2. Immersive Analytics: Time to Reconsider the Value of 3D for Information Visualisation

Kim Marriott[1], Jian Chen[2], Marcel Hlawatsch[3], Takayuki Itoh[4], Miguel A. Nacenta[5], Guido Reina[6], and Wolfgang Stuerzlinger[7]

[1] Monash University, Australia
Kim.Marriott@monash.edu
[2] The Ohio State University, USA
chen.8028@osu.edu
[3] University of Stuttgart (VISUS), Germany
Marcel.Hlawatsch@visus.uni-stuttgart.de
[4] Ochanomizu University, Japan
itot@is.ocha.ac.jp
[5] University of St Andrews, Scotland
mans@st-andrews.ac.uk
[6] University of Stuttgart (VISUS), Germany
guido.reina@visus.uni-stuttgart.de
[7] Simon Fraser University, Canada
w.s@sfu.ca

Abstract. Modern virtual reality display technologies engender spatial immersion by using a variety of depth cues such as perspective and head-tracked binocular presentation to create visually realistic 3D worlds. While 3D visualisations are common in scientific visualisation, they are much less common in information visualisation. In this chapter we explore whether immersive analytic applications should continue to use traditional 2D information visualisations or whether there are situations when 3D may offer benefits. We identify a number of potential applications of 3D depth cues for abstract data visualisation: using depth to show an additional data dimension, such as in 2.5D network layouts, views on non-flat surfaces and egocentric views in which the data is placed around the viewer, and visualising abstract data with a spatial embedding. Another important potential benefit is the ability to arrange multiple views in the 3D space around the user and to attach abstract visualisations to objects in the real world.

Keywords: immersive analytics, data visualisation, information visualisation, 3D

2.1. Introduction

In this chapter we are concerned with *spatial immersion*, in which the viewer is immersed within a virtual world that is perceptually convincing. We focus on

K. Marriott et al. (Eds.): Immersive Analytics, LNCS 11190, pp. 25–55, 2018.
https://doi.org/10.1007/978-3-030-01388-2_2

the visual aspects of this immersion and its relevance for data visualisation. An important component of spatial immersion in mixed reality systems[1] is the use of depth cues like occlusion and linear perspective, global lighting effects like shadows, texture rendering, as well as head tracking and binocular display technologies to simulate three-dimensional (3D) vision. Consequently these technologies support perceptually convincing rendering of data visualisations that make use of three dimensional space either in the visualisation itself or by allowing the visualisations to be placed anywhere in the viewer's environment.

Scientific visualisation researchers have been quick to adopt mixed reality display technologies, especially virtual reality systems, since most scientific data already lives in a 3D world. In contrast, information visualisation researchers have been very cautious about the use of 3D representations for abstract data and have therefore seen little benefit in the use of spatially immersive technologies. Most information visualisation researchers have continued to focus on flat 2D data visualisation designed for presentation on desktop computers, with a few forays into touch interfaces [51] and tangible graphics [52].

This conservatism is quite deliberate and is a response to the "unbridled enthusiasm" [71] of information visualisation researchers in the late 1980s and early 1990s for 3D representations. This was the time when the first 3D graphic workstations, such as those made by Silicon Graphics, were hitting the mass market and many information visualisation researchers were convinced that 3D visualisations utilising linear perspective, shading and shadows offered benefits over traditional 2D representations. For example, cone trees were introduced as an interactive way of visualising hierarchical data in 3D [78] (see Figure 1), the perspective wall [65] as a 3D focus+context way of viewing traditional 2D tables, while 3D desktops such as the data mountain [77] were introduced as a way of taking advantage of 3D spatial memory. The infamous 3D Pie or Bar Charts [80] that popular spreadsheet applications and business graphics showcased were also introduced at this time. Subsequent user studies, however, failed to find any benefits for these 3D representations over the traditional 2D representations in abstract data visualisation [16,19] and this early enthusiasm for 3D was replaced by a strong scepticism. Thus, in a recent popular text on data visualisation, Munzner [71] cautions against the "unjustified use" of 3D representations and immersive environments for representation of abstract data.

However, we believe that with the arrival of commodity immersive VR and AR devices such as the HTC Vive or Microsoft HoloLens it is timely to explore how to best visualise abstract data in such immersive environments and whether immersive analytic applications should continue to use traditional 2D information visualisations or whether there are situations in which 3D offers benefits. There are several reasons for doing so:

– **Potential of immersive displays.** Head-mounted mixed reality displays are almost certainly going to become much more common. As discussed

[1] We use the term mixed reality to refer to the continuum from pure virtual reality in which the user is totally immersed in the virtual environment, to augmented reality in which the physical environment is overlaid with virtual information [70].

in Chapter 1 these new platforms will free analytics applications from the desktop and encourage their use in all aspects of life. Furthermore these devices support situated analytics, embodied data exploration, collaboration as well as more engaging narrative visualisations. It is inevitable that mixed reality displays will be used for presentation of abstract data regardless of how well suited they are to this. We therefore need evidence-based guidelines on how to present abstract data on these devices, even if we find that the best way is to simply show abstract data on a flat 2D “billboard”.
- **Additional visual channel.** Allowing a third spatial dimension provides another visual channel for data visualisation. While prone to occlusion, depth disparity and foreshortening [71], as we shall discover, a number of studies demonstrate that there are benefits in using this channel for certain kinds of abstract data and tasks: we require a more nuanced understanding of when it is beneficial to do so.
- **Improving technology.** Furthermore, immersive display technologies have advanced considerably in the last decade. Modern displays have much higher resolution and less latency than the devices used in these earlier studies, as well as providing head-tracked binocular presentation[2] and supporting a wider range of interaction technologies. This means that modern immersive displays overcome at least some of the previously identified problems of viewing 3D visualisations on a flat desktop display.
- **Immersive work spaces.** Head-mounted mixed reality devices potentially allow the analyst to use the space around them as their workspace, placing data visualisations where they please, and in the case of augmented reality using the physical environment as the workspace. Such an immersive workspace is dramatically different from the traditional flat and workspace provided by desktop machines. As yet we do not know how to best use this extra freedom and what, if any, benefits it offers.
- **Beyond task effectiveness.** Information visualisation researchers have typically focussed on task effectiveness of data visualisations by measuring accuracy and speed. In immersive analytics we wish to understand effectiveness in a broader sense that includes other aspects of the user experience (Chapter 1). Does spatial immersion support deeper collaboration? Does it provide a more enjoyable, engaging, affective or memorable experience? Answers to these questions are vital if we want analytics applications to be used by the general public, not just professional analysts.

This chapter is intended to be a starting point for research into visual idioms for presenting and arranging abstract data in spatially immersive environments. In particular, it investigates the use of 3D for abstract data visualisation. However, a key problem in any discussion about the value of 3D visualisation is that the term 3D has been used to refer to very different things at different periods in time. In the 1980’s, 3D generally referred to graphics that rendered objects with three

[2] We avoid the use of the term stereoscopic, which derives from the greek $\sigma\pi\epsilon\rho\epsilon\omega\varsigma$, meaning “solid” (not “dual”, a common misconception). For an in-depth discussion see [98].

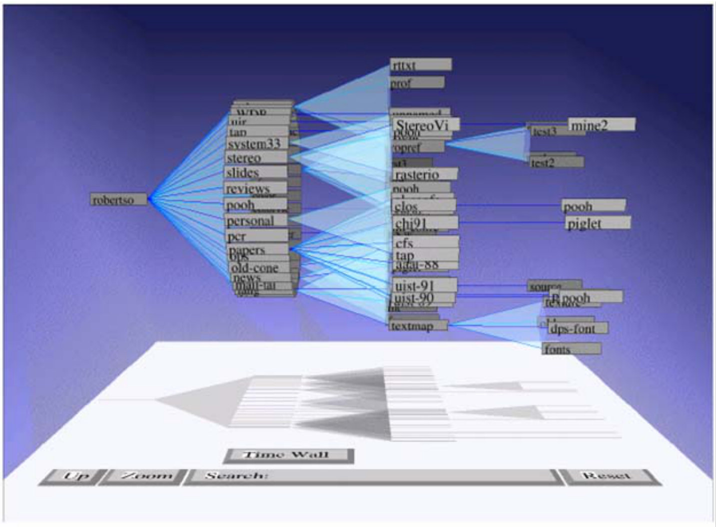

Fig. 1: Cone Trees and the horizontal variant Cam Trees were introduced by Robertson *et al.* in 1991 as a way to navigate large tree structures [80]. (Figure courtesy of S. Card, J. Mackinlay and G. Robertson and used with permission from Xerox.)

spatial dimensions (i. e., objects with volume) on to a flat (hence 2D) electronic screen while more recently 3D has become a synonym for (possibly head-tracked) binocular presentation. Thus it is important to clearly distinguish between the use of a 3D visualisation and the way in which this is presented to the viewer. We need to tease apart the effect of the different kinds of technologies and the depth cues they provide. As Ware argues, the use of 3D in data visualisation is really all about choosing the right depth cues [101]. We believe a major direction for immersive analytics research will be to investigate how to best make use of the depth cues provided in common spatially immersive environments for abstract data analysis. Our chapter provides the basis for this research by providing:

– A discussion of depth cues from a perceptual and technical perspective (*Section 2.2.*).
– A review of user studies evaluating the use of 3D in abstract data visualisation (*Section 2.3.*).
– Identification of applications and tasks for which spatially immersive visualisations may provide advantages over traditional 2D abstract data visualisation (*Section 2.4.*).

We finish the chapter with a number of open research questions.

2.2. Background: Perception and Presentation of 3D

As discussed in the introduction, one of the difficulties in understanding the value of 3D visualisation is that the term 3D has been used to refer to very different things depending upon sub-discipline and the time period. During the advent of the first powerful graphic workstations in the 1980's, 3D generally referred to graphics that rendered objects with three spatial dimensions (i. e., objects with volume) on to a flat (hence 2D) electronic screen.

More recently 3D has become a synonym for binocular presentation. Visualisation using this kind of 3D has been a popular area of research since the 80s (e. g., [76]), as sophisticated presentation technologies (e. g., binocular CAVEs, binocular head-mounted displays) were available to researchers before they were commercially viable for the general public. The assumption in this body of work, which we review in subsequent sections of this chapter, is that modes of presentation that are more faithful to how humans perceive the world will allow visualisations to take better advantage of the additional depth channel by overcoming the problems of the "flattened" 3D representations. For example, we might be able to overcome the inaccuracy in perceiving depth through the extra information contained in the differences between the images perceived by each eye, or the problem of occlusion [31] if we can move our heads (or ourselves) around a representation. This kind of thinking, in its limit, is exemplified by the relatively recent *physical visualisation* sub-area of data visualisation research [52], in which data is represented directly using real solid objects. Physical visualisation provides a representation that is not only faithful to real-world perception; it actually *is* real-world perception.

Obviously there are great differences between a 3D pie chart [80] printed on paper, a cone tree [78] seen through the Oculus Rift, or a bar chart printed with a volumetric 3D printer. For our purpose, a *3D visualisation* is any visualisation that *maps data to three different spatial dimensions*, independently of *how* these dimensions are presented to the viewer[3] and independently of whether these mappings are inherent in the semantics of the data (what is traditionally referred to as Scientific Visualisation) or they are abstract or arbitrary. A 3D visualisation may be good or bad depending on the device it is being viewed on, and devices will be different from each other depending on how they cater to the different ways of the human visual system to perceive depth dimension (depth cues).

We now provide a brief introduction into the perceptual and hardware terms that we will use in the rest of the chapter. First we summarise existing depth cues and attempts to rank them in order of strength, second we describe the most common presentation technologies and systems in connection to the cues that they support.

[3] This definition is related to the process called *spatialisation* [94] and resonates with the presentation-representation distinction made by Spence in his visualisation textbook [83], but we will not use this terminology since the distinction of representation and presentation is sometimes not clear and is intertwined with considerations of interactivity.

Depth perception The types of information that the visual system relies on for perceiving depth are called *depth cues*. Although there are multiple classifications and research is still ongoing (we do not know yet whether some of these cues actually provide information about depth to the visual system), we will use Cutting and Vishton's nomenclature [24], as well as mentioning some additional cues (or subcues) that are of relevance.

Occlusion refers to how objects that are closer in space prevent us seeing objects that are behind. This depth cue is ordinal, in the sense that it does not tell us how much further an object is if it is occluded, just that it is behind.

Relative size refers to the phenomenon that two objects of the same size that are at different distances from the observer will project a differently sized image in the eye. If we have an idea of the approximate size of an object, how big or small it appears in our field of view will give us information about how far it is.

Relative density relates to how spatial patterns of objects or visual features will appear more dense as the distance to the pattern increases. For example, if looking at a barley field, the individual spikes of the barley plants will be more concentrated and closer to each other in the retinal image of those parts of the field that are further away than those that are closer.

Height in visual field. If we assume that the space in front of us is approximately flat and that objects are bound to rest on the ground (due to the action of gravity), the position of their retinal images (or at least, their base) with respect to the horizon provides a proxy for their distance, since objects that are closer will be resting further from the horizon (or closer to our own feet).

The four cues described above, in combination, are often clustered as *linear perspective* [12] and are mostly a consequence of the projective properties of the eye as a sensor. Relative size, density and height in visual field are often also referred as *foreshortening*.

Aerial perspective refers to the change in colour properties (e. g., hue, saturation, lightness) of objects at large distances, caused by the scattering of light of the atmosphere in between the observer and the objects. This cue only works when distances are very large or when the atmosphere is optically dense (e. g., on a misty day).

Motion perspective is based on movement. If an object or the observer moves, the projection of the object in the retina changes, and the form and magnitude of these changes provide additional information about the 3D structure of the object [23] as well as its distance [79]. For example, when looking from a side window in a moving car, and assuming that most objects are static, we can perceive that objects that move faster in our field of view are closer to us, whereas objects in the horizon stay relatively static. Motion perspective is supported by head-tracking.

Binocular disparity and *stereopsis* are the cues most commonly associated with the modern popular understanding of 3D. Small differences in the images received by the left and the right eye (disparity), are processed in the brain to interpret depth and 3D shape of objects.

Accommodation refers to effects of dynamic physiological changes in the shape of each eye and the consequences on the retinal image. Specifically, accommodation is most evident in how objects that are not focused on appear proportionally blurry in relation to the distance to the current focus depth. This blur occurs because the eye lens works like a camera with a limited depth of field which adapts dynamically its focal length according to the object that is the center of attention. The amount of blur of the background and other objects provides information about their relative distances [66], although the amount of blur is also generally dependent on the lightness of the scene (a lighter scene causes the pupil to contract and, in consequence, the depth of field increase because the eye comes closer to a *pin hole camera*, which has almost infinite depth of field, i. e., all objects are in focus).

There are two additional possible sources of information from *accommodation*. The first is that the human nervous system controls the size of the pupil and the shape of the eye's lens, which could provide an additional input about the distance of an object (how much the body needs to stretch the lens for an object for it to be in focus is a function of its distance). The other is that different light frequencies are bent by the lens in slightly different ways. This results in subtle "halos" around objects that can be red or blue depending on which object is in focus and at which distance [46, 97], a manifestation of the chromatic aberration phenomenon.

Cast shadows can be an effective cue for judging the height of an object above the plane and act as an indirect depth cue by linking the depth of an object with that of the location on which its shadow falls [103].

Convergence refers to the change in rotation of the eyes that takes place to align the object or region of interest in the center of both eye's foveas. This is a reflex from the visual system [38]. Since the eye orientations are known by the perceptual system and are related to the distance of the object, the visual perceptual system can use this information to infer distance, although generally only in the short range (less than 3m) due to the fact that angle changes for far objects become increasingly small and therefore difficult to distinguish.

User interaction can provide further depth information by complementing or enhancing the above visual cues.

Controlled point of view refers to the ability of people to manipulate their point of view in a virtual space. This is common in desktop VR reality systems where the location of the point of view used to render a scene can be changed by the user through some input device, typically the mouse. This provides additional 3D information because the user knows what positional change she has triggered through her actions on the input device and therefore can expect different kinds of visual changes depending on the position and depth of the objects in the scene. Therefore this additional cue takes advantage of a closed feedback loop involving perception and motor action, and relies on the senses of touch and proprioception to complement mostly the *motion perspective* visual cue described above, but also triggering changes in most of the other cues (e. g., changes in vergence for binocular displays if an object comes close to the viewer).

Subjective motion is related to controlled point of view in that the viewer can change the perspective of the scene. However, in subjective motion the changes are triggered through the actual physical motion in space of the observer. These changes can be subtle (small changes in head position) or more dramatic if they involve walking. Subjective motion involves proprioception and motor action control as in the controlled point of view but, importantly, it provides additional information through the vestibular system (balance and movement detection). Equivalently to controlled point of view, this information is complementary to a range of visual cues [86].

Object manipulation through touch, using one's (tracked) hands, or input devices can change the position of objects with respect to the observer and therefore trigger motion perspective and changes in other cues. The change in relative position between the observer and the object of interest (virtual or physical) produces alterations of the information projected in the retina equivalent to what controlling the virtual point of view or moving oneself does, but it is generally limited to a single object (rather than the rest of the environment in the previous two), and it does not trigger vestibular system signals [72]. It still relies on touch (somatic), motor, and proprioceptive information.

3D display technologies By replicating all the cues listed above with electronic displays it is conceivable that we could recreate a (dynamic) experience of 3D perception that is very close to regular perception of physical objects in the real world. However, existing display technologies represent an additional bottleneck, since not all cues can be replicated simultaneously in a dynamic form. Therefore it is important to not only qualify what kind of depth perception cues our data representations will be depending on, but also which cues can actually be supported by the display itself.

Table 1 provides a summary of current display technologies and links it to the cues that they support. A technology can: a) clearly support a particular cue, b) might possibly support it, c) might only partially support it (or only in certain circumstances), and d) might intrinsically not support it. For more explanation of the technologies: Fishtank VR [104], Stereo CAVE [22], Accomodation Optics VR Headset [57], Holografika Holovizio see http://www.holografika.com, Optical 3D Displays [55], Volumetric 3D display [95] and Physical Displays [52]. Augmented reality (AR) provides a direct or indirect view of reality, augmented by computer-generated imagery. AR typically provides stereoscopic views, but monoscopic variants exist as well. Most AR systems are based on head-mounted devices, but some hybrid systems use large transparent displays in front of the scene to display the computer-generated content, e. g. [56].

Limitations of depth perception Knowing when and how best to use 3D for data visualisation requires us to understand the limitations of human depth perception and the additional limitations arising from current immersive display technologies. It is important to highlight that, although supported multi-modally by a plethora of cues as we have seen above, the perception of spatial 3D in

Table 1: Mapping between 3D display technologies and depth cues. "Y" (yes–Dark Green) indicates that existing systems do this. "P" (Possible–Light Green) indicates that existing systems could potentially do this. "D" (Depends/To Some Extent–Yellow) indicates that the specific property is achievable to some degree, although not completely, or needs to be simulated. "N" (No–Red) indicates that this is not currently possible.

	Linear Perspective	Aerial Perspective	Occlusion	Motion Perspective	Accommodation	Convergence	Binoc. Disparity and Stereopsis	User-controlled PoV	Subjective Motion	Interactive Content Manipulation	Examples
Regular photography or print	Y	Y	Y	N	N	N	N	N	N	N	
Desktop Computer Virtual Reality	Y	P	Y	Y	N	D	D	Y	N	D	[87]
Fishtank Virtual Reality	Y	P	Y	Y	N	D	D	P	Y	D	[104]
Non-disparity monocular/binocular viewing	Y	P	P	P	N	N	N	P	P	D	[99]
Head-mounted Binocular Displays	Y	P	Y	Y	N	Y	Y	P	Y	D	[90]
Multi-display Environments, Large Displays	Y	P	Y	Y	N	N	N	P	P	D	[43, 72, 73]
Binocular CAVEs	Y	P	Y	Y	N	Y	Y	P	Y	D	[22]
Gazer (Simulation of Accommodation)	Y	P	P	P	D	P	P	P	P	D	[66, 67]
Accommodation Optics VR Headset	Y	P	P	P	Y	Y	Y	P	Y	D	[81]
Multiview Autostereoscopic	Y	P	Y	Y	D	Y	Y	P	Y	D	[25, 49]
Volumetric 3D Displays	Y	N	N	Y	Y	Y	Y	P	Y	D	[36]
Optical Holographic 3D Displays	Y	N	N	Y	Y	Y	Y	P	Y	D	[55]
Augmented Reality (AR)	Y	P	D	Y	D	D	Y	N	Y	D	[5]
AR Hybrids	Y	P	Y	Y	D	D	Y	D	P	D	[56]
Physical Visualisations (reality)	Y	P	Y	Y	Y	Y	Y	N	Y	Y	[52]

humans is, strictly speaking, not *volumetric*, but instead layered and multiplexed in time. This is why the previous section describes the perception of depth (not 3D perception) and why Ware has argued that human visual perception is not 3D, not 2.5D, but instead closer to 2.05D [102].

Studies also indicate that more than 30% of the population may experience some form of binocular deficiency [47], and binocular acuity generally decreases with age [110]. Gracia *et al.* [39] identified several issues arising in 3D visualisation: line-of-sight ambiguity, occlusion and linear perspective distortion. Line-of-sight ambiguity refers to the fact that we can only see a single datapoint along each line of sight, whereas we can see many more across the field of view, text legibility, inappropriate view scale and movement distortion were also identified as potential issues. Text legibility is a particularly important issue as text and numbers are common in abstract data visualisations. Problems with legibility arise because current head-mounted AR and VR displays typically have lower resolution than desktop systems. This is compounded by foreshortening and 3D orientation, i. e., when the text is not directly facing the viewer [41, 62, 73].

An observer that has sufficient cues will be able to determine within a certain range of accuracy the 3D position of one particular object, but as the number of objects increases, or if the datum itself is continuous in the volume (e. g., the temperature of the water in a cube of $1km \times 1km \times 1km$) the dataset typically needs to be filtered and looked at a small number of layers at a time. With increasing numbers of objects, occlusion becomes a major limiting factor. The alternative of using transparency is limited by the fact that even the best 3D volume rendering techniques suffer from errors of 25% or more [33].

Apart from “true 3D” displays - volumetric display devices that show the graphics for a virtual object at the *actual depth* of the object, all other prevalent 3D stereo technologies, including stereo monitors and head-mounted displays, show the image on a different plane relative to where the eyes converge to in depth for objects away from the screen plane. Given that the eyes naturally focus onto the monitor image, this causes vergence-accomodation conflicts, e. g., [48], which affect the perception of 3D scenes. Moreover, if the disparity is too large, it results in double images (diplopia), which can also be effectively used by the perceptual system to infer depth [27].

Other limitations of spatially immersive data visualisation arise with interaction. These are discussed more fully in Chapter 4, but we point out that that navigation in 3D spaces is more challenging than 2D navigation, especially for abstract spaces, and 3D interaction is more difficult than 2D interaction as more degrees of freedom have to be controlled.

Following the review of depth perception and 3D presentation technologies, we now look at previous work that compare 2D and 3D representations.

2.3. Prior Research Comparing 2D with 3D Representations

In this section we review empirical studies comparing 2D with 3D representations and critically evaluate whether these support the orthodoxy that 2D is best for abstract data representations. As discussed earlier, as immersive technologies have improved what is implied when speaking of "3D" has changed so it is important to understand the depth cues as well as the kinds of interaction provided in the studies.

Cone trees and data mountains Cone trees [78] and the data mountain [77] were two very influential 3D visualisations from the 1990s which were claimed to have significant advantages over traditional 2D visualisations. Both visualisations were intended for use on a 3D graphics workstation with mouse interaction. There was no head tracking or binocular presentation, and by today's standards resolution was low. Later empirical studies critically evaluated the claim that cone trees and the data mountain offered benefits over 2D visualisations. As these significantly contributed to the subsequent scepticism of information visualisation researchers of the use of 3D in abstract data visualisation, we discuss these studies in some detail.

Robertson *et. al.* [78] argued that cone trees were more effective than traditional 2D trees because: (i) linear perspective provides a focus+context view of the tree; (ii) 3D cues of perspective, lighting and shadows help shape understanding; (iii) interactive animation reduces cognitive load by exploiting the visual system's object tracking; (iv) users enjoyed the visualisations because they were more "alive"; and (v) 3D allows more effective use of the display space by using depth which means that larger trees can be displayed.[4] Cockburn and McKenzie [16] conducted a user study comparing cone trees with a traditional indented list-like representation for navigating through a hierarchy to find particular nodes. They found that task performance was slower with cone trees and that users rated the cone interface as worse for seeing and for interacting with the hierarchy. They did not evaluate understanding of structure, but participant comments "indicated that cone trees may perform relatively better in [such] tasks". Caveats were that participants were much more familiar with the traditional indented list like representation, the cone tree was rendered with low fidelity because of implementation efficiency, and cone tree rotation was slower than scrolling in 2D.

The data mountain allowed the user to arrange documents on a virtual 3D desktop in front of them. The bottom of the billboard-like representations was always in contact with the desktop so it was really a 2.5D visualisation.[5] The

[4] Though not necessarily perceived because of linear perspective and occlusion.

[5] Unfortunately the term 2.5D visualisation is commonly used in several different ways. Here we are using it in the sense of Ware [101] in which the visualisation is essentially 2D but selected depth cues are used to provide some suggestion of 3D. In GeoVis 2.5D refers to showing a 2D continuous surface in 3D, while in network visualisation 2.5D refers to stacked 2D visualisations.

viewpoint of the user was fixed and occlusion, shadows and linear perspective were used to create the illusion of depth. Audio cues were also used. Motivations for using 3D were to: (i) allow more objects on the desktop (with linear perspective providing a focus+context view), (ii) a natural metaphor for grouping, and (iii) to leverage from 3D spatial memory. Robertson *et al.* [77] compared bookmark retrieval times and error rates with the Data Mountain and Internet Explorer (IE4). They found the data mountain led to increased efficiency and lower error rates. However since the two representations were quite different this did not directly address the question of whether 3D offered benefits over 2D. This was investigated by Cockburn and McKenzie [17] in a user study comparing a low fidelity implementation of the Data Mountain interface with the equivalent 2D interface in which thumbnails of web pages were organised on the screen much like in a standard window manager. There was no foreshortening and the user could bring thumbnails to the front or back. They found that participants were faster with the 2D interface than with 3D (though this was not statistically significant) but that the participants believed the 3D interface was more effective. They suspected that foreshortening (and low resolution) meant that in the 3D condition it was difficult to visually match thumbnails that were in the top half of the screen. In a subsequent study Cockburn [15] compared recall of non-overlapping 'cards' arranged on a desktop in horizontal rows. This compared a 2D presentation with a 3D presentation using foreshortening, shadows and proximity luminance covariance, i. e., cards at the front were lighter than things at the back. He found no difference between the two conditions. Another study compared the effectiveness of spatial memory in 2D, 2.5D (2D + linear perspective), and full 3D in both virtual and physical environments [18]. Flaws in the design and implementation mean it is difficult to draw conclusions from the physical implementation. While 2.5D outperformed 2D, which in turn outperformed 3D, the differences were not statistically significant.

So what can we take home from these evaluations of cone trees and the data mountain? It is fair to conclude that they show that 3D does not offer "magic" benefits over 2D for data visualisation because of its "naturalness". In hindsight, cone trees, because of occlusion and the inherent slowness of tree rotation, are a poor representation for hierarchical data. In the case of document management there is no difference between 2 and 2.5D. Indeed in some sense most window managers are 2.5D since they provide occlusion, the most powerful depth cue (hence the question: which depth cues are useful?). Thus, these studies do not rule out the use 2.5 or 3D in abstract data visualisation but rather suggest that depth cues need to be used carefully. We now look at the findings from other evaluations of 3D. These are organised by the kind of data being visualised.

Aviation Haskell and Wickens [45] compared 2D orthogonal views with 3D perspective for aviation displays showing current position and predicted flight path. The 3D display was better for lateral and altitude flight path tracking but worse for accurate measurement of airspeed. The display was not binocular or interactive. Van Orden and Broyles [95] compared a variety of 2D and 3D displays for judging aircraft altitude or speed, route planning (vectoring) and collision detection. Air traffic controllers were more accurate with 2D views except collision avoidance in which task the laser-based 3D volumetric display led to significantly more accurate performance than the other 2D or 3D displays.

3D shapes and landscapes St John *et al.* [84] compared a 3D orthographic perspective view with side-by-side 2D orthogonal views of complex block shapes. Shadows, motion and binocular cues were not used. 3D views were better for understanding the overall shape while 2D orthogonal views were better for precisely judging relative position.

Tory *et al.* [92] compared 2D, 3D, and combined visualisations for some basic tasks like position estimation and orientation. They suggest the combination of 2D and 3D views to improve precision. In further studies, Tory *et al.* [91] compared 2D orthogonal, 3D rotated orthographic projection, and combinations of 2D and 3D for determining relative position of a ball and complex block shape and for orienting a plane to cut a torus in half (Figure 2). They also compared two ways to orient and couple 2D planes with 3D. They concluded that 3D was effective for approximate navigation and relative positioning but combination 2D/3D displays were better for precise orientation and positioning.

In another study St John *et al.* [85] compared: (i) a 3D perspective view of a landscape; (ii) a topographic map with contours and (iii) a side-by-side combination of these two views. One task involved computing the best routing for antennas so that they were in line of sight. Participants were fastest with the side-by-side view, then the 2D view and slowest with the 3D perspective view. In a second study participants were asked to choose the best route from three given routes with either the 3D or the plan view. In this case the 3D view led to faster performance. These suggest that the 3D view is best for overall orientation and understanding, the 2D view is better for precise manipulation, and that it is beneficial to combine them using 3D for orientation and 2D for manipulation (“Orient and Operate").

Network visualisation Ware and Franck [100] report the results of a user study in which graph visualisation with 2D and different 3D displays were compared. According to their results, the 3D displays allowed participants to decide if there is a path between highlighted nodes with the same accuracy for graphs up to three times larger than in the 2D case. They found that both motion and binocular depth cues were beneficial. A binocular display was found to be 1.6 times more accurate than a 2D display when detecting paths of length two through the complex structure, and binocular presentation combined with head-coupled motion produced the best results.

The Ware and Frank study was replicated in [105] while Greffard *et al.* found benefits for binocular 3D over 2D for community detection in social networks [40]. Van Schooten [96] evaluated the impact of motion cues and binocular presentation on path following in 3D maze-like solid shaded structures (based on vascular structures). They found that the motion cue was more important than binocular

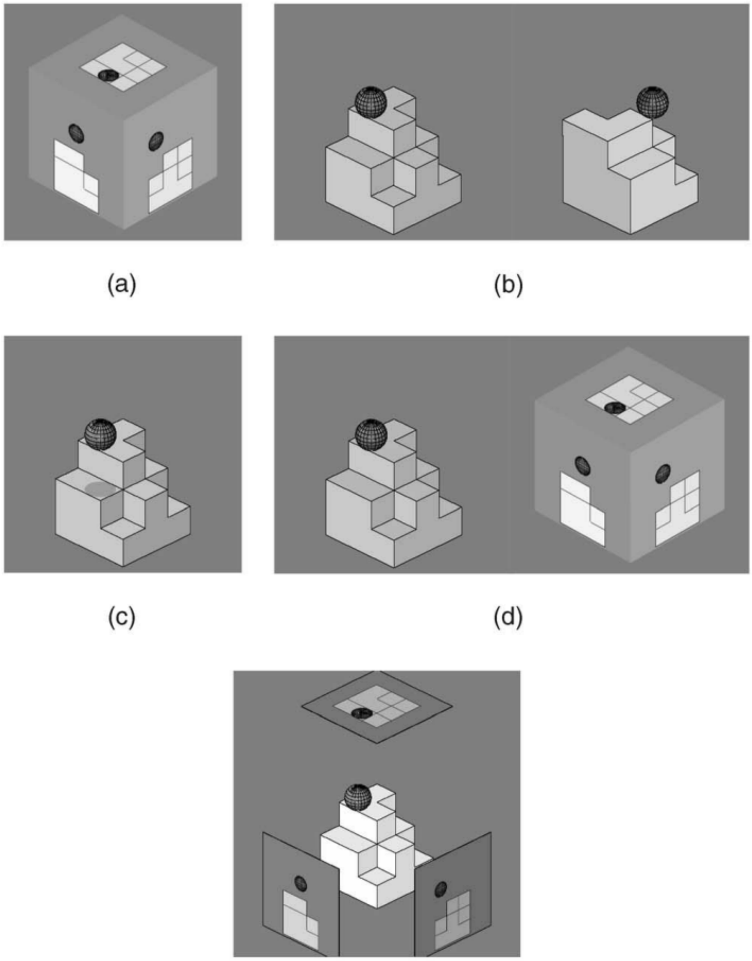

Fig. 2: Example stimuli from Tory *et al.* [91]. Position estimation of the ball relative to the block shape using the following visualisation techniques: (a) 2D, (b) 3D rotated, (c) 3D shadow, (e) Integrated 2D and 3D display. (Figure © 2006 IEEE. Reprinted, with permission, from [91])

presentation and that binocular presentation had little added benefit if motion cue was provided. Belcher *et al.* [5] found similar benefits to 3D and motion cues but not binocular views when viewing graphs using augmented reality.

Irani and Ware [50] found that 3D glyphs rendered using shading, surface texture and lighting as 3D cues where more easily recognised and remembered than flat 2D silhouettes of the glyphs. They studied this in the context of UML-like network diagrams. The diagrams themselves were laid out in a flat plane, only the glyphs utilised depth.

Kwon *et al.* recently proposed an immersive graph visualisation system [60,61]. Previous studies into 3D graph layout used exocentric layouts in which the user was outside the layout. They introduced an egocentric layout in which the graph is laid out in a sphere around the user's head and viewed using a binocular HMD VR environment. Their user study [60] found that their spherical layout outperformed a traditional 2D layout.

Alper *et al.* designed depth highlighting of 2D graph visualisation on binocular displays [1] (Figure 3). The technique makes use of depth cues to enable focus+context visualisation by overlaying a detailed image of a region of interest on the overall graph, which is visualized at a further depth with correspondingly less details. Their empirical study results show that binocular highlighting had about the same performance as the static visual highlighting and that performance was improved when binocular and static visual highlights were combined.

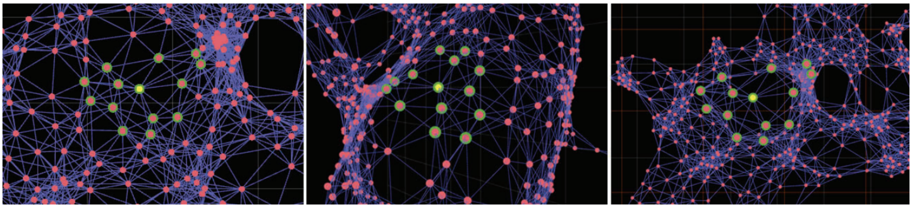

Fig. 3: Focus+context views from [1] for different visualisations illustrating how an increasingly large portion of graph nodes is kept in view from left to right, 2D, 3D, and 2.5D views. (Figure © 2011 IEEE. Reprinted, with permission, from [1])

Multivariate data visualisation An early paper by Lee *et al.* [63] ran two small user studies comparing 2D graphics with 3D graphics shown using a polarising binocular view on a standard monitor. They found that performance was more accurate with a 3D scatter plot of a three-dimensional data set than with three 2D scatter plots of the data, but that for another data set accuracy was the same with a 3D histogram as a multidimensional table. In a related study Wickens *et al.* [108] evaluated the use of 3D and a split-view orthogonal 2D representation of three-dimensional economic data. They found that participants using 3D were faster than those using 2D when answering integrative questions involving all three dimensions, while performance was similar for more focused questions involving fewer dimensions. In the 3D condition they found a positive effect for binocular depth cues while allowing user-controlled rotation or providing a mesh had no effect.

Wegman and Symanzik [106] reviewed the use of immersive technologies for visual exploration of multidimensional data (visual data mining). There are several empirical studies and the results are mixed. An early study by Nelson *et al.* [74] compared exploration of multidimensional data sets using brushing and a grand tour in 2D on a computer monitor and in 3D using binocular VR with head tracking. They found that 3D provided a large advantage in cluster identification and some advantage in identifying the shape of the cluster and similar performance when identifying the dimensionality of the data. More recently, Gracia *et al.* [39] evaluated the effectiveness of dimensionality reduction of a multidimensional dataset to a 3D scatterplot and to a 2D scatterplot. Both were shown in monocular 3D on a standard monitor. The extra dimension in the 3D condition meant that distances in the visualisation more closely matched the actual distances between points in the higher dimensional space. They found that with the 3D scatter plot users could more accurately compare distances between points (with respect to actual distance in multidimensional space) and more accurately detect outliers but had similar accuracy when classifying points. They were slower with 3D on the last two tasks. However, in both studies the number of participants was small and statistical significance was not reported. Westerman *et al.* [107] evaluated exploration of document spaces presented in two and three dimensions. They found that interaction and navigation was slower in 3D space. They used monocular 3D on a standard computer monitor.

Sedlmair *et al.* [82] evaluated two users ranking of cluster separation using a number of dimension reduction techniques for 2D scatterplots, scatterplot matrices (SPLOMs) and 3D scatterplots shown in monocular 3D on a standard monitor. They found that cluster separation was generally ranked to be the same or less with 3D scatterplots than with 2D scatterplots. This result is surprising given that that an extra dimension should allow the underlying structure to be revealed more clearly.

Tory *et al.* [93, 94] evaluated the usefulness of both 2D and 3D landscapes for enhancing understanding of clustering in 2D scatterplots. The first task was to identify the spatial region with the most values in a particular interval. Both 2D and 3D spatial encoding redundantly encoded the data value intervals. They found that the spatial landscape was detrimental and that the 3D landscape was worse than the equivalent 2D landscape. This may be the result of using the landscape color and/or height to encode attribute value range rather than density of data points which would have been more salient. They also evaluated the effect of 2D and 3D landscapes on recall. Recall without landscapes was more accurate, and 3D was generally better than 2D landscapes.

Spatial and spatio-temporal data visualisation Yalong *et al.* [109] compared task performance for three standard geographic tasks using different 2D and 3D representations for the Earth. It is one of the few studies to compare 2D and 3D representations in VR and used a head-tracked binocular HMD. They compared a 3D exocentric globe placed in front of the viewer, an egocentric 3D globe placed around the viewer, a flat map (rendered to a plane in VR) and a curved map, created by projecting the map onto a section of a sphere curved around the user. They found that (a) the exocentric globe was more accurate than the egocentric globe and the flat map for distances comparison, (b) for comparison of areas, more time is required with exocentric and egocentric globes than with flat and curved maps, and (c) that for direction estimation, the exocentric globe is more accurate and faster than the other visual presentations. There was a weak preference for the exocentric globe. Generally, the curved map had benefits over the flat map. In almost all cases the egocentric globe was found to be the least effective visualisation.

Kjellin *et al.* [59] compared 3D space-time cubes with 2D map + animation and a pure 2D representation in which tracks are drawn on a map and time is shown by annotating the tracks with orthogonal cross-lines spaced at regular time intervals. A Head-tracked binocular display was used for the 3D visualisation. It does not appear that the participants could interactively control their point of view or speed of the animation. They found that performance with the pure 2D representation was better than with the other two representations for predicting the point where 3 vehicles would meet in the future. In a final experiment they compared the pure 2D representation with the space-time cube and asked participants to determine the relative order 4 vehicles had passed through a particular point. For this task performance was better with the space-time cube. Kjellin *et al.* explained this in terms of Todd *et al.*'s theory [88,89] that tasks that can be solved by estimating properties of a viewed scene that are invariant under affine transformations are perceptually easier than those that are invariant only under Euclidean transformations. In another study Kjellin *et al.* [58] compared performance with monocular and head-tracked binocular presentation of a 3D space-time cube showing discrete spatiotemporal data. The task was to find the data set which had a cone like shape and the largest geographic spread. They found no difference between the conditions and believe this was because the task was essentially affine preserving and that there was little clutter, so little need for additional depth cues.

Summary It is clear from the studies that 3D representations are not generally better than 2D representations, but nor are 2D representations always better than 3D: which is better depends upon the kind of task. Previous studies suggest that 3D representations may show overall structure in multidimensional spaces more clearly: 3D shapes and terrain [84,85], multidimensional data [39,63,74,108], and networks [40,100,105]. On the other hand some of these studies also found that 2D representations were preferable for precise manipulation or accurate data value measurement or comparison and advocated the use of linked 2D and 3D representations [85,91].

It also seems that the choice of technology and depth cues makes a significant difference to the effectiveness of 3D visualisations. Ragan *et al.* [75] found that combining binocular presentation and head-tracking contributed to improved spatial judgment accuracy when participants need to distinguish between structural gaps and intersections between components of 3D models in a simulated underground systems. Two papers review user studies comparing binocular 3D with non-binocular 3D in wide variety of applications including scientific, medical and military [36, 69]. The most recent, McIntire *et al.* [69], summarised 180 experiments and found performance benefits for binocular 3D in 60% of the studies but little or no benefit in the others. They conclude that binocular 3D is beneficial for depth-related tasks including spatial understanding of complex scenes and spatial manipulation, especially when the objects were close and the tasks were difficult. This may help to explain the previously discussed mixed results on the effectiveness of 3D scatter plots when compared to 2D scatter plots. Studies providing binocular depth cues found benefits [63,74,108], while those using only monocular clues were mixed with one finding benefits [39] but others finding 2D was more effective [82,107,107]. It would be interesting to replicate these experiments using modern head-tracked binocular displays.

Finally, details of representation are important. For instance, Kjellin *et al.* [59] found that hanging the trajectory representation from a line to a stepped staircase of rectangles in the 3D space-time cube improved performance with the space-time cube considerably (though it was still worse than with the pure 2D representation).

2.4. Potential Benefits of Immersive Visualisation

In the preceding section we have reviewed user studies comparing 2D and 3D representations, In this section we present five reasons why immersive 3D display environments may offer advantages over traditional 2D visualisations on a desk display. Further information visualisation applications and concepts that might benefit from 3D representations or binocular displays are discussed by Brath [9] and McIntire and Liggett [68].

Using depth to show an additional abstract dimension In the last section we saw that two studies [39, 74] found benefits in using 3D scatter plots to understand structure in multi-dimensional data while three studies did not [82, 107]. Notably the three studies which did not find benefits did not use head-tracked binocular 3D while it was used in one of the studies finding benefits. If the data has more dimensions then projecting onto an additional dimension means that there will be less *stress* (error) in the multidimensional projection: the distance between data points in the projected 3D space will be closer to the distance between them in the original multi-dimensional space than when they are projected to a 2D space. Thus it seems plausible that clusters and other structure are likely to be more clearly shown in the 3D scatter plot. However for this benefit to be realised there need to be sufficient depth cues for the viewer to see the structure without the need for interaction beyond head movement.

Even if the primary visualisation is a 2D scatter plot, 3D may be used to show the relationship between different plots, for instance Elmqvist *et al.* [30] use an animated 3D cube to transition between different plots.

We have seen that there is empirical evidence that laying out node-link diagrams in 3D can benefit path following [100,105]. This may be because depth cues help to clarify edge crossings and resolve node-edge overlap. 3D network diagrams may also more clearly show structure [40] since (as discussed above) the extra dimension means that layouts can have less stress in the sense that the distance between nodes more closely reflects the graph theoretic distance between them and the edges have more uniform length.

Judicious use of depth cues may be beneficial even with essentially 2D representations of networks. For instance Alper *et al.* [1] found that using binocular “pop out” was beneficial for highlighting elements. In 2D node-link diagrams it is common to use occlusion to indicate that one edge crosses over another, so as to enforce the perception of the edges as separate objects. Other depth cues might also be beneficial. We also note that Irani and Ware [50] found benefits in using depth cues with glyphs in UML diagrams.

Several researchers have investigated different kinds of 2.5D layouts of network diagrams in which nodes in the network are laid out on 2D planes stacked on top of each other. For instance Brandes *et al.* [8] use stacked 2D layouts of metabolic pathways to compare pathways in different organisms while Eades and Feng used stacked layouts to show the hierarchical structure [29]. Another use of stacked 2D network representations is to show network changes over time. Each plane is a snapshot of the network at a particular point in time. Dwyer and Eades [28] used these to visualise trading data.

The use of the third dimension to show time is a successful idiom that has been employed with many kinds of data [4]. We have seen it used for showing trajectories and for dynamic network diagrams. It has also been used to show time series data from oscilloscopes organised in a layered eye diagram [64]. This uses an orthogonal projection to show time series wrapped in time with different slices clustered together behind one another.

Views on non-flat surfaces Another possible advantage of 3D is that 2D data with cyclic dimensions or without natural boundaries can be laid out on the surface of a sphere [10] or a cylinder so as to remove the visual illusion of a break in the dimension when the data is laid out in 2D. While it is not clear that this outweighs the problem of occlusion when viewed from outside the sphere or cylinder, Yalong *et al.* [109] found that an exocentric globe or curved map outperformed a standard flat map for some common geographic tasks. Non-flat viewing canvases also allow an egocentric view of the data from inside. Results comparing egocentric and exocentric views of abstract data are mixed. The study reported in Kwon *et al.* [60, 61] suggests that such an egocentric view is useful for network visualisation but Yalong *et al.* [109] found that for geographic tasks an exocentric globe outperformed an egocentric globe.

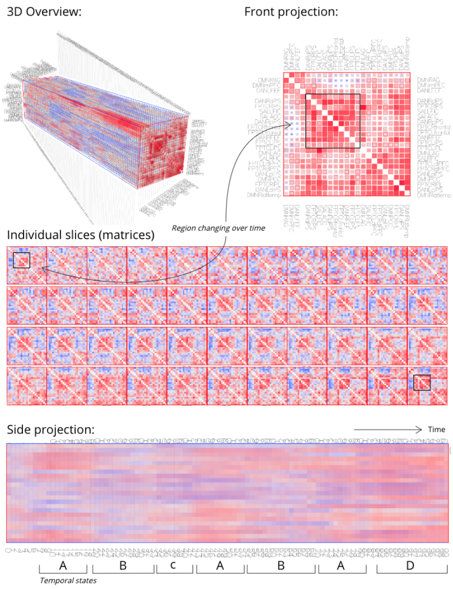

Fig. 4: Matrix cubes for visualising dynamic networks are an example of the use of the third dimension to show time [3]. (Figure courtesy Benjamin Bach.)

Visualising abstract data with a spatial embedding 3D has also been used quite successfully for visualising abstract data with a geographic embedding. Here it is natural to use two dimensions to show a map and use the third dimension to show a data attribute or to overlay data attributes in either 2D or 3D visualisations on a 3D landscape. Dübel *et al.* [26] classifies such visualisations based on whether the reference space (i. e., the map or surface) is shown in 2D or 3D and/or whether the abstract attribute is shown in 2D or 3D. Prism maps are widely used as an alternative to choropleth maps and allow more accurate comparison of data values so long as occlusion is not a problem. Vertical pins can be used to show magnitude of a data attribute at a particular location on a map rather than using proportional symbols such as a circle [9]. Another widely used example is the space-time cube. Introduced by Hägerstrand [42] in 1970 this shows trajectories across a map using time as a third (vertical) dimension. As we saw in the section it is well-suited to some tasks but for others a pure 2D representation may be more effective.

More generally 3D is potentially useful for showing abstract data with a 3D spatial embedding [54]. Bowman *et al.* defined the integration of spatial and non-spatial data as information-rich virtual environment [7]. The goal is to enhance spatial data with abstract information, which is the type of data that is not directly *perceptible* in the physical world but *added* to the 3D world (e. g., text labels to show velocity at every point in a complex vector field, trees, networks, and other multiple-dimensional information). The information-rich virtual environment defines the forms of the display, e. g., temperature in a room can be recorded spatially but cannot be directly perceived visually thus will need to be encoded (either with colour or texts). The follow-up lab-based experiment by Chen et al. [13] found that 2D-billboard style text display is better than 3D texts attached directly on spatial objects.

This might be by overlaying 2D or 3D abstract data representations on to 3D visualisations of say buildings or organisms, or by having separate 2D and 3D views linked in some way. For example, multi-dimensional scaling (MDS) techniques have been developed to to embed a 3D diffusion magnetic resonance imaging (DMRI) into a low-dimensional 2D representation [14]. Distance measurements (e. g., similarity) are also adopted to further cluster the 2D embedding [53]. Recently, Zhang *et al.* uses a topological approach to represent DTI using a contour tree to show brain water diffusion rate measured by fractional anisotropy in a brain [111].

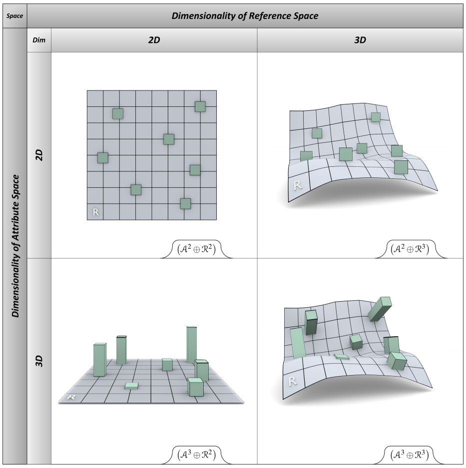

Fig. 5: Dübel *et al.* [26] categorise geo-visualisation techniques based on the dimensionality of the attribute space and reference space presentation. (Figure courtesy Steve Dübel).

Arranging multiple views in 3D space Another possible benefit of immersive 3D is that it allows the analyst to arrange their views in 3D space. Mixed mode displays, for instance, allow the use of virtual Powerwalls which provide unlimited display space and use linear perspective to provide an easily understood focus+context view [9].[6] Such a virtual Powerwall potentially provides the same benefits as actual physical Powerwalls for visual analytics [2].

[6] In a sense they are a modern equivalent of the perspective wall.

At present there is no agreed windowing metaphor in mixed mode applications and this is a research topic that warrants attention. Some researchers have suggested the use of flat 2D views whose position may be fixed with respect to the world or with respect to some part of the user's body, e.g., [34, 35, 37], while others have suggested more embodied three-dimensional data views, e.g., [11, 21], or blended views in which physical objects provide the view frame, e.g., [32].

Closely related to 2.5D layouts of networks are Collins and Carpendale's technique [20] of linking 2D representation laid out as views in 3D space by using arcs between elements in the different views to show connectivity with selected elements.

Engagement The final benefit of depth cues is that they may help to engage the reader. In many previous studies users indicated they preferred 3D representations even if they did not improve task performance. Of course this has to be done judiciously, adding gratuitous 3D depth cues to pie charts or bar charts is, and always will be, a bad idea. In general spatial immersiveness may aid user engagement. However, at present this has not been empirically investigated.

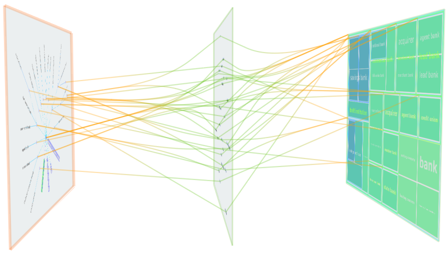

Fig. 6: Linking 2D representation in 3D space using arcs [20]. (Figure courtesy Christopher Collins.)

2.5. Research Questions and Issues

This chapter suggests a large number of different research directions and questions. One of the most obvious is to develop general, evidence based design rules and a portfolio of effective visual and interaction idioms for data visualisation in mixed reality environments. Some interesting questions are

– What should the data analyst's workspace look like in mixed reality applications? We have seen a variety of different suggestions, ranging from 2D windows similar to those found on a desktop to embodied scatter plots or parallel coordinates [11, 21, 34, 37]. At present there is no agreed approach and little empirical data. Closely related is the question of how to link different viewing canvases. Is brushing enough or should explicit links be drawn [20]? How do we guide/direct the analyst's attention in their immersive workspace, in particular how to communicate that there are important things behind them in the virtual space? (The same issues arise in video games and 360° immersive movies.)

- Some studies have advocated the use of tightly linked 2D and 3D representations [85, 91], with the 3D view providing overall orientation and understanding while the 2D view is better for precise manipulation and data comparison. Is this a general idiom that we should support in spatially immersive analytics applications?
- Egocentric data views potentially offer a more engaging user experience. However, studies comparing exocentric and egocentric views are mixed [60, 109]. When if ever, are egocentric views of data preferable to exocentric? Should the user be able to move between egocentric and exocentric views of the data analogously to the way the Magic Book allows the viewer to view actors in the story from outside or to transport themselves into the scene [6]? Is rapid motion in immersive egocentric data views unpleasant to the user and lead to motion sickness? If so, how best to change the position of the user in such views? One solution, used in video games, is teleportation of the user – is this also appropriate for immersive analytics?
- What are the differences between VR and AR environments and how do these affect the choice of immersive visualisation idiom?
- Parallel perspective–oblique, isometric, etc–is widely used in technical drawing as such projections have a uniform scale in each dimension and parallel lines remain parallel. An interesting question is when, if ever, is parallel perspective more appropriate than linear perspective for showing abstract data? It may for example have benefits when multiple users are looking at the same representation from slightly different viewpoints [44, 72].
- What is the influence of different aspects of spatial immersion—depth cues, realism of the rendering, integration of the environment in an AR setting (e. g., things are placed on the real table), consistency of multi-sensory output, egocentric vs. exocentric views, etc.—on all aspects of user experience and performance. In particular, how does spatial immersiveness affect longer-term engagement with the data/content (beyond the first "cool" impression)?
- To answer these questions we will need to devise and verify measures and ways of evaluating the effectiveness and usefulness of spatially immersive environments for data visualisation

One strong recommendation echoing that of McIntire [69], is that in all papers describing user studies involving some kind of 3D the technology and supported depth cues as well as available interactions and interaction devices are carefully described: In many existing papers this is unclear, yet it is now apparent that these significantly impact user performance and must be taken into account.

2.6. Conclusion

Mixed-reality display technologies utilise a variety of techniques such as depth cues (like occlusion, linear perspective or head-tracked binocular presentation) to enhance spatial immersion. In this chapter we have explored the potential use of these techniques to visualise abstract data in immersive analytics applications. Almost certainly more and more people will be using immersive environments as the technology matures and on occasions they will want to visualise data in these environments, not just on a desktop. We need to know the best way of supporting this.

While 3D visualisations are common in scientific visualisation, they are much less common in information visualisation. Indeed the accepted wisdom is that 3D effects are rarely useful for showing abstract data. In Section 2.3. we investigated the empirical evidence for this viewpoint in order to understand whether immersive analytics applications should continue to use traditional 2D information visualisations or whether there are situations when use of 3D depth cues may offer benefits. We found that previous user studies suggest that 3D representations more clearly show overall structure in higher dimensional datasets such as 3D terrain, networks or multidimensional data and are useful for providing orientation while 2D representations are preferable for precise manipulation or accurate data value measurement or comparison. Thus, it may be useful to provide linked 2D and 3D representations in immersive analytics applications.

In Section 2.4. we further explored the potential applications of depth cues for abstract data visualisation. These include: using depth to show an additional abstract dimension, such as in 2.5D network layouts, views on non-flat surfaces and egocentric views in which the data is placed around the viewer, visualising abstract data with a spatial embedding as well as arranging multiple views in 3D space.

Based on the discussions above, we feel that the current scepticism by information visualisation researchers for 3D visualisation of abstract data is too negative and that there are situations in which judicious use of depth cues to provide 2.5 or 3D views is warranted. We also believe that allowing the user to arrange views in the 3D space around them in mixed-reality settings offers benefits over traditional desktop windows management. We believe that further exploration of these topics should be a major focus of immersive analytics research.

Acknowledgements

We are thankful for the helpful suggestions of Benjamin Bach, Tim Dwyer and Mark Billinghurst. Marriott acknowledges the support of the Australian Research Council Discovery Scheme for this research through project DP180100755.

References

1. Alper, B., Hollerer, T., Kuchera-Morin, J., Forbes, A.: Stereoscopic highlighting: 2D graph visualization on stereo displays. IEEE Transactions on Visualization and Computer Graphics 17(12), 2325–2333 (2011)
2. 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. pp. 55–64. ACM (2010)
3. Bach, B., Pietriga, E., Fekete, J.D.: Visualizing dynamic networks with matrix cubes. In: the SIGCHI Conference on Human Factors in Computing Systems. pp. 877–886 (2014)
4. Bach, B., Dragicevic, P., Archambault, D., Hurter, C., Carpendale, S.: A review of temporal data visualizations based on space-time cube operations. In: Eurographics Conference on Visualization (EuroVis) (2014)
5. Belcher, D., Billinghurst, M., Hayes, S., Stiles, R.: Using augmented reality for visualizing complex graphs in three dimensions. In: Proceedings of the 2nd IEEE/ACM International Symposium on Mixed and Augmented Reality. p. 84. IEEE Computer Society (2003)
6. Billinghurst, M., Kato, H., Poupyrev, I.: The magicbook: a transitional ar interface. Computers & Graphics 25(5), 745–753 (2001)
7. Bowman, D.A., North, C., Chen, J., Polys, N.F., Pyla, P.S., Yilmaz, U.: Information-rich virtual environments: theory, tools, and research agenda. In: Proceedings of the ACM Symposium on Virtual Reality Software and Technology. pp. 81–90. ACM (2003)
8. Brandes, U., Dwyer, T., Schreiber, F.: Visual understanding of metabolic pathways across organisms using layout in two and a half dimensions. Journal of Integrative Bioinformatics 1(1) (2004)
9. Brath, R.: 3D infovis is here to stay: deal with it. In: 2014 IEEE VIS International Workshop on 3DVis. pp. 25–31 (2014)
10. Brath, R., MacMurchy, P.: Sphere-based information visualization: challenges and benefits. In: IV. pp. 1–6 (2012)
11. 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 SIGCHI Conference on Human Factors in Computing Systems. pp. 90:1–90:12 (2018)
12. Carlbom, I., Paciorek, J.: Planar geometric projections and viewing transformations. ACM Computing Surveys (CSUR) 10(4), 465–502 (1978)
13. Chen, J., Pyla, P.S., Bowman, D.A.: Testbed evaluation of navigation and text display techniques in an information-rich virtual environment. In: IEEE Proceedings of Virtual Reality. pp. 181–289. IEEE (2004)
14. Chen, W., Ding, Z., Zhang, S., MacKay-Brandt, A., Correia, S., Qu, H., Crow, J.A., Tate, D.F., Yan, Z., Peng, Q.: A novel interface for interactive exploration of DTI fibers. IEEE Transactions on Visualization and Computer Graphics 15(6), 1433–1440 (2009)
15. Cockburn, A.: Revisiting 2D vs 3D implications on spatial memory. In: Proceedings of the fifth conference on Australasian user interface-Volume 28. pp. 25–31. Australian Computer Society, Inc. (2004)
16. Cockburn, A., McKenzie, B.: An evaluation of cone trees. In: People and Computers XIV Usability or Else!, pp. 425–436. Springer (2000)

17. Cockburn, A., McKenzie, B.: 3D or not 3D?: evaluating the effect of the third dimension in a document management system. In: Proceedings of the SIGCHI conference on Human Factors in Computing Systems. pp. 434–441. ACM (2001)
18. Cockburn, A., McKenzie, B.: Evaluating the effectiveness of spatial memory in 2D and 3D physical and virtual environments. In: Proceedings of the SIGCHI conference on Human Factors in Computing Systems. pp. 203–210. ACM (2002)
19. Cockburn, A., McKenzie, B.: Evaluating spatial memory in two and three dimensions. International Journal of Human-Computer Studies 61(3), 359–373 (2004)
20. Collins, C., Carpendale, S.: VisLink: revealing relationships amongst visualizations. IEEE Transactions on Visualization and Computer Graphics 13(6), 1192–1199 (2007)
21. Cordeil, M., Cunningham, A., Dwyer, T., Thomas, B.H., Marriott, K.: ImAxes: Immersive axes as embodied affordances for interactive multivariate data visualisation. In: Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology. pp. 71–83. ACM (2017)
22. Cruz-Neira, C., Sandin, D.J., DeFanti, T.A., Kenyon, R.V., Hart, J.C.: The CAVE: audio visual experience automatic virtual environment. Commun. ACM 35(6), 64–72 (1992)
23. Cutting, J.E.: Rigidity in cinema seen from the front row, side aisle. Journal of Experimental Psychology: Human Perception and Performance 13(3), 323–334 (1987)
24. Cutting, J.E., Vishton, P.M.: Perceiving layout and knowing distances: The integration, relative potency, and contextual use of different information about depth. In: Perception of space and motion, pp. 69–117. Elsevier (1995)
25. Dodgson, N.A.: Autostereoscopic 3D Displays. Computer 38(8), 31–36 (2005)
26. Dübel, S., Röhlig, M., Schumann, H., Trapp, M.: 2D and 3D presentation of spatial data: a systematic review. In: IEEE VIS International Workshop on 3DVis. pp. 11–18. IEEE (2014)
27. Duwaer, A., Van Den Brink, G.: What is the diplopia threshold? Attention, Perception, & Psychophysics 29(4), 295–309 (1981)
28. Dwyer, T., Eades, P.: Visualising a fund manager flow graph with columns and worms. In: Proceedings Sixth International Conference on Information Visualisation. pp. 147–152. IEEE (2002)
29. Eades, P., Feng, Q.W.: Multilevel visualization of clustered graphs. In: International Symposium on Graph Drawing. pp. 101–112. Springer (1996)
30. Elmqvist, N., Dragicevic, P., Fekete, J.D.: Rolling the dice: multidimensional visual exploration using scatterplot matrix navigation. IEEE Transactions on Visualization and Computer Graphics 14(6), 1539–1148 (2008)
31. Elmqvist, N., Tsigas, P.: A taxonomy of 3D occlusion management for visualization. IEEE Transactions on Visualization and Computer Graphics 14(5), 1095–1109 (2008)
32. ElSayed, N.A., Smith, R.T., Marriott, K., Thomas, B.H.: Context-aware design pattern for situated analytics: Blended model view controller. Journal of Visual Languages & Computing 44 (2018)
33. Englund, R., Ropinski, T.: Evaluating the perception of semi-transparent structures in direct volume rendering techniques. In: SIGGRAPH ASIA 2016 Symposium on Visualization. pp. 9:1–9:8. ACM (2016)
34. Ens, B., Hincapié-Ramos, J.D., Irani, P.: Ethereal planes: a design framework for 2D information space in 3D mixed reality environments. In: Proceedings of the 2nd ACM Symposium on Spatial User Interaction. pp. 2–12. ACM (2014)

35. Ens, B.M., Finnegan, R., Irani, P.P.: The personal cockpit: a spatial interface for effective task switching on head-worn displays. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. pp. 3171–3180. ACM (2014)
36. Favalora, G.E.: Volumetric 3D displays and application infrastructure. Computer 38(8), 37–44 (2005)
37. Feiner, S., MacIntyre, B., Haupt, M., Solomon, E.: Windows on the world: 2D windows for 3D augmented reality. In: Proceedings of the 6th Annual ACM Symposium on User Interface Software and Technology. pp. 145–155. ACM (1993)
38. Fincham, E., Walton, J.: The reciprocal actions of accommodation and convergence. The Journal of Physiology 137(3), 488–508 (1957)
39. Gracia, A., González, S., Robles, V., Menasalvas, E., von Landesberger, T.: New insights into the suitability of the third dimension for visualizing multivariate/multidimensional data: a study based on loss of quality quantification. Information Visualization 15(1), 3–30 (2014)
40. Greffard, N., Picarougne, F., Kuntz, P.: Visual community detection: An evaluation of 2D, 3D perspective and 3D stereoscopic displays. In: Graph Drawing. pp. 215–225. Springer (2011)
41. Grossman, T., Wigdor, D., Balakrishnan, R.: Exploring and reducing the effects of orientation on text readability in volumetric displays. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. pp. 483–492. ACM (2007)
42. Hägerstrand, T.: What about people in regional science? Regional Science Association 24(1), 7–24 (1970)
43. Hancock, M., ten Cate, T., Carpendale, S., Isenberg, T.: Supporting sandtray therapy on an interactive tabletop. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. pp. 2133–2142. CHI '10, ACM (2010)
44. Hancock, M., Nacenta, M., Gutwin, C., Carpendale, S.: The effects of changing projection geometry on the interpretation of 3D orientation on tabletops. In: Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces. pp. 157–164. ACM (2009)
45. Haskell, I.D., Wickens, C.D.: Two-and three-dimensional displays for aviation: A theoretical and empirical comparison. The International Journal of Aviation Psychology 3(2), 87–109 (1993)
46. Held, R., Cooper, E., Banks, M.: Blur and disparity are complementary cues to depth. Current Biology 22(5), 426–431 (2012)
47. Hess, R.F., To, L., Zhou, J., Wang, G., Cooperstock, J.R.: Stereo vision: the haves and have-nots. i-Perception 6(3), 2041669515593028 (2015)
48. 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), 33–33 (2008)
49. Inc., H.: Holografika website. http://www.holografika.com, accessed: 2017-10-17
50. Irani, P., Ware, C.: Diagramming information structures using 3D perceptual primitives. ACM Transactions on Computer-Human Interaction (TOCHI) 10(1), 1–19 (2003)
51. Isenberg, P., Isenberg, T., Hesselmann, T., Lee, B., von Zadow, U., Tang, A.: Data visualization on interactive surfaces: a research agenda. IEEE Computer Graphics and Applications 33(2), 16–24 (2013)
52. 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 33rd Annual ACM Conference on Human Factors in Computing Systems. pp. 3227–3236. CHI 15, ACM (2015)

53. Jianu, R., Demiralp, C., Laidlaw, D.: Exploring 3D DTI fiber tracts with linked 2D representations. IEEE Transactions on Visualization and Computer Graphics 15(6), 1449–1456 (2009)
54. Kerren, A., Schreiber, F.: Why integrate InfoVis and SciVis?: an example from systems biology. IEEE ComputerGraphics & Application 34, 69–73 (2014)
55. Khan, J., Can, C., Greenaway, A., Underwood, I.: A real-space interactive holographic display based on a large-aperture hoe. In: Proc. SPIE. vol. 8644, p. 86440M (2013)
56. Kim, M., Lee, J., Stuerzlinger, W., Wohn, K.: Holostation: augmented visualization and presentation. In: SIGGRAPH Asia 2016 Symposium on Visualization. pp. 12:1–12:9. ACM (2016)
57. Kiyokawa, K., Kurata, Y., Ohno, H.: An optical see-through display for mutual occlusion with a real-time stereovision system. Computers & Graphics 25(5), 765–779 (2001)
58. Kjellin, A., Pettersson, L.W., Seipel, S., Lind, M.: Different levels of 3D: an evaluation of visualized discrete spatiotemporal data in space-time cubes. Information Visualization 9(2), 152–164 (2010)
59. Kjellin, A., Pettersson, L.W., Seipel, S., Lind, M.: Evaluating 2D and 3D visualizations of spatiotemporal information. ACM Transactions on Applied Perception (TAP) 7(3), 19:1–19:23 (2010)
60. Kwon, O.H., Muelder, C., Lee, K., Ma, K.L.: A study of layout, rendering, and interaction methods for immersive graph visualization. IEEE Transactions on Visualization and Computer Graphics 22(7), 1802–1815 (2016)
61. Kwon, O.H., Muelder, C., Lee, K., Ma, K.L.: Spherical layout and rendering methods for immersive graph visualization. In: 2015 IEEE Pacific Visualization Symposium (PacificVis). pp. 63–67 (2015)
62. Larson, K., van Dantzich, M., Czerwinski, M., Robertson, G.: Text in 3D: some legibility results. In: CHI'00 extended abstracts on Human Factors in Computing Systems. pp. 145–146. ACM (2000)
63. Lee, J.M., MacLachlan, J., Wallace, W.A.: The effects of 3D imagery on managerial data interpretation. MIS Quarterly pp. 257–269 (1986)
64. Lopez-Hernandez, R., Guilmaine, D., McGuffin, M.J., Barford, L.: A layer-oriented interface for visualizing time-series data from oscilloscopes. In: IEEE Pacific Visualization Symposium (PacificVis). pp. 41–48. IEEE (2010)
65. Mackinlay, J.D., Robertson, G.G., Card, S.K.: The perspective wall: detail and context smoothly integrated. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. pp. 173–176. ACM (1991)
66. Mauderer, M., Conte, S., Nacenta, M.A., Vishwanath, D.: Depth perception with gaze-contingent depth of field. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. pp. 217–226. ACM (2014)
67. Mauderer, M., Nacenta, M.A., Morrison, D.: Gazer: Application for gaze-contingent viewing of images(Apr 2018), https://github.com/MichaelMauderer/Gazer, original-date: 2015-06-16T12:05:34Z
68. McIntire, J.P., Liggett, K.K.: The (possible) utility of stereoscopic 3D displays for information visualization: The good, the bad, and the ugly. In: IEEE VIS International Workshop on 3DVis. pp. 1–9 (2014)
69. McIntire, J.P., Havig, P.R., Geiselman, E.E.: Stereoscopic 3D displays and human performance: A comprehensive review. Displays 35(1), 18–26 (2014)
70. Milgram, P., Takemura, H., Utsumi, A., Kishino, F.: Augmented reality: a class of displays on the reality-virtuality continuum. Proc. SPIE 2351, 282–292 (1995)

71. Munzner, T.: Visualization analysis and design. CRC Press (2014)
72. Nacenta, M.A., Hancock, M., Gutwin, C., Carpendale, S.: The effects of changing projection geometry on perception of 3D objects on and around tabletops. ACM Transactions on Computer-Human Interaction (TOCHI) 23(2), 11:1–11:54 (2016)
73. Nacenta, M.A., Sakurai, S., Yamaguchi, T., Miki, Y., Itoh, Y., Kitamura, Y., Subramanian, S., Gutwin, C.: E-conic: a perspective-aware interface for multi-display environments. In: Proceedings of the 20th Annual ACM Symposium on User Interface Software and Technology. pp. 279–288. ACM (2007)
74. Nelson, L., Cook, D., Cruz-Neira, C.: Xgobi vs the c2: results of an experiment comparing data visualization in a 3-D immersive virtual reality environment with a 2-D workstation display. Computational Statistics 14(1), 39–52 (1999)
75. Ragan, E.D., Kopper, R., Schuchardt, P., Bowman, D.A.: Studying the effects of stereo, head tracking, and field of regard on a small-scale spatial judgment task. IEEE Transactions on Visualization and Computer Graphics 19(5), 886–896 (2013)
76. Rekimoto, J., Green, M.: The information cube: Using transparency in 3D information visualization. In: Proceedings of the Third Annual Workshop on Information Technologies & Systems (WITS'93). pp. 125–132 (1993)
77. Robertson, G., Czerwinski, M., Larson, K., Robbins, D.C., Thiel, D., Van Dantzich, M.: Data mountain: using spatial memory for document management. In: Proceedings of the 11th Annual ACM Symposium on User Interface Software and Technology. pp. 153–162. ACM (1998)
78. Robertson, G.G., Mackinlay, J.D., Card, S.K.: Cone trees: animated 3D visualizations of hierarchical information. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. pp. 189–194. ACM (1991)
79. Rogers, B., Graham, M.: Motion parallax as an independent cue for depth perception. Perception 8(2), 125–134 (1979)
80. Schonlau, M., Peters, E.: Graph comprehension: An experiment in displaying data as bar charts, pie charts and tables with and without the gratuitous 3rd dimension. Social Science Research Network Working Paper Series pp. 1–16 (2008)
81. Schowengerdt, B.T., Seibel, E.J.: True three-dimensional displays that allow viewers to dynamically shift accommodation, bringing objects displayed at different viewing distances into and out of focus. CyberPsychology & Behavior 7(6), 610–620 (2004)
82. Sedlmair, M., Munzner, T., Tory, M.: Empirical guidance on scatterplot and dimension reduction technique choices. IEEE Transactions on Visualization and Computer Graphics 19(12), 2634–2643 (2013)
83. Spence, R.: Information visualization. Springer (2001)
84. St. John, M., Cowen, M.B., Smallman, H.S., Oonk, H.M.: The use of 2D and 3D displays for shape-understanding versus relative-position tasks. Human Factors: The Journal of the Human Factors and Ergonomics Society 43(1), 79–98 (2001)
85. St. John, M., Smallman, H.S., Bank, T.E., Cowen, M.B.: Tactical routing using two-dimensional and three-dimensional views of terrain. In: Proceedings of the Human Factors and Ergonomics Society Annual Meeting. vol. 45, pp. 1409–1413. SAGE Publications (2001)
86. Sun, H.J., Chan, G.S.W., Campos, J.L.: Active navigation and orientation-free spatial representations. Memory & Cognition 32(1), 51–71 (2004)
87. Tait, A.: Desktop virtual reality. In: IEE Colloquium on Using Virtual Worlds. pp. 5/1–5/5 (1992)

88. Todd, J.T., Norman, J.F.: The visual perception of 3-D shape from multiple cues: are observers capable of perceiving metric structure? Perception & Psychophysics 65(1), 31–47 (2003)
89. Todd, J.T., Oomes, A.H., Koenderink, J.J., Kappers, A.M.: On the affine structure of perceptual space. Psychological Science 12(3), 191–196 (2001)
90. Tomilin, M.G.: Head-mounted displays. Journal of Optical Technology 66, 528–533 (1999)
91. Tory, M., Kirkpatrick, A.E., Atkins, M.S., Möller, T.: Visualization task performance with 2D, 3D, and combination displays. IEEE Transactions on Visualization and Computer Graphics 12(1), 2–13 (2006)
92. Tory, M., Moller, T., Atkins, M.S., Kirkpatrick, A.E.: Combining 2D and 3D views for orientation and relative position tasks. In: Proceedings of the SIGCHI conference on Human Factors in Computing Systems. pp. 73–80. ACM (2004)
93. Tory, M., Sprague, D.W., Wu, F., So, W.Y., Munzner, T.: Spatialization design: comparing points and landscapes. IEEE Transactions on Visualization and Computer Graphics 13(6), 1262–1269 (2007)
94. Tory, M., Swindells, C., Dreezer, R.: Comparing dot and landscape spatializations for visual memory differences. IEEE Transactions on Visualization and Computer Graphics 15(6), 1033–1040 (2009)
95. Van Orden, K., Broyles, J.: Visuospatial task performance as a function of two-and three-dimensional display presentation techniques. Displays 21(1), 17–24 (2000)
96. Van Schooten, B.W., Van Dijk, E.M., Zudilova-Seinstra, E., Suinesiaputra, A., Reiber, J.H.: The effect of stereoscopy and motion cues on 3D interpretation task performance. In: Proceedings of the International Conference on Advanced Visual Interfaces. pp. 167–170. ACM (2010)
97. Vishwanath, D., Blaser, E.: Retinal blur and the perception of egocentric distance. Journal of Vision 10(10) (2010)
98. Vishwanath, D.: Toward a new theory of stereopsis. Psychological Review 121(2), 151–178 (2014)
99. Vishwanath, D., Hibbard, P.B.: Seeing in 3D with just one eye: stereopsis without binocular vision. Psychological Science 24(9), 1673–1685 (2013)
100. Ware, C., Franck, G.: Viewing a graph in a virtual reality display is three times as good as a 2D diagram. In: Proceedings of the IEEE Symposium on Visual Languages. pp. 182–183 (1994)
101. Ware, C.: Designing with a 2 1/2-D attitude. Information Design Journal 10(3), 258–258 (2001)
102. Ware, C.: Visual thinking: for design. Morgan Kaufmann (2010)
103. Ware, C.: Information Visualization: Perception for Design (3rd Ed.). Elsevier (2013)
104. Ware, C., Arthur, K., Booth, K.S.: Fish tank virtual reality. In: Proceedings of the INTERACT and CHI Conference on Human Factors in Computing Systems. pp. 37–42. CHI '93, ACM (1993)
105. Ware, C., Mitchell, P.: Reevaluating stereo and motion cues for visualizing graphs in three dimensions. In: Proceedings of the 2nd Symposium on Applied Perception in Graphics and Visualization. pp. 51–58. ACM (2005)
106. Wegman, E.J., Symanzik, J.: Immersive projection technology for visual data mining. Journal of Computational and Graphical Statistics 11(1), 163–188 (2002)
107. Westerman, S.J., Cribbin, T.: Mapping semantic information in virtual space: dimensions, variance and individual differences. International Journal of Human-Computer Studies 53(5), 765–787 (2000)

108. Wickens, C.D., Merwin, D.H., Lin, E.L.: Implications of graphics enhancements for the visualization of scientific data: Dimensional integrality, stereopsis, motion, and mesh. Human Factors: The Journal of the Human Factors and Ergonomics Society 36(1), 44–61 (1994)
109. Yang, Y., Jenny, B., Dwyer, T., Marriott, K., Chen, H., Cordeil, M.: Maps and globes in virtual reality. Computer Graphics Forum (2018)
110. 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), 891–900 (2003)
111. Zhang, G., Kochunov, P., Hong, E., Carr, H., Chen, J.: Towards visual mega-analysis of voxel-based measurement in brain cohorts. In: Proceedings of the Eurographics/IEEE VGTC Conference on Visualization: Short Papers. pp. 55–59. Eurographics Association (2016)