Gillet, A., Sanner, M., Stoffler, D., Olson, A.: Tangible interfaces for structural molecular biology. *Structure* 13(3), 483–491 (Mar 2005) doi: 10.1016/j.str.2005.01.009
54. Harrison, C., Sato, M., Poupyrev, I.: Capacitive fingerprinting: Exploring user differentiation by sensing electrical properties of the human body. In: Proceedings of the Annual ACM Symposium on User Interface Software and Technology (UIST). pp. 537–544. ACM, New York (2012) doi: 10.1145/2380116.2380183
55. Heer, J., Shneiderman, B.: Interactive dynamics for visual analysis. *Communications of the ACM* 55(4), 45–54 (Apr 2012) doi: 10.1145/2133806.2133821
56. Hess, R.F., To, L., Zhou, J., Wang, G., Cooperstock, J.R.: Stereo vision: The haves and have-nots. *i-Perception* 6(3) (Jun 2015) doi: 10.1177/2041669515593028
57. Hincapié-Ramos, J.D., Guo, X., Moghadasian, P., Irani, P.: Consumed endurance: a metric to quantify arm fatigue of mid-air interactions. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI). pp. 1063–1072. ACM, New York (2014) doi: 10.1145/2556288.2557130
58. Hinckley, K., Pausch, R., Goble, J.C., Kassell, N.F.: Passive real-world interface props for neurosurgical visualization. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI). pp. 452–458. ACM, New York (1994) doi: 10.1145/191666.191821
59. Hinckley, K., Yatani, K., Pahud, M., Coddington, N., Rodenhouse, J., Wilson, A., Benko, H., Buxton, B.: Pen + touch = new tools. In: Proceedings of the Annual ACM Symposium on User Interface Software and Technology (UIST). pp. 27–36. ACM, New York (2010) doi: 10.1145/1866029.1866036
60. Hoffman, D.M., Girshick, A.R., Akeley, K., Banks, M.S.: Vergence–accommodation conflicts hinder visual performance and cause visual fatigue. *Journal of Vision* 8(3), 33:1–33:30 (2008) doi: 10.1167/8.3.33
61. Horak, T., von Zadow, U., Kalms, M., Dachzelt, R.: Discussing the state of the art for “in the wild” mobile device localization. In: Proceedings of the ISS Workshop on Interacting with Multi-Device Ecologies “in the wild” (2016), http://cross-surface.com/papers/Cross-Surface_2016-2_paper_2.pdf

62. Huang, F.C., Chen, K., Wetzstein, G.: The light field stereoscope: immersive computer graphics via factored near-eye light field displays with focus cues. *ACM Transactions on Graphics* 34(4), 60:1–60:12 (2015) doi: 10.1145/2766922
63. Hutchins, E.L., Hollan, J.D., Norman, D.A.: Direct manipulation interfaces. *Human–Computer Interaction* 1(4), 311–338 (1985) doi: 10.1207/s15327051hci0104_2
64. Hutchinson, T.E., White, K.P., Martin, W.N., Reichert, K.C., Frey, L.A.: Human-computer interaction using eye-gaze input. *IEEE Transactions on Systems, Man, and Cybernetics* 19(6), 1527–1534 (1989) doi: 10.1109/21.44068
65. Isenberg, P., Dragicevic, P., Willett, W., Bezerianos, A., Fekete, J.D.: Hybrid-image visualization for large viewing environments. *IEEE Transactions on Visualization and Computer Graphics* 19(12), 2346–2355 (Dec 2013) doi: 10.1109/TVCG.2013.163
66. Isenberg, P., Isenberg, T.: Visualization on interactive surfaces: A research overview. *i-com* 12(3), 10–17 (Nov 2013) doi: 10.1524/icom.2013.0020
67. Isenberg, T.: Interactive exploration of three-dimensional scientific visualizations on large display surfaces. In: Anslow, C., Campos, P., Jorge, J. (eds.) *Collaboration Meets Interactive Spaces*, chap. 6, pp. 97–123. Springer, Berlin/Heidelberg (2016) doi: 10.1007/978-3-319-45853-3_6
68. Isenberg, T., Hancock, M.: *Gestures* vs. postures: ‘Gestural’ touch interaction in 3D environments. In: *Proceedings of the CHI Workshop on “The 3rd Dimension of CHI: Touching and Designing 3D User Interfaces” (3DCHI)*. pp. 53–61 (2012), <https://hal.inria.fr/hal-00781237>
69. Ishii, H., Ratti, C., Piper, B., Wang, Y., Biderman, A., Ben-Joseph, E.: Bringing clay and sand into digital design—Continuous tangible user interfaces. *BT Technology Journal* 22(4), 287–299 (Oct 2004) doi: 10.1023/B:BTTJ.0000047607.16164.16
70. Ishii, H.: The tangible user interface and its evolution. *Communications of the ACM* 51(6), 32–36 (Jun 2008) doi: 10.1145/1349026.1349034
71. Ishii, H., Ullmer, B.: Tangible bits: Towards seamless interfaces between people, bits and atoms. In: *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI)*. pp. 234–241. ACM, New York (1997) doi: 10.1145/258549.258715
72. Jackson, B., Schroeder, D., Keefe, D.F.: Nailing down multi-touch: Anchored above the surface interaction for 3D modeling and navigation. In: *Proceedings of Graphics Interface (GI)*. pp. 181–184. Canadian Information Processing Society, Toronto (2012) doi: 10.20380/GI2012.23
73. Jacob, R.J., Girouard, A., Hirshfield, L.M., Horn, M.S., Shaer, O., Solovey, E.T., Zigelbaum, J.: Reality-based interaction: A framework for post-WIMP interfaces. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*. pp. 201–210. ACM, New York (2008) doi: 10.1145/1357054.1357089
74. Jakobsen, M.R., Haile, Y.S., Knudsen, S., Hornbæk, K.: Information visualization and proxemics: Design opportunities and empirical findings. *IEEE Transactions on Visualization and Computer Graphics* 19(12), 2386–2395 (Dec 2013) doi: 10.1109/TVCG.2013.166
75. Jang, S., Stuerzlinger, W., Ambike, S., Ramani, K.: Modeling cumulative arm fatigue in mid-air interaction based on perceived exertion and kinetics of arm motion. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*. pp. 3328–3339. ACM, New York (2017) doi: 10.1145/3025453.3025523
76. Jankowski, J., Hachet, M.: Advances in interaction with 3D environments. *Computer Graphics Forum* 34(1), 152–190 (Jan 2015) doi: 10.1111/cgf.12466

77. Jansen, Y.: Physical and tangible information visualization. Ph.D. thesis, Université Paris Sud – Paris XI, France (Mar 2014), <https://tel.archives-ouvertes.fr/tel-00981521>
78. Jansen, Y., Dragicevic, P.: An interaction model for visualizations beyond the desktop. *IEEE Transactions on Visualization and Computer Graphics* 19(12), 2396–2405 (Dec 2013) doi: 10.1109/TVCG.2013.134
79. Jansen, Y., Dragicevic, P., Isenberg, P., Alexander, J., Karnik, A., Kildal, J., Subramanian, S., Hornbæk, K.: Opportunities and challenges for data physicalization. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*. pp. 3227–3236. ACM, New York (2015) doi: 10.1145/2702123.2702180
80. Johnson, R., O'Hara, K., Sellen, A., Cousins, C., Criminisi, A.: Exploring the potential for touchless interaction in image-guided interventional radiology. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*. pp. 3323–3332. ACM, New York (2011) doi: 10.1145/1978942.1979436
81. Jota, R., Nacenta, M.A., Jorge, J.A., Carpendale, S., Greenberg, S.: A comparison of ray pointing techniques for very large displays. In: *Proceedings of Graphics Interface (GI)*. pp. 269–276. Canadian Information Processing Society, Toronto (2010) doi: 10.20380/GI2010.36
82. Karam, M., schraefel, m.c.: A taxonomy of gestures in human computer interactions. Tech. Rep. 261149, University of Southampton (2005), <http://eprints.soton.ac.uk/261149/>, ISBN 0854328335
83. Keefe, D.F.: Integrating visualization and interaction research to improve scientific workflows. *IEEE Computer Graphics and Applications* 30(2), 8–13 (Mar/Apr 2010) doi: 10.1109/MCG.2010.30
84. Keefe, D.F., Isenberg, T.: Reimagining the scientific visualization interaction paradigm. *IEEE Computer* 46(5), 51–57 (May 2013) doi: 10.1109/MC.2013.178
85. Kerren, A., Schreiber, F.: Toward the role of interaction in visual analytics. In: *Proceedings of the Winter Simulation Conference (WSC)*. pp. 420:1–420:13. Winter Simulation Conference (2012), <http://dl.acm.org/citation.cfm?id=2429759.2430303> doi: 10.1109/WSC.2012.6465208
86. Kim, K., Elmqvist, N.: Embodied lenses for collaborative visual queries on tabletop displays. *Information Visualization* 11(4), 319–338 (Apr 2012) doi: 10.1177/1473871612441874
87. Kirmizibayrak, C., Radeva, N., Wakid, M., Philbeck, J., Sibert, J., Hahn, J.: Evaluation of gesture based interfaces for medical volume visualization tasks. In: *Proceedings of the International Conference on Virtual Reality Continuum and Its Applications in Industry (VRCAI)*. pp. 69–74. ACM, New York (2011) doi: 10.1145/2087756.2087764
88. Kister, U., Klamka, K., Tominski, C., Dachsel, R.: GRASP: Combining spatially-aware mobile devices and a display wall for graph visualization and interaction. *Computer Graphics Forum* 36(3), 503–514 (Jun 2017) doi: 10.1111/cgf.13206
89. Kister, U., Reipschläger, P., Dachsel, R.: MultiLens: Fluent interaction with multi-functional multi-touch lenses for information visualization. In: *Proceedings of the ACM Conference on Interactive Surfaces and Spaces (ISS)*. pp. 139–148. ACM, New York (2016) doi: 10.1145/2992154.2992168
90. Kister, U., Reipschläger, P., Matulic, F., Dachsel, R.: BodyLenses: Embodied magic lenses and personal territories for wall displays. In: *Proceedings of the International Conference on Interactive Tabletops & Surfaces (ITS)*. pp. 117–126. ACM, New York (2015) doi: 10.1145/2817721.2817726

91. Klamka, K., Siegel, A., Vogt, S., Göbel, F., Stellmach, S., Dachsel, R.: Look & pedal: Hands-free navigation in zoomable information spaces through gaze-supported foot input. In: *Proceedings of the International Conference on Multimodal Interaction (ICMI)*. pp. 123–130. ACM, New York (2015) doi: 10.1145/2818346.2820751
92. Klapperstuck, M., Czauderna, T., Goncu, C., Glowacki, J., Dwyer, T., Schreiber, F., Marriott, K.: ContextuWall: Peer collaboration using (large) displays. In: *Proceedings of the International Symposium on Big Data Visual Analytics (BDVA)*. pp. 1–8. IEEE, Red Hook, NY, USA (2016) doi: 10.1109/BDVA.2016.7787047
93. Klum, S., Isenberg, P., Langner, R., Fekete, J.D., Dachsel, R.: Stackables: Combining tangibles for faceted browsing. In: *Proceedings of the International Working Conference on Advanced Visual Interfaces*. pp. 241–248. AVI '12, ACM, New York, NY, USA (2012), <http://doi.acm.org/10.1145/2254556.2254600> doi: 10.1145/2254556.2254600
94. Konchada, V., Jackson, B., Le, T., Borazjani, I., Sotiropoulos, F., Keefe, D.F.: Supporting internal visualization of biomedical datasets via 3D rapid prototypes and sketch-based gestures. In: *Proceedings of the Symposium on Interactive 3D Graphics and Games (I3D)*. pp. 214–214. ACM, New York (2011) doi: 10.1145/1944745.1944794
95. Kruszyński, K.J., van Liere, R.: Tangible props for scientific visualization: Concept, requirements, application. *Virtual Reality* 13(4), 235–244 (Nov 2009) doi: 10.1007/s10055-009-0126-1
96. Kulik, A., Kunert, A., Beck, S., Reichel, R., Blach, R., Zink, A., Froehlich, B.: Clx6: a stereoscopic six-user display for co-located collaboration in shared virtual environments. *ACM Transactions on Graphics* 30(6), 188:1–188:12 (2011) doi: 10.1145/2070781.2024222
97. Kurzhals, K., Fisher, B., Burch, M., Weiskopf, D.: Evaluating visual analytics with eye tracking. In: *Proceedings of the Workshop on Beyond Time and Errors: Novel Evaluation Methods for Visualization (BELIV)*. pp. 61–69. ACM, New York (2014) doi: 10.1145/2669557.2669560
98. LaViola, J., Kruijff, E., Bowman, D., McMahan, R., Poupyrev, I.: 3D user interfaces: theory and practice. *Usability Series*, Pearson Education, Limited (2017)
99. Le Goc, M.: Supporting versatility in tangible user interfaces using collections of small actuated objects. Ph.D. thesis, Université Paris-Saclay, France (Dec 2016), <https://tel.archives-ouvertes.fr/tel-01453175>
100. Lee, B., Isenberg, P., Riche, N.H., Carpendale, S.: Beyond mouse and keyboard: Expanding design considerations for information visualization interactions. *IEEE Transactions on Visualization and Computer Graphics* 18(12), 2689–2698 (Dec 2012) doi: 10.1109/TVCG.2012.204
101. Lee, B., Kazi, R.H., Smith, G.: SketchStory: Telling more engaging stories with data through freeform sketching. *IEEE Transactions on Visualization and Computer Graphics* 19(12), 2416–2425 (2013) doi: 10.1109/TVCG.2013.191
102. Leithinger, D., Lakatos, D., DeVincenzi, A., Blackshaw, M., Ishii, H.: Direct and gestural interaction with relief: A 2.5D shape display. In: *Proceedings of the Annual ACM Symposium on User Interface Software and Technology (UIST)*. pp. 541–548. ACM, New York (2011) doi: 10.1145/2047196.2047268
103. López, D., Oehlberg, L., Doger, C., Isenberg, T.: Towards an understanding of mobile touch navigation in a stereoscopic viewing environment for 3D data exploration. *IEEE Transactions on Visualization and Computer Graphics* 22(5), 1616–1629 (May 2016) doi: 10.1109/TVCG.2015.2440233

104. Lucas, J., Bowman, D., Chen, J., Wingrave, C.: Design and evaluation of 3D multiple object selection techniques. In: ACM Interactive 3D graphics (2005)
105. Lundström, C., Rydell, T., Forsell, C., Persson, A., Ynnerman, A.: Multi-touch table system for medical visualization: Application to orthopedic surgery planning. *IEEE Transactions on Visualization and Computer Graphics* 17(12) (Dec 2011) doi: 10.1109/TVCG.2011.224
106. Lv, Z., Halawani, A., Feng, S., Li, H., Réhman, S.U.: Multimodal hand and foot gesture interaction for handheld devices. *ACM Transactions on Multimedia Computing, Communications, and Applications* 11(1s), 10:1–10:19 (Sep 2014) doi: 10.1145/2645860
107. MacKenzie, I.S.: Evaluating eye tracking systems for computer input. In: Majaranta, P., Aoki, H., Donegan, M., Hansen, D.W., Hansen, J.P., Hyrskykari, A., Rähkä, K.J. (eds.) *Gaze Interaction and Applications of Eye Tracking: Advances in Assistive Technologies: Advances in Assistive Technologies*, pp. 205–225. IGI Global, Hershey, PA, USA (2011) doi: 10.4018/978-1-61350-098-9.ch015
108. Majaranta, P., Bulling, A.: Eye tracking and eye-based human–computer interaction. In: Fairclough, S.H., Gilleade, K. (eds.) *Advances in Physiological Computing*, pp. 39–65. Springer, London (2014) doi: 10.1007/978-1-4471-6392-3_3
109. Malik, S., Ranjan, A., Balakrishnan, R.: Interacting with large displays from a distance with vision-tracked multi-finger gestural input. In: *Proceedings of the Annual ACM Symposium on User Interface Software and Technology (UIST)*. pp. 43–52. ACM, New York (2005) doi: 10.1145/1095034.1095042
110. McMahan, A.: Immersion, engagement, and presence: A method for analyzing 3-D video games. In: Wolf, M., Perron, B. (eds.) *The Video Game Theory Reader*, chap. 3, pp. 67–86. Routledge (2003), <http://www.alisonmcmahan.com/node/277>
111. McNeill, D.: *Hand and mind: What gestures reveal about thought*. University of Chicago Press (1992), <http://press.uchicago.edu/ucp/books/book/chicago/H/bo3641188.html>
112. Mohr, P., Kerbl, B., Donoser, M., Schmalstieg, D., Kalkofen, D.: Retargeting technical documentation to augmented reality. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*. pp. 3337–3346. ACM, New York (2015) doi: 10.1145/2702123.2702490
113. Myers, B.A.: A brief history of human-computer interaction technology. *ACM Interactions* 5(2), 44–54 (Mar/Apr 1998) doi: 10.1145/274430.274436
114. Ni, T., Bowman, D.A., Chen, J.: Increased display size and resolution improve task performance in information-rich virtual environments. In: *Proceedings of Graphics Interface*. pp. 139–146 (2006)
115. Nilsson, S., Gustafsson, T., Carleberg, P.: Hands free interaction with virtual information in a real environment: Eye gaze as an interaction tool in an augmented reality system. *PsychNology Journal* 7(2), 175–196 (2009)
116. Norman, D.A.: THE WAY I SEE IT: Signifiers, not affordances. *ACM Interactions* 15(6), 18–19 (Nov/Dec 2008) doi: 10.1145/1409040.1409044
117. Norman, D.A.: Natural user interfaces are not natural. *interactions* 17(3), 6–10 (May/Jun 2010) doi: 10.1145/1744161.1744163
118. Norman, D.A.: *The design of everyday things: Revised and expanded edition*. Basic books, New York (2013), <https://www.jnd.org/books/design-of-everyday-things-revised.html>
119. Piper, B., Ratti, C., Ishii, H.: Illuminating clay: A 3-D tangible interface for landscape analysis. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*. pp. 355–362. ACM, New York (2002) doi: 10.1145/503376.503439

120. Pirolli, P., Card, S.: Information foraging. *Psychological review* 106(4), 643–675 (Oct 1999) doi: 10.1037/0033-295X.106.4.643
121. Preim, B., Dachsel, R.: *Interaktive Systeme – Band 2: User Interface Engineering, 3D-Interaktion, Natural User Interfaces*, vol. 2. Springer/Vieweg, Berlin/Heidelberg (2015) doi: 10.1007/978-3-642-45247-5
122. Pretorius, A.J., Purchase, H.C., Stasko, J.T.: Tasks for multivariate network analysis. In: Kerren, A., Purchase, H.C., Ward, M.O. (eds.) *Multivariate Network Visualization: Dagstuhl Seminar #13201*, Dagstuhl Castle, Germany, May 12–17, 2013, Revised Discussions, pp. 77–95. Springer International Publishing, Cham, Switzerland (2014) doi: 10.1007/978-3-319-06793-3_5
123. Rädle, R., Jetter, H.C., Butscher, S., Reiterer, H.: The effect of egocentric body movements on users' navigation performance and spatial memory in zoomable user interfaces. In: *Proceedings of the International Conference on Interactive Tabletops and Surfaces (ITS)*. pp. 23–32. ACM, New York (2013) doi: 10.1145/2512349.2512811
124. Ragan, E.D., Endert, A., Sanyal, J., Chen, J.: Characterizing provenance in visualization and data analysis: An organizational framework of provenance types and purposes. *IEEE Transactions on Visualization and Computer Graphics* 22(1), 31–40 (Jan 2016), <https://doi.org/10.1109/TVCG.2015.2467551> doi: 10.1109/TVCG.2015.2467551
125. Rashid, U., Nacenta, M.A., Quigley, A.: The cost of display switching: A comparison of mobile, large display and hybrid UI configurations. In: *Proceedings of the International Working Conference on Advanced Visual Interfaces (AVI)*. pp. 99–106. ACM, New York (2012) doi: 10.1145/2254556.2254577
126. Rekimoto, J., Green, M.: The information cube: Using transparency in 3D information visualization. In: *Proceedings of the Annual Workshop on Information Technologies & Systems (WITS)*. pp. 125–132 (1993), <https://www.sonycsll.co.jp/person/rekimoto/cube/>
127. Renambot, L., Marrinan, T., Aurisano, J., Nishimoto, A., Mateevitsi, V., Bharadwaj, K., Long, L., Johnson, A., Brown, M., Leigh, J.: SAGE2: a collaboration portal for scalable resolution displays. *Future Generation Computer Systems* 54, 296–305 (2016) doi: 10.1016/j.future.2015.05.014
128. Roberts, J.C., Ritsos, P.D., Badam, S.K., Brodbeck, D., Kennedy, J., Elmqvist, N.: Visualization beyond the desktop—The next big thing. *IEEE Computer Graphics and Applications* 34(6), 26–34 (Nov/Dec 2014) doi: 10.1109/MCG.2014.82
129. Robles-De-La-Torre, G.: The importance of the sense of touch in virtual and real environments. *IEEE MultiMedia* 13(3), 24–30 (Jul 2006) doi: 10.1109/MMUL.2006.69
130. Roth, V., Schmidt, P., Güldenring, B.: The IR ring: Authenticating users' touches on a multi-touch display. In: *Proceedings of the Annual ACM Symposium on User Interface Software and Technology (UIST)*. pp. 259–262. ACM, New York (2010) doi: 10.1145/1866029.1866071
131. Samsung introduces 2007 LCD, plasma, DLP and CRT lineup (2007 (accessed 11th April 2017)), <https://www.engadget.com/2007/01/07/samsung-introduces-2007-lcd-plasma-dlp-and-crt-lineup/>
132. Scheurich, D., Stuerzlinger, W.: A one-handed multi-touch method for 3D rotations. In: *Human-Computer Interaction – INTERACT 2013*. pp. 56–69. Springer (2013) doi: 10.1007/978-3-642-40483-2
133. Shaer, O., Hornecker, E.: Tangible user interfaces: past, present, and future directions. *Foundations and Trends® in Human-Computer Interaction* 3(1–2), 4–137 (2010) doi: 10.1561/11000000026

134. Shneiderman, B.: The eyes have it: A task by data type taxonomy for information visualizations. In: *Proceedings of the IEEE Symposium on Visual Languages (VL)*. pp. 336–343. IEEE Computer Society, Los Alamitos (1996) doi: 10.1109/VL.1996.545307
135. Slater, M., Wilbur, S.: A framework for immersive virtual environments (FIVE): Speculations on the role of presence in virtual environments. *Presence: Teleoperators and virtual environments* 6(6), 603–616 (1997) doi: 10.1162/pres.1997.6.6.603
136. Facebook shows how it's gonna make virtual reality social, <https://www.cnet.com/au/news/facebook-mark-zuckerberg-shows-off-live-vr-virtual-reality-chat-with-oculus-rift/>
137. Spindler, M., Tominski, C., Schumann, H., Dachsel, R.: Tangible views for information visualization. In: *Proceedings of the International Conference on Interactive Tabletops and Surfaces (ITS)*. pp. 157–166. ACM, New York (2010) doi: 10.1145/1936652.1936684
138. Stasko, J., Görg, C., Liu, Z.: Jigsaw: Supporting investigative analysis through interactive visualization. *Information Visualization* 7(2), 118–132 (Apr 2008) doi: 10.1145/1466620.1466622
139. Stellmach, S., Dachsel, R.: Look & touch: Gaze-supported target acquisition. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*. pp. 2981–2990. ACM, New York (2012) doi: 10.1145/2207676.2208709
140. Stellmach, S., Stober, S., Nürnberger, A., Dachsel, R.: Designing gaze-supported multimodal interactions for the exploration of large image collections. In: *Proceedings of the Conference on Novel Gaze-Controlled Applications (NGCA)*. pp. 1:1–1:8. ACM, New York (2011) doi: 10.1145/1983302.1983303
141. Stuerzlinger, W., Wingrave, C.: The value of constraints for 3D user interfaces. In: *Virtual Realities: Dagstuhl Seminar 2008*. pp. 203–224. Springer (2011) doi: 10.1007/978-3-211-99178-7
142. Sun, J., Stuerzlinger, W., Shuralyov, D.: Shift-sliding and depth-pop for 3D positioning. In: *Proceedings of the 2016 Symposium on Spatial User Interaction*. pp. 69–78. ACM (2016)
143. Sutherland, I.E.: Sketchpad: A man-machine graphical communication system. In: *Proceedings of the Spring Joint Computer Conference (AFIPS, Spring)*. pp. 329–346. ACM, New York (1963) doi: 10.1145/1461551.1461591
144. Sutherland, I.E.: The ultimate display. In: *Proceedings of the IFIP Congress*. pp. 506–508 (1965)
145. Sutherland, I.E.: A head-mounted three dimensional display. In: *Proceedings of the Fall Joint Computer Conference (AFIPS, Fall, part I)*. pp. 757–764. ACM, New York (1968) doi: 10.1145/1476589.1476686
146. Teather, R.J., Stuerzlinger, W.: Pointing at 3D targets in a stereo head-tracked virtual environment. In: *Proceedings of the IEEE Symposium on 3D User Interfaces (3DUI)*. pp. 87–94. IEEE Computer Society, Los Alamitos (2011) doi: 10.1109/3DUI.2011.5759222
147. Cluster rendering, Unity user manual, <https://docs.unity3d.com/Manual/ClusterRendering.html>
148. Vertegaal, R.: Attentive user interfaces. *Communications of the ACM* 46(3), 30–33 (Mar 2003) doi: 10.1145/636772.636794
149. Wai, J., Lubinski, D., Benbow, C.P.: Spatial ability for STEM domains: aligning over 50 years of cumulative psychological knowledge solidifies its importance. *Journal of Educational Psychology* 101(4), 817 (2009) doi: 10.1037/a0016127

150. Walny, J., Carpendale, S., Henry Riche, N., Venolia, G., Fawcett, P.: Visual thinking in action: Visualizations as used on whiteboards. *IEEE Transactions on Visualization and Computer Graphics* 17(12), 2508–2517 (Dec 2011) doi: 10.1109/TVCG.2011.251
151. Walter, R., Bailly, G., Valkanova, N., Müller, J.: Cuenesics: Using mid-air gestures to select items on interactive public displays. In: *Proceedings of the International Conference on Human-computer Interaction with Mobile Devices & Services (MobileHCI)*. pp. 299–308. ACM, New York (2014) doi: 10.1145/2628363.2628368
152. Wigdor, D., Wixon, D.: *Brave NUI world: designing natural user interfaces for touch and gesture*. Elsevier/Morgan Kaufmann, Amsterdam (2011) doi: 10.1016/B978-0-12-382231-4.00037-X
153. Wither, J., DiVerdi, S., Höllerer, T.: Annotation in outdoor augmented reality. *Computers & Graphics* 33(6), 679–689 (Dec 2009) doi: 10.1016/j.cag.2009.06.001
154. Yi, J.S., Kang, Y.a., Stasko, J., Jacko, J.: Toward a deeper understanding of the role of interaction in information visualization. *IEEE Transactions on Visualization and Computer Graphics* 13(6), 1224–1231 (Nov/Dec 2007) doi: 10.1109/TVCG.2007.70515
155. Ynnerman, A., Rydell, T., Antoine, D., Hughes, D., Persson, A., Ljung, P.: Interactive visualization of 3D scanned mummies at public venues. *Communications of the ACM* 59(12), 72–81 (Dec 2016) doi: 10.1145/2950040
156. Yu, L., Efstathiou, K., Isenberg, P., Isenberg, T.: Efficient structure-aware selection techniques for 3D point cloud visualizations with 2DOF input. *IEEE Transactions on Visualization and Computer Graphics* 18(12), 2245–2254 (Dec 2012) doi: 10.1109/TVCG.2012.217
157. Yu, L., Svetachov, P., Isenberg, P., Everts, M.H., Isenberg, T.: FI3D: Direct-touch interaction for the exploration of 3D scientific visualization spaces. *IEEE Transactions on Visualization and Computer Graphics* 16(6), 1613–1622 (Nov/Dec 2010) doi: 10.1109/TVCG.2010.157
158. Brown University YURT homepage, <https://web1.ccv.brown.edu/viz-yurt>
159. von Zadow, U., Reipschläger, P., Bösel, D., Sellent, A., Dachselt, R.: YouTouch! Low-cost user identification at an interactive display wall. In: *Proceedings of the International Working Conference on Advanced Visual Interfaces (AVI)*. pp. 144–151. ACM, New York (2016) doi: 10.1145/2909132.2909258
160. Zaroff, C.M., Knutelska, M., Frumkes, T.E.: Variation in stereoacuity: Normative description, fixation disparity, and the roles of aging and gender. *Investigative Ophthalmology & Visual Science* 44(2), 891 (Feb 2003) doi: 10.1167/iovs.02-0361